Chapter 4
Landmark Sliding for 3D Shape Correspondence

Pahal Dalal
University of South Carolina, USA

Song Wang
University of South Carolina, USA

ABSTRACT

Shape correspondence, which aims at accurately identifying corresponding landmarks from a given population of shape instances, is a very challenging step in constructing a statistical shape model such as the Point Distribution Model. Many shape correspondence methods are primarily focused on closed-surface shape correspondence. The authors of this chapter discuss the 3D Landmark Sliding method of shape correspondence, which is able to identify accurately corresponding landmarks on 3D closed-surfaces and open-surfaces (Dalal 2007, 2009). In particular, they introduce a shape correspondence measure based on Thin-plate splines and the concept of explicit topology consistency on the identified landmarks to ensure that they form a simple, consistent triangle mesh to more accurately model the correspondence of the underlying continuous shape instances. The authors also discuss issues such as correspondence of boundary landmarks for open-surface shapes and different strategies to obtain an initial estimate of correspondence before performing landmark sliding.

BACKGROUND

The Point Distribution Model (Cootes, 1995) has become a very popular tool for statistical shape analysis and has been widely used in various computer-vision and medical-imaging applications such as image segmentation and shape based diagnosis. The major challenge in constructing a Point Distribution Model (PDM), especially in 3D, is the step of landmark-based shape correspondence. Shape correspondence aims at identifying a set of accurately corresponding landmarks from a population of given shape instances. The non-linearity of the shape description and shape

DOI: 10.4018/978-1-4666-1806-0.ch004

Figure 1. Illustration of each representation of a shape instance: (a) Point cloud S_P representing the surface and the surface boundary S_B. (b) Landmark-based triangle mesh S_T where each vertex is a landmark in S_L. (c) Discrete triangle mesh S_M to approximate the surface S. Note that S_M and S_T are not the same.

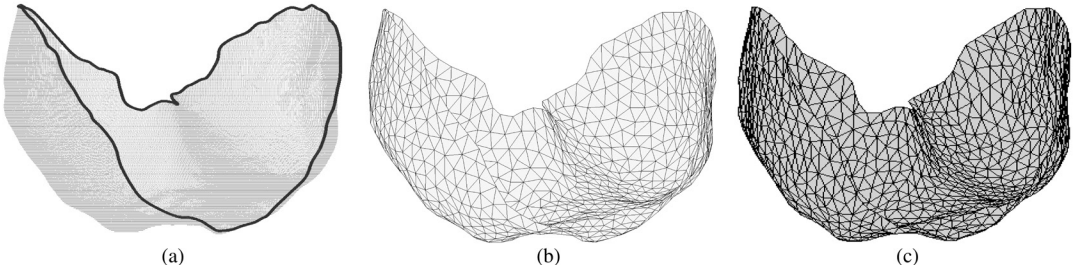

variation for most anatomical structures leads to a problem where it is very difficult to find an optimal solution.

Various 3D shape correspondence methods have been proposed for PDM construction. However, most of these methods are aimed at closed-surface shape correspondence. For example, both the *Minimum Description Length* (Davies, 2002; Heimann, 2005) and *Spherical Harmonics* (Brechbuhler, 1995; Gerig, 2001) methods map each shape instance to a sphere and reduce the shape correspondence problem to that of parameterizing the sphere. It is usually difficult to apply such a sphere-mapping step to open-surface shapes. Hence, we require a method that can perform 3D shape correspondence for both closed-surface and open-surface shapes.

PROBLEM FORMULATION

The aim of shape correspondence is to obtain a set of corresponding landmarks on a population of shape instances. As shown in Figure 1, we can represent each shape instance S as:

S_P, a dense point cloud defining the entire closed or open surface;
S_B, the subset of S_P that describe the closed boundary of the surface if S_P is an open surface (S_B is empty if S_P is a closed surface);
S_M, a triangle mesh constructed to approximate S;
S_L, the set of landmarks identified by the shape correspondence method;
S_T, the triangle mesh on S_L.

With correspondence across all shape instances, the triangle mesh S_T can be used to ensure topology consistency among all instances in the population. Different from S_L and S_T, S_M is constructed independently and may contain different number of vertices and triangles for each shape instance. Further, S_B is required to ensure that the landmarks along the boundary of an open-surface shape instance are correctly corresponded with points along the boundary of another instance in the population.

First we consider a simplified form of the shape correspondence where the population consists of only two instances: a *template* U and a *target* V. If we identify a set of landmarks U_L on the template, the problem then is to identify target landmarks V_L. Following this same method for a population of N instances, we can first select an instance as the template and identify template landmarks U_L. Then, we can construct corresponding landmarks on each target V_i in a pair-wise manner. In this way, corresponding landmarks can be identified on the entire population of shape instances. Note that these landmarks may not coincide with any anatomically significant locations or features.

LANDMARK SLIDING

Shape Correspondence Error

For shape correspondence, the major problem then is to define a shape correspondence error φ between U_L and V_L. This shape correspondence error must describe the underlying non-rigid deformation between the two shape instances represented by these landmarks. Further, the shape correspondence error must well-represent the entire target surface V and its properties. For example, we must ensure that landmarks along the boundary of U have corresponding landmarks along the boundary of V for open-surface shapes. To achieve these two goals, we can define the shape correspondence error as

$$\varphi(U_L, V_L) = d(U_L, V_L) + R(V_L)$$

where $d(U_L, V_L)$ quantifies the non-rigid shape deformation between U_L and V_L, and $R(V_L)$ is a regularization term to reflect the surface properties.

