Posture Invariant Correspondence Of Triangular Meshes In Shape Space

Stefanie Wuhrer Chang Shu
National Research Council of Canada
Ottawa, Canada

Prosenjit Bose
Carleton University
Ottawa, Canada

Abstract

We find dense point-to-point correspondences between two surfaces corresponding to different postures of the same articulated object in a fully automatic way. The approach requires no prior knowledge about the shapes being registered. Furthermore, the approach does not require any user-specified parameters. We register possibly incomplete triangular meshes. We model the deformations of an object as isometries and solve the correspondence problem by aligning the intrinsic geometries of the manifolds in a suitable space. We apply the technique to segment the surface into near-rigid components.

1. Introduction

We aim to find *dense point-to-point correspondences* between two articulated surfaces $S^{(0)}$ and $S^{(1)}$. The correspondence problem is defined as follows. Given a position $x^{(0)}$ on $S^{(0)}$, we aim to find the position $x^{(1)}$ on $S^{(1)}$ that corresponds to the same intrinsic location on $S^{(1)}$ as does $x^{(0)}$ on $S^{(0)}$. If the position $x^{(1)}$ is absent on $S^{(1)}$ due to incomplete data, no correspondence is found for $x^{(0)}$. We assume that $S^{(0)}$ and $S^{(1)}$ are triangular manifolds.

Applications that require knowledge of dense point-to-point correspondences between two articulated surfaces include mesh deformation [3], shape registration [23], object recognition [16], and mesh parameterization [20, 27].

This type of precessing is often applied to real-world data captured using 3D laser-range scanners or image-based reconstruction. The reconstructed surfaces are often noisy or incomplete due to occlusions. This makes the problem of automatically computing the dense point-to-point correspondences challenging.

Finding point-to-point correspondences is a difficult problem because local regions on the surface are often not distinctive. This results in searching a large set of candidate correspondences when computing the correspondence for each object point. Previous methods for finding correspondences for articulated surfaces often restrict the search space using prior knowledge about the objects being registered [1] or use probabilistic methods to solve the problem [2]. These approaches may produce inaccurate correspondences. Two methods were proposed to find the correspondences in a fully automatic way by minimizing a deformation cost [15, 33]. These methods do not assume knowledge of markers or template shapes. However, both methods rely heavily on non-intuitive user-specified parameters.

We propose a new method to solve the correspondence problem in a fully automatic way that does not require any user-specified parameters. We find candidate correspondences using multi-dimensional scaling as in Jain et al. [17] and Wuhrer et al. [31]. We then find the best alignment by choosing the candidate that has the shortest deformation distance. The deformation distance is measured in a shape space that represents near-isometric morphs of triangular meshes [7].

2. Related Work

This section reviews previous work on finding correspondences between two articulated surfaces. Allen et al. [1] use a given template model of a human body and a set of marker positions to solve the correspondence problem in an automatic way. The approach deforms the template mesh to fit a range scan of a human using a smooth deformation. This method requires prior knowledge about the shape being registered and about marker positions.

Recently, a number of markerless registration algorithms were developed [2, 9, 17, 31]. Anguelov et al. [2] solve the registration problem using loopy belief propagation on a Markov network. This approach maximizes a joint probabilistic function over all possible correspondences. The method ensures that close-by points on one surface map to close-by points on the other surface and that geodesics are preserved. This method fails when registering surfaces of a human body due symmetric misalignments of the front and the back of the body.

Jain et al. [17], Wuhrer et al. [31], and Bronstein et al. [9] solve the non-rigid correspondence problem by matching the intrinsic geometries of the surfaces. First, the approaches embed the intrinsic geometry of the surface into a Euclidean space using multi-dimensional scaling (MDS).

Second, the approaches match the embeddings. These approaches register shapes with nearly symmetric canonical forms erroneously. This paper presents an approach that overcomes the drawback by aligning the shapes using a more stable procedure. The alignment procedure can be viewed as minimizing a deformation distance.

Bronstein et al. [8] also solve the correspondence problem by matching the intrinsic geometries of the surfaces. However, instead of embedding the intrinsic geometry of a shape S in Euclidean space, Bronstein et al. embed the intrinsic geometry of S into a triangular surface Q using generalized MDS. Generalized MDS aims to embed the points on S into the surface Q, such that the geodesic distance on S is approximated well by the geodesic distance of the corresponding points on Q. This method avoids the large embedding errors caused by embedding into Euclidean space. However, Bronstein et al. require a template mesh when the aim is to register two incomplete surfaces.

Zeng et al. [32] embed the intrinsic geometry of a surface with arbitrary topology in the plane using a conformal mapping. They then solve the registration problem by matching two conformal mappings in the plane.

