

A Fast 3D Correspondence Method for Statistical Shape Modeling

Pahal Dalal, Brent C. Munsell, Song Wang*, Jijun Tang and Kenton Oliver
Department of Computer Science and Engineering
University of South Carolina, Columbia, SC 29208, USA

Hiroaki Ninomiya, Xiangrong Zhou and Hiroshi Fujita
Department of Intelligent Image Information
Division of Regeneration and Advanced Medical Sciences
Graduate School of Medicine, Gifu University, Japan

Abstract

Accurately identifying corresponded landmarks from a population of shape instances is the major challenge in constructing statistical shape models. In this paper, we address this landmark-based shape-correspondence problem for 3D cases by developing a highly efficient landmark-sliding algorithm. This algorithm is able to quickly refine all the landmarks in a parallel fashion by sliding them on the 3D shape surfaces. We use 3D thin-plate splines to model the shape-correspondence error so that the proposed algorithm is invariant to affine transformations and more accurately reflects the nonrigid biological shape deformations between different shape instances. In addition, the proposed algorithm can handle both open- and closed-surface shape, while most of the current 3D shape-correspondence methods can only handle genus-0 closed surfaces. We conduct experiments on 3D hippocampus data and compare the performance of the proposed algorithm to the state-of-the-art MDL and SPHARM methods. We find that, while the proposed algorithm produces a shape correspondence with a better or comparable quality to the other two, it takes substantially less CPU time. We also apply the proposed algorithm to correspond 3D diaphragm data which have an open-surface shape.

1. Introduction

Statistical shape modeling provides an effective way to quantitatively describe various shape structures and their possible variations. It has been widely used in many important computer-vision applications, especially in medical-imaging. For example, in [2, 3], statistical shape models are constructed to accurately locate subtle differences of

corpus-callosum shapes between schizophrenia patients and normal controls. Certain shape structures are also related to other diseases by constructing and comparing the statistical shape models [6, 14, 22]. In [7, 25], statistical shape models are successfully used to guide image segmentation by detecting the structures with desirable shapes. While earlier research efforts focused on the 2D case, where each shape instance is a contour, current efforts primarily focus on 3D surface shape instances, because most anatomic structures bear a 3D shape [28]. This paper focuses on the 3D case and therefore, we formally define a shape instance to be a *shape surface* that may be closed or open.

To construct a statistical shape model from a collection of training shape instances, the main difficulty is the required step of (*landmark-based*) *shape correspondence*, which identifies a set of corresponded landmark points across the given training shape instances. This way, each shape instance is represented by a certain number of landmarks, which facilitate the construction of statistical shape models. Note that, as in most prior work [2, 3, 7], the identified corresponded landmarks may not be coincident with any anatomically significant point of a structure. Prior research [3, 10] reveals that a small error in shape correspondence may substantially influence the accuracy of the resultant statistical shape models. Shape correspondence is a well known challenging problem, especially in the 3D case, because of the non-linear shape description and non-rigid shape variation.

Many methods have been developed for shape correspondence in recent years [1, 3, 5, 8, 9, 10, 18, 19, 20, 21, 23, 26, 27, 29, 30, 32, 33, 34]. Among them, the *minimum description length method (MDL)* [10] and the *spherical harmonics descriptors method (SPHARM)* [4, 15] are treated by many researchers as two of the state-of-the-art methods that have been successfully used for 3D shape correspondence. Specifically, in MDL, the shape-

*Corresponding author. Email: songwang@cse.sc.edu.

correspondence error is measured by the required bit-length to transmit the resulting statistical shape model and all the training shape instances used to construct this model. Given the high complexity of this shape-correspondence error, random optimization algorithms, such as the genetic algorithm [10, 13, 24, 29, 30], are usually used to solve this problem. Recently, Heimann *et al.* [17] adapted this MDL-based shape-correspondence error and suggest a gradient-descent algorithm to find a local-optimal solution. In SPHARM, using the first order ellipsoid from the spherical harmonic coefficients, the spherical parametrization are aligned to establish the correspondence across 3D shape instances [4]. However, both MDL [17] and SPHARM can only handle genus-0 closed-surface shape since they need to first apply a conformal mapping to map each shape surface to a sphere. Furthermore, both methods are computationally expensive given the complexity of the conformal mapping.

