Volumetric Segmentation via 3D Active Shape Models

Molly M. Dickens
Texas Tech University
molly.dickens@coe.ttu.edu

Shaun S. Gleason
Oak Ridge National Laboratory
gleasonss@ornl.gov

Hamed Sari-Sarraf
Texas Tech University
hamed.sari-sarraf@coe.ttu.edu

Abstract

A volumetric image segmentation algorithm has been developed and implemented by extending a 2D algorithm based on Active Shape Models. The new technique allows segmentation of 3D objects that are embedded within volumetric image data. The extension from 2D involved four components: landmarking, shape modeling, gray-level modeling, and segmentation. Algorithms and software tools have been implemented to allow a user to efficiently landmark a 3D object training set. Additional tools were built that subsequently generate models of 3D object shape and gray-level appearance based on this training data. An object segmentation strategy was implemented that optimizes these models to segment a previously unseen instance of the object. Results of this new 3D segmentation algorithm have been generated for a synthetic volumetric data set.

1. Introduction

The increasing quality and availability of imaging instruments such as magnetic resonance imaging and x-ray computed tomography (CT), has resulted in more prevalent use of these instruments for nondestructive industrial evaluation, patient diagnostics, and medical research. Employing these scanning instruments results in a tremendous amount of volumetric data. Analyzing these large data sets is both challenging and time consuming for the end users. Therefore, 3D data analysis tools are in increasing demand. Object segmentation is an important aspect of such data analysis.

Segmentation is one of the most challenging and valuable tasks in the field of computer vision and objects that display wide variation in shape or appearance pose a particular challenge. While progress in algorithm development for 2D image segmentation techniques has been considerable [1-3], the progress in successful 3D segmentation has been minimal by comparison.

Statistical shape modeling techniques, such as Active Shape Models [4] (ASM), rely on the statistics of an object's shape and gray-level appearance as gathered from a training set of manually landmarked instances of the object. Once the model is derived, image segmentation is performed by allowing the model to deform until it fits a previously unseen instance of the object. The deformation of the model occurs as it searches for nearby edges in the image but is limited to shapes that are consistent with those seen in the training set. ASMs are tolerant of faint or missing edges and have been successfully used to segment 2D objects that exhibit wide shape variation [5-7]. Because of the qualities of ASM, it has excellent potential as a method for segmenting 3D objects from image volumes. Therefore, the focus of this work is the extension of ASM to 3D, involving four components: landmarking, shape modeling, gray-level modeling, and segmentation.

1.1. Synthetic data generation

For the purposes of developing and testing the 3D ASM implementation, a synthetic data set was used to simulate simple shape variation but avoid the difficulties of real data such as noise and indistinguishable features. Since existing synthetic data is difficult to find in a voxelized form, a procedure was devised to generate voxelized volumes from a 3D model. First, a simple 3D model of a kidney was created in AutoCAD™. Next, this model was deformed by stretching the object by varying amounts in each of the 3 axes. A public domain software tool was obtained to convert the deformed AutoCAD™ models to 3D voxelized volumes. Using this procedure, a set of 19 voxelized 3D kidneys was generated (Figure 1).

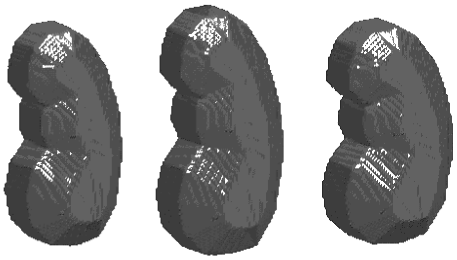

Figure 1. Samples of the synthetic kidney models

2. Generating a landmarked training set

To obtain the necessary statistics from the ASM training set, there must be point-to-point correspondence of the landmark points (LPs) that delineate the surface of the object under study. Although straightforward for 2D objects, the landmarking process is especially challenging for volumetric data because of the increased number of points and the intractable problem of having to directly landmark an object that is embedded within another image volume. Therefore, a very challenging and important extension task is the development of a volumetric landmarking technique. Previously, this problem has been addressed by methods that require an initial manual segmentation of the training set and then use automatic algorithms to orient those shapes and find corresponding point pairs [8, 9]. Our approach, first introduced in [10], is novel in that object orientation is established first, so that the subsequent segmentation directly produces corresponding LPs. This distributes the need for user input throughout the process to maximize the efficiency of the user's contribution rather than concentrating it on a tedious a priori manual segmentation. It also leaves the responsibility of establishing correspondence to the user, who is often better equipped than existing algorithms.