Recently Huang et al. [15] and Zhang et al. [33] independently developed two similar approaches to solve the non-rigid correspondence problem. Both approaches find the correspondence by minimizing a deformation energy. Huang et al. [15] propose a technique that solves the correspondence problem iteratively by alternating between a correspondence optimization and a deformation optimization. The approach can be viewed as an extension of the Iterative Closest Point algorithm (ICP) [5] that is often used to solve the rigid correspondence problem. The method is shown to perform well if the two meshes are initially well aligned. If the alignment is poor, the method fails. The main drawback of this method is that it relies heavily on non-intuitive user-defined parameters. This makes the method impractical. Zhang et al. [33] propose a technique that solves the correspondence problem by finding a small set of features and by choosing the best feature correspondence as the one that minimizes a deformation energy. To improve the efficiency of the algorithm, the tree of all matching features is pruned if the features are too dissimilar. Nonetheless, the algorithm is not as efficient as the algorithm of Huang et al. [15]. Once the feature correspondences are computed, the full correspondence is found by deforming the full mesh based on the feature points. The main drawback of this method is the computational inefficiency. Results are only demonstrated for models with less than 4000 vertices. Furthermore, like the method of Huang et al., the tree pruning relies heavily on non-intuitive user-defined parameters. The authors leave finding parameters automatically for future work.

In this work, we find the best alignment of the shapes in a shape space representing the intrinsic geometry of a triangular manifold. Shape spaces that represent the intrinsic geometry of shapes are used in morphing, where the goal is to smoothly transform a source shape into a target shape [30, 18, 7]. By finding the best alignment in shape space, we minimize a deformation energy.

3. Overview of Approach

This section gives an overview of our algorithm. The approach starts by finding a set of candidate correspondences. The correct correspondence is found as the candidate correspondence that minimizes a deformation energy.

The candidate correspondences are found using multi-dimensional scaling. This approach has previously been used to solve the correspondence problem [17, 31]. However, the previous approaches have the drawback that near-symmetric shapes are registered erroneously because the wrong candidate correspondence is accepted as the correct solution to the correspondence problem. For instance, the back of a person may be registered to the front although the back and front of a person are not entirely symmetric to each other. Section 4 reviews how to efficiently find candidate correspondences.

To find the correct candidate correspondence, we evaluate a deformation cost for each candidate correspondence. The candidate correspondence with the smallest deformation cost is chosen as the final correspondence. To evaluate the deformation cost, we compute a deformation distance between $S^{(0)}$ and $S^{(1)}$ in a shape space that was previously used to morph between near-isometric objects [7]. Section 5 outlines how to compute the deformation cost for a given candidate correspondence.

4. Finding Candidate Correspondences Via Multi-Dimensional Scaling

This section describes how to find a set of candidate correspondences.

4.1. Finding Dense Correspondences

We model the deformations of a shape as isometries and solve the correspondence problem by matching the intrinsic geometries of two shapes. To match the intrinsic geometries, we map the intrinsic geometries of the surfaces into a low-dimensional Euclidean space via multi-dimensional scaling. Elad and Kimmel [13] introduced the embedding of the intrinsic geometry of a surface into a low-dimensional Euclidean space and denoted it the *canonical form* of a surface. If shapes are deformed isometrically, then canonical forms are posture-invariant shapes that represent the intrinsic geometry of a manifold. Note that in this application, multi-dimensional scaling is used to embed the intrinsic geometry of a shape in Euclidean space without reducing the dimensionality of the data.

To find the candidate correspondences, we employ a coarse-to-fine strategy. We compute sample sets $P^{(r)}$ containing $n'^{(r)}$ vertices from $S^{(r)}$ for $r = 0, 1$ using Voronoi sampling. We compute the geodesic distance $\delta_{i,j}^{(r)}$ between every pair of samples p_i and p_j using fast marching [19]. To account for incomplete models, we compute confidence values

$$\omega_{i,j}^{(r)} = 1 - \frac{m_{i,j}^h}{m_{i,j}},$$

where $m_{i,j}$ is the number of edges on the geodesic path computed by the fast marching technique from p_i to p_j and where $m_{i,j}^h$ is the number of edges tracing a hole of $S^{(r)}$ on the geodesic path from p_i to p_j. We use $\delta_{i,j}$ as dissimilarities and $\omega_{i,j}$ as weights to embed the samples $P^{(r)}$ of the manifold $S^{(r)}$ into $\mathbb{R}^k$ via least-squares multi-dimensional scaling (MDS) [6, p.146-155]. In the following, $k = 3$.

