Non-parametric Iterative Model Constraint Graph min-cut for Automatic Kidney Segmentation

M. Freiman[1], A. Kronman[1], S.J. Esses[2,3], L. Joskowicz[1], and J. Sosna[3]

[1] School of Eng. and Computer Science, The Hebrew Univ. of Jerusalem, Israel
[2] Mount Sinai School of Medicine, New York, NY
[3] Dept. of Radiology, Hadassah Hebrew University Medical Center, Jerusalem, Israel
freiman@cs.huji.ac.il

Abstract. We present a new non-parametric model constraint graph min-cut algorithm for automatic kidney segmentation in CT images. The segmentation is formulated as a maximum a-posteriori estimation of a model-driven Markov random field. A non-parametric hybrid shape and intensity model is treated as a latent variable in the energy functional. The latent model and labeling map that minimize the energy functional are then simultaneously computed with an expectation maximization approach. The main advantages of our method are that it does not assume a fixed parametric prior model, which is subjective to inter-patient variability and registration errors, and that it combines both the model and the image information into a unified graph min-cut based segmentation framework. We evaluated our method on 20 kidneys from 10 CT datasets with and without contrast agent for which ground-truth segmentations were generated by averaging three manual segmentations. Our method yields an average volumetric overlap error of 10.95%, and average symmetric surface distance of 0.79mm. These results indicate that our method is accurate and robust for kidney segmentation.

1 Introduction

Kidney segmentation and volumetric measurement from Computed Tomography (CT) datasets has been proven to be an effective and accurate indicator for renal function in many clinical situations. These include urological treatment decision-making, radiotherapy planning, and estimation of the glomerular filtration rate of living donors [1,2,3]. CT imaging is widely used for kidney analysis and diagnosis since it provides essential anatomical information, including kidney morphology and renal vessel characteristics.

Automatic kidney segmentation is a challenging task. The main factors are unclear borders between the kidney, the liver and the spleen, image acquisition artifacts, image noise, and various pathologies, such as tumors and nephrolithiasis. The interactive graph min-cut based segmentation method provides globally optimal segmentation based on both weighted voxel adjacencies and prior models of the object and the background [4]. However, extensive user interaction

T. Jiang et al. (Eds.): MICCAI 2010, Part III, LNCS 6363, pp. 73–80, 2010.
© Springer-Verlag Berlin Heidelberg 2010

is required to provide estimates of the prior intensity models, and to prevent "segmentation leaks" with predefined spatial constraints. This user interaction limits its routine clinical use for kidney segmentation [5].

A variety of methods have been proposed to incorporate fixed parametric shape information to spatially constrain the graph-min cut optimization. Slabaugh and Unal [6] constrain the optimization to a narrow band of a predefined ellipse. Freedman and Zhang [7] use a level-set representation of the shape prior to compute the probability of each voxel to belong to the object class. Ali et al. [8] use the Poisson distribution and distance maps to compute the shape term of the graph for the segmentation of 2D kidney slices from DCE-MRI. Freiman et al. [9] use a local tubular descriptor to adapt the graph-cut segmentation to vascular structures. All these methods require the explicit formulation of a specific shape model which is inadequate as a general solution for kidney segmentation due to the large inter-patient shape variability. In addition, an interactive initialization is often required to properly position the shape model in the image domain [6,7]. Although Freiman et al. [9] use an automatic initialization, their method is designed specifically for carotid arteries segmentation, and cannot be directly applied to the segmentation of other organs.

Recently, Kumar et al. [10] and Malcolm et al. [11] proposed to use adaptive parametric shape models to constrain the min-cut optimization. In their method, the model parameters are considered as a latent variable, and an iterative approach is used to simultaneously estimate the shape model parameters and to compute the segmentation. The drawback of parametric shape models is that they are less suitable for medical images in which the inter-patient organ shape variability is relatively large, and where various types of pathologies are present [12,13].