2.1. Object reorientation

Since 3D objects cannot be precisely oriented during imaging, especially when concealed within another body, this technique first establishes the actual orientation of the object within the volume based on features from the transaxial slices. To that end, a cylindrical coordinate system is located on each volume by presenting the user with appropriate image slices and annotation tools, as dictated by the data under study. This is demonstrated for the synthetic kidney data set in Figures 2 through 4. The user first marks a reference plane through the volume that will establish the reference angle in the cylindrical coordinate system. In this case, the line of symmetry of the kidney is marked on 3 slices (Figure 2) and the position of the reference plane is a best-fit through these lines (Figure 3). Note that the reference plane may be rotated in the transaxial plane as well as in the dimension orthogonal to it, as indicated by a translation of the lines from slice to slice. Next, the axis of the cylindrical coordinate system is marked on the reference plane and the volumes are resliced proportionally along their new coordinate axis (Figure 4) to form new volumes each made up of m corresponding and aligned image slices. Thus, the landmarking task has been reduced to m training sets of 2D images and tools to streamline 2D landmarking can be utilized.

2.2. Landmarking

To perform the 2D landmarking, a software tool was implemented to allow a user to interactively place and adjust LPs. This tool allows the user to efficiently and accurately landmark several members of each training set manually. A second software tool was implemented that builds a 2D ASM from these initial members and uses it to automatically segment the remaining training set images. During this process, the user may adjust the automatically placed points on any subsequently landmarked image and add it to the model, hence, improving its accuracy. A separate model can be built for each of the m sets of slices or similar slices can be grouped together, further streamlining the process.

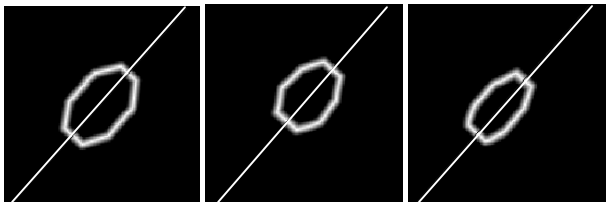

Figure 2. Three original transaxial slices; lines indicate the user's determination of the reference plane for each slice

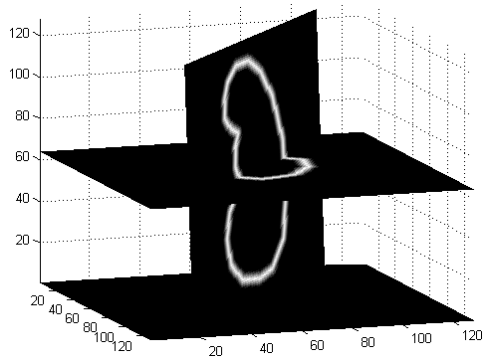

Figure 3. The reference plane shown in the 3D space with two of the transaxial slices for reference

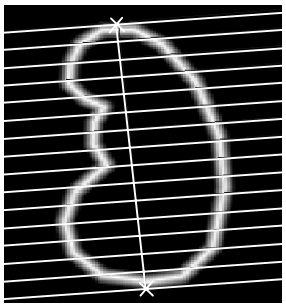

Figure 4. : Reference plane with 16 new slices equidistantly along, and perpendicular to, the reference axis

3. Extending ASM

The training process for 3D ASM requires the extension of shape model (SM) training and gray-level model (GLM) training. In addition to the training, the segmentation process must also be extended.

Note that in a 2D application, a shape is represented by a one-dimensional (1D) vector whose elements are the collection of x- and y-coordinates of the LPs, thus

$$\boldsymbol{x} = [x_1, x_2, \ldots, x_n, y_1, y_2, \ldots, y_n]^T \quad (1)$$

where n is the number of LPs. In the extension of ASM to 3D, the z-coordinates are simply appended to this same shape vector, thus

$$\boldsymbol{x} = [x_1, x_2, \ldots, x_n, y_1, y_2, \ldots, y_n, z_1, z_2, \ldots, z_n]^T. \quad (2)$$

3.1 Shape modeling

The first step in creating a 3D SM is to align all of the shapes gathered in the aforementioned training process. Alignment of any two shapes in the training set is accomplished by finding the scale and rotation that minimize the sum-of-squared distances between their corresponding LPs. In 2D, this involves only one rotation and there is a simple analytic solution for the gradient of the distance function. In 3D, alignment involves minimization with respect to scale and three different rotations (about the x-, y-, and z-axes). The function was minimized numerically in this case. Once the alignment is complete, the mean 3D shape can be calculated.