To solve the correspondence problem, we compute the rigid correspondence between the canonical forms. The canonical forms $X^{(r)}$ are invariant with respect to rotation, translation, and reflection [11]. Hence, we need to consider multiple alignments of $X^{(0)}$ and $X^{(1)}$. We follow the common approach to align both $X^{(0)}$ and $X^{(1)}$ by the eigenvectors of their respective data covariance matrices [17, 31]. For each possible sign alignment of the eigenvectors, we compute an optimal rigid correspondence using the Hungarian method [26]. The Hungarian method solves the following problem: given a sign alignment, find an assignment function $a(i)$ that assigns exactly one point $X_{a(i)}^{(1)}$ to every point $X_i^{(0)}$, such that $E_H = \sum_{i=1}^{n'^{(0)}} d(X_i^{(0)}, X_{a(i)}^{(1)})$ is minimized, where d denotes the Euclidean distance in $\mathbb{R}^k$. This results in 2^k different rigid candidate correspondences.

4.2. Eliminating Erroneous Matchings

To eliminate erroneous matchings that assign a point $X_{a(i)}^{(1)}$ to a point $X_i^{(0)}$, we wish to only accept correspondences that form a graph $G^{(0)}$ on $S^{(0)}$ that is approximately isometric to a graph $G^{(1)}$ on $S^{(1)}$. We can find these correspondences using a kernel extraction method as proposed by Leordeanu and Hebert [22] and used by Huang et al. [15]. The approach proceeds as follows. We first compute a consistency matrix C with entries

$$c_{i,j} = \min\left(\frac{\delta_{i,j}^{(0)}}{\delta_{a(i),a(j)}^{(1)}}, \frac{\delta_{a(i),a(j)}^{(1)}}{\delta_{i,j}^{(0)}}\right).$$

Note that $0 \leq c_{i,j} \leq 1$. The matrix C measures if two pairs of correspondences are consistent. They are consistent if and only if $\delta_{i,j}^{(0)}$ and $\delta_{a(i),a(j)}^{(1)}$ are almost identical. In case of consistency, $c_{i,j}$ is close to 1. To allow for small non-isometric deformations, we do not use C to extract the valid correspondences, but define a matrix M to find the kernel [15]. The entries of M are

$$m_{i,j} = \begin{cases} \left(\frac{c_{i,j}-c_0}{1-c_0}\right)^2 \text{ if } c_{i,j} > c_0, \\ 0 \text{ otherwise,} \end{cases}$$

where c_0 is a threshold. The threshold c_0 describes how much non-isometric deformation is acceptable. Unlike Huang et al. [15], we do not set the threshold empirically. Instead, to compute c_0 in a fully automatic manner, we analyzed the distribution of the entries of the matrix C over all 2^k candidate correspondences. We found that the distribution resembles part of a Gaussian distribution. The distribution for one of the Alien models discussed in Section 6 is shown in Figure 1. The figure shows a histogram of the values $c_{i,j}$. We use the 2^k matrices C (one matrix per candidate correspondence) to learn the underlying normal distribution (μ, σ^2) via maximum likelihood estimation. We then set c_0 to $\mu - 3 * \sigma$. This way, 99.7% of the correspondences are expected to be accepted.

Finally, we project the vertices $S^{(r)} \setminus P^{(r)}$ to the embedding space. The approach finds the correspondence of the projected vertices by evaluating an approximating thin-plate spline (TPS) mapping the embedding $X^{(0)}$ to $X^{(1)}$ [12].

A coarse-to-fine strategy is implemented by performing the least-squares MDS and the Hungarian assignment on a sample set $P^{(r)}$ of vertices computed using Voronoi sampling [14]. To find dense correspondences, all vertices of $S^{(r)} \setminus P^{(r)}$ are projected to the embedding space and registered using the TPS mapping and a nearest neighbor search using a kd-tree. We only accept correspondences that are consistent with the sample correspondences according to the learned parameter c_0.

The running time of computing all 2^k candidate correspondences is

$$O(n'n \log n + n'^2 t + 2^k((n'+k)^3 + n(tk + \sqrt{n}))),$$

where $n = max(n^{(0)}, n^{(1)}), n' = max(n'^{(0)}, n'^{(1)}), k$ is the dimension of the embedding space, and t is the maximum number of iterations used to compute the embedding of the points. The reason is that initially, the algorithm computes the set $P^{(r)}$ via Voronoi sampling in $O(n'n \log n)$ time and the canonical forms of $P^{(r)}$ in $O(n'^2 t)$ time. Next, for each of the 2^k possible alignments, the algorithm computes the Hungarian matching and the TPS mapping in $O((n'+k)^3)$ time. Finally, the algorithm computes all nearest neighbors in $O(n(tk+\sqrt{n}))$ time using a kd-tree. Since $k = 3$ and since on average n' and t are much smaller than n, the average running time of the algorithm is $O(n\sqrt{n})$.