In this paper, we develop a new *landmark-sliding* based algorithm to achieve a faster 3D correspondence for both closed- and open-surface shapes. Specifically, we first construct a rough initial correspondence between two shape surfaces by removing the possible translation, rotation and scaling transformations. We then slide these roughly corresponded landmarks on the respective shape surface to minimize a specified shape-correspondence error. In this paper, we use 3D thin-plate splines to model the shape-correspondence error so that the proposed algorithm is invariant to affine transformations and more accurately reflects the nonrigid biological shape deformations between different shape instances. We conduct experiments on 41 instances of closed-surface hippocampus shape and compare the performance of the proposed method to that of both MDL and SPHARM. We also apply the proposed method to correspond 18 instances of open-surface diaphragm shape.

2. Surface Representation and Problem Description

In this paper, we consider three different representations for each shape surface S : (a) *point cloud* S_P , extracted or labelled directly from medical images, (b) *triangle mesh* S_T , constructed from the point cloud S_P , and (c) *landmarks* S_L , identified in the shape correspondence. Among them, S_P is the original data form of the shape surface and consists of a dense set of *surface points*. S_T can be constructed from these surface points by a mesh-generation tool. In this paper, we use the COCONE package developed by Dey and Giesen [11] to generate the triangle mesh S_T . The finally identified landmarks S_L may not be contained in S_P , but they are located on (or very close to) the underlying continuous surface S modelled by S_P so that they can represent well this surface.

For simplicity, let's consider the case of corresponding

a *target* shape surface V to a *template* shape surface U . Furthermore, we assume that U_L , or the landmarks in the template U , are known and fixed: denote them as $\mathbf{u}_i = (u_{ix}, u_{iy}, u_{iz})$, $i = 1, 2, \dots, n$. The goal is then to identify n corresponded landmarks V_L , or $\mathbf{v}_i = (v_{ix}, v_{iy}, v_{iz})$, $i = 1, 2, \dots, n$, on the target surface V . If we have an algorithm to accomplish this two-instance shape correspondence, we can easily extend it to handle multiple-instance shape correspondence. Specifically, to correspond more than two shape instances, we can pick one of them as the template, construct the landmarks in this template, and then repeat the two-instance shape correspondence algorithm to correspond each of the remaining shape instances to this template.

The major problem is then to define a shape-correspondence error ϕ between U_L and V_L . We wish this error to well describe the underlying nonrigid deformation between two shape surfaces. Usually, it is desirable that this error is invariant to the affine transformations of rotation, translation, and scaling. In addition, we wish V_L to well represent the shape surface V instead of being aggregated in a small local area on the surface V . We denote this shape-representation error to be $R(V_L)$ and suggest a shape-correspondence error in the form of

$$\phi(U_L, V_L) = d(U_L, V_L) + \lambda \cdot R(V_L), \quad (1)$$

where $d(U_L, V_L)$ is a measure of the nonrigid shape deformation between U_L and V_L and $\lambda > 0$ is a balance factor.

2.1. Definition of $d(U_L, V_L)$

In this paper, we use the 3D thin-plate splines [2, 12, 31] to model the nonrigid deformation between the template and target shape surfaces and use the thin-plate bending energy as $d(U_L, V_L)$. In particular, the thin-plate spline finds a mapping $\mathbf{t} = (t_x, t_y, t_z)$ from U_L to V_L , i.e., $v_{ix} = t_x(\mathbf{u}_i)$, $v_{iy} = t_y(\mathbf{u}_i)$, and $v_{iz} = t_z(\mathbf{u}_i)$, $i = 1, 2, \dots, n$. The thin-plate bending energy, which measures the required energy to deform a volume to match these two sets of landmarks, is then characterized by

$$d(U_L, V_L) = \iiint_{-\infty}^{\infty} (L(t_x) + L(t_y) + L(t_z)) dx dy dz,$$

where $L(\cdot) = \left(\frac{\partial^2}{\partial x^2}\right)^2 + \left(\frac{\partial^2}{\partial y^2}\right)^2 + \left(\frac{\partial^2}{\partial z^2}\right)^2 + 2\left(\frac{\partial^2}{\partial x \partial y}\right)^2 + 2\left(\frac{\partial^2}{\partial y \partial z}\right)^2 + 2\left(\frac{\partial^2}{\partial z \partial x}\right)^2$. It can be further simplified to a quadratic form

$$d(U_L, V_L) = \mathbf{v}_x^T \mathbf{L} \mathbf{v}_x + \mathbf{v}_y^T \mathbf{L} \mathbf{v}_y + \mathbf{v}_z^T \mathbf{L} \mathbf{v}_z, \quad (2)$$

where \mathbf{v}_x , \mathbf{v}_y and \mathbf{v}_z are the columnized vectors that contain the x , y and z coordinates of the landmarks in V_L , respectively. \mathbf{L} is the $n \times n$ upper left submatrix of