The principal component analysis (PCA) [4] is used in exactly the same way to reduce the shape vector dimensionality down to the t most significant modes of variation. The SM consists of the mean shape $\bar{\boldsymbol{x}}$, the PCA transformation matrix Φ, and the variances corresponding to each shape mode λ_i ($i = 1, 2, \ldots, t$). A shape vector $\boldsymbol{x}$ is transformed into the t-dimensional PCA space via

$$\boldsymbol{b} = \Phi^T (\boldsymbol{x} - \bar{\boldsymbol{x}}) \quad (3)$$

where the value of each b_i determines the variance of the shape in the i^{th} shape mode.

3.2 Gray-level modeling

The extension of GLM to 3D requires a new profile sampling strategy. In the 2D case, GLM for each LP is formed by extracting gray-level samples from all images in the training set along the profile that passes through each LP and is normal to the object boundary. This same idea has been extended for the 3D case in that the GLMs are formed by sampling all training *volumes* along the profile that is normal to the object *surface*. Once these profiles are extracted, the GLM for each LP consists of the objective function

$$f(\boldsymbol{g}_s) = (\boldsymbol{g}_s - \bar{\boldsymbol{g}}_d)^T S_d^{-1} (\boldsymbol{g}_s - \bar{\boldsymbol{g}}_d) \quad (4)$$

where $\bar{\mathbf{g}}_d$ and S_d are the mean profile and covariance matrix of all the profiles across the training set. The value of this function is the Mahalanobis distance between a candidate profile $\mathbf{g}_s$ and the mean profile $\bar{\mathbf{g}}_d$.

There were two significant new developments in the gray-level modeling process. Since the objects were landmarked after the reorientation process, the LPs had to be mapped back to the original 3D space before the gray-level profiles could be obtained. This mapping involved inverting the x, y, and z translations and the rotations in each of the three dimensions.

The second new development was that of estimating the normals to the surface. In 2D, the normal through an LP is defined to be perpendicular to a line joining the two neighboring points. In 3D however, the neighbors of an LP are not automatically known. The eight nearest neighbors for each LP were found and ordered so that a triangulation of the points provided an approximation to the surface of the object. The normal was taken as the mean of the normals to each triangle in the surface.

3.3 Segmentation

After SM and GLMs are formulated, they work together to guide the process of searching for the surface of a new 3D object within a volumetric data set. The role of SM during segmentation is to constrain the deformation of the model within the shape limits imposed by the training set. The role of GLM for each LP is to guide the search for a new position in the volume that most likely corresponds to the surface of the object of interest. As is the case with 2D ASM, the 3D shape vectors are represented as 1D vectors in PCA space and the gray-level model is based on a 1D profile; hence, the deformation and search scheme implemented here is very similar to that used in the 2D application.

To implement the model deformation and search process, software was written to perform the following functions:

- Given an initial location of 3D LPs within the new volume, update the position of all LPs by minimizing the objective function (Equation 4) generated during GLM training.
- Adjust the set of LPs to comply, in a statistical sense, with the SM generated during training by adjusting the shape parameters in the PCA space.

During segmentation, these two functions are called iteratively until the set of LPs (i.e. shape) settles into a consistent position within the volume. This process will be described in more detail as it pertains to the results presented in Section 4.3.

The development of visualization tools was crucial to validate the segmentation results. Note that visualization is not an issue for 2D images – the entire image and

segmentation can be viewed at once. However, in image volumes the segmentation must be evaluated relative to the image data by viewing 2D slices through the volume. This was made possible by the development of a software tool that allows the user to scroll through the image slices in each of the orthogonal directions while a boundary approximation based on the locations of the LPs is superimposed on each slice.

4. Results

The synthetic kidney volumetric image data described in Section 1.1 was used to develop and verify each component of the algorithm. Using the leave-one-out method, the completed algorithm was applied to the members of this training set. This section discusses the resulting 3D model and its ability to deform to fit the shape variations in this training set.

4.1. Landmarking the data

The 19-volume synthetic kidney training set was oriented and landmarked via the new procedure described in Section 2. The orientation and reslicing was performed, as shown in Figures 2 through 4, using 16 slices per volume. Landmarking was performed slice-by-slice using 2D ASM to help guide and accelerate the process. Generally, 4 to 5 manually landmarked instances were needed before the 2D model could sufficiently segment the subsequent slices. In two cases, a model derived from one set of slices was used to segment/landmark the next set of slices. A total of 224 LPs (4 to 18 per slice) were used to delineate the synthetic kidney in each image volume. Figure 5 shows the complete sets of 224 LPs for three of the volumes.

