A Novel Application of Principal Surfaces to Segmentation in 4D-CT for Radiation Treatment Planning

Sheng You, Esra Ataer-Cansizoglu, Deniz Erdogmus
Cognitive Systems Laboratory
Northeastern University
Boston, MA
Email: erdogmus@ece.neu.edu

James Tanyi, Jayashree Kalpathy-Cramer
Oregon Health & Science University
Portland, OR
Email: kalpathy@ohsu.edu

***Abstract*—Radiation therapy is one of the most effective options used in the treatment of about half of all people with cancer. A critical goal in radiation therapy is to deliver optimal radiation doses to the observed tumor while sparing the surrounding healthy tissues. Radiation oncologists typically manually delineate normal and diseased structures on three-dimensional computed tomography (3D-CT) scans. Manual delineation is a labor intensive, tedious and time-consuming task. In recent years, concerns about respiration induced motion have led to the popularity of four-dimensional computed tomography (4D-CT) for the tracking of tumors and deformation of organs. However, as manually contouring in all phases would be prohibitively expensive, the development of fast, robust, and automatic segmentation tools has been an active area of research in 4D radiotherapy.**

In this paper, we describe a novel application of principal surfaces for the propagation of contours in 4D-CT studies. Regions of interest (ROIs) are manually delineated slice-by-slice in the reference 3D-CT scans. Edges are detected on all of the slices of the target 3D-CT phase. A kernel density estimation (KDE) based on the detected edges is then calculated. The principal surface algorithm is applied to find the ridges of the edge KDE to provide the object contours. Manually drawn contours from the reference phase are used as an initialization. Contours of ROIs are propagated recursively in all consecutive phases to complete a respiration cycle. Results are provided for a phantom data set of simulated tumor motion as well as on a de-identified data set of the lung of a patient. Evaluation of the efficacy of automatic segmentation in organs and tumors are based on the comparison between manually drawn contours and automatically delineated contours. The Dice coefficients are approximately 0.97 for the lung tumor on the phantom data sets and 0.95 for the patient data sets. The centroid distances between manually delineated lung volume and automatically segmented lung volume in each CT direction are $< 1\ mm$ for both phantom data sets and patient data sets.

***Keywords*-radiation therapy; contour propagation; segmentation; principal surface;**

I. Introduction

Technological advances in radiation oncology have instigated to significant changes in the delivery of radiotherapy over the last few decades. These changes have led to the replacement of conventional radiotherapy that utilizes simple rectangular treatment apertures with increasingly conformal radiotherapy techniques, such as three-dimensional conformal radiation therapy (3D-CRT) and intensity-modulated radiation therapy (IMRT). Conformal radiotherapy techniques facilitate the delivery of tumoricidal doses of radiation to tumor-bearing tissues while simultaneously minimizing the dose to adjacent organs-at-risk (OARs), hence improving the therapeutic ratio. The radiotherapy dose is typically prescribed to a planning target volume (PTV). This volume includes a gross tumor volume (GTV) plus areas of microscopic spread (clinical target volume, CTV) and a margin around it that accounts for the systematic and random errors plus physiological organ changes that occur during the treatment planning and delivery process [1, 2]. Using a conformal radiotherapy technique to deliver a radically high dose to the PTV could result in a significant dose to the OARs. However, there is robust evidence from several tumor sites indicating that employing dose-escalation and/or altered fractionation treatments results in improved outcomes [3]-[6]. Reducing the dose to the OARs using techniques such as IMRT and reducing the size of the PTV using image guided radiotherapy (IGRT) enables radiation dose-escalation to be effective.