We use the 3D thin-plate splines to model the non-rigid deformation between the template and target 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 $v_{iz} = t_z(\mathbf{u}_i), i = 1, 2, \ldots, n$. The thin-plate bending energy, which measures the energy required to deform a volume to match these two sets of landmarks, is then characterized by

$$d(U_L, V_L) = \int\int\int_{-\infty}^{+\infty} (\mathbf{L}(t_x) + \mathbf{L}(t_y) + \mathbf{L}(t_z))dxdydz$$

where

$$\mathbf{L}(\cdot) = \left(\frac{\partial^2}{\partial x^2}\right)^2 + \left(\frac{\partial^2}{\partial y^2}\right)^2 + \left(\frac{\partial^2}{\partial z^2}\right)^2 + 2\left(\frac{\partial^2}{\partial x \partial y}\right)^2 + 2\left(\frac{\partial^2}{\partial y \partial z}\right)^2 + 2\left(\frac{\partial^2}{\partial z \partial x}\right)^2.$$

Expressed in a quadratic form, the bending energy can be defined as

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

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 sub-matrix of

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

where the $n \times n$ matrix $\mathbf{K}$ has element $k_{ij} = -\frac{1}{8\pi}\|\mathbf{u}_i - \mathbf{v}_i\|$, $n \times 4$ matrix $\mathbf{D} = \left[1_{n\times 1}, \mathbf{u}_x, \mathbf{u}_y, \mathbf{u}_z\right]$ 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.

For the regularization term $R(V_L)$ we discuss two strategies: (a) a strategy based on consistency in the distribution of landmarks between the template and target, and (b) a strategy for modeling the correspondence between boundaries of open-surface shapes. First, we note that the thin-plate spline based shape deformation measure $d(U_L, V_L)$ does not prevent the aggregation of target landmarks in a small region of the surface V. Since we wish to construct landmarks that

represent the entire surface, we will ensure that the selected template landmarks U_L are uniformly (or nearly uniformly) distributed over U. Certainly, target landmarks V_L aggregated within a small region on V do not truly reflect the correspondence in the underlying surfaces U and V. The major reason is that the numeric value of $d(U_L, V_L)$ is dependent on the size of V_L. If all the coordinates of V_L are scaled to half their original values, $d(U_L, V_L)$ will be reduced to one-fourth of its original value. Therefore, if all the target landmarks are aggregated within a small region of V the value of $d(U_L, V_L)$ may become smaller. One way is to construct an initial estimate of V_L by co-aligning the template and target after removing their translation, rotation, and scaling difference, and then we can impose a distribution consistency between U_L and V_L by setting

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

where $\lambda_i > 0$ can be used to balance the terms $d(U_L, V_L)$ and $R(V_L)$ by setting $\lambda_1 = \lambda_2 = ... = \lambda_n$. On the other hand, the regularization term may be used to reflect a specific characteristic of the surface. For example, we can use $R(V_L)$ to reflect that boundaries of open-surface shapes must correspond to each other. More specifically, we can first consider only the template boundary landmarks $U_{BL} \subseteq U_L$ and use a contour correspondence method to obtain the *expected* target boundary landmarks. Then we can define

$$R(V_L) = \sum_{i=1}^{n} \lambda_i \left\| \mathbf{v}_i - \mathbf{v}_i^b \right\|$$

where $\mathbf{v}_i^b$ is the *expected* position of $\mathbf{v}_i$ based on the correspondence of only the boundary contour and $\lambda_i > 0$ if $\mathbf{v}_i$ is a boundary landmark, $\lambda_i = 0$ otherwise. In this way, the regularization term can be applied to selected landmarks to better reflect the underlying characteristics of the shape.

Topology Consistency

Using the thin-plate spline bending energy as the shape correspondence error does not guarantee the preservation of topology consistency between the template and the target. For example, in Figure 2, the landmarks V_L shown in (b) have a finite bending energy with the template landmarks U_L in (a). However, when we connect the landmarks V_L by following the same vertex connectivity as the template triangle mesh, we obtain a non-simple triangle mesh V_T with self intersections. Landmark topology consistency reflects the geometric homeomorphism between the template and the target shape instances. Hence, landmark topology consistency should be included as a critical constraint in shape correspondence and statistical shape modeling. Including such a constraint explicitly into the shape correspondence error may not be tractable. Instead, we may algorithmically ensure that topology consistency is maintained at each step in the construction of template and target landmarks. Note that we cannot redefine the connection order in V_T to remove self intersections because this connection order is defined by the template mesh U_T. In our experiments, we notice that topology consistency is generally not violated for closed-surface shapes but can be a very serious problem in open-surface shapes.

Template Landmarks and Initial Estimate of Target Landmarks

We construct the template and target landmarks in three steps. First, we construct the template landmarks, U_L independently of all other target instances. Second, we construct an initial estimate of target landmarks V_L. Finally, we apply a refinement step to minimize the shape correspondence error $\varphi(U_L, V_L)$ through landmark sliding. The

Figure 2. Illustration of inconsistency in topology of template landmarks U_L and target landmarks V_L: (a) Triangulation of template landmarks $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \mathbf{u}_4, \mathbf{u}_5, \mathbf{u}_6$. (b) Triangulation of corresponding target landmarks $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4, \mathbf{v}_5, \mathbf{v}_6$ showing self-intersection.

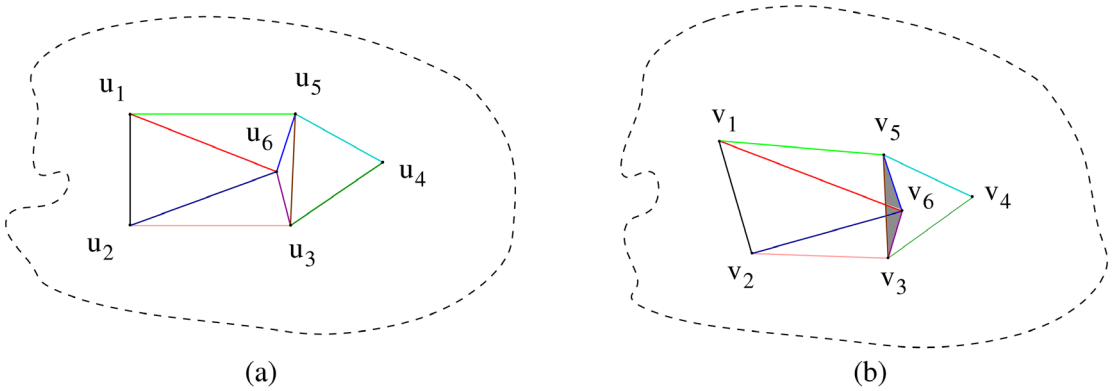

landmarks U_L must well-represent the entire template surface and hence be well-distributed on the surface U. We may impose a uniformly divided 3D grid on the template instance and select the point in U_P closest to the center of a cell as a template landmark. By tuning the size and number of cells in this 3D grid, we can change the number of template landmarks.

