Semi-supervised Segmentation using Non-parametric Snakes for 3D-CT Applications in Radiation Oncology

Jayashree Kalpathy-Cramer[1], *Member, IEEE,* Umut Ozertem[4], *Student Member, IEEE,* William Hersh[1], Martin Fuss[3], Deniz Erdogmus[2], *Senior Member, IEEE*

Departments of [1]Medical Informatics and Clinical Epidemiology, [2]Radiation Medicine, [3]Computer Science and Electrical Engineering, Oregon Health & Science University, Portland, OR, [4]Yahoo! Inc. Santa Clara CA

***Abstract*—We present a semi-supervised protocol for segmentation of tumors and normal anatomy for applications in Radiation Oncology. A primary goal in radiation therapy in oncology is to deliver high radiation dose to the perceived tumor while sparing the surrounding non-diseased organs. Consequently, a critical task in the workflow of radiation oncologists is the manual delineation of normal and diseased structures on 3D-CT scans. In this paper, we compare the results using a non-parametric snake technique with a gold standard consisting of manually delineated structures. Structures include tumors as well as normal organs including lungs, liver and kidneys. This technique provides fast segmentation that is robust with respect to noisy edges. In addition, this algorithm does not require the user to optimize a variety of parameters unlike many segmentation algorithms. We provide results that show the improvement in overlap between the manually delineated gold standard and the output of the segmentation algorithms using the user input.**

I. Introduction

CANCER is the second most frequent cause of death in the United States. Radiation therapy is one of the most effective treatments for many types of cancer, and is used in the care of about half of all people being treated for cancer. Ionizing radiation can be used to kill cancer cells and shrink tumors, thereby either treating the cancer in curative intent or proving relief of symptoms (palliation) [1]. A primary goal in radiation therapy is to deliver high radiation doses to the perceived tumor while sparing the surrounding healthy tissues. Recent technological advances such as Intensity-modulated Radiation Therapy (IMRT) [1, 2] or 3-dimensional conformal radiation therapy allow radiation oncologists to deliver high doses of radiation to the tumor while minimizing the dose delivered to nearby areas, resulting in fewer side effects and reducing toxicity.

A critical aspect of the treatment is the planning stage where the radiation oncologist determines the area to be treated and the dose profile that will be delivered. For radiation therapy planning, multi-slice, three-dimensional computed tomography (CT) scans are obtained from the patient. Typically, the radiation oncologist manually delineates (or "contours") the tumor and the neighboring organs on each slice of the CT volume using a mouse or other pointing device.

Although there are many algorithms available for automatic and semi-automatic contouring, radiation oncologists often adjust the contours on the slices using their knowledge of general anatomy, prior knowledge about the particular patient and data from other sources like Positron Emission Tomography (PET) scans. In addition, when tumors lie on the walls of organs or are present in organs with similar density, the edge of the object can be occluded or hard to discern. These situations can degrade the performance of algorithms that rely on clear boundaries or sharp differences in pixel intensity.

Segmentation, especially for medical applications, is an active area of research. Common algorithms include those based on clustering, histograms, edge detection, region growing, and level sets. Gradient vector flow (GVF) snakes [3] are a popular class of algorithms in the active contours (snakes) [4] category. The segmentation problem for snakes is traditionally defined as a parametric energy minimization problem. Traditional snakes are sensitive to noise and capture range issues. Their initialization must be close to boundary desired and their convergence can depend on the initial position. GVF snakes solve many of these issues and are good at detecting shapes with boundary concavities. However, they are sensitive to the parameters of the algorithm and may need to be tuned for each application. They can also be quite slow as extracting the GVF field can be computationally expensive. Both these issues can be problematic for medical applications where the clinicians do not wish to tune algorithmic parameters and real-time segmentation is required. In this paper, we apply a fast, robust algorithm for segmentation in multi-slice volumes that does not require many parameters to be optimized. The segmentation is performed in near real-time. The algorithm can take into account the user's input in one or more slices and propagate that to other slices. Non-parametric snakes [5] are used as the basis for this procedure. This semi-supervised method incorporates the knowledge from the clinician along with the features of the image to provide a robust algorithm that can deal with noisy or somewhat

This work was supported in part by the National Library of Medicine (NLM) Training grant 2T15LM007088 (JKC) and NSF grants ECS-0524835 and ECS-0622239 (UO, DE)

occluded edges.