$$\begin{bmatrix} \mathbf{K} & \mathbf{D} \\ \mathbf{D}^T & 0 \end{bmatrix}^{-1},$$

where the $n \times n$ matrix \mathbf{K} has element $k_{ij} = -\frac{1}{8\pi} \|\mathbf{u}_i - \mathbf{u}_j\|$ and the $n \times 4$ matrix $\mathbf{D} = [\mathbf{1}_{n \times 1}, \mathbf{u}_x, \mathbf{u}_y, \mathbf{u}_z]$ with $\mathbf{u}_x, \mathbf{u}_y$ and \mathbf{u}_z being the columnized vectors that contain the x, y and z coordinates of the landmarks in U_L , respectively. The thin-plate bending energy is invariant to any affine transformation, i.e., if the mapping \mathbf{t} is affine, the resulting bending energy is always zero.

2.2. Definition of $R(V_L)$

Ideally, we may need to check the fitting error between V_L and the continuous surface V underlying the point cloud V_P to define this representation error. However, this may lead to a very complex function that precludes efficient algorithms for optimization. In this paper, we adopt an alternative way: we first co-align the template U and the target V by removing their relative rotation, translation and scaling transformations and then require a spatial-distribution consistency between U_L and V_L . Here, we simply set

$$R(V_L) = \sum_{i=1}^n \|\mathbf{u}_i - \mathbf{v}_i\|^2$$

after removing the affine transformations between U and V . This can effectively prevent all the target landmarks V_L from aggregating in a small area on the surface V .

Note that, the thin-plate-spline based shape-deformation measure $d(U_L, V_L)$ may not be sufficient to prevent the aggregation of the target landmarks in the shape correspondence. In fact, from Eq. (2), we can see that, if all the coordinates of the target landmarks V_L are scaled to half of their original values, $d(U_L, V_L)$ will be reduced to one-fourth of its original value. Therefore, when all the target points are moved closer to each other, the value of $d(U_L, V_L)$ may get smaller. The introduction of the term $R(V_L)$ can effectively avoid this problem. In the following, we develop a shape-correspondence algorithm to identify target landmarks V_L that minimizes the shape-correspondence error (1).

3. Proposed Method

The proposed shape correspondence algorithm consists of three steps: (a) efficiently aligning U and V by removing translation, rotation, and scaling transformations, (b) constructing the template landmarks U_L and an initial estimate of V_L , and (c) refining V_L by a landmark-sliding algorithm that minimizes the shape-correspondence error (1).

3.1. Aligning the Target V and the Template U

This step is operated on U_T and V_T , the triangle mesh representations of the template and the target. We first remove the location and scaling differences between U_T and V_T by moving their centers of mass to the origin and normalizing their sizes to be the same.

We then remove the rotations between the template and the target by aligning their principal axes. Consider the template shape surface U_T as an example. Suppose the template triangle mesh U_T contains m triangle faces with areas W_1, W_2, \dots, W_m , respectively. For each triangle, we calculate its centroid as the average of its three vertices. Denote the centroid of these m triangles as $\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_m$, respectively. Denote $\mathbf{r}_i = \frac{W_i \mathbf{c}_i}{\sum_{j=1}^m W_j}$ and each $\mathbf{r}_i = (r_{ix}, r_{iy}, r_{iz})^T$ is a 3D vector. We can then create a 3×3 covariance matrix

$$\begin{bmatrix} \sum_{i=1}^m r_{ix}^2 & \sum_{i=1}^m r_{ix} r_{iy} & \sum_{i=1}^m r_{ix} r_{iz} \\ \sum_{i=1}^m r_{ix} r_{iy} & \sum_{i=1}^m r_{iy}^2 & \sum_{i=1}^m r_{iy} r_{iz} \\ \sum_{i=1}^m r_{ix} r_{iz} & \sum_{i=1}^m r_{iy} r_{iz} & \sum_{i=1}^m r_{iz}^2 \end{bmatrix}. \quad (3)$$

Applying an eigenvalue decomposition to this matrix yields three orthogonal principal vectors $\mathbf{e}_1, \mathbf{e}_2$, and \mathbf{e}_3 with three corresponding eigenvalues being sorted in a decreasing order. We then rotate the whole template shape surface so that \mathbf{e}_1 and \mathbf{e}_2 are aligned with x - and y - axes, respectively.