However, creating treatment plans with steep dose gradients (as is typical of IMRT) and decreasing target volume margins raises concern about the effect of inter- and/or intrafraction motion on target coverage and the importance of properly defining the PTV. Respiratory-correlated or 4D computed tomography is increasingly being used to assess motion in the thoracic and abdominal regions. By including respiration-induced motion into radiotherapy, 4D-CT ensures improvements in the accuracy of target definition and localization, ensures dose coverage of the target throughout the breathing cycle and calculation of the dose distributions for targets and OARs, and reduces the occurrence of breathing artifacts. In order to incorporate 4D-CT information into 4D radiotherapy, delineation of tumors and organs in these data sets becomes critically important. As manually contouring on all phases would be prohibitively time/labor-intensive, delineating 4D ROIs manually is a challenging task in radiotherapy. The ability to accurately propagate 3D surface information would be useful to reduce the clinical workflow.

Deformable registration-based techniques are widely used to propagate contours from a reference phase to other remaining phases of a 4D-CT data set. Deformable registration provides displacement maps between the reference phase image and the target phase image. 3D contours on a reference phase are then deformed to the target phase according to the displacement maps to generate contours on that phase [7]. The procedure for obtaining displacement maps is based on the entire voxel-voxel matching between the reference phase and the target phase. Propagating object boundaries from the reference phase to other phases of 4D-CT data sets without performing full blown 3D-3D deformable registration is expected to be more computationally efficient and less time consuming.

In this paper, we demonstrate a novel application of the KDE-based principal surface algorithm for volumetric segmentation and contour propagation of tumors or organs. Rather than implementing deformable registration, our method incorporates the user inputs from the reference phase as an initialization and edge information from all the target phases. The initialized 3D contours are propagated consecutively and recursively through all the phases by applying the principal surface algorithm to converge to the ridges of the edge maps. Automatically propagated contours after a breathing cycle are quantitatively validated through a comparison with manual contours. The principal surface algorithm had been applied for image segmentation before [8,9]. Kalpathy-Cramer et al. presented a semi-supervised protocol for segmentation of tumors and non-affected tissues by using non-parametric snakes with the modified principal curves to propagate boundaries between CT slices. A given two-dimensional slice for 3D-CT scans was segmented first. The final segmentation on that slice was then used as a reference contour for segmenting neighboring slices. This process was repeated for all the 2D slices to form a continuous 3D surface model. In this paper, the principal surface algorithm is applied for directly segmenting the objects in three dimensions thereby permitting 4D contour propagation and analysis. Direct 3D segmentation without the need to optimize parameters makes this approach faster and less computationally complex.

II. Methodology

The overall algorithm starts with preprocessing of images. Edge maps are then detected on all the slices of the target 3D-CT phase. Then, a KDE based on the detected edges is calculated. Based on the edge maps, the principal surface algorithm is applied to find the ridges of the edge KDE to provide the object contours on the target phase using manually drawn contours from the reference phase as an initialization. Finally, some postprocessing methods are used to optimize the output of the principal surface algorithm.

A. Edge Detection

Let I be the 3D-CT such that $\mathbf{p} = [x, y, z]^T$ is the location of a voxel, and $I(x, y, z)$ is the intensity at that location. Let $\mathbf{E}(\mathbf{p})$ be the edge image. The edge maps can be obtained using a suitable edge detector, and the edge values can be binary or continuous [10].

1) Binary edge voxels obtained using the Canny edge detector with appropriate thresholds can be used in this situation.

$$\mathbf{E}(\mathbf{p}) = \begin{cases} \mathbf{E}(\mathbf{p}_i) = 1; & \mathbf{p}_i \text{ is an edge voxel} \\ \mathbf{E}(\mathbf{p}_i) = 0; & \text{otherwise} \end{cases}$$

A KDE of the edge distribution is constructed using the edge image $\mathbf{E}(\mathbf{p})$ via a summation of the kernel functions centered at every voxel in the image.

$$p_e(\mathbf{p}) = \frac{1}{N} \sum_{i=1}^{N} \mathbf{E}(\mathbf{p}_i) k_{\mathbf{\Sigma}_i}(\mathbf{p} - \mathbf{p}_i)$$

where N is the number of edge voxels, k is a kernel.