The following section outlines how we choose the best candidate correspondence.

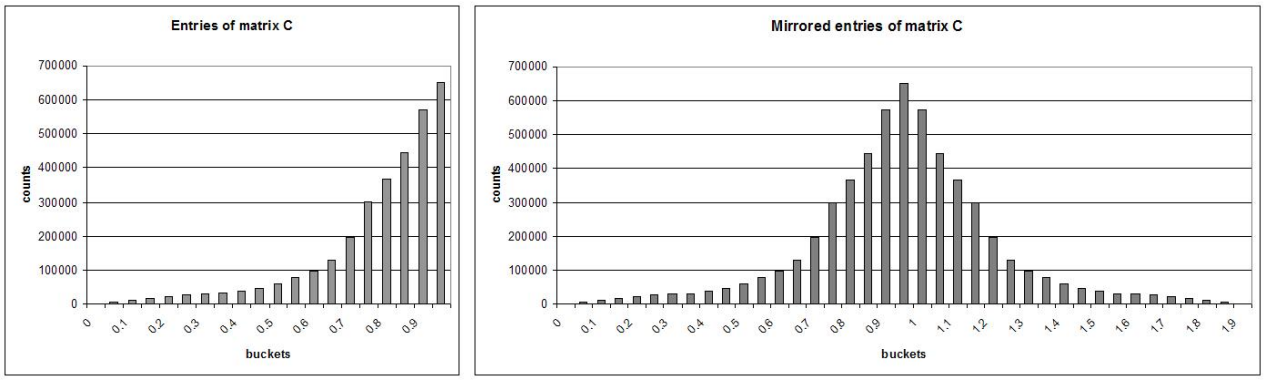

Figure 1. *The distribution of the entries of the matrix C over all 2^k candidate correspondences for one of the Alien models. The left side shows the distribution. The right side shows the Gaussian obtained by mirroring the values along $x = 1$.*

5. Choosing a Candidate Correspondence

Recall that we assume that $S^{(0)}$ and $S^{(1)}$ are near-isometric. Hence, we wish to accept the correspondence that deforms $S^{(0)}$ to $S^{(1)}$ in the most isometric way. Given a candidate correspondence, we know two sets of ordered vertex coordinates: a set of vertex coordinates $V^{(0)}$ on $S^{(0)}$ and a set of vertex coordinates $V^{(1)}$ on $S^{(1)}$. Note that $V^{(r)}$ is a subset of the vertices of $S^{(r)}$ because some vertices have no correspondence. To choose the correspondence that deforms the shape in the most isometric way, we need to compute a common mesh structure M such that embedding M into $\mathbb{R}^3$ with $V^{(0)}$ as vertices approximates $S^{(0)}$ and such that embedding M into $\mathbb{R}^3$ with $V^{(1)}$ as vertices approximates $S^{(1)}$. Once M is known, we need to evaluate a cost function such that minimizing the cost function over all candidate alignments yields the correspondence that deforms $S^{(0)}$ to $S^{(1)}$ in the most isometric way.

5.1. Computing a Common Mesh M

Recall that we know the underlying mesh structures $M^{(0)}$ of $S^{(0)}$ and $M^{(1)}$ of $S^{(1)}$. Note that $M^{(0)}$ and $M^{(1)}$ do not have the same topology. Furthermore, $M^{(0)}$ and $M^{(1)}$ may be incomplete or even disconnected.

We compute two mesh structures $M'^{(r)}$ starting from $M^{(r)}$. The goal for $M'^{(r)}$ is to only contain vertices that are in the set $V^{(r)}$. We compute $M'^{(r)}$ from $M^{(r)}$ using half-edge collapses. We collapse all the half-edges until only vertices in $V^{(r)}$ remain. We always collapse the half-edge that results in the least change in volume. Lindstrom and Turk [24] explain how to find this half-edge. This results in two mesh structures on $V^{(r)}$: $M'^{(0)}$ and $M'^{(1)}$. Both $M'^{(0)}$ and $M'^{(1)}$ have the property that embedding them into $\mathbb{R}^3$ with $V^{(0)}$ as vertices approximates $S^{(0)}$ and that embedding them into $\mathbb{R}^3$ with $V^{(1)}$ as vertices approximates $S^{(1)}$.

