A Simultaneous Segmentation and Regularization Framework for Vessel Extraction in CT Images

Gang Song[1], Alonso Ramirez-Manzanares[1,2] and James C. Gee[1]

[1] Penn Image Computing and Science Laboratory (PICSL), Department of Radiology, University of Pennsylvania, Philadelphia, PA 19104, USA
http://www.picsl.upenn.edu
[2] Univ. de Guanajuato, Facultad de Matematicas, Callejon Jalisco S/N, Valenciana, Guanajuato, Gto. Mexico. C.P. 36240

Abstract. A new framework is proposed for segmenting the pulmonary vessel tree while simultaneously estimating vessel orientations from lung Computed Tomography (CT) images. The problem is formulated as a joint optimization for both the segmentation and the orientation of the vessel tree. We propose to use a histogram vector to describe vessel orientation, which avoids explicit classification of branching points from a vessel tree. The objective function encodes the orientation information by defining a neighboring relationship between voxels, and is solved iteratively by alternately optimizing the segmentation and regularizing the orientation. The validation on manually labeled datasets suggests the potential value of our algorithm.

1 Introduction

Pulmonary vessel extraction is an important step in performing a quantitative analysis for lung CT images. Blood vessels inside lungs have a salient tree structure over the entire lung. The radii of vessels get decrease as they extend from the center to the periphery of the lung. The vessels typically have a higher intensity than the lung parenchyma. However, it is not sufficient to classify the vessels from the rest of the lung volume only by using the voxel intensities. The imaging process adds noise to the data volume and under current limitations of scanning resolution, the vessel structure in the images may become disconnected. These factors make vessel extraction a challenging problem.

In order to achieve a robust vessel segmentation, researchers have proposed various approaches. One is to preprocess the volume by enhancing the vesselness. For example, Shikata et al ([1]) first enhanced the vessels by a line-filter based on the Hessian matrix, after which the vessel tree was segmented by thresholding the vesselness and filling local gaps. Another method turns the problem into a tracking problem on the vessel tree after getting the initial segmentation using some heuristic thresholding. In [2], the initial estimation of the vessels was obtained from maximum intensity projection, and a rule-based scheme was adopted

to separate the vessel tree by extracting the center lines. Wu et al ([3]) used the method in a more principled way. They proposed a regulated morphology approach to produce a set of fuzzy spheres; a tracking algorithm subsequently generated a set of connected trees based on constraints such as for instance collinearity and size. There are also level-set based approaches such as those in [4][5], which optimize the vessel boundary as the zero level set of a deformable curve model.

In terms of optimization, the variable that all these methods try to optimize is only the segmentation label. Other information, especially the orientation of the vessel tree, is fixed after precomputation from either the Hessian matrix ([1]) or the morphology ([2][3]). In contrast to these previous approaches, we exploit the orientation information in a new way. An accurate estimation of orientation can provide better tracking in segmentation. Meanwhile, a good segmentation eliminates noise and improves the orientation information. In this paper we propose a new method of vessel extraction to compute both the segmentation and the orientation simultaneously.

In our framework, the orientation is regarded as a variable to be optimized rather than a precomputed feature for segmentation. We formulate our approach as an optimization problem and propose an iterative solution. First, a new description of orientation is applied for vessel extraction. We employ a histogram to accommodate multiple directions at tree branching points. This is different from using the direction vector, which is not suitable for describing bifurcation; for instance, in previously reported methods ([1][2]) special tracking rules have to be defined for the branching points. To initialize the orientation response in multiple directions, we design a bank of elongated second-order filters, which can detect more than one dominant direction.

Second, we propose a new formulation to alternately optimize both the orientation information and the segmentation label. We apply the graph cut method in segmentation, which has become a popular method for medical image segmentation ([6][7][8]). This method guarantees the global optimal solution and is computationally efficient. It can model both the data likelihood and the neighboring relationship between two voxels. We introduce a novel term defined in terms of the orientation of neighboring voxels. This leads us to a new way to integrate the tracking and segmentation procedures.

This paper is organized as follows: Section 2 presents our simultaneous segmentation and regularization framework. Section 3 shows both quantitative and qualitative results on High-Resolution Computed Tomography ($HRCT$) ([9]) data. Finally, we give our conclusion in Section 4.

