Open-Curve Shape Correspondence Without Endpoint Correspondence

Theodor Richardson and Song Wang

Department of Computer Science and Engineering,
University of South Carolina, Columbia, SC 29208, USA
richa268@cse.sc.edu, songwang@cse.sc.edu

Abstract. Shape correspondence is the foundation for accurate statistical shape analysis; this is usually accomplished by identifying a set of sparsely sampled and well-corresponded landmark points across a population of shape instances. However, most available shape correspondence methods can only effectively deal with complete-shape correspondence, where a one-to-one mapping is assumed between any two shape instances. In this paper, we present a novel algorithm to correspond 2D open-curve partial-shape instances where one shape instance may only be mapped to part of the other, i.e., the endpoints of these open-curve shape instances are not presumably corresponded. In this algorithm, some initially identified landmarks, including the ones at or near the endpoints of the shape instances, are refined by allowing them to slide freely along the shape contour to minimize the shape-correspondence error. To avoid being trapped into local optima, we develop a simple method to construct a better initialization of the landmarks and introduce some additional constraints to the landmark sliding. We evaluate the proposed algorithm on 32 femur shape instances in comparison to some current methods.

1 Introduction

It is well known that the performance of statistical shape analysis [1, 3, 10, 7, 8] is highly dependent upon the performance of shape correspondence [4, 2, 11, 9], which identifies a set of sparsely sampled and well-corresponded landmark points across a population of shape instances. However, accurate shape correspondence is a very challenging problem given the strong nonlinearity of shape geometry and the complex nonrigid deformation occurring among the considered shape instances.

Many new models and methods have been developed recently to achieve more accurate landmark-based shape correspondence. Each method has its own merits and difficulties. Davies et al [4] develop a Minimum Description Length (MDL) algorithm where the correspondence error is measured by the required bit-length to transmit these shape instances as well as the template shape. Both random-searching [4] and gradient-descent algorithms [6] have been developed to minimize this nonlinear MDL-based measure. Xie and Heng [14] develop a medial-axis-based shape-correspondence method, where the medial axes of all

R. Larsen, M. Nielsen, and J. Sporring (Eds.): MICCAI 2006, LNCS 4190, pp. 17–24, 2006.

the shape instances are assumed to be of the same topology and then the shape-correspondence problem is decomposed into a set of simpler problems of corresponding some short curve segments with very few high-curvature points. With some roughly-corresponded landmarks, Bookstein [2] develops an algorithm to move these landmarks along the tangent directions of the shape contour to achieve a minimum landmark-correspondence error that is defined by the thin-plate bending energy. However, the resultant landmarks may not be located on the underlying shape contour. Wang, Kubota, and Richardson [13, 9] address this problem by adding a step of projecting the landmarks back to the shape contour.

However, all these methods can only address the complete-shape correspondence where a one-to-one mapping is assumed between each pair of the shape instances. In 2D cases, each shape instance is in the form of a continuous curve. For convenience, we sometimes refer to this ground-truth continuous form of a shape instance as a *shape contour*. These shape contours may be open or closed (the starting and ending point of the contour is in the same physical location). When the shape contour is open, the endpoints are always corresponded landmarks across all the shape instances. In many real applications, however, the assumption of the complete-correspondence does not usually hold for open-curve shape contours. An shown in the femur contours in Fig. 1, we can see that it is difficult to guarantee that the initially extracted open femur boundaries have exact endpoint correspondence and one boundary may only be mapped to a portion of the other. Such examples are common in practice since we may only be interested in portions of a structure and it is difficult, even manually, to accurately extract open-curve shape contours with consistently corresponded endpoints. Furthermore, many medical images suffer from cropping and occlusion within a whole structure. In this paper, we develop an algorithm to address such an open-curve partial shape correspondence problem by extending and adapting the landmark-sliding algorithms [2, 13, 9]. This algorithm is tested on 32 femur shape instances and the correspondence performance is quantitatively compared with the results of a popular implementation of the MDL algorithm [12] and the most recent landmark-sliding algorithm [9].