5.2. Evaluating a Cost Function

Recall that M denotes a mesh structure such that embedding M into $\mathbb{R}^3$ with $V^{(0)}$ as vertices approximates $S^{(0)}$ and such that embedding M into $\mathbb{R}^3$ with $V^{(1)}$ as vertices approximates $S^{(1)}$. Given M, $V^{(0)}$, and $V^{(1)}$, we can evaluate how isometric the two deformed shapes are by computing the energy

$$E(M) = \sum_{(i,j)\in E} \left(\left\| v_i^{(0)} - v_j^{(0)} \right\| - \left\| v_i^{(1)} - v_j^{(1)} \right\| \right)^2,$$

where E is the edge set of M and $v_i^{(r)}$ is the i-th vertex in the ordered set $V^{(r)}$ for $r = 0, 1$. The energy $E(M)$ measures how isometric the shapes are because a deformation of a shape represented by a triangular mesh is isometric if and only if all triangle edge lengths are preserved during the deformation [18].

The energy $E(M)$ is a deformation distance in the shape space $\mathcal{S}$ that captures the intrinsic geometry of the deformed mesh [7]. Kilian et al. [18] note that we can find the true point-to-point correspondences between two deformed triangular meshes over a set of possible candidate correspondence alignments by computing the distance in $\mathcal{S}$ for each alignment. The correct alignment is found as the alignment with minimum distance. This observation holds because the correct correspondence is the shortest geodesic path in shape space $\mathcal{S}$. Note that the shape space used by Kilian et al. simply encodes the extrinsic geometry of the deformed shapes. We use a different shape space introduced by Bose et al. [7] because this shape space encodes the intrinsic geometry of the deformed mesh. It is therefore more suitable to measure the deformation distance.

To obtain a symmetric deformation energy, we evaluate the energy twice and choose the cost as the maximum of the two results. That is, the deformation cost for a candidate correspondence is computed as

$$E_D = \max(E(M'^{(0)}), E(M'^{(1)})).$$

6. Results

This section presents experiments using the algorithm presented in this paper. The experiments were conducted using an implementation in C++ on an Intel Pentium D with 3.5 GB of RAM. OpenMP was used to improve the efficiency of the algorithms. To minimize the energy when computing the canonical form, a quasi-Newton method is used. The quasi-Newton method used is the limited-memory Broyden-Fletcher-Goldfarb-Shanno scheme [25]. For implementation details of this approach, see Wuhrer et al. [31].

We first show some results of the correspondence algorithm. We then show how the computed correspondence can be used to find near-rigid components.

6.1. Correspondence Results

The first experiment evaluates the quality of our approach. We compute the correspondences between four poses of an alien model. Each of the poses contains 6858 vertices. The models were obtained by animating a model from the Princeton Shape Benchmark [28] using the approach by Baran and Popović [4]. The experiment is shown in Figure 2. We found the correspondences between posture (a) and postures (b),(c), and (d) in this experiment. Vertices on the arms, legs, and torso are assigned a color in posture (a). The corresponding points in postures (b),(c), and (d) are then displayed using the same color. Note that some vertices do not obtain a color in postures (b),(c), and (d) because no correspondence was found for those vertices. We can see that a visually pleasing correspondence is found for all postures. Our implementation takes about 35 minutes to find each of the correspondences.

We compare the correspondences found by our algorithm to the ground truth by computing the geodesic distances between the correspondence found by the algorithm and the true correspondence for each vertex. We measure the error in correspondence as the number of edges along the shortest path between the correspondence found by the algorithm and the true correspondence. Since the algorithm rejects erroneous matchings automatically, some points do not obtain a correspondence. We do not assign an error to rejected correspondences. When registering posture (a) and posture (b), 5764 correspondences are rejected as erroneous. When registering posture (a) and posture (c), 4710 correspondences are rejected as erroneous. When registering posture (a) and posture (d), 4930 correspondences are rejected as erroneous. A histogram of the error encountered is shown in Figure 3. The histogram shows three different data sets: the set of errors when corresponding posture (a) to posture (b) is shown in light grey, the set of errors when corresponding posture (a) to posture (c) is shown in dark grey, and the set of errors when corresponding posture (a) to posture (d) is shown in white. Nearly all of the correspondences found by our algorithm between posture (a) and postures (c) and (d) are accurate within a distance of two edge lengths. This shows that the presented approach yields correspondences of high quality. In case of corresponding postures (a) and (b), the best candidate correspondence is found. That is, no front-to-back mismatching occurs. The large correspondence error is a result of non-isometric models.

Note that unlike in the previous approach by Wuhrer et al. [31], no symmetry alignment problems occur. However, our approach accepts fewer correspondences.