In this paper we present a non-parametric global shape constrained graph min-cut approach for the automatic segmentation of kidneys from CT images. Our approach defines a non-parametric hybrid model that couples both shape and patient specific intensity information as a latent variable. An Expectation Maximization (EM) algorithm is then used to simultaneously estimate the latent model and to produce the kidney segmentation. The main advantages of our method are that: 1) it iteratively refines both the shape and the intensity models during the segmentation to increase its accuracy; 2) it uses a non-parametric shape representation that allows weighting of the training instances, and; 3) it uses the graph min-cut framework to perform global optimization on the entire volume in each iteration, as opposed to level-sets [13] or Maximum A-Posteriori (MAP) estimation [12], which perform only local optimizations. Experimental evaluation of our method on 20 kidneys from both contrasted and non-contrasted CT datasets yield an average volumetric overlap error of 10.95% (std: 3.7%).

2 Method

The input of our method is a set of training images $\{I_i\}$ with their corresponding segmentations $\{M_i\}$,$i \in \{1 \dots N\}$ and the image I to segment. The output is the kidney segmentation M. The EM approach is applied as follows:

1. Initially, each training dataset is registered to the current image.
2. **E-step**: Non-parametric shape and intensity models are computed with adaptive weighting of each training dataset.
3. **M-step**: The kidney segmentation is computed with the graph min-cut technique with edges weights derived from the kidney model.

The kidney model is refined with respect to the current image by iterating over steps 2 and 3 until convergence. Next, we formalize our framework and describe each step of the algorithm in detail.

The graph min-cut approach [4] defines segmentation as a voxel labeling problem. The optimal labeling is defined as the Maximum A-Posteriori (MAP) estimation of a Markov Random Field (MRF) M that minimizes an energy function $E(M)$ incorporating both a prior model and voxel neighborhood information:

$$E(M) = \sum_{\boldsymbol{x}} \left(\phi(\boldsymbol{x}|m(\boldsymbol{x})) + \sum_{\boldsymbol{y}} \psi(m(\boldsymbol{x}), m(\boldsymbol{y})) \right) \tag{1}$$

where $\boldsymbol{x}$ is the voxel coordinates vector, $\boldsymbol{y}$ is a neighboring voxel of $\boldsymbol{x}$, and $m(\boldsymbol{x}) \in \{m_{Obj}, m_{Bkg}\}$ are the object and background voxel labels, respectively. The prior model likelihood term $\phi(\boldsymbol{x}|m(\boldsymbol{x}))$ is defined as:

$$\phi(\boldsymbol{x}|m(\boldsymbol{x})) = \begin{cases} -log(p(\boldsymbol{x} \in \text{Obj}|\Theta_{Obj})) & \text{if } m(\boldsymbol{x}) = m_{Obj}, \\ -log(p(\boldsymbol{x} \in \text{Bkg}|\Theta_{Bkg})) & \text{if } m(\boldsymbol{x}) = m_{Bkg}. \end{cases} \tag{2}$$

where $p(\boldsymbol{x} \in \text{Obj}|\Theta_{Obj})$ is the likelihood of voxel $\boldsymbol{x}$ to belong to the object or background classes given a prior model Θ. In the original graph min-cut approach [4], this likelihood term is computed using a fixed prior intensity model, which is usually acquired via user interaction. The voxel neighborhood term $\psi(m(\boldsymbol{x}), m(\boldsymbol{y}))$ penalizes neighboring voxels with different labels. It is usually proportional to a contrast-based term that reduces the penalty when neighboring voxels have large intensity differences.

Instead of using a fixed intensity model, we define the prior model Θ as a combination of both shape and intensity information. This prior model is pre-computed from the training datasets. We refine the initial estimate Θ by considering it to be a latent variable in the energy functional, which now takes the form of $E(M, \Theta)$. An EM approach is then used to simultaneously estimate the unknown model Θ and to find the segmentation M [10,11]. We describe the EM steps and the iterative process in detail next.