2 Problem Formulation and Landmark-Sliding

In this paper, we consider 2D partial-shape correspondence, where each shape instance is in the form of a continuous open-shape contour, as shown in Fig. 1. Denote the given set of open-curve shape instances to be $S = \{S_1, S_2, \ldots, S_n\}$. Each shape instance S_i is in the form of an arc-length parameterized curve $\mathbf{s}_i(t_i) = (x_i(t_i), y_i(t_i))$, $0 \leq t_i \leq L_i$, where L_i is the perimeter length of S_i and t_i is the traversed arc-length along the curve from $\mathbf{s}_i(0)$ to $\mathbf{s}_i(t_i)$. Note that $\mathbf{s}_i(0)$ and $\mathbf{s}_i(L_i)$ are the two endpoints of the shape contour S_i. In practice, we obtain these shape instances by applying the Catmull-Rom spline interpolation to approximate some structural boundaries that are manually extracted from the medical images.

To address the partial shape correspondence, we need to determine a continuous subcontour P_i of each shape contour S_i such that a complete correspondence exists across P_i, $i = 1, \ldots, n$. As a subcontour of S_i, we can denote the two endpoints of P_i as $P_{i1} = \mathbf{s}_i(t_{i1})$ to $P_{iK} = \mathbf{s}_i(t_{iK})$, where $t_{i1} \geq 0$ and $t_{iK} \leq L_i$ are the traversed curve lengths from $\mathbf{s}_i(0)$ to these two endpoints of P_i. Without loss of generality, we let $t_{i1} < t_{iK}$. The problem formulation of the open-curve complete-shape correspondence is well known for nonrigid shape correspondence: identifying K nonfixed landmarks along P_i, $i = 1, \ldots, n$ such that a predefined correspondence error among them is minimum and the two endpoints of each shape contour are always selected as corresponded landmarks, i.e., P_{i1}, $i = 1, \ldots, n$ are priorly known to be corresponded and P_{iK}, $i = 1, \ldots, n$ are also priorly known to be corresponded. This way, we can define the open-curve partial shape correspondence as finding parameters $0 \leq t_{i1} \leq \ldots \leq t_{iK} \leq L_i$, $i = 1, \ldots, n$ such that the sampled landmarks $\mathbf{s}(t_{ik})$, $k = 1, \ldots, K$, $i = 1, \ldots, n$ minimize a selected correspondence error.

As in [2] and [13, 9], we define the landmark shape-correspondence error using the thin-plate bending energy [5], which has been widely used to describe nonrigid shape deformation. Since corresponding n shape instances is usually treated by corresponding each shape instance separately to a template shape instance, we can focus on the correspondence between two shape instances: the template V in which the K sequentially sampled landmarks $\mathbf{v}_1, \ldots, \mathbf{v}_K$ have been priorly given and the target U in which we want to find the K corresponded landmarks $\mathbf{u}_1, \ldots, \mathbf{u}_K$. The correspondence error is defined to be the thin-plate bending energy

$$\beta(V \to U) = \frac{1}{8\pi}(\mathbf{x}_u^T \mathbf{M}_v \mathbf{x}_u + \mathbf{y}_u^T \mathbf{M}_v \mathbf{y}_u),$$

where $\mathbf{x}_u$ and $\mathbf{y}_u$ are columnized vectors of x- and y-coordinates of K landmarks along the target U, and $\mathbf{M}_v$ is the thin-plate bending matrix calculated from K landmarks in the template V [5]. Specifically, the bending matrix $\mathbf{M}_v$ is the upper-left $K \times K$ submatrix of

$$\begin{pmatrix} \mathbf{K}_v & \mathbf{P}_v \\ \mathbf{P}_v^T & \mathbf{0} \end{pmatrix}^{-1},$$

where $\mathbf{K}_v$ is $K \times K$ kernel matrix with element $k_{ij} = \mathcal{K}(\mathbf{v}_i, \mathbf{v}_j) = \|\mathbf{v}_i - \mathbf{v}_j\|^2 \log \|\mathbf{v}_i - \mathbf{v}_j\|$ and $\mathbf{P}_v = (\mathbf{x}_v, \mathbf{y}_v, \mathbf{1})$.