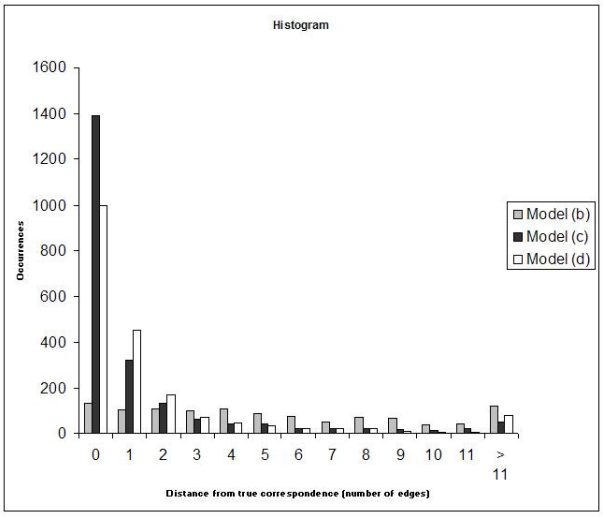

Figure 3. *Histogram of errors.*

Second, we evaluated the proposed approach using the four models with large non-rigid deformation shown in Figure 4. Each row shows the correspondences between the leftmost model in the row with each of the other models in the same row. The correspondences are visualized as follows. For the experiments shown in the first and third rows, we manually segmented the leftmost models. Each point that belongs to a segment is assigned a color in the leftmost model. For the remaining experiments, each point that has a corresponding point in another posture is assigned a unique color in the leftmost models. The corresponding points in the other postures are then displayed using the same color.

The cat and horse models were created and used by Sumner et al. [29]. The cat models contain 7207 vertices. The horse models contain 8431 vertices. In both cases, we use 2000 samples to compute the canonical forms. The gorilla and centaur models were created by Bronstein et al. [9]. The gorilla models contain between 2028 and 2046 vertices. The centaur models contain 3400 vertices. In both cases, we use 1000 samples to compute the canonical forms.

The experiments show that the proposed approach is suitable to compute pairwise correspondences between

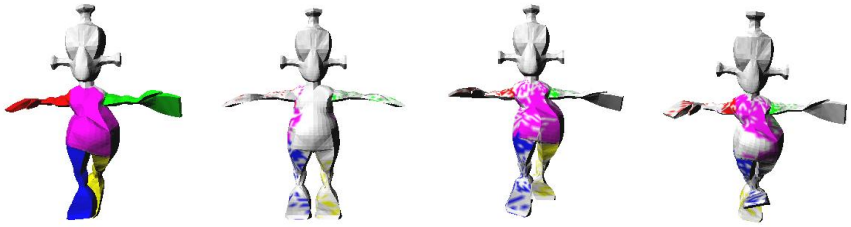

(a) (b) (c) (d)

Figure 2. *Correspondence for models of alien in postures (a) to (d) with known ground truth.*

models with large deformation. Our implementation computes each of the correspondences in less than 10 minutes.

Third, we evaluated the performance of the proposed approach for models with incomplete underlying meshes. We visualize the correspondence using the same approach as in the previous experiments. The models shown in Figure 5 contain many small holes. The models are from the McGill 3D shape benchmark [34] and contain 21338 and 25658 vertices. We use 2050 samples to compute the canonical forms. Note that the correct correspondence is found. No symmetric alignment problems occur. This shows that the proposed approach is suitable for incomplete models with small holes. Our implementation computes the correspondences in about 80 minutes.

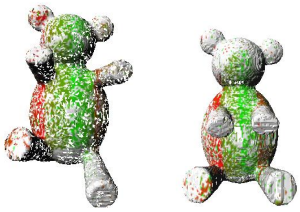

(a) (b)

Figure 5. *Correspondence for incomplete models of a bear.*

The proposed approach has limitations. One limitation we encountered in our experiments is due to the assumption that models deform approximately isometrically. If the models deform non-isometrically, the incorrect correspondence is found. An example of this limitation is shown in Figure 6. The models were created by Bronstein et al. [9]. Poses (e) and (f) are registered correctly. However, the remaining poses register incorrectly. Note that none of the incorrect poses is simply a symmetric misalignment. Instead, parts of the object are registered correctly and other parts of the object are registered incorrectly. For instance, in pose (b) the arms are registered correctly while the legs are registered incorrectly. Note that in those cases, none of the 8 candidate correspondences we compute is correct. This problem does not arise because we choose the incorrect candidate correspondence, but because none of the candidate correspondences is correct. This problem occurs although the embedding error is small when computing the canonical forms. This shows that the models are not isometric.

6.2. Application to Finding Near-Rigid Components

An application of the correspondence problem is the problem of segmenting a mesh into near-rigid components based on a set of deformed input meshes. This section applies the proposed algorithm to the problem of segmenting a mesh into near-rigid components. To compute the segmentation, we first need to find a complete corresponding mesh. Note that the proposed algorithm only computes point-to-point correspondences for a subset of the vertices because erroneous correspondences are rejected using the kernel technique described above.