However, we must determine whether to align \mathbf{e}_1 and \mathbf{e}_2 to positive or negative directions of the x - and y -axes, so that they are consistent when processing both template and target surfaces. In this paper, we calculate $\sum_{i=1}^m W_i (\mathbf{e}_1 \cdot \mathbf{c}_i)$, where \cdot indicates the operation of dot product. If its value is positive, we align \mathbf{e}_1 to the positive direction of x -axis. Otherwise, we align \mathbf{e}_1 to the negative direction of x -axis. Similarly, we can use the same strategy to determine the unique rotation to align \mathbf{e}_2 to the y -axis. For the target shape surface, we perform the same eigenvalue-decomposition and axis rotations. This will make the template and target to have the same orientations. There are rare cases where such eigenvalue-decomposition and axis rotation strategy fail to work i.e. when any two eigenvalues of the covariance matrix are too close to each other and/or the value of $\sum_{i=1}^m W_i (\mathbf{e}_1 \cdot \mathbf{c}_i)$ (or $\sum_{i=1}^m W_i (\mathbf{e}_2 \cdot \mathbf{c}_i)$) is very close to zero. However, for most shapes with certain amount of complexity, we find that this strategy can effectively remove the rotation transformations between two shape instances.

In some applications, the target V and the template U may be *a priori* known to be aligned (or partially aligned) by having no translation, rotation, and/or scaling transformations. For example, the diaphragm shape data used in Section 4 have no rotation transformation among them since they are extracted from medical images that are known to be aligned. In such cases, we can turn off the corresponding aligning operations to save computational time.

3.2. Constructing U_L and an Initial Estimate of V_L

By selecting one shape instance as the template U , constructing the template landmarks U_L is independent of any other training shape instances. The basic consideration is

that U_L is sufficiently dense to represent the surface U and sufficiently sparse for the construction of a compact statistical shape model. In this paper, we uniformly divide the 3D space into equal-size cubic grid cells. For each cell containing some surface point in U_P , we pick the one closest to the center of the this cell as a landmark in U_L . This process is shown in Fig. 1, where the cells are shown in 2D for convenience. We can tune the size of the grid cell to control the number or the density of the template landmarks U_L .

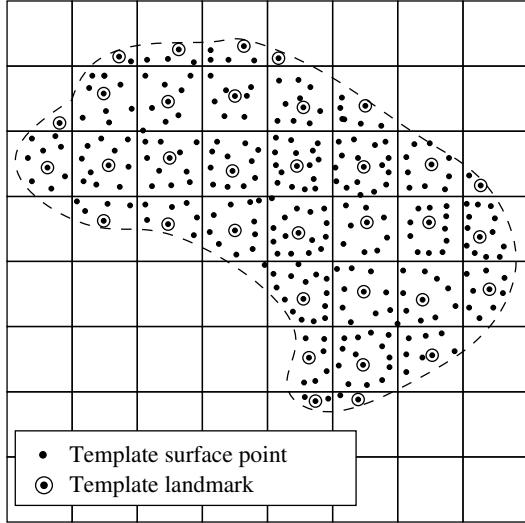


Figure 1. An illustration of constructing the template landmarks U_L from point clouds U_P . Note that, we show 2D grid cells here for illustration.

By removing the possible rotation, translation and scaling transformations between the template and target, we can construct an initial estimate of the target landmarks V_L . We construct V_L by finding, from the target point cloud V_P , the points with small distance to the template landmarks U_L . Specifically, for each template landmark \mathbf{u}_i , we find from all the surface points in V_P the one with the smallest Euclidean distance as the initial estimate of the target landmark \mathbf{v}_i , $i = 1, 2, \dots, n$, as shown in Fig. 2. To avoid the possible problem of finding the same target surface point for two different template landmarks, we exclude a target surface point from the search space if it has been included in V_L in previous searches. Since the template and the target shape surfaces are pre-aligned, this simple algorithm is able to find V_L that correspond roughly to U_L .

3.3. Refining the Target Landmarks by Iterative Landmark Sliding

With the template landmarks U_L , we can iteratively refine the initially estimated V_L by sliding V_L on the surface V . We achieve this goal by sliding landmarks \mathbf{v}_i on the tangent plane passing through \mathbf{v}_i . Each step of sliding aims at minimizing the shape-correspondence error (1). We

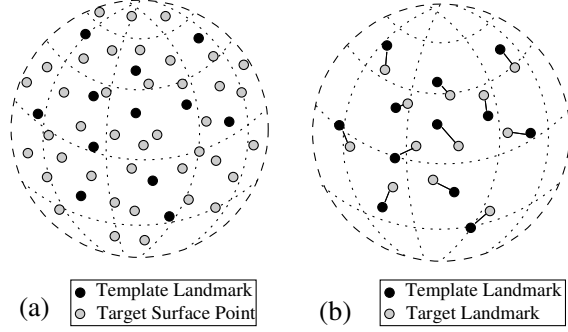


Figure 2. An illustration of finding an initial estimate of the target landmarks V_L .