2 Method

The goal of our method is two-fold: to segment the pulmonary vascular vessels, and to estimate the vessel orientation at each voxel. We propose a new framework for optimizing both the segmentation label and the orientation estimation. Given the image volume I, for each voxel i with coordinates $\boldsymbol{s}_i$, we want to assign a

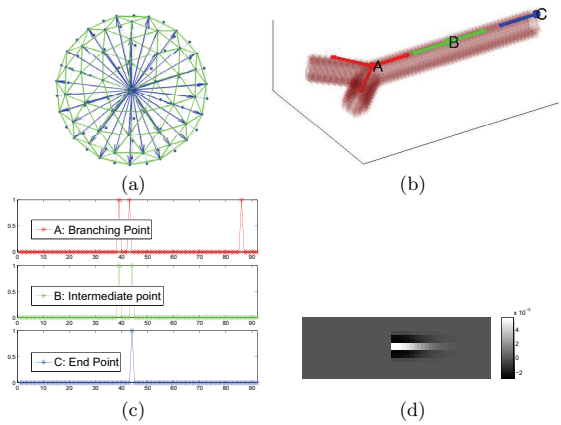

Fig. 1. (a) The 92 predefined directions ψ_k illustrated on a 3D sphere. (b) Three types of points on a Y-junction. (c) The orientation vectors o of for A) branching points, B) intermediate points and C) end points. (d) 2D view of the orientation filter (Equ. (2)).

label l_i indicating whether it belongs to the vessel tree ($l_i = 1$) or the background ($l_i = 0$). Moreover we want to estimate the orientation of the vessel voxel at $\boldsymbol{s}_i$.

The orientation is described by a histogram vector $\boldsymbol{o}_i = \langle o_i^1, \ldots, o_i^K \rangle$. Each bin in the histogram corresponds to one of the predefined K unit normal vectors, $\{\boldsymbol{\psi}_1, \ldots, \boldsymbol{\psi}_K\}$, $\|\boldsymbol{\psi}_k\| = 1$. Figure 1(a) shows the 92 directions used in the paper. The value of bin o_i^k is a continuous value in the range of $[0, 1]$, representing the likelihood to the k-th direction of the vessel. Such a histogram vector is capable of describing multiple dominant vessel directions at the same location; thus it can generalize the description of the three types of points in the vascular tree: the branching points with three non-zero bins, the intermediate points with two, and the end points with one. In general $\boldsymbol{o}_i$ is a sparse vector $\in [0, 1]^K$ (see Figure 1(b)(c) for an illustration on a Y-junction).

Such a vector description is different from the one used by most previous methods ([1][2]), in which the orientation is only defined on the intermediate points, and the branching points need to be classified differently from other points in order to initiate vessel tracking. In contrast, our histogram representation is capable of describing orientation information in all parts of the vascular tree with the same format. Therefore, there is no need for classifying and tracking branching points like in [3]. This description is more robust to local noise and broken vessels.

The data term, D_i^S, denotes the negative log-likelihood of assigning label l_i to the pixel i by the intensity at site $\boldsymbol{s}_i$. $W_{i,j}^S$ is the neighboring connectivity

term referring to the compatibility between i and j. $\mathcal{N}$ is the second order spatial neighborhood of 26-connectivity. D_i^R and $W_{i,j}^R$ are the terms for orientation. W^S and W^R ensure that the segmentation labels and the orientation of the vessels change smoothly along the tree. Our algorithm computes the minimum of the objective function:

$$C(L, \boldsymbol{O}) = \sum_{i \in I} D_i^S + \lambda_S^W \sum_{i \in I} \sum_{l_j \neq l_i | j \in \mathcal{N}_i} W_{i,j}^S + \sum_{i \in I} D_i^R + \lambda_R^W \sum_{i | l_i = 1} \sum_{l_j = 1 | j \in \mathcal{N}_i} W_{i,j}^R. \quad (1)$$

Each term is a function of the segmentation label field $L = \{l_i\}$, and the orientation field $\boldsymbol{O} = \{\boldsymbol{o_i}\}$. The first half of the formulation, $\sum_i D_i^S + \sum_i \sum_j W_{i,j}^S$, is the popular segmentation energy function, which can be efficiently optimized by the min-cut/max-flow algorithm ([6]). In contrast, the variable $\boldsymbol{o}_i$ in the second half, $\sum_i D_i^R + \sum_i \sum_j W_{i,j}^R$, is a continuous variable, which leads us to a different regularization approach as explained in Section 2.3.