Given the nonlinearity of the continuous shape contour and the inability to slide landmarks along the contour directly, we adopt the landmark sliding method of [13], which starts from some initialized K landmarks along the target shape U, moves the landmarks along the tangent directions to reduce the correspondence error, and then projects the updated landmarks back to the shape contour. The sliding and projection operations are iteratively performed until convergence. Specifically, letting α_k be the sliding distance of $\mathbf{u}_k$, $k = 1, \ldots, K$, the landmark-sliding algorithm aims to minimize the correspondence error

$$\phi(\alpha_1, \ldots, \alpha_K) = (\mathbf{x}_u + \mathbf{P}'_x \boldsymbol{\alpha})^T \mathbf{M}_v (\mathbf{x}_u + \mathbf{P}'_x \boldsymbol{\alpha}) + (\mathbf{y}_u + \mathbf{P}'_y \boldsymbol{\alpha})^T \mathbf{M}_v (\mathbf{y}_u + \mathbf{P}'_y \boldsymbol{\alpha}) \quad (1)$$

where $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_K)^T$ and $\mathbf{P}'_x$ and $\mathbf{P}'_y$ are diagonal matrices of the x- and y-components of the estimated tangent vectors at the current K landmarks in the target shape U. The projection is performed by simply traversing the current landmarks along the shape contour by the sliding distance α_k, $k = 1, \ldots, K$. To preserve the shape topology, each landmark is not allowed to slide across its two neighbors; in [13], it is shown that this can be described by some linear constraints [13, 9]. Therefore, the above cost function can be optimized efficiently by the quadratic-programming algorithm. The major problem of the landmark sliding algorithm is its vulnerability to being trapped into local minima. For complete-shape correspondence with known corresponded endpoints, good initialization of the K target landmarks can be easily constructed, which, however, is not the case in the partial-shape correspondence discussed in this paper.

3 Landmark-Sliding for Open-Curve Partial Shape Correspondence

There are several important issues in extending the landmark-sliding algorithm to the open-curve partial shape correspondence. First, the landmark initialization becomes a very important issue without endpoint correspondence. Second, by moving the first and last landmark points, we actually resize the target shape to be mapped to the template. Reducing the size of the target shape, although it is an affine transformation, can decrease the bending energy. Therefore, without further considerations, we can not directly use the bending energy to decide the desirable correspondence results. Third, we also cannot allow the landmarks to move beyond the original endpoints of the open shape contours.

To simplify the algorithm description, we assume that the template shape contour V corresponds to part of the target shape U, i.e., the endpoints of V can be included as landmarks and their correspondence in U is located along U. Then the third issue above is easy to handle: we need only to add a constraint on the sliding distance α_1 and α_K so that they are smaller than the respective remaining external curve length. These are still linear constraints and, therefore, the quadratic programming algorithm still applies. In the following, we extend the landmark-sliding algorithms to address the other two important issues.

3.1 Initial Landmark Construction

The landmark-sliding algorithm is vulnerable to being trapped into a local optima, especially when the considered shape instances contain many high-curvature points, motivating an accurate initialization method. In the open-curve partial shape correspondence, it can be assumed that the shape contours extracted from medical images contain a similar subject of interest that can be extracted within a certain variance expressed as a percentage. For completeness, this variance is not assumed to be *a priori* knowledge though it may be known by an expert. For the cases considered herein, it is reasonable to assume that the variance of shape information at the endpoints is not greater than 50 percent.

We therefore iteratively test the bending energy by clipping an equal percentage ω from both sides of target shape instance U and distributing the K landmarks as described above over the remainder of the curve as follows:

$$\mathbf{t}_u = \mathbf{t}_v \frac{(1-\omega)L_u}{L_v} + \frac{\omega L_u}{2} \tag{2}$$

where $\mathbf{t}_u$ is the set of arc-length parameters for the K landmarks of shape instance U and L_u is the perimeter length of shape instance U ($\mathbf{t}_v$ and L_v are similarly defined for template shape instance V). With each distribution, we can then test the initial bending energy to determine the optimal ω clip length that should be used to establish the final correspondence with V. Because the landmarks of each U are not fixed, the step length is flexible and, within a certain range, the sliding will converge to the same result; for our purposes, we varied the clip length from 0 to 50 percent with a step length of 0.5 percent. The optimal clip length for each shape instance is determined independently to account for greater variability. As an added constraint to prevent undue collapse of the shape contour, endpoints are not allowed to move inward more than $\frac{\omega L_u}{2}$ for each shape instance U; however, the endpoints are allowed to move outward to recapture length that was previously assumed to be removed.