II. Image Segmentation using Non-Parametric Snakes

A. Proposed protocol

We envision the following flow for the semi-supervised segmentation in clinical practice.

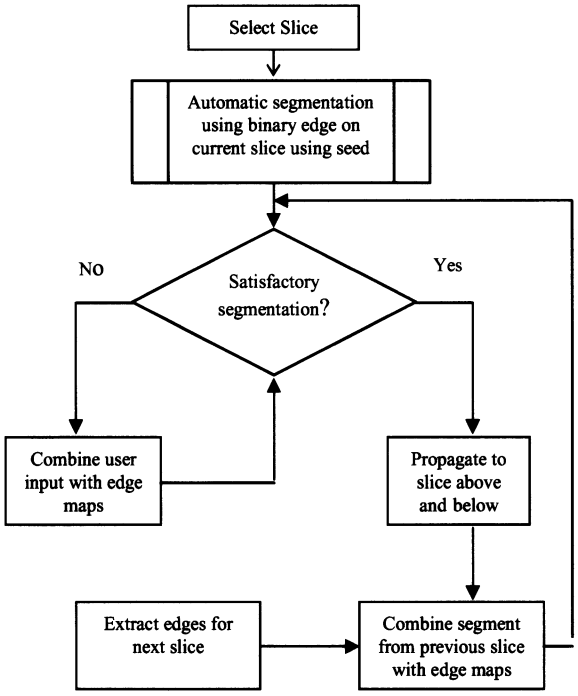

B. Non-parametric Snakes

Non-parametric snakes are a recently proposed [5] kernel density estimation (KDE) based segmentation technique that is fast and relatively robust to noise. Let I be the CT image volume in 3 dimensions such that $[x,y,z]^T$ is the location of a pixel in space and $I(x,y,z)$ is the intensity of the image at that location. We consider a given slice, S, chosen by the user. We want to segment the object in this slice contained in the image $I(x,y)$. In the implementation used in this paper, we use a binary edge detector (Canny) with thresholds adjusted for the organ of interest. Let $E(\mathbf{s})$ be the edge image, where $\mathbf{s}$ is the vector of pixel coordinates x, y, and z.

$$E(\mathbf{s}) = \begin{cases} E(\mathbf{s}) = 1 & s \text{ is an edge pixel} \\ E(\mathbf{s}) = 0 & otherwise \end{cases}$$

In the general case, as described in [5], the KDE of the edge map is constructed using the edge image E(s) via a summation over the kernel functions centered at every point in the image.

$$p_{edge}(\mathbf{s}) = \frac{1}{N_{edge}} \sum_{i=1}^{N} E(\mathbf{s}_i) K_{\Sigma_i}(\mathbf{s} - \mathbf{s}_i)$$

where N_{edge} is the number of pixels and

$$N_{edge} = \sum_{i=1}^{N} E(\mathbf{s}_i)$$

However, in this application, we used a fixed-width isotropic Gaussian kernel. We used the fixed-point algorithm as described in [5] to achieve fast convergence. Specifically:

$$\mathbf{s} \leftarrow \frac{\sum_{i=1}^{N} E(\mathbf{s}_i) K_{\sigma^2 \mathbf{I}}(\mathbf{s} - \mathbf{s}_i)\mathbf{s}_i}{\sum_{i=1}^{N} E(\mathbf{s}_i) K_{\sigma^2 \mathbf{I}}(\mathbf{s} - \mathbf{s}_i)} = \mathbf{m}(\mathbf{s})$$