estimate the tangent plane by fitting all the surface points in V_P that are located in a small neighboring area centered at \mathbf{v}_i . This tangent plane can be described by two linearly-independent, unit-length tangent vectors \mathbf{p}_i and \mathbf{q}_i in this plane. After a step of sliding, the target landmark \mathbf{v}_i will be moved to $\mathbf{v}_i + \alpha_i \mathbf{p}_i + \beta_i \mathbf{q}_i$, where α_i and β_i are the sliding distances along two tangent vectors. The problem is then to locate the optimal sliding distances α_i and β_i , $i = 1, 2, \dots, n$ so that after this sliding, the shape-correspondence error (1) will be minimized.

First, the shape-correspondence error (1) after a step of sliding can be written as

$$\begin{aligned} \phi(U_L, V_L; \boldsymbol{\alpha}, \boldsymbol{\beta}) = & \sum_{\star \in \{x, y, z\}} (\mathbf{v}_\star + \mathbf{P}_\star \boldsymbol{\alpha} + \mathbf{Q}_\star \boldsymbol{\beta})^T \mathbf{L} (\mathbf{v}_\star + \mathbf{P}_\star \boldsymbol{\alpha} + \mathbf{Q}_\star \boldsymbol{\beta}) \\ & + \lambda \sum_{i=1}^n \|\mathbf{v}_i + \alpha_i \mathbf{p}_i + \beta_i \mathbf{q}_i - \mathbf{u}_i\|^2, \end{aligned}$$

where

$$\begin{aligned} \boldsymbol{\alpha} &= (\alpha_1, \alpha_2, \dots, \alpha_n)^T \\ \boldsymbol{\beta} &= (\beta_1, \beta_2, \dots, \beta_n)^T \\ \mathbf{P}_\star &= \text{diag}(p_{1\star}, p_{2\star}, \dots, p_{n\star}) \\ \mathbf{Q}_\star &= \text{diag}(q_{1\star}, q_{2\star}, \dots, q_{n\star}), \star \in \{x, y, z\}, \end{aligned}$$

and $\mathbf{p}_i = (p_{ix}, p_{iy}, p_{iz})$ and $\mathbf{q}_i = (q_{ix}, q_{iy}, q_{iz})$.

Another consideration is that we should prevent the target landmarks V_L from moving away from the target surface V since we expect the resulting corresponded landmarks V_L to well represent the surface V . The best way to solve this problem is to restrict the maximum distance for each step of the sliding and to project the target landmarks after each step of sliding back to the surface V . For example, the projection can be achieved by finding, from V_P , the surface points with the smallest Euclidean distance to the target landmarks after the sliding. In practice, we find that it is usually sufficient by only setting the maximum distance for each step of the sliding to a small value ϵ . The main

reason is that, in addition to the target landmarks V_L , their association with the tangent planes and tangent vectors will also be updated after each step of the landmark sliding. This update involves a set of surface points in V_P around the target landmarks. Therefore, we simply impose the following constraints in every step of landmark sliding.

$$\begin{aligned} |\alpha_i| &\leq \epsilon \\ |\beta_i| &\leq \epsilon, i = 1, 2, \dots, n \end{aligned} \quad (4)$$

We can see that finding the sliding distances α and β to minimize the shape-correspondence error $\phi(U_L, V_L; \alpha, \beta)$, subject to constraints (4), is a quadratic programming problem that can be solved very efficiently. In our experiments, we use the OQP quadratic-programming solver developed by Gertz and Wright [16]. For the stop condition, we iteratively run this landmark-sliding algorithm until the change of the shape-correspondence error $\phi(U_L, V_L; \alpha, \beta)$ between two iterations becomes insignificant or a preset number of iterations is reached.

4. Experiments

We implemented the proposed method in C++ and used two 3D data sets for testing its performance. The first data set contains 41 instances of closed-surface hippocampus shapes and the second contains 18 instances of open-surface human diaphragm shapes. Example hippocampus and diaphragm shape instances are shown in Fig. 3. For the proposed method, we set $\lambda = 0.001$ in (1) and $\epsilon = 0.05$ in (4), given that all the shape instances are normalized as introduced in Section 3.1. Landmark sliding is set to stop after 20 iterations. The CPU time reported in this section is acquired from Linux workstations equipped with Intel Xeon 3.4GHz processors.