3.2 An Adapted Correspondence Error

The remaining problem is to choose the desirable correspondence result out of the ones that are resulting from different clip lengths. It is well known that the thin-plate bending energy is invariant to the affine transformation, i.e., if the map between V and U can be described by an affine transformation, the resulting bending energy is always zero. However, if the map between V and U is not an affine transformation, then some additional affine transformations on the target landmarks U would change the bending energy that is calculated without consideration of this additional affine transformation. As an example, a scaling factor of 0.75 applied to the target landmarks U in Eq. (2) would decrease the bending energy $\beta(V \to U)$ to $(0.75)^2\beta(V \to U)$. A scaling in either x- or y-coordinates alone will decrease the bending energy in a similar way. In a degenerate case where all target landmarks are moved to the same point, the size of the target shape contour that finally corresponds to the template is zero and, in this case, the bending energy, or the shape-correspondence error, is also zero. Therefore, the direct comparison of the bending energy may not help us decide the desirable correspondence result out of the ones from different clip lengths.

To address this problem, we introduce a normalized bending energy to help decide the desirable correspondence result. Specifically, for the correspondence result from each clip length, we normalize the resulting U by removing the affine transformation between V and U and then calculating their bending energy as the final shape-correspondence error. Then, for the correspondence results from different clip lengths, we pick the one with the smallest such normalized correspondence error as the desirable correspondence. In the next section, we will show that, given the flexility of the landmark-sliding algorithm, the desirable correspondence with smallest normalized correspondence error can be obtained from a certain range of the clip length.

4 Experiments and Analysis

To test the proposed method, we used a fixed template of 41 landmark points and a data set of 32 open-curve femur shapes extracted from medical images. We test the effectiveness of the method by comparing the initial normalized bending energy to the final normalized bending energy (after applying the landmark-sliding algorithm) for each clip length. We can see that the desirable correspondence with minimum normalized bending energy can be obtained from a wide range of clip lengths between 8% and 20%. Therefore, the optimal clip length is not unique and can take a range of values. A sample of the visual results for the optimal clip length can be seen in Fig. 1(c-g) along with (b), which shows the polyline interpolated landmarks in the template shape.

Evaluation and comparison among shape correspondence methods is a difficult problem in which the ground truth is often unknown. As such, we focus on known metrics with which to compare the performance of the proposed method to extant methods for shape correspondence, namely a popular implementation of MDL capable of operating with *nonfixed endpoints* in which the landmarks at the endpoints of the contour are allowed to vary to the same extent as any other landmark in the set[12] and the most recent incarnation of the landmark sliding algorithm [9]. The same template shape was used for all of the methods under consideration to allow for a fair comparison. A brief explanation of the metrics used for comparison follows: (a) $E(\beta)$ and $std(\beta)$: the mean and standard deviation of the thin-plate bending energy between the template and all the shape instances according to the identified landmarks; (b) λ_1, λ_2, and λ_3: the three principal eigenvalues of the covariance matrix of $U_i, i = 1, 2, \ldots, n$. These are the principal zeros of the characteristic polynomial describing the covariance of the data set, which should be as small as possible if the regularity of the shape has been extracted successfully. In calculating the covariance matrix, the Procrustes analysis [1] is applied to normalize the size and orientation of all the shape instances. (c) the sum of the eigenvalues of the covariance matrix of $U_i, i = 1, 2, \ldots, n$; (d) G, the generality of the resulting shape model, calculated according to the leave-one-out cross-validation measure common to statistical processing [11]; and (e) the total CPU time used for processing each data set, based on the adopted specific implementations. In general, with a similar representation error, a good shape correspondence is expected to have small $E(\beta)$, λ_1, λ_2, λ_3, $\sum \lambda$, and G. The statistical improvement gained by this method can be seen in Table 1 as it clearly outperforms the sliding method at a small cost of CPU time and produces results that are comparable to or better than the MDL implementation using free endpoints.

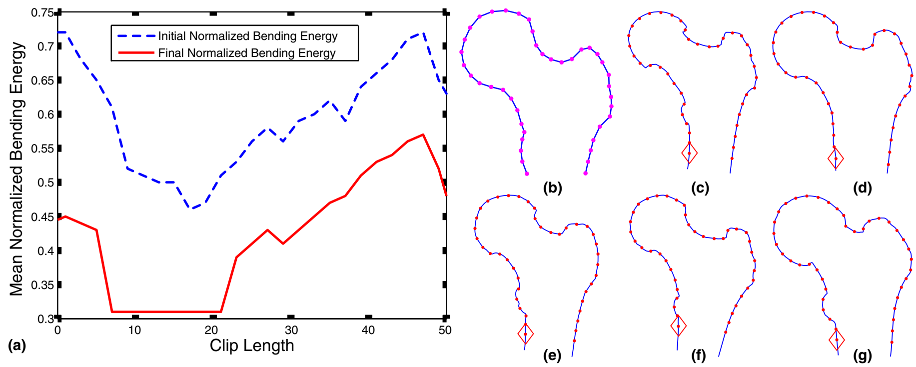

Fig. 1. (a) Comparison results for clip length showing the initial and final normalized bending energy. This bending energy shown here is the mean value on all 32 shape instances. (b) Template shape and (c-g) identified target landmarks on five sample shape instances using the proposed method.