This is similar to the mean shift described in [6] and can avoid spurious maxima by a suitable choice of kernel size. For continuous edgemaps, the algorithm becomes mean shift with a weighted KDE model. Consequently, the algorithm primarily seeks to converge to local maxima in the smoothed edge distribution. Because the convergence speed of each eigenmode in the vicinity of a local maximum is proportional to the eigenvalue magnitude, if a ridge exists, the algorithm practically converges to the ridge first and then climbs up the ridge towards the peak. Another consequence of this behavior is, as seen in some of the experiments in the next section, the snake may not progress into low edge density regions along the object boundary and also into boundary concavities depending on the kernel size and the distance of the initialization curve from the ridges. This drawback is addressed by generating additional sample points using the principal surface approach described below in section E.

C. User Intervention:

The segmentation on a given slice is reviewed and corrected by the clinician. The algorithm allows the combination of output of the edge detector and the user-defined segmentation to create a new KDE for that slice. The respective weights are user defined variables. This enhances the segmentation where the edges are not discernable and where the clinician has indicated the existence of an edge based on prior knowledge from a different information source like a different imaging modality or anatomical knowledge.

The new edge used for the non-parametric snake is a weighted sum of these two edges, where the weights are user determined as:

$$E_{new}(\mathbf{s}) = w_{edge} E_{edge}(\mathbf{s}) + w_{user} E_{user}(\mathbf{s})$$

where $w_{edge} + w_{user} = 1$

E_{edge} are the edge pixels resulting from the edge detector (Canny, in this case) and E_{user} be the user determined edge, w_{edge} and w_{user} are the respective weights. Note that the user does not need to specify an edge for the entire region of interest but can chose to provide edge information in the case of occluded or obscure edges.

D. Propagation

The final segmentation on a given slice either as automatically detected or in a semi-supervised fashion in

conjunction with the clinician, is propagated to slices above and below that slice. This is used in two ways. Not surprisingly, it is used as the initialization for the slices immediately adjacent. However, it is also used for the KDE for the next slice. The KDE on a given slice is the weighted sums of the edges detected on that slice and the contours propagated from neighboring slices, with lower weights for the slices that are farther away along z orthogonal to the slice under consideration. This incorporates the users' edges without requiring user input on all slices. By contouring a few selected slices in a volume, a three dimensional object can be segmented in a robust and fast manner using non-parametric snakes.

$$E_{thisSlice}(\mathbf{s}) = w_{thisEdge}E_{thisEdge}(\mathbf{s}) + \sum_{i=-N}^{N} w_i E_i(\mathbf{s})$$

where 2N is the number of neighboring slices that are used and $w_{edge} + \sum_i w_i = 1$

E. Principal Curves

We expanded used the principal curves approach as described in [7] to add points in sparse areas resulting from the output of the non-parametric snakes. We used a Gaussian kernel for the KDE using the updated edge map. This updated edge map is the weighted sum of the output of the edge detector, the user input (if applicable) and the contour from neighboring slices. For each nearest neighbor pair of points in the nonparametric snake, the midpoint $\mathbf{s}_0$ of the line segment connecting each pair is taken as an initial point. Let $\mathbf{v}_0$ be the major eigenvector of the local covariance of $p_{edge}(\mathbf{s})$ at $\mathbf{s}_0$ and $\mathbf{g}_0$ be the gradient at this point. Let $\mathbf{m}(\mathbf{s})$ be the mean-shift update direction as defined earlier. Starting from $\mathbf{s}_0$, we iterate the following:

$$\mathbf{s} \leftarrow sign(\mathbf{v}_0^T\mathbf{g}_0)\mathbf{v}_0\mathbf{v}_0^T\mathbf{m}(\mathbf{s})$$

This iteration has been shown to converge to the ridge of the edge distribution, therefore is an interpolating sample from the desired boundary. The process is repeated until a satisfactory number of samples are obtained from the boundary. This may be determined, for instance by considering the distribution of nearest neighbor distances between the samples in the nonparametric snake.