2) Continuous edge density map can be obtained using prethreshold edge strength values leading to a weighted edge KDE.

Suppose we have a continuous edge map $\mathbf{E}_c(\mathbf{p}) \in [0, \infty)$, obtained for instance by employing a suitable rule such as:

$$\mathbf{E}_c(\mathbf{p}) = \begin{cases} (i) : \mathbf{E}_b(\mathbf{p}) * \mathbf{G}(\mathbf{p}) \\ (ii) : ||\nabla I(\mathbf{p})||^2 \\ (iii) : (ii) * \mathbf{G}(\mathbf{p}) \end{cases}$$

where $\mathbf{E}_b(\mathbf{p})$ is a binary edge map, $\mathbf{G}(\mathbf{p})$ is a gaussian filter, and $\nabla I(\mathbf{p})$ is the magnitude of the gradient field of the image.

The KDE based on continuous edge maps can be expressed as

$$p_e(\mathbf{p}) = \sum_{i=1}^{N} w_i k_{\mathbf{\Sigma}_i}(\mathbf{p} - \mathbf{p}_i)$$

where the weights w_i are the normalization factor with

$$w_i = \frac{\mathbf{E}_c(\mathbf{p}_i)}{\sum_{i=1}^{N} \mathbf{E}_c(\mathbf{p}_i)}$$

.

B. Principal Curve/Surface

A self-consistent principal curve or surface, defined by Hastie [11], is a smooth differentiable curve that passes through the middle of the data. In Hastie's definition, every point on the curve or surface is the conditional mean of the points projecting onto this point. Locally defined principal curves and surfaces are presented by Erdogmus and Ozertem [12,13]. They are inspired by Hastie, but utilizes local first and second order derivatives of the underlying density function instead. According to this definition of the principal curve/surface, a point is on the principal curve/surface iff it is the local maximum in the orthogonal subspace, instead of the expected value.

Given a vector $\mathbf{x} \in \mathbf{R}^n$, let $q(\mathbf{x})$ be its pdf, $\mathbf{g}(\mathbf{x})$ be the transposition of the local gradient, and $\mathbf{H}(\mathbf{x})$ be the local Hessian of the pdf [14].

Assume $q(\mathbf{x}) > 0$ for all $\mathbf{x}$, and is at least twice differentiable. Local covariance is defined as: $\mathbf{\Sigma}^{-1}(\mathbf{x}) = -q(\mathbf{x})^{-1}\mathbf{H}(\mathbf{x}) + q(\mathbf{x})^{-2}\mathbf{g}(\mathbf{x})\mathbf{g}^T(\mathbf{x})$. $((\lambda_1(\mathbf{x}), \mathbf{v}_1(\mathbf{x})), ..., (\lambda_n(\mathbf{x}), \mathbf{v}_n(\mathbf{x}))$ are the eigenvalue-eigenvector pairs of local covariance, where the eigenvalues are sorted such that $\lambda_1(\mathbf{x}) > \lambda_2(\mathbf{x}) > ... > \lambda_n(\mathbf{x})$ and $\lambda_i \neq 0$. In general, a point is on the d dimensional principal surface iff the local gradient is in the span of d eigenvectors of the local Hessian and the corresponding $(n-d)$ eigenvalues are negative. Using the KDE, selecting the orthogonally constrained subspace spanned by corresponding $(n-d)$ eigenvectors of the local covariance can constrain the mean-shift iterations into the subspace by d leading eigenvectors of the local covariance.

