3D shape isometric correspondence by spectral assignment

Xiang Pan [1] Linda Shapiro [2]
[1] College of Computer, Zhejiang University of Technology, 310014 Hangzhou, P. R. China
[2] Computer Science and Engineering, University of Washington, 98195 Seattle, USA
panx@zjut.edu.cn, shapiro@cs.washington.edu

Abstract

Finding correspondences between two 3D shapes is common both in computer vision and computer graphics. In this paper, we propose a general framework that shows how to build correspondences by utilizing the isometric property. We show that the problem of finding such correspondences can be reduced to the problem of spectral assignment, which can be solved by finding the principal eigenvector of the pairwise correspondence matrix. The proposed framework consists of four main steps. First, it obtains initial candidate pairs by performing a preliminary matching using local shape features. Second, it constructs a pairwise correspondence matrix using geodesic distance and these initial pairs. Next, the principal eigenvector of the matrix is computed. Finally, the final correspondence is obtained from the maximal elements of the principal eigenvector. In our experiments, we show that the proposed method is robust under a variety of poses. Furthermore, our results show a great improvement over the best related method in the literature.

1. Introduction

In computer vision and computer graphics, building correspondences between two 3D shapes is an important problem that allows the semantic structure of shape parts to be determined. The correspondence results can be used for numerous applications, such as object recognition, time-varying surface reconstruction, mesh parameterization, and statistical shape modeling. As a consequence, the problem of finding 3D shape correspondence has attracted the attention of many researchers.

Isometry is a very fundamental characteristic of 3D shape correspondences. An isometric transformation is one that is invariant with distance. Thus we would like to find an optimal correspondence that minimizes the distance error between two shapes. If the distance is defined to be geodesic distance, this correspondence is robust under pose-variation and large deformations. However, finding such an optimal correspondence is a very typical NP-hard problem. Therefore, a general solution that minimizes isometric error is the key point to building good correspondences.

Addressing this problem, we proposed a general framework of minimizing the isometric error by spectral assignment. The idea of spectral assignment is to build a pairwise correspondence matrix using geometric properties and geodesic distance. Based on Raleigh's ratio theorem, the problem of minimizing isometric error can be obtained by finding pairs that correspond to the maximal elements of the principal eigenvector of a pairwise correspondence matrix. This proposed framework has two distinctive advantages. First, it works directly on the original Euclidean space instead of a transformed space and will therefore not produce an approximation error. Secondly, this method will not force every source point to be mapped to a destination point, a characteristic that is important for partial correspondences. To further verify the proposed framework, this paper also gives an implementation, which uses multidimensional scaling (MDS) with least squares error for the initial matching. Experiments show our method is robust under pose variation and shows a great improvement over the most related method in the literature.

2. Related work

The problem of 3D shape correspondence can be solved by different mapping methods, such as distance-preserving mappings, conformal mappings, and heat mappings [1]. Our method is based on distance-preserving mappings. Therefore, this section gives a brief review of this kind of method.

Correspondence methods by distance-preserving mapping can be further classified into two categories. The first category uses multidimensional scaling

(MDS), proposed by Elad *et al* [2], who define bending invariant signatures by approximating Euclidean distance by geodesic distance over the surface. Samuel *et al* [3] fit 3D humans to a template, iteratively transforming the template to the input model using MDS. Wuhree *et al* [4] uses Markov network to learn the spatial relations among a human's feature points and builds correspondences by maximizing a joint probability over all possible configurations of 3D human poses. Jain *et al* [5] transform two input shapes into spectral space using geodesic distance, apply principal component analysis to get an initial alignment, and use thin-plated splines to obtain the correspondence. To further improve Jain's algorithm, Sahillioğlu *et al* [6] perform an additional step after spectral correspondence to minimize the isometric cost using an iterative greedy algorithm. Ruggeri *et al* [7] construct a point-based statistical descriptor by combining geodesic shape distribution and other geometric information.