2.1 E-Step: Non-parametric Hybrid Model Estimation

Given a set of training images $\{I_i\}$ with their corresponding segmentations $\{M_i\}$,$i \in \{1 \ldots N\}$ and a new image to segment I, our goal is to construct a prior model Θ that combines both shape S_{Obj} and intensity I_{Obj} information.

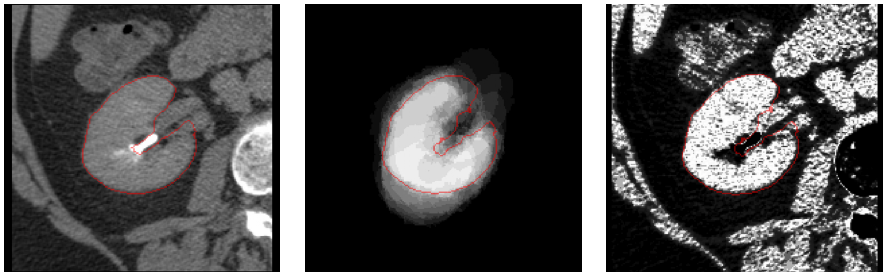

(a) Original CTA (b) Initial shape model (c) Initial likelihood map

Fig. 1. Model components. The red contour denotes the ground-truth segmentation. (a) Axial slice from the kidney CT dataset to segment, (b) The shape model, and; (c) Intensity based Likelihood image. Bright voxels represent high probability to belong to the object class based on the shape/intensity information. Note that using each component solely, does not provide enough information to obtain accurate segmentation. The presented model components were taken from a first iteration of the algorithm.

The shape information is computed from the object segmentation in the training datasets $\{I_i\}$. The probability of voxel $\boldsymbol{x}$ to belong to the object class based on the shape model is defined as:

$$p(\boldsymbol{x}|S_{Obj}) = \frac{1}{N}\sum_{i=1}^{N}\alpha_i\chi(\boldsymbol{x}, M_i) \tag{3}$$

where N is the number of training datasets, α_i is a weighting parameter, and the function χ is:

$$\chi(\boldsymbol{x}, M_i) = \begin{cases} 1 & \text{if } M_i(\Phi_i(\boldsymbol{x})) \in Obj, \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

where Φ_i is a geometrical transformation that maps training dataset I_i to the current image I. The transformations Φ_i are computed using an intensity-based B-Spline registration algorithm [14]. The parameter α_i indicates how close is dataset i to the current patient image I. In the first iteration, α_i is proportional to the Mutual Information (MI) between $M_i(\Phi_i)$ and I. Subsequently, α_i is proportional to the Dice coefficient between current segmentation and the training segmentation M_i.

The intensity model is defined as a patient-specific non-parametric estimation of the Intensity Probability Distribution Function (IPDF), which is estimated using intensity histogram of the voxels in current image I, that belong to the object. In the first iteration, only voxels with high confidence to belong to the object based on the shape model are included in the histogram computation. In subsequent iterations, all the voxels that were labeled as object in the previous iteration are used. The intensity-based background likelihood is set to be its complement, i.e., $1 - p(I(\boldsymbol{x})|Obj)$. Fig. 1 illustrates the different components of our model.

2.2 M-Step: Graph min-cut MAP-MRF Optimization

Given an estimation of the prior model Θ, the goal now is to compute the MAP-MRF that best model the given image I. The graph min-cut formulation for MAP-MRF estimation is as follows.