Another approach is to sample the points U_P using a method such as *k-mean clustering*. Using K-means clustering is certainly more flexible than the grid based method, since the number of landmarks can be directly specified. While the traditional K-means clustering method is suitable for closed-surface shapes, it may lead to the boundary of open-surface shapes not being represented sufficiently as shown in Figure 3(a). The major reason is that there may be no cluster centers along the boundary of the open-surface shape instance. To address this problem, we adapt the K-means clustering algorithm to a *constrained* K-means algorithm which identifies sufficient cluster centers from the surface boundary as shown in Figure 3(b). The approach is similar to the traditional K-means algorithm where clusters are iteratively identified and their centers are calculated by averaging the points in each cluster. For the constrained K-means clustering, we identify whether a cluster contains any points in the template surface boundary U_B and if so, we calculate the center of such a cluster as the average of only the points that are in U_B. In this way, we can obtain well-distributed template landmarks U_L which include sufficient landmarks representing the boundary of the template surface.

To obtain the initial estimate of target landmarks V_L, we propose two strategies: the first is a general method applicable to both closed-surface and open-surface shapes, the second is aimed towards a better initial estimate of V_L for open-surface shapes. In the first strategy, the basic idea is to remove the location, scaling and rotation differences between the template and target represented by their triangle meshes U_M and V_M. We remove the location and scaling differences between U_M and V_M 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 surface U_M as an example. Suppose the template triangle mesh U_M contains m triangle faces with areas $w_1, w_2, \ldots, w_m$, respectively. For each triangle, we calculate its centroid as the average of its three vertices. Denote the centroid of these m triangles as $\mathbf{c}_1, \mathbf{c}_2, \ldots, \mathbf{c}_m$ respectively. Denote $\mathbf{r}_i = \frac{w_i \mathbf{c}_i}{\sum_{j=1}^{m} w_j}$ and each

Figure 3. A simple 2D illustration of the constrained K-means clustering algorithm: (a) Traditional K-means clustering of S_P leads to well distributed cluster centers, but surface boundary points are not usually included as cluster centers. (b) Constrained K-means clustering with cluster centers along the boundary.

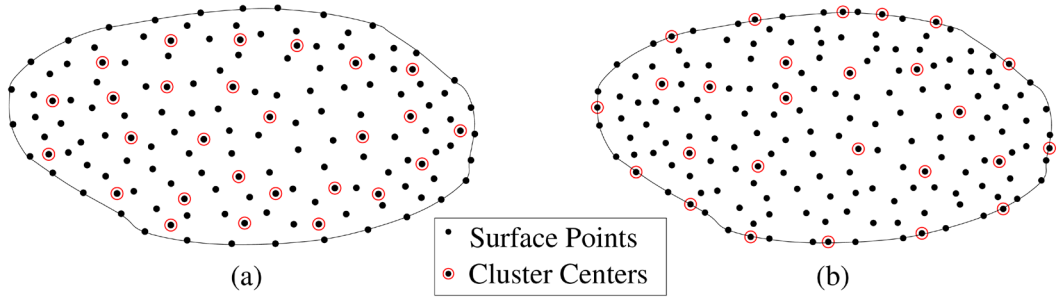

$\mathbf{r}_i = (r_{ix}, r_{iy}, r_{iz})^T$ is a 3D vector. We can then create a 3×3 covariance matrix

$$\begin{bmatrix} \sum_{i=1}^{m} r_{ix}^2 & \sum_{i=1}^{m} r_{ix} r_{iy} & \sum_{i=1}^{m} r_{ix} r_{iz} \\ \sum_{i=1}^{m} r_{ix} r_{iy} & \sum_{i=1}^{m} r_{iy}^2 & \sum_{i=1}^{m} r_{iy} r_{iz} \\ \sum_{i=1}^{m} r_{ix} r_{iz} & \sum_{i=1}^{m} r_{iy} r_{iz} & \sum_{i=1}^{m} r_{iz}^2 \end{bmatrix}$$

Applying 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 $\mathbf{e}_1$, $\mathbf{e}_2$ are aligned with the *x* and *y* axes, respectively. However, we must determine whether to align $\mathbf{e}_1$, $\mathbf{e}_2$ to positive or negative direction of the *x* and *y* axes, so that they are consistent when processing both template and target surfaces. In order to address this problem, we calculate the sum of dot products $\sum_{i=1}^{m} w_i(\mathbf{e}_1 \cdot \mathbf{c}_i)$. 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. However, 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.

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 smallest 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 $\mathbf{v}_i, i = 1, 2, \ldots, n$ as shown in Figure 4. 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

Figure 4. An illustration of finding an initial estimate of the target landmarks V_L by co-aligning the template and target surfaces

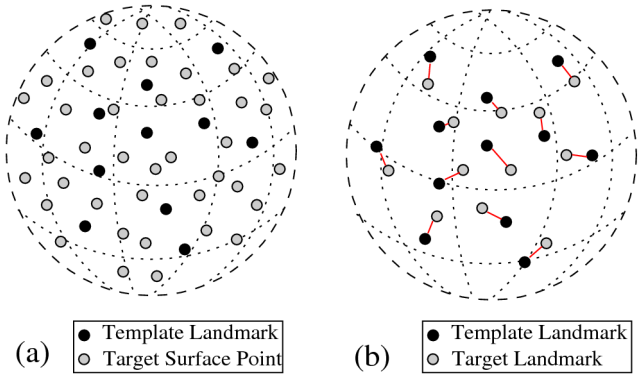

template and the target shape surfaces are pre-aligned, this algorithm is able to find V_L that correspond roughly to U_L.