4.2. Building the models

The process of building the SM first involved aligning the 18 training shapes (where the 19th shape has been held out for testing) and calculating the mean shape. PCA is used to reduce the shape vector dimensionality down to the most significant modes of variation. For this data, the shape vectors were reduced from 672 dimensions down to 8. This number of shape modes was chosen to capture 98% of the total variance in the training set. During segmentation, the shape of the model will be constrained to fall within ±3 standard deviations of the mean for each shape mode. The mean shape is shown in Figure 6 (center) along with new shapes created by varying the first mode by ±3 standard deviations from the mean (right and left, respectively).

A GLM was formed for each of the 224 LPs by sampling all 18 training volumes along a normal profile through the LP. Twelve samples were taken in each direction (at one pixel intervals) forming profiles of 25 intensity values.

4.3. Segmenting new objects

The mean shape establishes the initial locations of the LPs in the iterative ASM segmentation process. The first step in each iteration is to fit the GLM. For this process, testing profiles, sampled with 10 additional values in each direction, were taken from the test volume. By shifting the training profile along the longer testing profile, the Mahalanobis distance could be calculated at 21 positions. Thus, fitting the GLM involved moving the LP to the position along the testing profile for which the objective function (Equation 4) was minimized. Figure 7 shows an example of one of the LPs in the first iteration being adjusted to fit its GLM.

Completing an iteration of the segmentation process, the new candidate LP locations are then sent to the SM where they are adjusted to constrain the shape. This is done by transforming the candidate shape into the SM's PCA space and adjusting any of the 8 mode parameters that fall outside of the allowable range.

Figure 8 shows the progress of the segmentation process on three orthogonal image planes through the test volume. The intersection of each plane and a surface formed by the current LPs is shown by a white boundary. Since the surface is formed by a simple triangulation of the LPs, this boundary is just an approximation to the actual boundary of the object between LPs. Therefore, only the LPs (shown by white dots) accurately represent the location of the model. The 1st row of Figure 8 shows views of the test volume with the initial location of the model superimposed. Note that the volume has been altered by adding noise and translating it 5 pixels in each dimension to further test the performance of the model. The 2nd and 3rd rows of Figure 8 show the results of the 1st and 5th iterations, respectively. The LPs along the outer boundary of the object on these image planes indicate that the model has successfully deformed to segment the synthetic kidney in this volume.

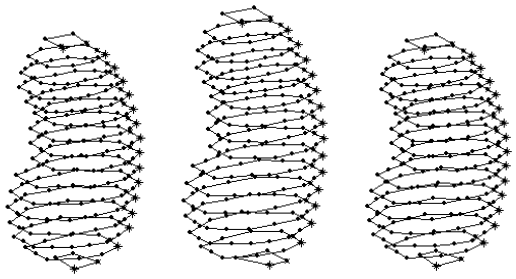

Figure 5. Complete LP sets for the objects in Figure 1

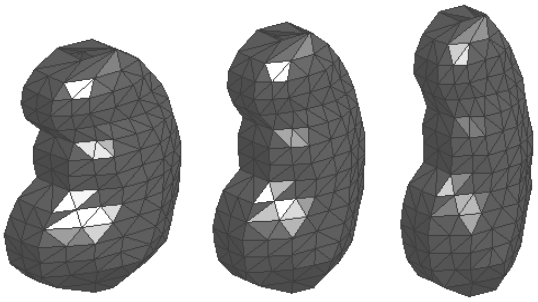

Figure 6. Variation of the first shape mode; (from left to right) - 3σ, mean shape, + 3σ

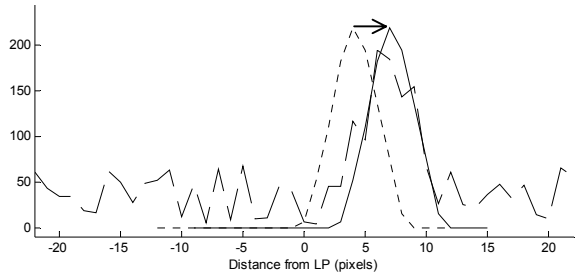

Figure 7. Example of GLM fit: test profile (dashed line), training profile (dotted line), and training profile shifted to indicate new location of LP (solid line)