Consider the data samples $\{\mathbf{x}_i\}_{i=1}^N$, where $\mathbf{x}_i \in \mathbf{R}^n$. The KDE is defined as $q(\mathbf{x}) = \frac{1}{N}\sum_{i=1}^N k_h(\mathbf{x} - \mathbf{x}_i)$, where k is a kernel, and h is the bandwidth of the kernel. The most commonly used kernels are Gaussian kernels. The Gaussian kernel based KDE of this data set is

$$q(\mathbf{x}) = \sum_{i=1}^N w(\mathbf{x}_i) k_{\mathbf{\Sigma}_i}(\mathbf{x} - \mathbf{x}_i)$$

where $w(\mathbf{x}_i)$ is the weight ($= 1/N$ with binary edge voxels) and $\mathbf{\Sigma}_i$ ($= \sigma^2\mathbf{I}$ in this case for computational simplicity) is the kernel covariance for $\mathbf{x}_i$; $k_{\mathbf{\Sigma}_i}(\mathbf{x} - \mathbf{x}_i) = \frac{1}{(2\pi)^{d/2}|\mathbf{\Sigma}_i|^{1/2}} e^{-\frac{1}{2}(\mathbf{x}-\mathbf{x}_i)^T \mathbf{\Sigma}_i^{-1}(\mathbf{x}-\mathbf{x}_i)}$. The gradient and the Hessian of the KDE are

$$\mathbf{g}(\mathbf{x}) = -\sum_{i=1}^N w(\mathbf{x}_i) c_i \mathbf{u}_i$$

$$\mathbf{H}(\mathbf{x}) = \sum_{i=1}^N w(\mathbf{x}_i) c_i (\mathbf{u}_i \mathbf{u}_i^T - \mathbf{\Sigma}_i)^{-1}$$

$$\mathbf{\Sigma}^{-1}(\mathbf{x}) = -q(\mathbf{x})^{-1}\mathbf{H}(\mathbf{x}) + q(\mathbf{x})^{-2}\mathbf{g}(\mathbf{x})\mathbf{g}^{\mathbf{T}}(\mathbf{x})$$

where $c_i = k_{\mathbf{\Sigma}_i}(\mathbf{x} - \mathbf{x}_i)$, $\mathbf{u}_i = \mathbf{\Sigma}_i^{-1}(\mathbf{x} - \mathbf{x}_i)$.

At the mode, the gradient becomes zero:

$$\mathbf{g}(\mathbf{x}) = -\sum_{i=1}^N w(\mathbf{x}_i) k_{\mathbf{\Sigma}_i}(\mathbf{x} - \mathbf{x}_i)\mathbf{\Sigma}_i^{-1}(\mathbf{x} - \mathbf{x}_i) = 0$$

Reorganizing it and solving for $\mathbf{x}$, one obtains the mean-shift update

$$\mathbf{x} \leftarrow \mathbf{m}(\mathbf{x}) = (\sum_{i=1}^N k_{\mathbf{\Sigma}_i}(\mathbf{x}-\mathbf{x}_i)\mathbf{\Sigma}_i^{-1})^{-1} \sum_{i=1}^N k_{\mathbf{\Sigma}_i}(\mathbf{x}-\mathbf{x}_i)\mathbf{\Sigma}_i^{-1}\mathbf{x}_i$$

The eigen decomposition of $\mathbf{\Sigma}^{-1}(\mathbf{x}) = \mathbf{V\Gamma V}$. Let $\mathbf{V}_\perp = [\mathbf{v}_1 ... \mathbf{v}_{n-d}]$ be the $(n-d)$ largest eigenvectors of $\mathbf{\Sigma}^{-1}$. Then, $\mathbf{x}$ is iteratively forced to converge to the principal surface in the constrained space $\mathbf{V}_\perp \mathbf{V}_\perp^{\mathbf{T}} \mathbf{m}(\mathbf{x})$ through the subspace mean-shift update. If the gradient is orthogonal to the subspace spanned by the selected $(n-d)$ eigenvectors when projecting the data from n to d dimensions, the mean-shift iterations stop.