2.1 Initial Estimation of Orientation

To locate the multiple dominant directions, we apply a bank of directional filters tuned along the set of directions $\{\boldsymbol{\psi}_k\}$. This is different from using the Hessian matrix ([1]), which is suitable for only one dominant direction. To detect thin structures, we generalize the elongated second-order derivative filter in [10][11] to 3D as:

$$F(\alpha, \beta, \gamma) = \delta_{\tilde{x} \geq 0} \frac{1}{C} \frac{\partial^2}{\partial \tilde{y}^2} \exp\left(\frac{\tilde{x}^2}{\lambda_x \sigma^2} + \frac{\tilde{y}^2}{\sigma^2} + \frac{\tilde{z}^2}{\lambda_z \sigma^2}\right)\Bigg|_{(\tilde{x}, \tilde{y}, \tilde{z})^T = \mathbf{R}(\alpha, \beta, \gamma)(x, y, z)^T}. \quad (2)$$

The direction of the filter is controlled by the Euler angle (α, β, γ). The predefined direction $\boldsymbol{\psi}_k$ determines α and β. The third Euler angle γ comes from n evenly distributed angles from 0 to π. We empirically choose $n = 4$ as a tradeoff between accuracy and running time. $\delta_{\tilde{x} \geq 0}$ makes the filter respond only to the forward direction. σ controls the scale of the filter, λ_x controls the elongated scale along x-axis, and λ_z controls the thickness of the thin structure (see Figure 1(d)). We set $\sigma = 1$, $\lambda_x = 16$ and $\lambda_z = 1$ in our experiments. $\mathbf{R}$ is the rotation matrix defined by (α, β, γ). By rotating the filter in 3D we get a bank of filters $\{F(\alpha, \beta, \gamma)\}$ for different directions. So an image volume needs to be convoluted with $K \times n$ filters. The initial estimation of the orientation bin $\hat{\boldsymbol{o}}^k$ is the maximum response out of the n filters of $\boldsymbol{\psi}_k$, normalized to the range $[0, 1]$ divided by the maximum response of the whole volume.

2.2 Segmenting the Frontier Band

The vascular vessels are typically brighter than the rest of the lung region. However, in the regions towards the boundary of the lung, the vessels become as thin as only one or two voxels. Small motion blur and imaging noise make it difficult to identify these vessels. In contrast to voxel intensity, the orientation information is relatively more robust because the orientation filter typically covers a

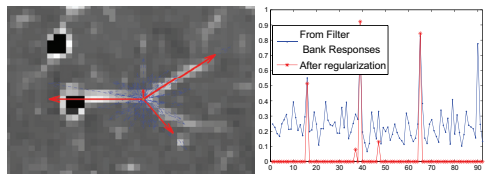

(a) (b)

Fig. 2. (a) A patch from HRCT lung volume. *Blue arrows*: the initial orientation $\hat{o}_i$ estimated from convolution with the filter bank. *Red arrows*: the orientation after regularization, o_i. (b) The histogram vector of $\hat{o}_i$ and o_i.

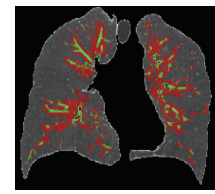

Fig. 3. The active band (red region) and the segmentation (green region) for the first iteration.