In the second strategy, which is applied only to open-surface shapes, we use shape matching based on geometric information along with conformal mapping to construct an initial estimate for the target landmarks V_L. First, we flatten the template triangle mesh U_T and target triangle mesh V_M to their 2D conformal representations $\tilde{U}_T$ and $\tilde{V}_M$ respectively, as shown in Figure 5(a) and (b). With the 2D conformal maps, we can easily build a 2D parameterization for the 3D surface. Second, we find a set of corresponding landmarks on the 3D boundary of U and V using a contour correspondence algorithm (Wang, 2004). By projecting these landmarks to their respective conformal mappings, we obtain $\tilde{U}_{BL}$ and $\tilde{V}_{BL}$ which are the 2D, conformal projections of the template and target boundary landmarks $U_{BL} \subseteq U_L$ and $V_{BL} \subseteq V_L$. We use $\tilde{U}_{BL}$ and $\tilde{V}_{BL}$ to eliminate the translation, rotation and scaling differences between the template and the target using *Procrustes* analysis on the 2D plane. We then use the shape-context method (Belongie, 2002) to find a matching between all the template landmarks $\tilde{U}_L$ (on the 2D plane) and $\tilde{V}_D$, a down-sampled subset of points on the target in the 2D plane. We use the constrained K-means clustering method to preserve the boundary while constructing $\tilde{V}_D$. We only require that the number of points in $\tilde{V}_D$ be larger than or equal to the number of landmarks in $\tilde{U}_L$ so that each template landmark is assigned a corresponding target landmark. To enforce the requirement that template boundary landmarks $\tilde{U}_{BL}$ are corresponded to a point along $\tilde{V}_B$, we set the cost of matching boundary landmarks with non-boundary points to a very large number. After finding the matches of $\tilde{U}_L$ in $\tilde{V}_D$, we perform inverse conformal mapping to get a set of target landmarks on the 3D target surface. However, we cannot directly use this set of 3D landmarks on the target as the initial estimate of V_L. It may contain mismatches, because shape context only considers geometry information but not spatial topology, as shown in Figure 5(c). If we construct the triangle mesh V_T by connecting V_L using the same connection information as U_T, the mesh V_T may not be a simple mesh without self-intersections. We use statistical regression and develop a process of elimination to remove such pairs $(\mathbf{u}_i, \mathbf{v}_i)$. We then use the remaining matched pairs to establish a topologically consistent initial estimate for V_L.

Specifically, we define the regression cost function as

Figure 5. An illustration of conformal mapping and shape-context descriptor: (a) 3D mesh of the human diaphragm. (b) The conformal map of the 3D mesh. (c) The shape-context descriptor of a landmark $\mathbf{u}_i$.

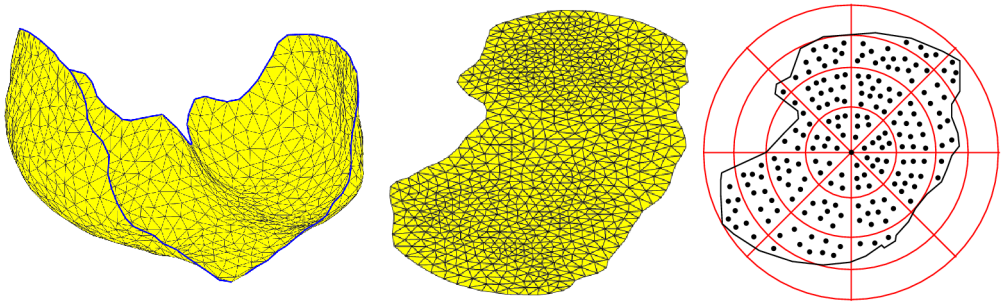

(a) (b) (c)

$$\psi = \sum_{i=1}^{n} \left\| \mathbf{v}_i - \hat{\mathbf{v}}_i \right\|^2 + v_x^T \mathbf{L} v_x + v_y^T \mathbf{L} v_y + v_z^T \mathbf{L} v_z$$

where $\mathbf{L}$ is the same thin-plate spline bending matrix as before. We calculate the optimal values of $\hat{\mathbf{v}}$ to minimize this cost by equating to 0 the partial derivative of ψ with respect to $\hat{\mathbf{v}}$. To identify the mismatched pairs $(\mathbf{u}_i, \mathbf{v}_i)$, we calculate the Euclidean distance between each $\mathbf{v}_i$ and $\hat{\mathbf{v}}_i$. Then, we sort the list of pairs $(\mathbf{u}_i, \mathbf{v}_i)$ by decreasing value of this Euclidean distance (worst matched pair to best) and discard the pair $(\mathbf{u}_i, \mathbf{v}_i)$ if it satisfies the following conditions:

The pair violates the topology consistency requirement. For example, in Figure 2 we discard the pair $(\mathbf{u}_6, \mathbf{v}_6)$.

They are not boundary landmarks. However, neighboring boundary landmarks may be switched to ensure topology consistency.

Worst 30% pairs of matches $(\mathbf{u}_i, \mathbf{v}_i)$ are considered as mismatches, but we do not remove all the landmarks on the same triangle in U_T.

By this process of elimination, we get U_R and V_R which are topologically consistent and representative landmarks. Finally, we calculate the 3D thin-plate transform between U_R and V_R, such that $V_R = \mathbf{t}(U_R)$ and apply this transform $\mathbf{t}$ to all landmarks in U_L to get the initial estimate for V_L i.e. $V_L = \mathbf{t}(U_L)$. We now connect the landmarks in V_L using the same triangulation information as U_T to get V_T.

Landmark Refinement

The initial estimate V_L described above is constructed using only the geometry information, whether by co-aligning the template and target or using shape-context based matching for open-surface shapes. While it preserves the landmark-topology consistency, it may not minimize the 3D shape correspondence error. Hence, we adopt an iterative refinement procedure for V_L such that the shape correspondence error is minimized. In this algorithm, all landmarks in V_L are simultaneously and iteratively moved on the surface V to minimize the correspondence error. In an iteration, we first move each landmark on its tangent plane and then project the new landmarks back onto the surface V. More specifically, each landmark $\mathbf{v}_i \in V_L$ is moved to $\mathbf{v}'_i = \mathbf{v}_i + \alpha_i \mathbf{p}_i + \beta_i \mathbf{q}_i$ on the tangent plane, where $\mathbf{p}_i$ and $\mathbf{q}_i$ are linearly independent unit tangent vectors and α_i and β_i are the sliding distances along these two tangent vectors. The optimal sliding distances can be found by solving the following quadratic-programming problem