C. Data

The principal surface algorithm to segmentation is implemented in two studies. The first study is a phantom study to simulate the motion of lung tumors. All images are acquired on a dedicated 16-slice helical big-bore simulator (Philips Medical Systems, Cleveland, OH, USA). For respiratory-correlated acquisitions, the scanner is operated in helical cine mode with a table pitch of 0.09 $mm\ s^{-1}$ and gantry rotation time of 0.5 s. Respiratory cycle signals are monitored using the Real-time Position Management respiratory gating system (RPM, Varian Medical Systems, Palo Alto, CA, USA) for precise temporal correlation to CT data acquisition. In brief, an infrared camera mounted to the CT couch records the position of a small, reflective block placed on the upper abdomen of patients to continually measure breathing cycles throughout the cine scan. The breathing trace is then used to retrospectively correlate the cine CT images with 10 equally-spaced phases of the respiratory cycle. All CT scans are reconstructed at 3 mm thickness with pixel size in the transverse direction of 1 mm. All reconstructed CT series are transferred to a 4D treatment planning platform (Eclipse v8.6, Varian Medical Systems, Palo Alto, CA, USA) for manual segmentation. Because the reconstructed 3D-CT series have the same coordinate system (DICOM frame of reference), Eclipse recognizes the images as being registered to each other and creates a 4D image object. The size and resolution of the CT images are $512 \times 512 \times 54$ and $0.668 \times 0.668 \times 3\ mm^3$, respectively.

The second study is a 4D-CT patient study to propagate the lung contours from one of the reference phase to all target phases and compare back with the initial manual contours. These 4D-CT data sets, in the DICOM format, cover one full breathing cycle with 10 3D-CT scan phases numbered phase 0 through phase 9. Each 3D-CT scan has a slice thickness of 3 mm with 116 slices, the size for each slice 512×512, and the pixel resolution for each slice $1.17188 \times 1.17188\ mm^2$. Phase 6 corresponds to the end-of-expiration phase. Manual contours of all ROIs are drawn by a physician on that phase. All ROIs are stored as separate binary mask files with the same dimensions as each 3D-CT scan. A binary ROI mask file is a three-dimensional matrix with ones at the voxels of ROIs and zeros everywhere else.

D. Evaluation Metrics

Two methods are used in quantitative validation:

1) Dice similarity coefficient index
In order to measure the overlap between the automatically propagated contours achieved by the principal surface algorithm and the manually drawn contours by a physician, the Dice coefficient is applied.

$$d = 2\frac{|A \cap B|}{|A| + |B|}$$

where A and B indicate the volume of objects. The output of the principal surface algorithm is the 3D coordinates of object surfaces. A binary 3D volume mask should first be created in order to calculate the Dice coefficients.

2) Absolute centroid distance
In order to measure the distances between the volumes defined by the surfaces, the absolute centroid distance is applied. The absolute centroid distance is the Euclidean distance between the centroids of two volumes.

III. Experiment Results

A. Phantom Study

We test the principal surface algorithm on each phase of the phantom data for tumor segmentation. In the preprocessing step, each slice of the 3D-CT scan is down sampled by 3/0.668 while keeping the same resolution in the z dimension to yield $3 \times 3 \times 3\ mm^3$ resolution. The edge maps are detected by the Canny edge detector. A Gaussian kernel is used with the kernel covariance $\sigma^2 I$, where $\sigma = 1$. The resampled contours after the principal surface algorithm are 1-4 pixels inside the manually drawn contours. Those contours are converted into a binary volume with a pixel value of 1 inside the contours and 0 outside. Morphological operation is then used to grow the contours by n pixels, where n is determined based on the data to maximize the Dice coefficient. Here we choose n=2.

Figure 1 shows the results of the proposed algorithm on one slice of a 3D-CT scan with its associated automatic contours (blue lines). The manual contours are included for comparison (red lines).

Figure 2 shows the Dice coefficient index between the automatic contours by the principal surface algorithm and manually drawn contours of tumors in each 3D-CT phantom data set. The mean Dice coefficient index for this 4D phantom data is 0.9666 with a standard deviation of 0.0069.