III. Results

A. Automatic Segmentation on a Slice

Fig. 1 shows a section of a lung tumor that has been contoured by a radiation oncologist. The output of the Canny edge detector [8], using the appropriate threshold for lung tumors, can be seen in Fig. 2. A crude initialization of the nonparametric snake is also displayed. The probability density estimate based on this edge map is shown in Fig. 3. In just 2 iterations, the non-parametric snake has segmented the tumor as seen in Fig. 4. However, as can be observed, the user has chosen a larger area for the segmentation, albeit by one to two pixels.

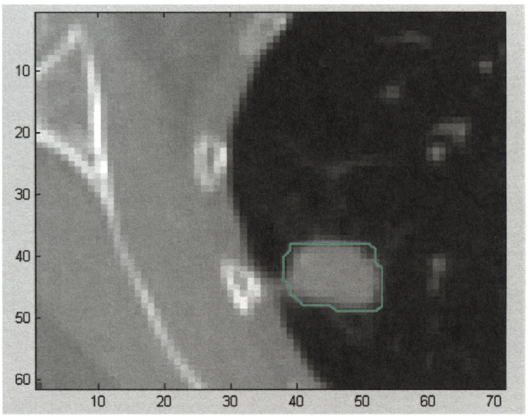

Fig. 1. Lung tumor contoured by a radiation oncologist (in green), considered the gold standard.

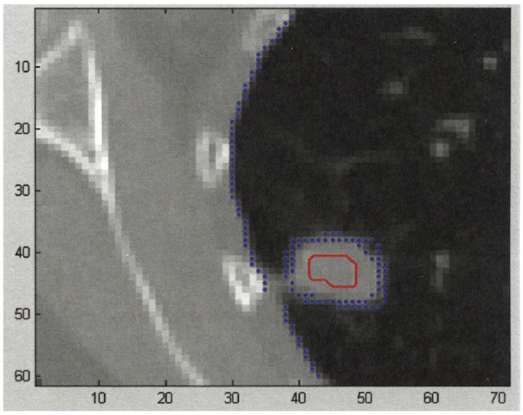

Fig. 2. Canny edge detection (in blue) and initialization (in red)

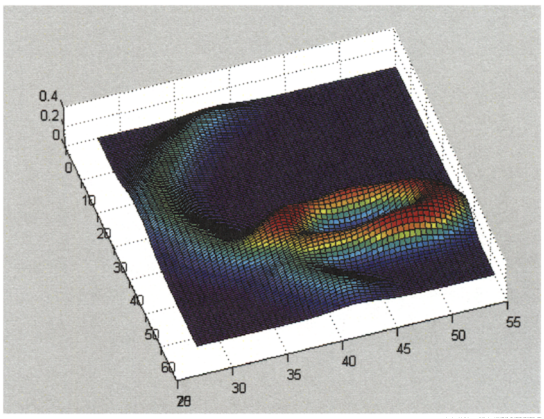

Fig. 3 Probability density estimate of the edge map of the lung tumor in Fig. 2.

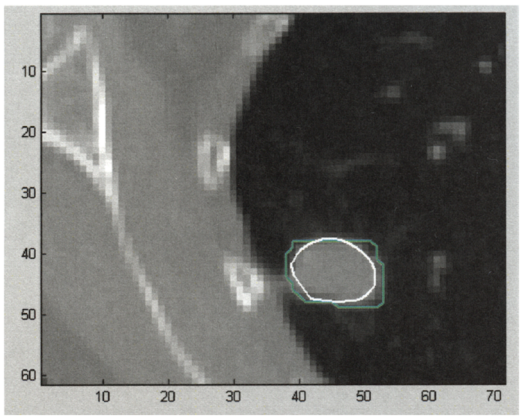

Fig. 4. Final result after two iterations of the non-parametric snakes algorithm (in white) compared to hand-drawn contour (in green).