$$\min_{\alpha,\beta} \varphi(U_L, V_L; \alpha, \beta)$$
$$= \sum_{*\in\{x,y,z\}} (v_* + \mathbf{P}_*\alpha + \mathbf{Q}_*\beta)^T \mathbf{L}(v_* + \mathbf{P}_*\alpha + \mathbf{Q}_*\beta)$$
$$+\sum_{i=1}^{n} \lambda_i \left\| \mathbf{v}_i + \alpha_i \mathbf{p}_i + \beta_i \mathbf{q}_i - \mathbf{v}_i^e \right\|^2$$

subject to constraints

$$|\alpha_i| \le \varepsilon, |\beta_i| \le \varepsilon$$

where,

$$\alpha = (\alpha_1, \alpha_2, ..., \alpha_n)^T, \beta = (\beta_1, \beta_2, ..., \beta_n)$$
$$\mathbf{P}_* = diag(\mathbf{p}_{1*}, \mathbf{p}_{2*}, ..., \mathbf{p}_{n*})$$
$$\mathbf{Q}_* = diag(\mathbf{q}_{1*}, \mathbf{q}_{2*}, ..., \mathbf{q}_{n*})$$

For closed-surface shapes, we may set $\mathbf{v}_i^e = \mathbf{u}_i$ where the regularization term $R(V_L)$ is used to ensure consistency in distribution of landmarks. For open-surface shapes we may set $\mathbf{v}_i^e = \mathbf{v}_i^b$ where $\mathbf{v}_i^b$ is the expected position of a boundary landmark based on its boundary contour correspondence. In the case of closed-surface shapes, we can then set $\lambda_1 = \lambda_2 = ... = \lambda_n$. On the other hand for open-surface shapes, we can set $\lambda_i > 0$ if $\mathbf{v}_i$ is a boundary landmark or $\lambda_i = 0$ otherwise. In this way, the second term of the cost function only considers the expected boundary landmarks since $\lambda_i = 0$ for all other landmarks in the case of open-surface shapes.

Usually ε may be set to a small value to avoid the landmarks from sliding too far from the surface. The major concern is that topology consistency may not be maintained for large values of ε. However, if we can ensure topology consistency, we may adopt an aggressive landmark sliding approach. First, we set ε to a large value and perform a step of landmark sliding. If the topology consistency is broken by the result of this step of sliding, we set $\varepsilon = \frac{\varepsilon}{2}$ and redo the step of landmark sliding. We continue this process until a value for ε is reached such that the landmark sliding can be performed without violating the topology consistency or the value of ε becomes too small to perform any landmark sliding. For projecting $\mathbf{v}'_i$ back onto the surface V, we find, from the input dense point cloud V_P, a point that has the smallest distance to $\mathbf{v}'_i$. The resulting closest point is then used as $\mathbf{v}_i$ for the next iteration of landmark sliding.

Generally, it is very difficult to directly check the topology consistency of landmarks in 3D. We may use the 2D parameterization rendered by the conformal mapping for open-surface shapes and a conformal mapping to a canonical sphere for closed-surface shapes. Specifically, we find from triangle mesh V_M, the closest point to $\mathbf{v}'_i$. This point denoted as $\mathbf{v}''_i$, may be located within a triangle or along the sides of a triangle. Since $\mathbf{v}''_i, i = 1, 2, ..., n$ are on the triangle mesh V_M, we can apply the conformal mapping to find their mapping $\tilde{\mathbf{v}}''_i$ (to flat plane for open-surface shapes and sphere for closed-surface shapes). If the triangle mesh built on the point-set $\tilde{\mathbf{v}}''_i, i = 1, 2, ..., n$ using the same connection order as in U_T has any self-intersections; we know the landmark topology is inconsistent with the template.

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 26 instances of open-surface human diaphragm shapes.

For the hippocampus data we set each $\lambda_i = 0.001, i = 1, 2, ..., n$ and $\varepsilon = 0.05$ toperform a version of landmark sliding referred to as SLIDE. Specifically, this method uses co-alignment for initial estimate of V_L and does not explicitly check

Figure 6. Quantitative evaluation of MDL, SPHARM and SLIDE methods 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.

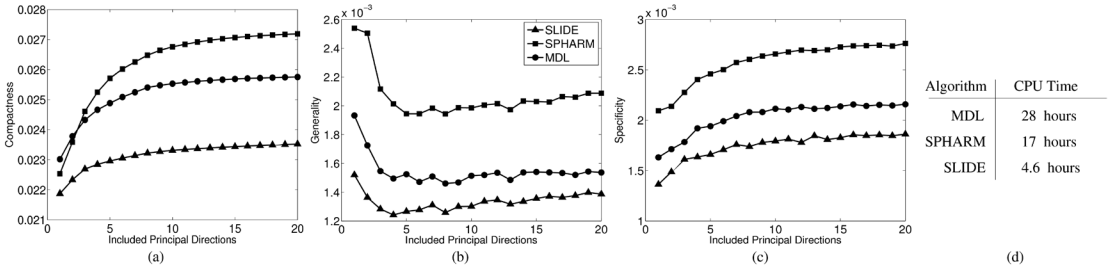

Algorithm	CPU Time
MDL	28 hours
SPHARM	17 hours
SLIDE	4.6 hours

for topology consistency since ε is set to a very small value. We identify the initial estimate of landmark correspondence by eliminating translation, rotation, and scaling differences between the template and target. After performing the landmark sliding to refine this initial estimate, we construct a point distribution model (PDM) and use the *compactness*, *generality*, and *specificity* measures (Styner, 2003) 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. Note that, the smaller these measures the better the shape correspondence. 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 MDL method using the implementation based on gradient-descent algorithm (Heimann, 2005) and SPHARM methods (Brechbuhler, 1995; Gerig, 2001). For a fair comparison we identify 642 corresponding landmarks on this population of hippocampus instances using each of the three methods. The resulting PDMs for each method are shown in Figure 7 by considering the first two principal directions, where SLIDE indicates the result of the proposed method. Further, the compactness, generality, and specificity measures are shown in Figure 6.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.