Figure 3. Examples of hippocampus (left) and diaphragm (right)

For performance evaluation, we construct a point distribution model (PDM) [7] that defines the mean shape and its associated covariance matrix from the correspondence result. Using the PDM, Styner *et al.* [28] suggest three quantitative measures: *compactness*, *generality* and *specificity* to evaluate the performance of shape correspondence. Particularly, compactness evaluates the amount of variance in the PDM. Generality uses a “leave one out” test to evaluate a PDM’s capability to describe unseen shape instances outside of the training set. Specificity evaluates a PDM’s capability to represent only valid shapes. According to [28],

the smaller these measures, the better the shape correspondence. Additionally, a qualitative method for evaluating shape correspondence is to vary the PDM mean shape by intentionally changing the projection lengths along the first several principal directions between $\pm 3\sigma$ standard deviations [6, 10], where σ^2 represents the eigenvalues. By varying the PDM we can visually check the resulting shape deformation from the mean shape to determine the validity of the PDM shape space. In this section, we adopt both these quantitative and qualitative methods for performance evaluation.

4.1. Experimental Results on the Hippocampus Data

We test the proposed method on 41 hippocampus instances. Each hippocampus shape instance contains 8,000 to 10,000 surface points. We compare the performance of the proposed method with the state-of-the-art MDL [10] and SPHARM methods [4, 15]. Specifically, we use an MDL implementation based on the gradient-descent algorithm [17]. The number of the identified landmarks on each shape instance can not be arbitrarily set in MDL and SPHARM. For both MDL and SPHARM, we select $n = 642$ landmarks for each shape instance in this experiment. To get a fair comparison, we tune the cell size (Section 3.2, Fig. 2) so that the proposed method also selects $n = 642$ landmarks in the template U_L .

The quantitative-evaluation results for all three methods are shown in Fig. 4, where “SLIDE” indicates the result of the proposed method. These results show that the proposed method outperforms both MDL and SPHARM in terms of compactness, specificity, and generality. In addition, the proposed method requires much less CPU time than MDL and SPHARM. Qualitative evaluation results are shown in Fig. 5. We can see that all three methods have similar shape deformation when changing the projection lengths along the first two principal directions. These results indicate that the proposed method produces a shape correspondence that is of a better or comparable quality to the state-of-the-art MDL and SPHARM methods, but takes much less CPU time.

4.2. Experimental Results on the Diaphragm Data

The human diaphragm is an open-surface tissue located beneath the lung and above the liver, with its edge affixed to the corresponding bone frame (rib, xiphoid, lumbar). Identifying the human diaphragm from torso CT images is a very important task in computer aided diagnosis [35]. We collect 18 diaphragm shape instances to test the open-surface shape correspondence. Each of the 18 diaphragm shape instances is defined by approximately 11,000 to 12,500 surface points. We vary the cell size as discussed in Section 3.2 to construct template U_L with




Figure 4. Quantitative evaluation of MDL, SPHARM and the proposed “SLIDE” method on the hippocampus data: (a) compactness measure, (b) generality measure, (c) specificity measure, and (d) total CPU time taken by each method for corresponding all 41 hippocampus shape instances. Note that, for the measures shown in (a), (b) and (c), the smaller the better.




Figure 5. Qualitative evaluation of MDL, SPHARM and the “SLIDE” method on the hippocampus data. “0” indicates the mean shape.

$n = 305$, $n = 600$ and $n = 910$ landmarks, respectively. For each of these three templates, the quantitative and qualitative evaluation of the proposed method is shown in Figs. 6 and 7, respectively. Since neither MDL nor SPHARM can correspond open-surface shapes, we do not include SPHARM and MDL comparison results for the open surface diaphragm shape. From the qualitative evaluation result shown in Fig. 7, we can see that the resulting shape deforms reasonably along its principal directions.

5. Conclusion

In this paper, we developed a new method for 3D landmark-based shape correspondence that is important for statistical shape modeling. By an initial estimate of the corresponded landmarks, this method can quickly refine all the landmarks in a parallel fashion by sliding all the landmarks on the shape surface. In this method, we applied 3D thin-plate bending energy to model the shape-correspondence error so that the proposed algorithm well reflects the nonrigid biological shape deformations between different shape in-




Figure 6. Quantitative evaluation of the proposed method on the diaphragm data: (a) compactness measure, (b) generality measure, (c) specificity measure, and (d) total CPU time for corresponding all 18 diaphragm shape instances with different number of template landmarks.

stances. Unlike many other available shape-correspondence methods, the proposed algorithm can handle both open- and closed-surface shapes. Both quantitative and qualitative evaluation on the 3D hippocampus data showed that the proposed method produces a shape correspondence that is of a better or comparable quality to the state-of-the-art MDL and SPHARM methods, but takes much less CPU time. We also applied the proposed method to correspond the open-surface diaphragm instances, with promising results.