Fig. 5 shows the results of the non-parametric curves on the right lung. Again, without any user input or tuning, the non-parametric snake provides good segmentation in just two fixed-point iterations. However, we can see in Fig. 5 that there is a slight difference between the user's contour and the output of the non-parametric snake based segmentation, especially in concavities. Poor initialization and clustering due to the mean shift behavior can result in sparse points in the concavities. This can be improved by interpolating using the principal curves algorithm, as seen in part C below.

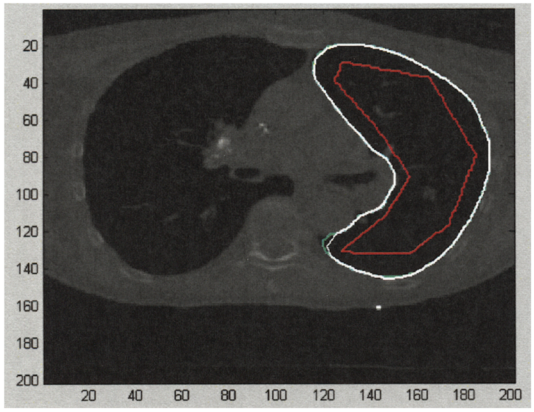

Fig. 5. Initialization (in red), automatic segmentation (white), user-drawn contour (green)

We compared the performance of this automatic segmentation algorithm on 100 slices from four different patients. Overall, we achieved over 90% overlap with the gold standard on lungs and over 85% overlap on lung tumors, where a margin of one to two pixels can have a large impact on the overlap percentage. However, this is dependent on the initialization. Methods to reduce this dependence on the initialization are presented in sections B, C and D.

B. Modified KDE Incorporating User Input

Fig. 6 demonstrates the impact of using the user-defined contour in estimating the KDE on a given slice. In this case, the Canny edge and the user-specified contour were given equal weights. We can see the improvement in meeting the gold standard compared to the completely automatic segmentation.

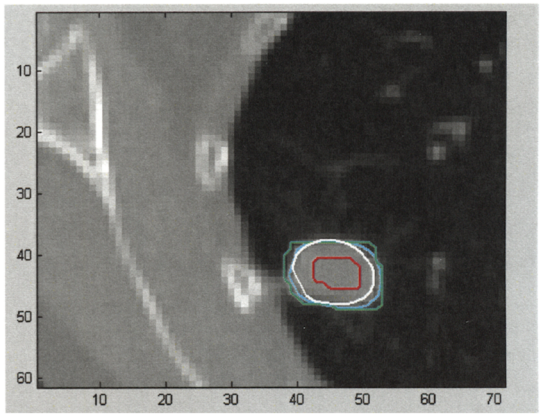

Fig. 6. Improved segmentation by incorporating user input: Initialization (in red), automatic segmentation (white), user-drawn contour (green), segmentation incorporating user input (cyan)

C. Propagation of Curves

The propagation of the contours to adjacent slices for initialization improves the segmentation of neighboring slices by providing an initialization that is very close to the expected output. This reduces the requirements normally necessary for the initializations. For instance, Fig. 7 shows a lung segmented starting with a poor initialization. Although, the non-parametric snake matches the user-drawn contour along most of the perimeter, we can see the problems caused the initial curve being far away from the edge at the concavities.