For the diaphragm data, we perform experiments on 26 instances, each described by approximately 100,000 to 250,000 points. Since both MDL and SPHARM are not able to correspond open-surface shapes, we compare the correspondence results of the landmark sliding method without any consideration for open-surfaces or enforcement of topology consistency (SLIDE), and a version that ensures boundaries are corresponded and topology consistency is maintained between template and target (SLIDE-T). In SLIDE-T, we set $\varepsilon = 50$ and $\lambda_i = 1$ for boundary landmarks and $\lambda_i = 0$ for all others. Using these settings, we identify 842 corresponding landmarks to construct and evaluate the point distribution model (PDM) derived from the correspondence results for SLIDE-T. To perform a fair comparison, we identify 842 landmarks on the same dataset using SLIDE. In Figure 8, we show a few significant instances from the deformable shape space described by the PDMs constructed using SLIDE and SLIDE-T. Particularly,

Figure 7. Qualitative evaluation of MDL, SPHARM and SLIDE method on the hippocampus data. "0" indicates the mean shape.

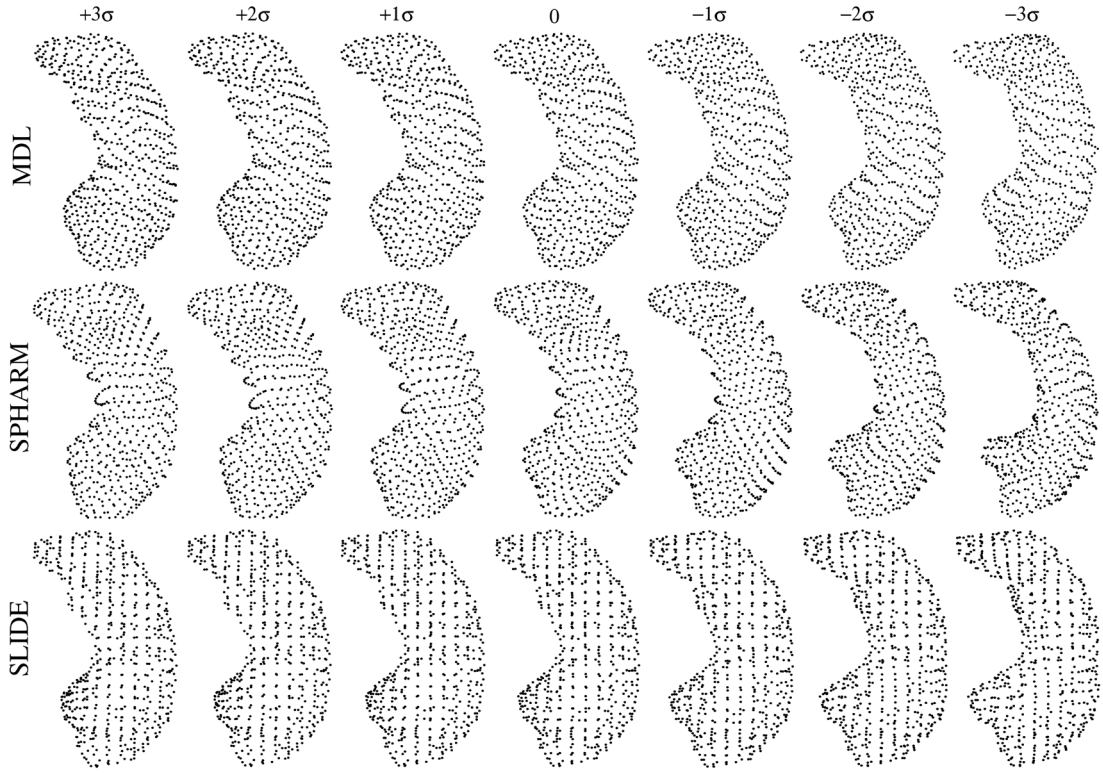

we show the shape instances after deforming the mean shapes along their first two principal directions, respectively. It is clear that the two PDMs describe similar shape spaces, but differ in two aspects. First, the shape space described by the PDM of SLIDE allows shape instances with self-intersecting surfaces and unnatural folding in some regions, while SLIDE-T always leads to a simple triangle mesh with no self intersections. Second, it can be observed that the boundary of the diaphragm shape is not well described by the PDM constructed using SLIDE and shows some folding, while SLIDE-T leads to a smooth boundary. We introduce two new measures to check whether the identified landmarks can well represent the topology and the boundary of each diaphragm shape instance. Without correct topology and good representation of the boundary, the identified landmarks cannot well represent the shape instance. First, we define the *topology error* as:

Δ_t = Number of triangles in V_T that are not topologically consistent with U_T.

This measure evaluates the ability of a method to correctly represent the topological homeomorphism of the underlying shape. Second, we define the *boundary error* as:

$$\Delta_b = \sum\nolimits_{\mathbf{v}_i \in V_{BL}} \left\| \mathbf{v}_i - \mathbf{v}^{(i)} \right\|$$

where $\mathbf{v}^{(i)}$ is the closest point along V_B to $\mathbf{v}_i$. This measure evaluates the ability of a shape correspondence method to represent the boundary of an open-surface shape. We apply these two new

Figure 8. Comparison of SLIDE and SLIDE-T: (First Row, M=1) Shape space of the PDM created using SLIDE (top) and SLIDE-T (bottom) for M=1. (Second Row, M=2) Shape space of the PDM created using SLIDE (top) and SLIDE-T (bottom) for M=2. M = 1, 2 are the two principal directions, μ is the mean shape and σ is the standard deviation along a principal direction.