Table I illustrates the centroid distances of between the automatically segmented tumor volumes and manually delineated tumor volumes in each 3D-CT phantom data set based on the absolute centroid distances in Medial-Lateral, Superior-Inferior, Anterior-Posterior directions. The results demonstrate that the automatically segmented contours are consistent with the manually drawn contours of tumors.

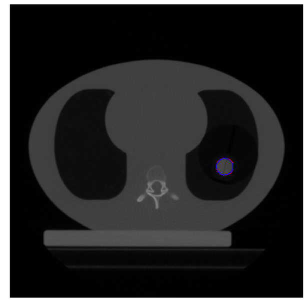

Figure 1: Automatically segmented contours (blue lines) and manual contours (red lines) on one slice of one 3D-CT phase out of 4D-CT phantom data

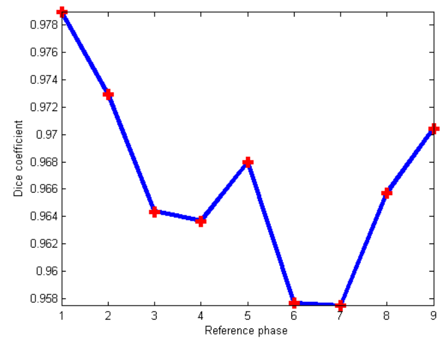

Figure 2: Dice coefficients for each 3D-CT phase

B. Patient Study

We also apply the principal surface identification algorithm to a patient 4D-CT data set for lung segmentation. 10 phase 3D-CT scans form a whole breathing cycle in 4D-CT data sets. In the preprocessing step, each slice is down sampled by 3/1.17188 while keeping the same resolution in the z dimension to yield $3 \times 3 \times 3\ mm^3$ resolution. The edge maps are also detected by the Canny edge detector. A gaussian kernel is used with the kernel covariance $\sigma^2 I$, where $\sigma = 2$. Lung contours are only manually delineated on phase 6. Manual contours of lung from phase 6, therefore, are used as an initialization surface to find the lung surface of the consecutive phase. The output is then used as the initialization for the next phase. This procedure is executed recursively in all consecutive phases to complete one respiration cycle. The final output contour after a complete propagation is resampled back to the original 3D-CT scan. The Dice coefficient between the 10-phase recursively iterated principal surface (to coincide with the initial reference phase) and the manually delineated lung volume in this reference phase is 0.946.

Figure 3 shows the triangulated surfaces of the initial manually segmented lung and the corresponding automatically

Table I: Absolute centroid distances of tumors versus respiration phase in each direction

Phase	ML (mm)	SI (mm)	AP (mm)
1	0.0361	0.0028	0.0114
2	0.0917	0.0352	0.0015
3	0.0819	0.0426	0.0118
4	0.0679	0.0639	0.0326
5	0.0118	0.0939	0.0158
6	0.0654	0.0113	0.0074
7	0.0432	0.0075	0.0638
8	0.0273	0.0021	0.0704
9	0.0632	0.0402	0.0359

propagated contours for the lung.

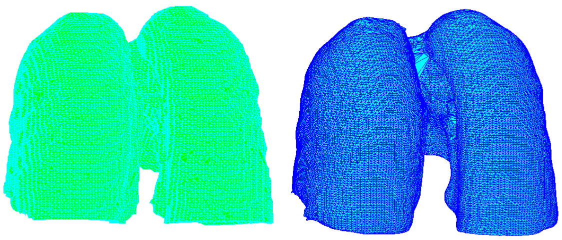

Figure 3: Left figure: Triangulated surfaces of the lung by physicians (green); Right figure: Triangulated surfaces of the lung segmented by the principal surface algorithm (blue)

Figure 4 compares the automatic contours with the manual contours in transversal, sagittal, and coronal views. Slicer is used as the visualization tool. The original manually drawn contours for the lung are displayed as solid golden lines, and the automatically segmented contours are shown as red shades. The manual contours are only drawn on the transversal view, and the visualization tool does some approximations to display it on the sagittal and coronal views. These comparisons show that the consistency between the manual contours and automatic contours.