Let $G = (V, E)$ be the image graph. Graph nodes $V = \{v_1, \dots v_n, v_s, v_t\}$ are defined such that node $v_{\boldsymbol{x}}$ corresponds to voxel $\boldsymbol{x}$ and terminal nodes v_s and v_t correspond to the object and background classes. Graph edges $E = \{(v_{\boldsymbol{x}}, v_s), (v_{\boldsymbol{x}}, v_t), (v_{\boldsymbol{x}}, v_{\boldsymbol{y}})\}$ consist of three groups: 1) edges $(v_{\boldsymbol{x}}, v_s)$ from voxels to the object terminal node; 2) edges $(v_{\boldsymbol{y}}, v_t)$ from voxels to the background terminal node, and; 3) edges $(v_{\boldsymbol{x}}, v_{\boldsymbol{y}})$ between adjacent voxels. The cost of a cut that divides the graph into the object class and the background class is defined as the sum of the cut edges' weights. The segmentation is the bipartite graph partition that minimizes the cut cost.

Edge weights are assigned as follows. Edge weights $w(v_{\boldsymbol{x}}, v_s)$ represent the likelihood of voxel $\boldsymbol{x}$ to belong to the kidney (object) based on hybrid intensity and geometric model Θ:

$$w(v_{\boldsymbol{x}}, v_s) = -\log(p(\boldsymbol{x} \in \text{Obj}|\Theta_{Obj})) = -\log(p(I(\boldsymbol{x})|I_{Obj}) \cdot p(\boldsymbol{x}|S_{Obj})) \tag{5}$$

Edge weights $w(v_{\boldsymbol{x}}, v_t)$ represent the likelihood of each voxel to belong to the background class. We define it as the complement of the object prior model:

$$w(v_{\boldsymbol{x}}, v_t) = -\log(1 - (p(\boldsymbol{x} \in \text{Obj}|\Theta_{Obj}))) \tag{6}$$

Edge weights $w(v_{\boldsymbol{x}}, v_{\boldsymbol{y}})$ penalize for nearby voxels $\boldsymbol{x}, \boldsymbol{y}$ that have different labels. This edge weight is a combination of the local intensity difference and the spatial location of the edge with respect to the prior shape model S_{Obj}:

$$w(v_{\boldsymbol{x}}, v_{\boldsymbol{y}}) = \exp\left(-\frac{(I(\boldsymbol{x}) - I(\boldsymbol{y}))^2}{\sigma}\right) \cdot \psi(m(\boldsymbol{x}), m(\boldsymbol{y})|S_{Obj}) \tag{7}$$

where σ is a normalization constant that represents the standard deviation of the intensity values inside the object, and $\psi(m(\boldsymbol{x}), m(\boldsymbol{y})|S_{Obj})$ is computed using Eq. 3. The minimal cut of the graph is then computed as described in [4].

2.3 The Iterative Solution

The iterative process starts with the computation of the pairwise registration transformations Φ_i that map each training dataset to the current image domain. Since the pairwise registrations are independent, they can be computed in parallel by spanning simultaneous registration processes on computer clusters [15]. The method described in [16] can be used as an alternative, to reduce the number of required registrations. Although this initialization step may be computationally intensive, it prevents the accumulation of individual registrations errors that pervade in existing parametric methods [12].

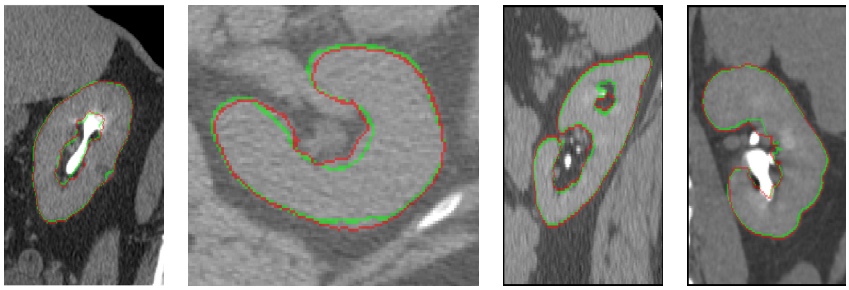

(a) (b) (c) (d)