Our work is most similar to Sahillioğlu 's work. Both methods directly work on the original 3D Euclidean space. However, Sahillioğlu's method requires many correct correspondences to be found in the initial matching. Our method only obtains a set of possible correspondences for the initial matching. Our condition is less strict than Sahillioğlu's method, and therefore, it is more general. In addition, we improved his initial matching method to find better correspondences.

3. Problem formulation

We cast the problem of 3D shape correspondence as a point-to-point correspondence problem. Therefore, we assume there are N and M points ($N \leq M$) evenly distributed over the surface S and T, respectively. The correspondence task is to build a point-to-point mapping between two sets S: $\{s_1, s_2, \cdots s_i, \cdots s_N\}$ and T: $\{t_1, t_2, \cdots t_i, \cdots t_M\}$. For any two points x_i and x_j , their geodesic distance is denoted by $geod(x_i, x_j)$. Therefore, the problem of minimizing isometric error can be defined by the following equation:

$$D_{iso} = \operatorname{argmin} \sum_{i,j=1}^{N} \left| geod(s_i, s_j) - geod\left(\varphi(s_i), \varphi(s_j)\right) \right| \quad (1)$$

where $\varphi(s_i)$ denotes a correspondence for point s_i. To generate such an isometric correspondence, we first obtain a set of possible correspondences

$$P = \{(s_1, T_1), \cdots (s_i, T_i), \cdots, (s_N, T_N)\}$$

where T_i is a set of K possible corresponding points in T for point s_i in S. This process can be performed by some typical shape features, for example SIFT features, MeshDOG features and Euclidean distance of two aligned shapes. Second, we compute a pairwise correspondence matrix M, each of whose rows and each of whose columns represent possible correspondences (s_i, t_j) between a point $s_i \in S$ and a point $t_j \in T$. Thus the matrix has $N \times K$ rows and also $N \times K$ columns. The non-diagonal elements of the matrix represent the compatibility of a correspondence (s, t) with another correspondence (s', t') . The compatibility is given by:

$$c - |geod(s, s') - geod(t, t')| \quad (2)$$

where c is a constant to convert the equation (1) into a maximum problem. The compatibility is set to be zero if $s = s'$ and $t \neq t'$ or $s \neq s'$ and $t = t'$. The diagonal elements of the matrix represent the dissimilarity of a single correspondence (s, t) and are computed by shape features as mentioned above. In our implementation, we only use local features to build the initial matching set P . Therefore, the value of the diagonal elements is initialized to 0. Given P and M, the isometric correspondence can be obtained by maximizing the following equation:

$$X^* = \operatorname{argmax}(X^T M X) \quad (3)$$

where X is a characteristic vector denoting the final correspondences. The length of X is equal to $|P|$. For each element, its value is 1 if it is a good correspondence and 0 otherwise. Based on Raleigh's ratio theorem, the above maximum problem can be solved by computing the principal vector of matrix M. Section 4.2 will give more details about how to obtain X by greedy algorithm.

4. Implementation

This section gives an implementation for the above framework. We first perform an even sampling over two shapes. Then the initial matching is performed using Euclidean distance, since two corresponding points are close to each other in two aligned shapes. Finally, isometric correspondence can be obtained by solving equation (3).

4.1 Even sampling

The process evenly samples a subset of vertices over the surface of each shape. After even sampling, each 3D shape is divided into a set of almost equal patches. For each patch, its center is called a base vertex. The correspondences are only built for the base vertices.

The whole sampling process can be iteratively performed by geodesic distance. First, all vertices are set to be unmarked. In each iteration, one unmarked vertex is selected as the seed point and marked. Those vertices whose geodesic distances to the seed point are less than a predefined radius r are also marked. The iterations continue until all vertices are marked. Based on these patches, we can construct base vertices, which are at least a distance r apart from each other. Figure 1 shows a result of even sampling. These base vertices are evenly distributed over the 3D shape. Therefore, correspondences of base vertices can be easily extended to build dense correspondences.