The absolute centroid distances between the volumes of the automatically segmented lung and the manually segmented lung are 0.3420 mm, 0.1026 mm, and 0.3153 mm, respectively in Medial-Lateral, Superior-Inferior, Anterior-Posterior directions. It can be noted that the outputs of our principal surface algorithm are consistent with the manually drawn contours of ROIs by physicians.

IV. Discussions

Modern technological advances have significantly improved the ability to target tumors while sparing nearby OARs. However, the steep dose gradients applied necessitate the closer monitoring of breathing induced motion to ensure tumor coverage. The use of 4D-CT imaging is gaining popularity in radiation therapy as a means of measuring and

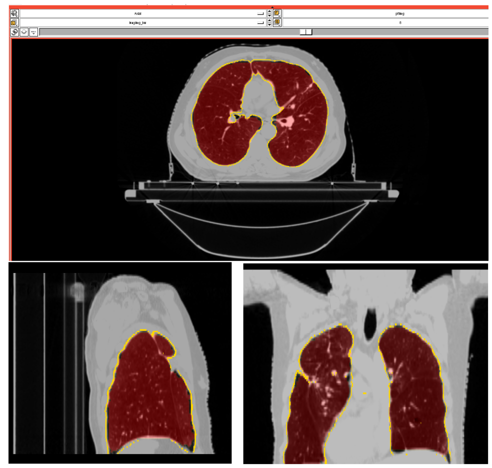

Figure 4: Transversal, sagittal, coronal views of CT images comparing the originally manually delineated lung contours with the contours obtained by the principal surface algorithm

adapting to inter- and/or intrafraction motion. As manual contouring of critical structures in all phases is not realistic in the routine workflow in the clinic, registration is typically used to propagate manually drawn contours from one phase to all subsequent phases. Rigid or deformable registration is first applied using the reference phase as the fixed image and other phases as moving images. The contours are deformed using the resulting deformation fields.

However, as deformable registration can be computationally expensive if large volumes are used, global rigid registration followed by local deformable registration is commonly used. As the current implementation of the principal surface-based algorithm uses only binary edge points detected by the edge detector, the computational complexity and speed are better than many deformable registration algorithms. Finding the accurate edge information for tumors or organs is a very crucial step. Consequently, in our approach the success of the algorithm may be limited if edges based on intensity or texture are not easily discernible. The algorithm generally performs well on the lungs and simulated tumors in phantom data because of the clear boundaries detected. Some organs do not have significant contrast against the background, such as the heart, the liver, the prostate, and the kidney. Tumors often have similar intensity and texture as the surrounding soft tissues. Thus, our proposed algorithm might have limited success in segmenting those ROIs. Physicians delineate these ROIs based on their previous knowledge of the shape and relative positions of the ROIs. A more accurate and robust segmentation would incorporate prior shape information. The prior knowledge of ROIs is critical to successfully segment those ROIs with lower contrast to the background or similar intensity values to the surrounding tissues.

V. Conclusions and Future Work

We have presented preliminary results on a novel application of the principal surface algorithm for the propagation of contours in 4D-CT for applications in radiation therapy. Our experimental results demonstrate the feasibility of automatic segmentation of the ROIs in 4D-CT. We are currently performing a more extensive evaluation, comparing this algorithm to commonly used deformable and rigid registration-based contour propagations. In the future, we will incorporate prior shape information into the KDE procedure to make principal surfaces smoother and more consistent with the prior shapes. We are also evaluating the use of weighted edges instead of purely binary edges to improve the robustness of the algorithm.

References

[1] International Commission on Radiation Units, "Prescribing, Recording and Reporting Photon Beam Therapy," ICRU Report 50. Bethesda International Commission on Radiation Units and Measurement, MD, ICRU, 1993.