Acknowledgements

We would like to thank Martin Styner from University of North Carolina at Chapel Hill for providing the 3D hippocampus data and the MDL and SPHARM packages. We thank Ipek Oguz for the help on using the MDL package.

This work was funded, in part, by NSF-EIA-0312861 and AFOSR FA9550-07-1-0250.

References

- [1] A. Baumberg and D. Hogg. Learning flexible models from image sequences. In *European Conference on Computer Vision*, pages 299–308, 1994.
- [2] 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.
- [3] F. Bookstein. Landmark methods for forms without landmarks: Morphometrics of group differences in outline shape. *Medical Image Analysis*, 1(3):225–243, 1997.
- [4] C. Brechbuhler, G. Gerig, and O. Kubler. Parametrization of closed surfaces for 3-D shape description. *CVGIP*, 61:154–170, 1995.
- [5] A. Brett and C. Taylor. A framework for automated landmark generation for automated 3D statistical model construction. In *Information Processing in Medical Imaging Conference*, pages 376–381, 1999.
- [6] T. Cootes, A. Hill, C. Taylor, and J. Haslam. The use of active shape models for locating structures in medical images. *Image and Vision Computing*, 12:355–366, 1994.
- [7] 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.
- [8] R. Davies, T. Cootes, and C. Taylor. A minimum description length approach to statistical shape modelling. In *Information Processing in Medical Imaging Conference*, pages 50–63, 2001.
- [9] R. Davies, T. Cootes, C. Twining, and C. Taylor. An information theoretic approach to statistical shape modelling. In *British Machine Vision Conference*, pages 3–11, 2001.
- [10] 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.
- [11] T. Dey and J. Giesen. Detecting undersampling in surface reconstruction. *Proc. 17th ACM Sympos. Comput. Geom.*, pages 257–263, 2001.
- [12] J. Duchon. Splines minimizing rotation-invariant seminorms in Sobolev space. In *Constructive Theory of Functions of Several Variables, Lecture Notes in Mathematics*, 571, pages 85–100, 1977.
- [13] A. Ericsson and K. Astrom. Minimizing the description length using steepest descent. In *British Machine Vision Conference*, pages 93–102, 2003.
- [14] B. Geiger. Three-dimensional modelling of human organs and its application to diagnosis and surgical planning. In *Technical Report: INRIA France*, 1993.
- [15] G. Gerig, M. Styner, D. Jones, D. Weinberger, and J. Lieberman. Shape analysis of brain ventricles using spharm. In *MMBIA*, pages 171–178, 2001.
- [16] M. Gertz and S. Wright. Object-oriented software for quadratic programming, 2001.
- [17] T. Heimann, I. Wolf, T. Williams, and H.-P. Meinzer. 3D active shape models using gradient descent optimization of description length. In *Information Processing in Medical Imaging Conference*, 2005.
- [18] A. Hill and C. Taylor. A method of non-rigid correspondence for automatic landmark identification. In *British Machine Vision Conference*, pages 323–332, 1996.
- [19] A. Hill and C. Taylor. A framework for automatic landmark identification using a new method of nonrigid corre-

spondence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(3):241–251, March 2000.