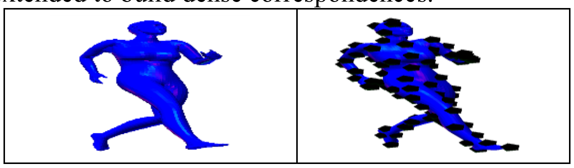

Figure 1. 3D shape and its even sampling

4.2 Initial matching

Using these base vertices, we can perform initial matching. As shown in Section 3, our framework is to select good correspondences from the results of the initial matching. Therefore, the quality of initial matching plays a very important role for the final correspondences. Here the initial matching is performed by combining MDS and least squares error. First, for each 3D shape, MDS is performed to build a pose-invariant representation. Second, two transformed 3D shapes are aligned by minimizing the least square error. Figure 2 shows an example of two aligned shapes. It can be seen that each base vertex in the source shape and its correct correspondence in the destination shape are very close. Therefore, we can use Euclidean distance to build the set P of possible correspondences. For each base vertex of the source shape, the top K nearest neighbors are used to construct possible correspondences. (We use K=5.)

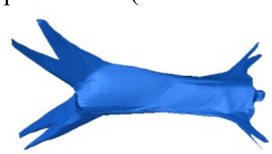

Figure 2. Alignment of two 3D shapes

In a practical implementation, the process of minimizing least squares error is replaced by Principal Components Analysis (PCA), which can greatly reduce the complexity. First, two transformed shapes are aligned by PCA. Due to the ambiguity of PCA, we need to flip the sign of each axis to see which direction has the least squared error between the two shapes. Finally the direction with the least error is selected as the final alignment. In this way, we only test 8 possible directions to get the initial alignment. Therefore, it is very efficient compared to the traditional process of minimizing least squares error.

4.3 Correspondences by greedy optimization

Now the problem of correspondences can be solved by computing the eigenvectors of matrix M. We first obtain the principal eigenvector of matrix M. Notice that each component of the principal eigenvector is mapped to a possible correspondence, while, its scalar value denotes the possibility of a correct correspondence. In particular, the larger the scalar value of one component, the better the correspondence pair. Therefore, we can formulate the problem of finding final correspondences as a greedy optimization as summarized in Table 1.

Table 1. Greedy best first search algorithm

1. Initialize the vector X as a zero vector.
2. Compute principal eigenvector $PrinV$ of M.
3. Find maximum element V_{max} of vector $PrinV$. Its corresponding element in X is set to be 1.
4. Remove maximum element V_{max} from $PrinV$. At the same time, remove all potential assignments in confliction with V_{max}.
5. If $PrinV$ is empty, the process returns the solution X. Otherwise go back to step 3.

5. Experimental analysis

In our experiments, we verify our proposed framework using a public 3D shape database called TOSCA. The database contains 80 models classified into 9 categories. For each category, we select one model as the source shape. Then correspondences between this source shape and the other shapes in the same category are performed to verify the proposed framework. Figure 3 shows some results, with only 5% of matching pairs displayed for better visualization. For two mapped points of a pair, we find the position of a source point is very similar to the position of its destination point regardless of the varying poses.

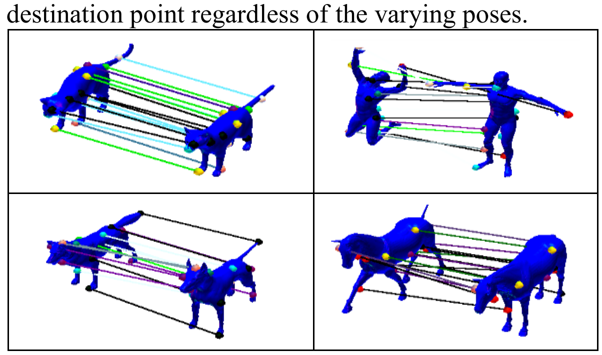

Figure 3. Some examples of correspondences

Next, we perform a statistical analysis of the corresponding results. There are 68 pairs in the test database. In addition, we accept symmetric flips in isometric correspondence. Table 2 lists the number of incorrect pairs for each category. As a comparison, we also show the number of incorrect pairs by Sahillioğlu's algorithm [6].