Fig. 2. Representative results from four kidneys with comparison to the ground truth. Our algorithm result (red contour) and the ground truth segmentation (green contour) are overlaid on the CT slices. (a) Coronal view of a left kidney from CTA, (b) Axial view of a right kidney from CT, (c) Sagittal view of a right kidney from CTA, and (d) Coronal view of a right kidney from CTA. Additional images and 3D movies are available on: `http://www.cs.huji.ac.il/~freiman/kidney_seg`

In the E-step, the shape and intensity model Θ is computed. The shape model is first computed with Eq. 3 and then the intensity model is computed as described in Sec. 2.1. In the M-step, the updated estimation of the segmentation M is computed based on these models. Both steps are iteratively repeated until Θ and M remain stable.

3 Experimental Results

We evaluated our method on 20 kidneys from 10 CT datasets of size $512 \times 512 \times 350 - 500$ voxels, $0.5 - 1.0 \times 0.5 - 1.0 \times 1.0 - 1.5\text{mm}^3$, with and without contrast agent administration. The datasets were acquired on a 64-row CT scanner (Brilliance 64 - Phillips Healthcare, Cleveland, OH) and were chosen randomly from the hospital archive to represent wide variety of patients with different ages and pathologies. Three different observers annotated both the left and right kidney on each dataset. The STAPLE algorithm [17] was used to estimate the ground-truth from the three manual segmentations. In addition we measured the intra-observer performance with respect to the estimated ground-truth.

We evaluated our method using the Leave-One-Out (LOO) approach, where all datasets except the one tested were used for training. Fig. 2 presents representative results of kidney segmentation.

Both volumetric and surface based measures were computed. Table 1 summarizes the average (std) results for the evaluated metrics as compared to the human observers performance. Our method yields an average volumetric overlap error of 10.95% (std: 3.7%) which is better than the previously published error of 17% for semi-automatic kidney segmentation from CT images [2] . The average volumetric overlap error of the observers was 4.67% (std: 0.28%). The average running time using computer cluster to perform the registrations in parallel was 13:54 min (std: 0:27 min). Although our automatic method did not achieve

Table 1. Comparison metrics for kidney segmentation. The first column is the observer number. The second column is the Absolute Volume Difference (AVD) from the ground-truth in %. The third column is the Average Symmetric Surface Distance (ASSD) in mm. The fourth is the Root Mean Square Symmetric Surface Distance (RMS SSD) in mm. The fifth column is the Maximal Symmetric Surface Distance (MSSD) in mm. The sixth column is the Volumetric Overlap Error (VOE) in %. The first row presents the results of our method. The additional three rows present the observers performance compared to the ground-truth.

	AVD (%)		**ASSD (mm)**		**RMS (mm)**		**MSSD (mm)**		**VOE (%)**	
	mean	std	mean	std	mean	std	mean	std	mean	std
Our	6.2	4.09	0.79	0.34	1.52	0.73	10.46	3.96	10.95	3.7
Observer 1	2.72	2.53	0.32	0.18	0.72	0.28	6.76	2.19	4.88	2.11
Observer 2	2.37	2.13	0.27	0.17	0.66	0.38	5.87	3.31	4.36	2.15
Observer 3	2.45	2.11	0.3	0.18	0.71	0.39	6.29	3.35	4.78	2.39

the human performance, it produces clinically acceptable results for volumetric measurements and surgical planning with no need for user interaction.

4 Conclusion

We have presented a new iterative non-parametric model-based graph min-cut approach for kidney segmentation in CT images. Given a set of training images with their segmentations, our method computes the kidney segmentation based on MAP-MRF estimation of the current image. The algorithm treats the model parameters as a latent variable in a discrete energy functional. An Expectation-Maximization approach is then used to iteratively estimate the model and to obtain the MAP-MRF estimation using the graph min-cut technique. We evaluated the performance of our method on 20 kidneys. Our results show that the proposed method is accurate, robust, easy to use, and provides relevant clinical measurements for many applications. In the future we plan to apply the proposed method to other organ segmentations from various imaging modalities.