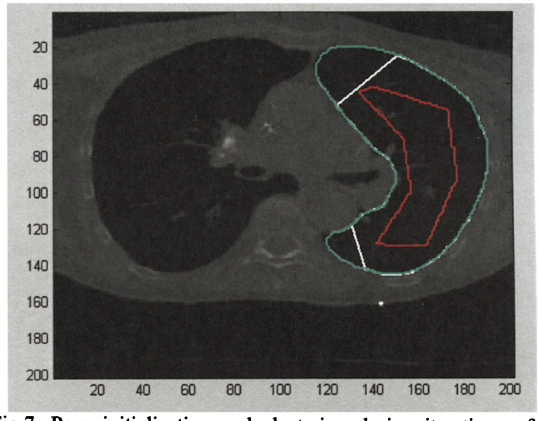

Fig.7. Poor initialization and clustering during iterations of the non-parametric snakes can cause sparse points in concavities in lung: initialization (red), final contour (white).

Fig. 8 demonstrates the effectiveness of the propagation

method. In this case, the output of the automatic segmentation from one slice was used as the initialization for the slice above. The edge points for the slice of interest were a weighted combination of the edge points as detected by the Canny detector and the final contour from the previous slice.

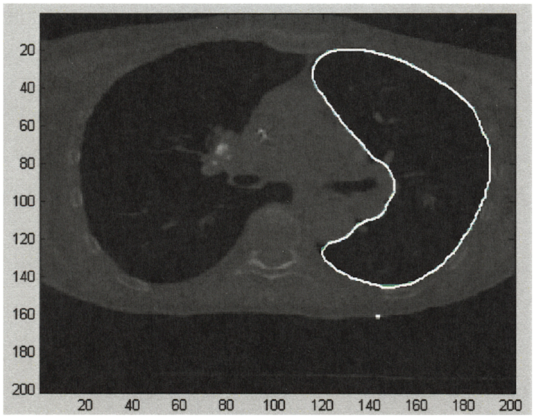

Fig. 8. Propagation of contours from adjacent slices greatly improves segmentation: User-defined contour edge (green), final segmentation (white)

Overall, we achieved over 95% overlap with the gold standard on lungs and over 92% overlap on lung tumors by propagating a contour to the slice immediately adjacent.

D. Principal curves

The non-parametric snakes algorithm in [5], being essentially a mean-shift approach using weighted KDE, can cause the points to cluster in regions of high density if the edges are non-uniformly sampled along the ridge. We explored the use of principal curves as a means to generate additional points in areas of sparse data. In Fig. 9, the final non-parametric snake has few points in the concavity. This can occur if the initialization points are far from the KDE of the edge and the kernel size is small.

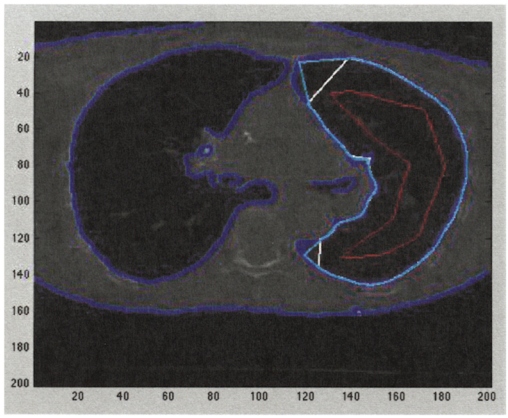

Fig. 9 Additional points using principal curves improve conformity to object boundary: Canny edge (blue), final segmentation (white), interpolation using principal curves (cyan).

Principal curves as defined in [7] can be used to generate additional points along the ridge as detailed above. This interpolation procedure can be repeated iteratively until the minimum distance between adjacent points is below the desired threshold. Figure 9 shows the effect of adding a just one additional point between a pair of points whose distance exceeded the threshold. This greatly improves the overlap with the user-defined contour. Using principal curves reduces the dependence of the initialization.

E. Comparison to other methods

We performed limited comparison of the automatic segmentation using non-parametric snakes with a publicly available implementation of the level set methods [8]. We would like to note that only very limited optimization of the parameters was performed. The parameters used for the tumor example for the level set algorithm were: $\lambda = 5.0$, $\mu = 0.001$, $\nu = 3.0$, and time step $\tau = 200$. The curve evolution used 300 iterations.