Table 2. Comparison of two algorithms

Category	Our algorithm	Sahillioğlu's
Cat	0	1
Centaur	0	2
David	0	1
Dog	0	0
Gorilla	0	1
Horse	0	0
Michael	0	7
Victoria	0	9
Wolf	0	0

As shown in the above table, our method has a great improvement over his. Since both methods are actually a refinement of initial matching, the final correspondences are highly depending on that initial matching. Sahillioğlu's algorithm uses spectral embedding for initial matching, and each eigenvector is independently scaled by the corresponding eigen-value. However, this scaling strategy will seriously destroy the topological features of the 3D shape, causing erroneous correspondences. Figure 4 gives such an example, where the head of the human is wrongly aligned to the feet of the other human (Figure 4(a)). On the contrary, our algorithm outputs correct results (Figure 4(b)). Finally, our algorithm is much more general than Sahillioğlu's method.

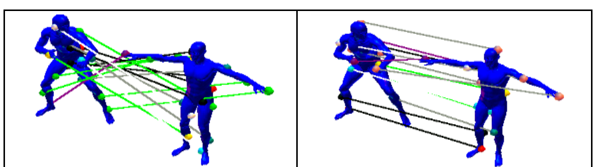

(a) Sahillioğlu's algorithm

(b) our algorithm

Figure 4. Comparison to Sahillioğlu's algorithm

Like other related algorithms, the main limitation of our current implementation is still the problem of symmetric flip, which cannot be distinguished in initial matching, as shown in Figure 5. For this problem, it is possible to add some constraints by matching local features.

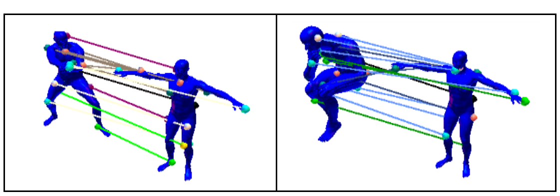

Figure 5. The problem of symmetric flip

6. Conclusion

This paper gives a general framework for finding isometric correspondence by spectral assignment which is performed by computing the principal eigenvector of the pair-wise correspondence matrix and is stable under pose-variation. There is much future work to be done. First, our current implementation of initial matching only works well when two shapes can be aligned correctly. Therefore, we need to improve the initial matching for better results. Secondly, better local shape features will be considered to support partial correspondence.

Acknowledgement

This work was partly supported by Natural Science Foundation of Zhejiang Province, China (Grant No: Y1110780, Y1100824) and partly supported by NIH/NIDCR grant No. U01DE020050. The authors thank Yusuf Sahillioğlu for providing his source code.

References

[1] Kaick, V., et al., A survey on shape correspondence. Computer Graphics Forum, 2010. 30(6): p. 1681–1707.

[2] Elad, A. and R. Kimmel, On Bending Invariant Signatures for Surfaces. IEEE Trans on Pattern Analysis and Machine Intelligence, 2003. 25: p. 1285–1295

[3] Samuel, S.M.L., C.C.L. Wang, and K.C. Hui, Bending-invariant correspondence matching on 3D human bodies for feature point extraction. IEEE Transactions on Automation Science and Engineering, 2011. 8(4): p. 805 - 814.

[4] Wuhrer, S., C. Shu, and P. Xi, Landmark-Free Posture Invariant Human Shape Correspondence. The Visual Computer, 2011. 27(9): p. 843-852.

[5] Jain, V. and H. Zhang. Robust 3D Shape Correspondence in the Spectral Domain. in Shape Modeling and Applications. 2006.

[6] Sahillioğlu, Y. and Y. Yemez. 3D Shape Correspondence by Isometry-Driven Greedy Optimization. in Computer Vision and Pattern Recognition. 2010.

[7] Ruggeri, M.R., et al., Spectral-Driven Isometry-Invariant Matching of 3D Shapes. International Journal of Computer Vision, 2010. 89(3): p. 248-265.