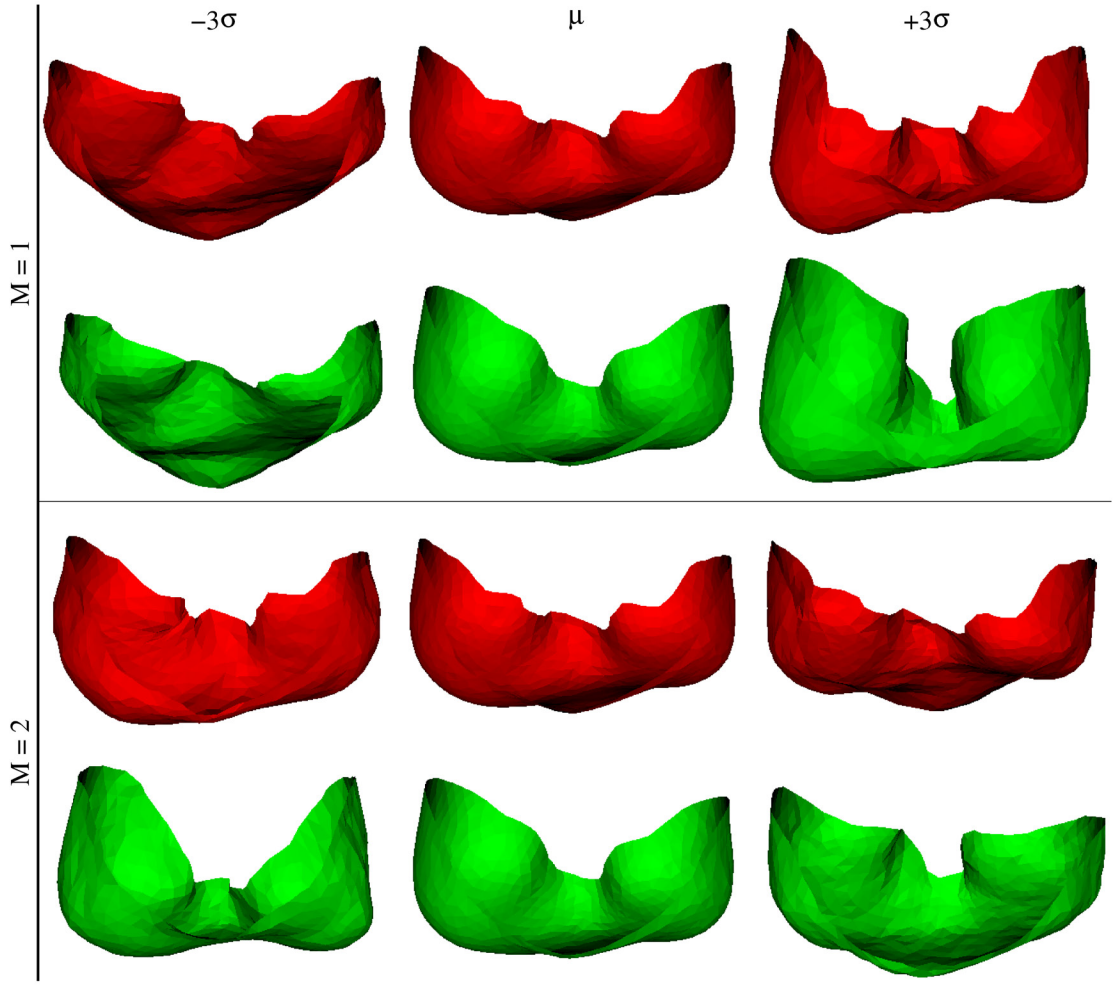

measures to all 26 shape instances and show them in Figure 9. It is clear that SLIDE-T performs better than SLIDE.

In addition to these new measures of shape correspondence evaluation, we find the compactness, generality and specificity measure for SLIDE and SLIDE-T. As shown in Figure 10, SLIDE leads to better performance on the compactness, generality and specificity measures compared to SLIDE-T. This is in contrast to what we obtained above. The major reason is that, without a reference to any ground truth, these three measures only evaluate the properties of the resulting PDM and do not check whether the resulting PDM can really represent the underlying shape. In other words, these three measures are only meaningful when the identified landmarks can well represent the original shape instances, both topologically and geometrically. More specifically, we can clearly see from Figure 11(a) that landmarks identified by SLIDE do not correctly model the whole diaphragm by missing the area near the

Figure 9. Comparison of SLIDE and SLIDE-T: Measures of topology error and boundary error. $\Delta_t = 0$ for all instances in SLIDE-T which is represented by the horizontal axis. The horizontal axis represents the 26 shape instances.

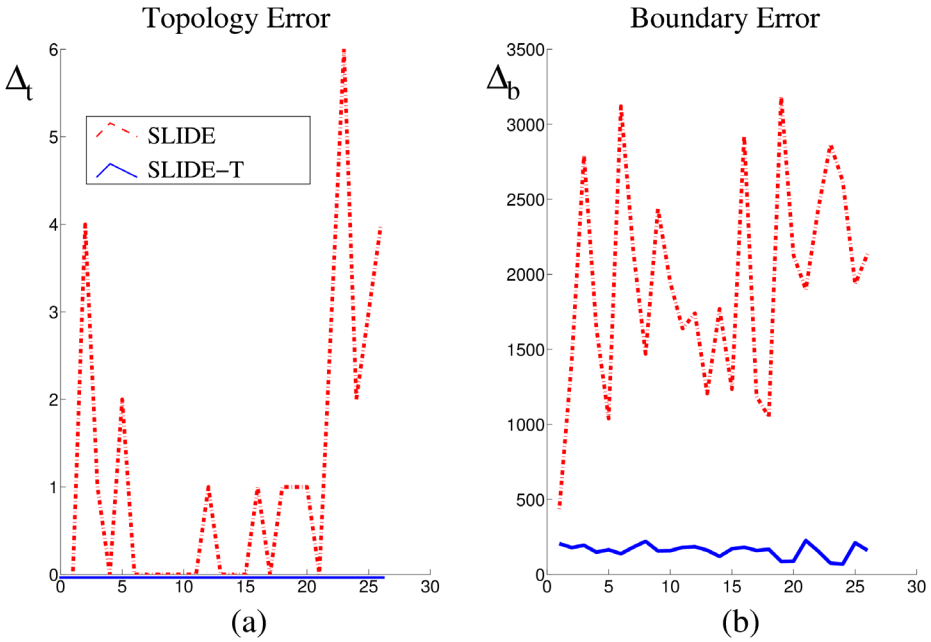

Figure 10. Compactness, generality, and specificity measures for SLIDE and SLIDE-T