As seen in figure 10, the resulting segmentation was very similar between the non-parametric snakes and the level sets. The same initialization was used for both methods. The non-parametric snake algorithm, requiring only two iterations, was significantly faster than the level set, as implemented in [9]. On a Windows XP machine with 4GB of RAM, running MATLAB version 7.0.1.15 R14, the average time for segmenting lung tumors using the level set algorithm was 5.85 +/-0.3 seconds while the average time for the non-parametric snakes was 0.64 +/-0.14 seconds.

The results on the lung were similar in that the level set algorithm took almost 10 times as long as the non-parametric snakes and the results were highly dependent on the initialization and the parameters. The non-parametric snakes took 1.53 +/-0.21 seconds.

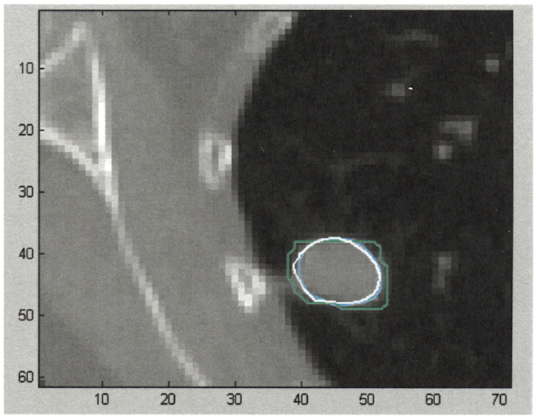

Fig. 10. Comparison of level set (cyan) and non-parametric snakes (white) with user-defined contour (green).

IV. Conclusion

Non-parametric approaches to snakes are an attractive alternative for robust and fast segmentation for medical applications. They do not require optimization of many parameters, unlike many other segmentation algorithms. The kernel width can be optimized automatically using leave-one-out cross-validation if desired. They allow user input to be combined with edge features for segmentation of noisy images or in the case of missing edges. Propagation of contours improves segmentation by improving the initialization and providing a means for user input on one slice to be applied to neighboring slices.

These preliminary results indicate that a semi-supervised 3D object segmentation algorithm is feasible and only information from a user about the desired object indicated in a few slices in the form of reasonably accurate contours is sufficient to carry these priors to other slices. Future work will involve improving algorithm robustness and accuracy by incorporating better edge features and experimentation with data from multiple subjects. Due to the speed of the algorithms, our future plans also include an interactive cooperative-segmentation procedure that will not only benefit radiation oncologists, but will have impact on general object segmentation in multidimensional images and videos.

Acknowledgment

We would like to acknowledge Dr. Apiradee Srisuthep, Research Fellow in the Department of Radiation Medicine, OHSU, Portland, OR for her help in acquiring the manually labeled data.

References

[1] "Radiation Therapy for Cancer: Q & A - National Cancer Institute"; http://www.cancer.gov/cancertopics/factsheet/therapy/radiation.

[2] "Intensity-Modulated Radiation Therapy"; http://www.radiologyinfo.org/en/info.cfm?pg=imrt&bhcp=1.

[3] Chenyang Xu and J. Prince, "Snakes, shapes, and gradient vector flow," *Image Processing, IEEE Transactions on*, vol. 7, 1998, pp. 359-369.

[4] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *International Journal of Computer Vision*, vol. 1, Jan. 1988, pp. 321-331.

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

[6] U. Ozertem, D. Erdogmus, and Tian Lan, "Mean Shift Spectral Clustering for Perceptual Image Segmentation," *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on*, 2006, p. II.

[7] D. Erdogmus and U. Ozertem, "Self-consistent locally defined principal surfaces," *IEEE Workshop on Machine Learning in Signal Processing, 2008 (submitted)*.

[8] J. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 8, Nov. 1986, pp. 679-698.

[9] Chunming Li et al., "Level set evolution without re-initialization: a new variational formulation," *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, 2005, pp. 430-436 vol. 1.