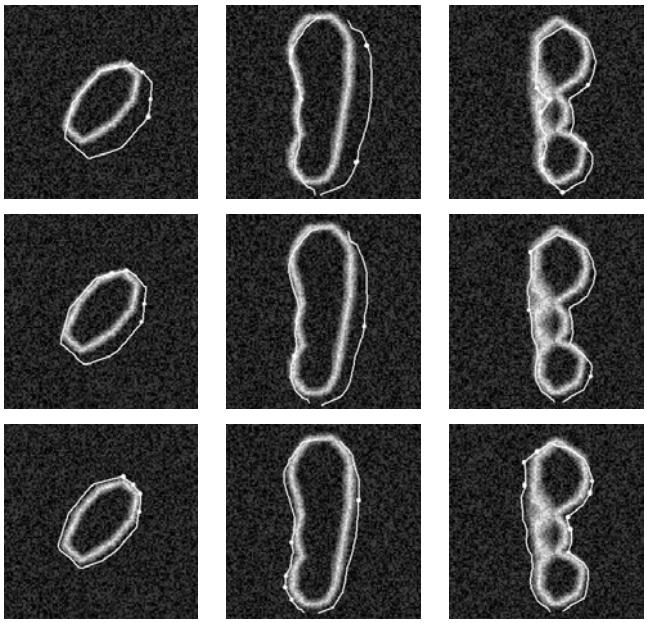

Figure 8. Results of the segmentation process: model initial position (1st row), result of first iteration (2nd row), result of 5th iteration (3rd row)

5. Conclusions

A volumetric segmentation algorithm has been developed and implemented by extending the 2D ASM methodology. A unique approach was taken to generate the necessary training sets and several new developments were required in the modeling and segmentation processes. For the computer-generated synthetic image volumes, the new algorithm successfully segments the object of interest. This paper is one of the first to present final segmentation results of a 3D ASM implementation.

Current work is focused on applying this algorithm to real image volumes. A set of CT image volumes of laboratory mice has been obtained and the process of training a 3D kidney model for that data is underway. There are two notable limitations in the current algorithm that will be addressed in future research. First, the landmarking scheme can only handle a limited class of simple shapes. A solution under consideration is to apply the scheme to individual components of more complex shapes. Second, if the model is not well initialized, it will diverge. As in the 2D case, a multi-resolution implementation is needed to prevent this by allowing the entire model to travel farther at course resolutions and make smaller adjustments at fine resolutions.

6. References

[1] T. McInerney and D. Terzopoulos, "Medical image segmentation using topologically adaptable surfaces," *Proc. CVRMed '97*, pp. 1-10, March 1997.

[2] M. Atkins and B. Mackiewich, "Fully automatic segmentation of the brain in MRI," *IEEE Trans. Med. Imag.*, vol. 17, n. 1, pp. 98-107, Feb. 1998.

[3] F. Lefebvre, G. Berger, and P. Laugier, "Automatic detection of the boundary of the calcaneus from ultrasound parametric images using an active contour model; clinical assessment," *IEEE Trans. Med. Imag.*, vol. 17, n. 1, pp. 45-52, Feb. 1998.

[4] T. Cootes, C. Taylor, D. Cooper, and J. Graham, "Active shape models - their training and application," *Comp. Vis. Image Understanding*, vol. 61, n. 1, pp. 38-59, Jan. 1995.

[5] N. Duta and M. Sonka, "Segmentation and interpretation of MR brain images: An improved active shape model," *IEEE Trans. Med. Imag.*, vol. 17, no. 6, pp. 1049-1062, Dec. 1998.

[6] Y. Wang and L. Staib, "Boundary finding with prior shape and smoothness models," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 22, no. 7, pp. 738-743, July 2000.

[7] P. Smyth, C. Taylor, and J. Adams, "Automatic measurement of vertebral shape using active shape models," *Proc. BMVC96*, pp. 176-180, 1996.

[8] Y. Wang, B.S. Peterson, and L.H. Staib, "Shape-based 3D surface correspondence using geodesics and local geometry," *Comp. Vis. Patt. Recog.*, vol. II, pp. 644-651, June 2000.

[9] C. Lorenz and N. Krahnstover, "Generation of point-based 3D statistical shape models for anatomical objects," *Comp. Vis. Image Understanding*, vol. 77, pp. 175-191, 2000.

[10] M. Dickens, H. Sari-Sarraf, and S. Gleason, "A streamlined volumetric landmark placement method for building three-dimensional active shape models," *Proc. SPIE Med. Imag.*, vol. 4322, pp. 269-280, 2001.