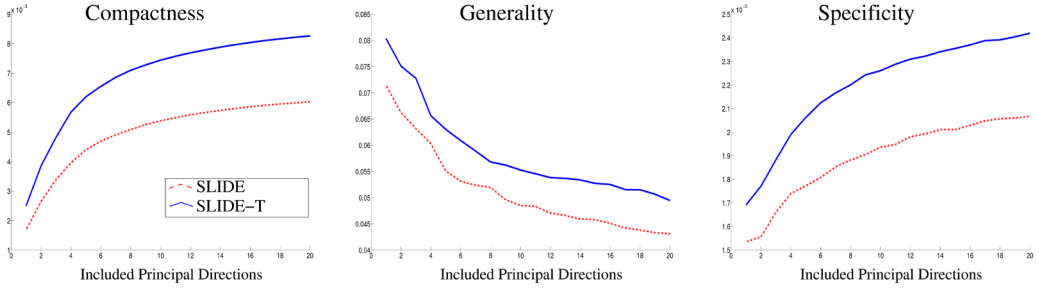

boundary; leading to a PDM with lesser variation and hence a better compactness measure. This also makes the size of the mean shape resulting from SLIDE smaller than that from SLIDE-T along the vertical direction, as shown in Figure 11(b), which gives SLIDE an undesirable advantage in terms of these three measures. As mentioned above, in this case, the correspondence results from SLIDE are in fact incorrect and the results from these three measures are misleading.

CONCLUSION

In this chapter, we have described a landmark sliding method to address the important problem of 3D shape correspondence for both closed-surface shapes and open-surface shapes. This method can identify corresponding landmarks on a population of shape instances with high accuracy such that they are topologically consistent. In this method, the 3D thin-plate spline bending energy is used to

Figure 11. (a) The PDM resulting from SLIDE is unable to model the shape boundary correctly. (b) The mean shape resulting from SLIDE (red) and SLIDE-T (green) have the same size in the horizontal and depth directions but not the vertical direction; giving SLIDE an advantage in the values of compactness, generality, and specificity.

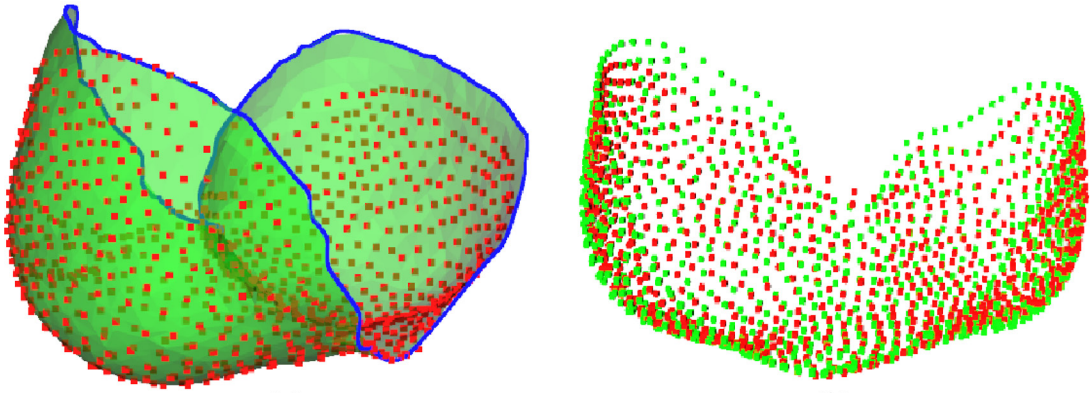

(a) (b)

model the landmark correspondence error and an efficient, iterative landmark refinement process is developed. For the closed-surface shapes, experiments reveal that the landmark sliding method leads to better shape correspondence results compared to MDL and SPHARM methods. Further, the introduction of topology consistency significantly improves the shape correspondence results of the landmark sliding method for open-surface shapes.

REFERENCES

Belongie, S., Malik, J., & Puzicha, J. (2002). Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *4*(24), 509–522. doi:10.1109/34.993558

Brechbuhler, C., Gerig, G., & Kubler, O. (1995). Parametrization of closed surfaces for 3-D shape description. *Computer Vision Graphics and Image Processing*, *61*, 154–170. doi:10.1006/cviu.1995.1013

Cootes, T., Taylor, C., Cooper, D., & Graham, J. (1995). Active shape models - Their training and application. *Computer Vision and Image Understanding*, *61*(1), 38–59. doi:10.1006/cviu.1995.1004

Dalal, P., Ju, L., McLaughlin, M., Zhou, X., Fujita, H., & Wang, S. (2009). 3D open-surface shape correspondence for statistical shape modeling: Identifying topologically consistent landmarks. In *IEEE International Conference on Computer Vision,* (pp. 1857-1864).

Dalal, P., Munsell, B., Wang, S., Tang, J., Oliver, K., & Ninomiya, H. … Fujita, H. (2007) A fast 3D correspondence method for statistical shape modeling. In *IEEE Conference on Computer Vision and Pattern Recognition*, (pp. 1-8).

Davies, R., Twining, C., Cootes, T., Waterton, J., & Taylor, C. (2002). A minimum description length approach to statistical shape modeling. *IEEE Transactions on Medical Imaging*, *21*(5), 525–537. doi:10.1109/TMI.2002.1009388

Gerig, G., Styner, M., Jones, D., Weinberger, D., & Lieberman, J. (2001) Shape analysis of brain ventricles using spharm. In *Mathematical Methods in Biomedical Image Analysis*, (pp. 171–178).

Heimann, T., Wolf, I., Williams, T., & Meinzer, H.-P. (2005). 3D active shape models using gradient descent optimization of description length. In *Information Processing in Medical Imaging, LNCS 3565*.

Styner, M., Rajamani, K., Nolte, L.-P., Zsemlye, G., Szekely, G., Taylor, C., & Davies, R. (2003). 3D correspondence methods for model building. In *Proceedings of Information Processing in Medical Imaging* (pp. 63–75). Evaluation of. doi:10.1007/978-3-540-45087-0_6

Wang, S., Kubota, T., & Richardson, T. (2004). Shape correspondence through landmark sliding. In *IEEE Conference on Computer Vision and Pattern Recognition*, (pp. 143–150).