To compute two complete meshes in correspondence from one of the reduced meshes $M'^{(r)}$ for $r = 0, 1$, we use the mean-value geometry encoding introduced by Kraevoy and Sheffer [21]. With the mean-value encoding, we can find a deformation of $M^{(r)}$ that interpolates the points with known correspondence in $M^{((r+1) \mod 2)}$. To compute this deformation, an energy needs to be minimized. We minimize this energy using numerical derivatives as in Kraevoy and Sheffer. We make the deformed mesh as isometric as possible to $M^{(r)}$ by minimizing the isometric energy introduced by Kilian et al. [18]. We slightly modify the energy to ensure that the deformed mesh interpolates the points with known correspondence in $M^{((r+1) \mod 2)}$.

Once the correspondence is known, the near-rigid components are found using a clustering technique in the shape space encoding the intrinsic geometries of the meshes. The clustering algorithm proceeds in two steps. First, clusters are computed in the dual domain of the mesh. Two triangles are in the same cluster if they move almost rigidly to each other. Second, small clusters are merged to the closest adjacent cluster to avoid over segmentation of the model. The closest adjacent cluster is the one that moves the most rigid with respect to the current cluster.

Figure 7 shows the near-rigid segmentation obtained for the models shown in the top row of Figure 4. Note that the overall skeletal structure is captured. However, the face of the cat is over-segmented due to incorrect correspondences.

Figure 4. *Correspondence for cat, horse, gorilla, and centaur models with large non-rigid deformation.*

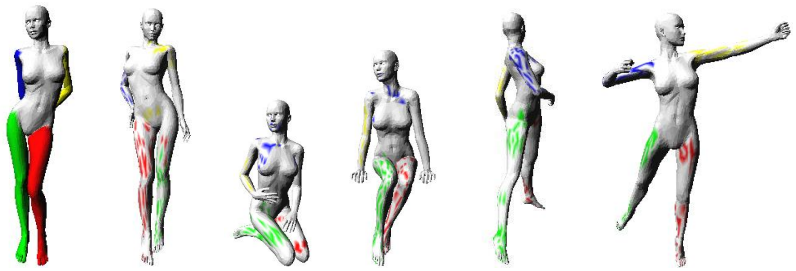

(a) (b) (c) (d) (e) (f)

Figure 6. *Correspondence for models of a dancer that does not deform isometrically.*

Figure 7. *Near-rigid segmentation of the cat model.*

7. Conclusion

This section summarizes the contribution of this work by comparing our algorithm to the recent related work by Zhang et al. [33] and Huang et al. [15]. Furthermore, we summarize limitations of our algorithm.

We find the correct correspondence by minimizing a deformation cost. This is similar to the recent approaches by Zhang et al. and Huang et al. Like the algorithms by Zhang et al. and Huang et al., our algorithm works well for isometric surfaces. No prior knowledge about the objects being registered is required. The algorithm by Huang et al. outperforms our algorithm in terms of running time. The running time of Zhang et al.'s algorithm is comparable to the running time of our algorithm.

The main advantage of this work compared to the work by Zhang et al. and Huang et al. is the presentation of a parameter free method that solves the correspondence problem. This is an important contribution because the large number of non-intuitive user-specified parameters required by the algorithms by Zhang et al. and Huang et al. renders their algorithms impractical.

Finally, we summarize some limitations of our approach that should be addressed in the future:

- Surfaces that are not isometric cannot be registered reliably using this algorithm, as discussed in Section 6.
- Surfaces with large holes cannot be registered reliably using this algorithm, since large holes alter the global shape of the canonical embedding of the sample points.
- Surfaces that cannot be represented well in Euclidean spaces cannot be registered reliably using this algorithm. Although the gorilla models have large embedding error, we never encountered this problem in our experiments. This problem is intrinsic to approaches that use canonical forms in $\mathbb{R}^k$ and can be avoided by computing canonical forms on template shapes [9].
- Surfaces with many significant outliers cannot be registered reliably using this algorithm, because MDS is not robust with respect to outliers [11].

References

[1] B. Allen, B. Curless, and Z. Popović. The space of human body shapes: reconstruction and parameterization from range scans. *ACM TOG*, 22(3):587–594, 2003.

[2] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, H.-C. Pang, and J. Davis. The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *Neural Inf. Proc. Systems*, 2004.

[3] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis. Scape: shape completion and animation of people. *ACM TOG*, 24(3):408–416, 2005.

[4] I. Baran and J. Popović. Automatic rigging and animation of 3d characters. *ACM TOG*, 26(3), 2007.