large neighborhood. We introduce the orientation into the connectivity term as follows:

$$\tilde{W}^S_{i,j} = \sum_{k=1}^{K} w^k_{ij} o^k_i, \quad w^k_{ij} = \begin{cases} \langle \boldsymbol{\psi}_k, \boldsymbol{s}_j - \boldsymbol{s}_i \rangle & \text{for} \quad \langle \boldsymbol{\psi}_k, \boldsymbol{s}_j - \boldsymbol{s}_i \rangle \geq t \\ 0 & \text{otherwise} \end{cases}, \tag{3}$$

where $\langle \cdot, \cdot \rangle$ is the vector dot product and t is a threshold ($t = 0.9$ for our experiments). In this definition, the spatial regularization weight w^k_{ij} is high if the voxel j is along any of the detected directions at voxel i. $W^S_{i,j}$ in Equ. (1) is defined to be symmetric: $W^S_{i,j} = \max(\tilde{W}^S_{i,j}, \tilde{W}^S_{j,i})$. We set $\lambda^W_S = 3/7$ in Equ. (1).

The data likelihood term is defined as $D^S_i = -\log \Pr(I_i|l_i)$. This likelihood is computed by the local Parzen-window nonparametric density estimation. A small neighborhood around each voxel is sampled by an isotropic Gaussian PDF. The implementation details for optimizing the segmentation labels L given Equ. (3) are illustrated in [8] by using the min-cut/max flow algorithm.

2.3 Regularization of Orientation Vectors

The initial estimation of vessel orientation depicted in Section 2.1 is very noisy. We desire a sparse solution for vector $\boldsymbol{o}_i$, in which each vector should have strong responses only in no more than three bins. Also the estimation should be spatially coherent. For this purpose we apply the multi-modal regularization framework introduced in [12]. The value of $\boldsymbol{o}_i$ is regularized from the initial estimation $\hat{\boldsymbol{o}}_i$. The data term in Equ. (1) is defined as:

$$D^R_i = \|\boldsymbol{o}_i - \hat{\boldsymbol{o}}_i\|^2 + \kappa(c\,\overline{o}^2_i - \sum_{k=1}^{K}(o^k_i)^2), \quad \text{in which } \overline{o}_i = \frac{1}{K}\sum_{k=1}^{K} o^k_i. \tag{4}$$

The first term restricts the regularization results $\boldsymbol{o}_i$ to be close to the input $\hat{\boldsymbol{o}}_i$. The second term enhances predominant orientations and attenuates the spurious ones, while c controls the number of non–zero peaks (see details in [12]).

We set $c = 6$ and $\kappa = 0.5$ in this paper to favor 1 to 3 non-zero different peaks. For the spatial regularization term, $W^R_{i,j}$ in Equ. (1) is defined as: $W^R_{i,j} = \sum_{k=1}^{K} w^k_{ij}(o^k_i - o^k_j)^2$. As normalized bin values in $\boldsymbol{o}_i$ are non-negative real-valued, we minimize this stage by solving a constrained system of linear equations. The details for optimizing the orientation $\{\boldsymbol{o}\}$ with the Gauss-Seidel algorithm are given [12]. We set λ^W_R to be 6. Figure 2(a)(b) shows an example of the regularization results.

2.4 Iterative Optimization of the Objective Function

We propose an iterative approach to alternately optimize the first and the second half of the objective function Equ. (1). In each iteration, we maintain a region called *Active Band* (*AB*). The segmentation and regularization is only applied sequentially inside this region, and the labels outside AB are fixed.

In the initialization stage, AB is obtained by thresholding the response of multiple orientation filters (see Figure 3). Next, orientation is regularized within AB. After the regularization, segmentation labels within AB are computed by the min-cut/max flow algorithm. At the end of each iteration, AB is updated by dilating current vessel tree labels (we use a dilation radius of 2).

The tubular-like structure of AB prevents the segmentation from leaking into the lung parenchyma. This improves computation efficiency without sacrificing accuracy. Such an idea is also adopted in [7]; but in contrast, we do not use Laplacian pyramids to obtain the active band. The convergence point of iteration gives a sub-optimal solution to the objective function of both L and $\boldsymbol{O}$.

3 Results

The two datasets were acquired from one patient with no known pulmonary disease. Using a volumetric expiratory HRCT protocol ([9]), the patient was scanned once at end-inspiration and again at maximal end-expiration with a 4-detector CT scanner (GE Lightspeed, 2.5 mm collimation, 120kVp, 240mA, 0.5s gantry rotation time, 15mm per rotation). Images were reconstructed to 1.25 mm-thick slices with a 512×512 matrix of 0.63-mm in-plane resolution and were further downsampled by half-resolution to 256×256×200. Before vessel extraction, a rough mask of the whole lung was estimated via semi-automatic level-set segmentation ([13]) to remove extra parts of the bones, heart, and large airways.