- [20] M. Jones and T. Poggio. Multidimensional morphable models. In *International Conference on Computer Vision*, pages 683–688, 1998.
- [21] M. Jones and T. Poggio. Multidimensional morphable models : A framework for representing and matching object classes. *International Journal on Computer Vision*, 2(29):107–131, 1998.
- [22] A. Kelemen, G. Szekely, and G. Gerig. Elastic model-based segmentation of 3-d neuroradiological data sets. *IEEE Transactions on Medical Imaging*, 18(10):828–839, 1999.
- [23] A. Kotcheff and C. Taylor. Automatic construction of eigen-shape models by genetic algorithm. In *Information Processing in Medical Imaging Conference*, pages 1–14, 1997.
- [24] F. D. la Torre. Automatic learning of appearance face models. In *Recognition Analysis and Tracking of Faces and Gestures in Realtime Systems*, pages 32–39, 2001.
- [25] 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.
- [26] A. Pitiot, H. Delingette, A. Toga, and P. Thompson. Learning object correspondences with the observed transport shape measures. In *Information Processing in Medical Imaging Conference*, pages 25–37, 2003.
- [27] C. Shelton. Morphable surface models. *International Journal on Computer Vision*, 1(38):75–91, 2000.
- [28] 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 in Medical Imaging Conference*, 2003.
- [29] H. Thodberg. Adding curvature to minimum description length shape models. In *British Machine Vision Conference*, volume 2, pages 251–260, 2003.
- [30] H. Thodberg. Minimum description length shape and appearance models. In *Information Processing in Medical Imaging Conference*, pages 51–62, 2003.
- [31] G. Wahba. *Spline Models for Observational Data*. Society Industrial and Applied Mathematics, 1990.
- [32] K. Walker, T. Cootes, and C. Taylor. Automatically building appearance models from image sequences using salient features. In *British Machine Vision Conference*, pages 463–562, 1999.
- [33] Y. Wang, B. Peterson, and L. Staib. Shape-based 3D surface correspondence using geodesics and local geometry. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages (II)644–651, 2000.
- [34] 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.
- [35] X. Zhou, H. Ninomiya, T. Hara, H. Fujita, R. Yokoyama, H. Chen, M. Kanematsu, and H. Hoshi. Automated identification of diaphragm in non-contrast torso ct images and its application to computer-aided diagnosis systems. *International Journal of Computer Assisted Radiology and Surgery*, 1:366–367, 2006.

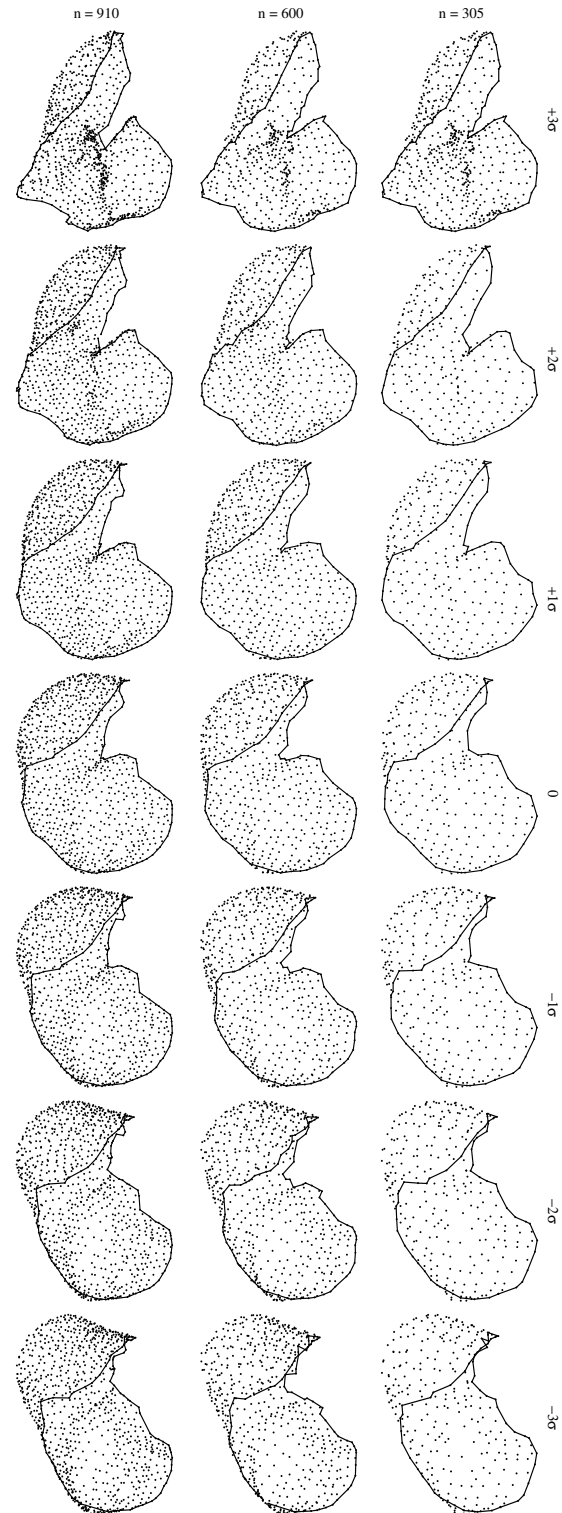


Figure 7. Qualitative evaluation of the proposed method on the diaphragm data. We rotate all the images by 90 degrees (clockwise) for saving space. The depicted contour is not a product of “SLIDE”, but is used for clarity of presentation.