[2] International Commission on Radiation Units, "Prescribing, Recording and Reporting Photon Beam Therapy," Supplement to ICRU Report 50. Bethseda International Commission on Radiation Units and Measurement, MD, ICRU, 1999.

[3] S. Bhide, M. T Guerrero Urbano, C. Clark, V. Hansen, E. Adams, E. Miles, H. McNair, A. Warrington, K Harrington, and C Nutting, "Results of Intensity Modulated Radiotherapy (IMRT) in Laryngeal and Hypopharyngeal Cancer: a Dose Escalation Study," Radiother Oncol, vol. 82, 2007, pp. S74-S75.

[4] S. T. Peeters , W. D. Heemsbergen, P. C. Koper, W. L. van Putten, A. Slot, M. F. Dielwart, J. M. Bonfrer, L. Incrocci, and J. V. Lebesque, "Dose-response in Radiotherapy for Localized Prostate Cancer: Results of the Dutch Multicenter Randomized Phase III Trial cCmparing 68 Gy of Radiotherapy with 78 Gy," J Clin Oncol, vol. 24, no. 13, 2006, pp. 1990-1996.

[5] T. G. Urbano, C. H. Clark, V. N. Hansen, E. J. Adams, E. A. Miles, H. Mc Nair, A. M. Bidmead, J. Warrington, D. P. Dearnaley, C. Harmer, K. J. Harrington, and C. M. Nutting, "Intensity Modulated Radiotherapy (IMRT) in Locally Advanced Thyroid Cancer: Acute Toxicity Results of a Phase I Study," Radiother Oncol, vol. 85, no. 1, 2007, pp. 58-63.

[6] M. J. Zelefsky, Z. Fuks, M. Hunt, H. J. Lee, D. Lombardi, C. C. Ling, V. E. Reuter, E. S. Venkatraman, and S. A. Leibel, "High Dose Radiation Delivered by Intensity Modulated Conformal Radiotherapy Improves the Outcome of Localized Prostate Cancer," J Urol, vol. 166, no. 3, 2001, pp. 876-881.

[7] M. Peroni, M. F. Spadea, M. Riboldi, G. Baroni, G. T. Y. Chen, and G. C. Sharp, "Validation of an Automatic Contour Propagation Method for Lung Cancer 4D Adaptive Radiation Therapy," IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI '09), 2009. pp. 1071-1074.

[8] J. Kalpathy-Cramer, U. Ozertem, W. Hersh, M. Fuss, and D. Erdogmus, "Semi-supervised Segmentation using Non-parametric Snakes for 3D-CT Applications in Radiation Oncology," Machine Learning for Signal Processing (MLSP 2008), IEEE Workshop on, 2008, pp. 109-114.

[9] J. Kalpathy-Cramer, U. Ozertem, W. Hersh, M. Fuss, and D. Erdogmus, "Robust Segmentation using Non-parametric Snakes with Multiple Cues for Applications in Radiation Oncology," SPIE Medical Imaging, 2009

[10] U. Ozertem, and D. Erdogmus, "Nonparametric Snakes," IEEE Transactions on Image Processing, vol. 16, 2007, 2361-2368.

[11] T. Hastie, W. Stuetzle, "Principal Curves," Jour. Am. Statistical Assoc., vol. 84, no. 406, 1989, pp. 502-516.

[12] D. Erdogmus, and U. Ozertem, "Self-consistent Locally Defined Principal Surfaces," IEEE Internatinal Conference on Acoustics, Speech and Signal Processing(ICASSP 2007), pages II-549-II-552, April, 2007.

[13] U. Ozertem, and D. Erdogmus, "Local Conditions for Critical and Principal Manifolds," Proceedings of ICASSP'08, 2008, pp. 1893-1896.

[14] U. Ozertem, "Locally Defined Principal Curves and Surfaces," Ph.D. dissertation, OHSU, 2008