[5] P. Besl and N. McKay. A method for registration of 3-d shapes. *IEEE TPAMI*, 14(2):239–256, 1992.

[6] I. Borg and P. Groenen. *Modern Multidimensional Scaling Theory and Applications*. Springer, 1997.

[7] P. Bose, J. O'Rourke, C. Shu, and S. Wuhrer. Isometric morphing of triangular meshes. *CCCG*, pages 55–58, 2008.

[8] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS*, 103(5):1168–1172, 2006.

[9] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Calculus of non-rigid surfaces for geometry and texture manipulation. *IEEE TVCG*, 13(5):902–913, 2007.

[10] K. L. Clarkson. A randomized algorithm for closest-point queries. *SIAM Journal on Comput.*, 17(4):830–847, 1988.

[11] T. Cox and M. Cox. *Multidimensional Scaling, Second Edition*. Chapman & Hall CRC, 2001.

[12] I. Dryden and K. Mardia. *Statistical Shape Analysis*. Wiley, 2002.

[13] A. Elad and R. Kimmel. On bending invariant signatures for surfaces. *IEEE TPAMI*, 25(10):1285–1295, 2003.

[14] Y. Eldar, M. Lindenbaum, M. Porat, and Y. Zeevi. The farthest point strategy for progressive image sampling. *IEEE TIP*, 6(9):1305–1315, 1997.

[15] Q. Huang, B. Adams, M. Wicke, and L. J. Guibas. Non-rigid registration under isometric deformations. *CGF*, 27(5), 2008.

[16] V. Jain and H. Zhang. A spectral approach to shape-based retrieval of articulated 3d models. *CAD*, 39(5):398–407, 2007.

[17] V. Jain, H. Zhang, and O. van Kaick. Non-rigid spectral correspondence of triangle meshes. *IJSM*, 13(1):101–124, 2007.

[18] M. Kilian, N. J. Mitra, and H. Pottmann. Geometric modeling in shape space. *ACM TOG*, 26(3), 2007.

[19] R. Kimmel and J. Sethian. Computing geodesic paths on manifolds. *PNAS*, 95:8431–8435, 1998.

[20] V. Kraevoy, A. Sheffer, and C. Gotsman. Matchmaker: Constructing constrained texture maps. *ACM TOG*, 22(3):326–333, 2003.

[21] V. Kraevoy, A. Sheffer, and C. Gotsman. Mean-value geometry encoding. *IJSM*, 12(1):29–46, 2006.

[22] M. Leordeanu and M. Hebert. A spectral technique for correspondence problems using pairwise constraints. *IEEE ICCV*, 2005.

[23] X. Li and I. Guskov. Multi-scale features for approximate alignment of point-based surfaces. In *SGP*, 2005.

[24] P. Lindstrom and G. Turk. Fast and memory efficient polygonal simplification. In *IEEE Vis.*, pages 279–286, 1998.

[25] D. C. Liu and J. Nocedal. On the limited memory method for large scale optimization. *Math. Programming*, 45:503–528, 1989.

[26] J. Munkres. Algorithms for the assignment and transportation problems. *Journal of the Soc. of Industrial and Applied Math.*, 5(1):32–38, 1957.

[27] J. Schreiner, A. Asirvatham, E. Praun, and H. Hoppe. Inter-surface mapping. *ACM TOG*, 23(3):870–877, 2004.

[28] P. Shilane, P. Min, M. Kazhdan, and T. Funkhouser. The princeton shape benchmark. *SMI*, 2004.

[29] R. W. Sumner and J. Popović. Deformation transfer for triangle meshes. *ACM TOG*, 23(3):399–405, 2004.

[30] Y. M. Sun, W. Wang, and F. Chin. Interpolating polyhedral models using intrinsic shape parameters. *The Journal of Vis. and Computer Animation*, 8(2):81–96, 1997.

[31] S. Wuhrer, C. Shu, Z. B. Azouz, and P. Bose. Posture invariant correspondence of incomplete triangular manifolds. *IJSM*, 13(2):139–157, 2007.

[32] W. Zeng, Y. Zeng, Y. Wang, X. Yin, Xianfeng Gu, D. Samaras, 3D Non-rigid Surface Matching and Registration Based on Holomorphic Differentials. *ECCV*, 2008.

[33] H. Zhang, A. Sheffer, D. Cohen-Or, Q. Zhou, O. van Kaick, and A. Tagliasacchi. Deformation-driven shape correspondence. *CGF*, 27(5), 2008.

[34] J. Zhangand,K. Siddiqi, D. Macrini, A. Shokoufandeh, and S. Dickinson. Retrieving articulated 3-d models using medial surfaces and their graph spectra. *EMMCVPR*, 2005.