The datasets that were used in this work are now publicly available for further evaluation and comparison with other algorithms on: `http://www.cs.huji.ac.il/~caslab/kidneyEval`.

Acknowledgment

This research is supported in part by MAGNETON grant 38652 from the Israeli Ministry of Trade and Industry. M. Freiman is also supported in part by the Hebrew University Hoffman Leadership and Responsibility Fellowship Program.

References

1. Herts, B.R., et al.: Estimating glomerular filtration rate in kidney donors: a model constructed with renal volume measurements from donor CT scans. Radiology 252(1), 109–116 (2009)
2. Rao, M., et al.: Comparison of human and automatic segmentations of kidneys from CT images. Int. J. Rad. Onc. Bio. Phys. 61(3), 954–960 (2005)
3. Hidas, G., et al.: Estimating relative renal function from relative parenchymal volume-a feasibility study. J. Endourol. 22(11), 2527–2530 (2008)
4. Boykov, Y., Funka-Lea, G.: Graph Cuts and Efficient N-D Image Segmentation. Int. J. of Comp. Vis. 70(2), 109–131 (2006)
5. Shim, H., et al.: Semiautomated segmentation of kidney from high-resolution multidetector computed tomography images using a graph-cuts technique. J. Comput. Assist. Tomogr. 33(6), 893–901 (2009)
6. Slabaugh, G., Unal, G.: Graph cuts segmentation using an elliptical shape prior. In: IEEE ICIP 2005, vol. 2, pp. 1222–1225 (2005)
7. Freedman, D., Zhang, T.: Interactive graph cut based segmentation with shape priors. In: IEEE CVPR 2005, vol. 1, pp. 755–762 (2005)
8. Ali, A., Farag, A., El-Baz, A.: Graph Cuts Framework for Kidney Segmentation with Prior Shape Constraints. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part I. LNCS, vol. 4791, pp. 384–392. Springer, Heidelberg (2007)
9. Freiman, M., et al.: Vessels-Cut: a graph based approach to carotid arteries patient-specific modeling. In: Magnenat-Thalmann, N. (ed.) Modelling the Physiological Human. LNCS, vol. 5903, pp. 1–12. Springer, Heidelberg (2009)
10. Kumar, M., Torr, P., Zisserman, A.: OBJ CUT. In: IEEE CVPR 2005, pp. 18–25 (2005)
11. Malcolm, J., Rathi, Y., Tannenbaum, A.: Graph Cut Segmentation with Nonlinear Shape Priors. In: IEEE ICIP 2007, pp. 365–368 (2007)
12. Sabuncu, M.R., et al.: Supervised Nonparametric Image Parcellation. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) MICCAI 2009. LNCS, vol. 5762, pp. 1075–1083. Springer, Heidelberg (2009)
13. Wimmer, A., Soza, G., Hornegger, J.: A Generic Probabilistic Active Shape Model for Organ Segmentation. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) MICCAI 2009. LNCS, vol. 5762, pp. 26–33. Springer, Heidelberg (2009)
14. Klein, S., et al.: elastix: a toolbox for intensity based medical image registration. IEEE Trans. Med. Imaging 29(1), 196–205 (2010), `http://elastix.isi.uu.nl/`
15. Barak, A., Shiloh, A.: The MOSIX2 Management System for Linux Clusters, Multi-Clusters and Clouds, `http://www.mosix.org/pub/MOSIX2_wp.pdf`
16. Commowick, O., Malandain, G.: Efficient Selection of the Most Similar Image in a Database for Critical Structures Segmentation. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part II. LNCS, vol. 4792, pp. 203–210. Springer, Heidelberg (2007)
17. Warfield, S., Zou, K., Wells, W.: Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation. IEEE Trans. Med. Imaging. 23(7), 903–921 (2004)