Table 1. Statistical comparison of correspondence results

Measures	MDL [12]	Landmark Sliding developed in [9]	Proposed Method
$E(\beta)$	1.2304	0.4628	0.3149
$std(\beta)$	0.3373	0.1125	0.0824
λ_1	0.0024	0.0030	0.0026
λ_2	0.0008	0.0007	0.0007
λ_3	0.0005	0.0006	0.0005
$\sum \lambda$	0.0051	0.0059	0.0046
$G(\cdot)$	0.8722	0.6337	0.4714
CPU time(s)	418.7	124.7	160.6

5 Conclusion

Because most of the current shape correspondence methods can only effectively deal with complete-shape correspondence, where a one-to-one mapping is assumed between any two shape instances, we present a novel algorithm to correspond open-curve shape instances where a one-to-one correspondence does not exist across all shape instances. In this algorithm, we first introduce an initialization by allowing some portion along the open-curve shape contours to be clipped. We then apply the landmark-sliding algorithm to find the correspondence from each starting clip length. Finally, we introduce a normalized shape-correspondence error to decide the desirable correspondence. The proposed algorithm is tested on 32 femur shape instances and compared to the current landmark-sliding algorithm [9] and an implementation of the MDL method [12]. In this comparison, there is significant statistical improvement using the proposed algorithm compared to the other current methods with only a small CPU time increase over the current landmark-sliding algorithm.

Acknowledgement

This work was funded, in part, by NSF-EIA-0312861 (S. Wang) and an NSF graduate fellowship (T. Richardson).

References

1. F. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585, June 1989.
2. F. Bookstein. Landmark methods for forms without landmarks: Morphometrics of group differences in outline shape. *Medical Image Analysis*, 1(3):225–243, 1997.
3. T. Cootes, C. Taylor, D. Cooper, and J. Graham. Active shape models - their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, Jan. 1995.
4. R. Davies, C. Twining, T. Cootes, J. Waterton, and C. Taylor. A minimum description length approach to statistical shape modeling. *IEEE Transactions on Medical Imaging*, 21(5):525–537, May 2002.
5. J. Duchon. Splines minimizing rotation-invariant semi-norms in Sobolev space. In *Constructive Theory of Functions of Several Variables, Lecture Notes in Mathematics, 571*, pages 85–100, 1977.
6. T. Heimann, I. Wolf, T. Williams, and H.-P. Meinzer. 3d active shape models using gradient descent optimization of description length. In *Information Processing for Medical Imaging Conference*, 2005.
7. D. Kendall, D. Barden, T. Carne, and H. Le. *Shape and Shape Theory*. John Wiley & Sons, LTD, 1999.
8. M. Leventon, E. Grimson, and O. Faugeras. Statistical shape influence in geodesic active contours. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 316–323, 2000.
9. T. Richardson and S. Wang. Nonrigid shape correspondence using landmark sliding, insertion and deletion. In *International Conference on Medical Image Computing and Computer Assisted Intervention*, pages II–435–442, 2005.
10. C. Small. *The Statistical Theory of Shape*. Springer-Verlag, 1996.
11. M. Styner, K. Rajamani, L.-P. Nolte, G. Zsemlye, G. Szekely, C. Taylor, and R. Davies. Evaluation of 3D correspondence methods for model building. In *Information Processing for Medical Imaging Conference*, 2003.
12. H. Thodberg. Minimum description length shape and appearance models. In *IPMI*, pages 51–62, 2003.
13. S. Wang, T. Kubota, and T. Richardson. Shape correspondence through landmark sliding. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages I–143–150, 2004.
14. J. Xie and P. Heng. Shape modeling using automatic landmarking. In *International Conference on Medical Image Computing and Computer Assisted Intervention*, pages II–709–716, 2005.