Three iterations were run for each volume to obtain convergence. Figure 4(a) shows the final vessel 3D mask on one volume. Most of the extracted vessels are thin (i.e. not over segmented) and connected without manual initialization. We evaluated our segmentation results using a similar approach to that used in [1]. About 3000 points uniformly distributed within vessels were manually labeled for each volume. For verification, the lung region is divided into five distinct regions (see Figure 4(b)) depending on the distance from two seed points manually placed on the border of the left and right lung around the hilum, where most vessels enter the lung region. The peripheral region E contains mostly thin

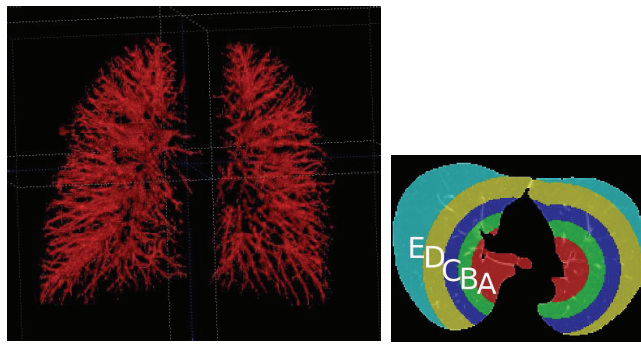

(a) (b)

Fig. 4. (a) The extracted 3D vessel mask from the 256×256×200 volume, with no leakage apparent from visual inspection. (b) The five divided regions according to the distances from the two seed points.

region	A	B	C	D	E	all regions
case 1	127/129	300/307	491/503	794/821	557/585	96.8%
case 2	321/327	384/393	544/564	554/577	201/215	96.5%
total	98.3%	97.7%	97.0%	96.4%	94.8%	96.7%

Table 1. Validation on different regions. The left in -/- is the number of points correctly extracted; the right is the number of points available for validation in each region. The *true negative* is mainly from thin vessels around the lung periphery due to the restriction of the resolution. Note that *false positive* ratios are, however, unable to be reported, as in [1], due to the absence of the ground truth of the entire vessel segmentation.

vessels, while thick vessels exist in region A. For each region we counted the number of correctly classified samples. Table 1 lists the true positive rates averaged over different regions and different volumes. While our results are comparable to those reported in [1] for regions A, B and C, our accuracy decreases from the hilum to the periphery of the lungs (region D and E), where the vessels are almost one pixel thin and get blurred by both imaging noise and low resolution.

Due to the absence of the ground truth of orientation, we only provide qualitative results for the estimated orientation. Figure 5(a) shows an example of the initial orientation estimated from the multi-orientation filter bank in a small region. Figure 5(b) is the final regularized orientation, which is more smooth and sparse, i.e. the noise was eliminated.

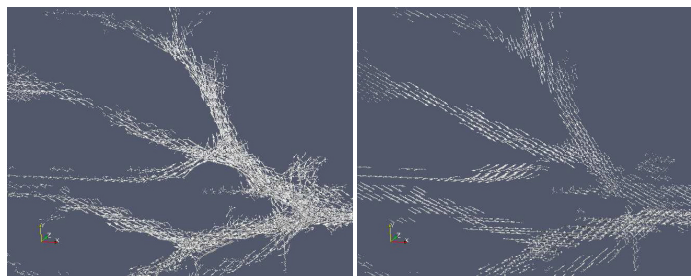

(a) (b)

Fig. 5. A partial region from the entire mask. (a) The initial orientation from the filters. (b) The final orientation, which is much more regularized and has less noise than (a).

4 Conclusion

In this paper a new method is proposed for simultaneously segmenting the pulmonary vessel trees and estimating the vessel orientation from lung CT images. The orientation information is represented by histogram vectors, which unifies the representation for branching, intermediate and end points in a vessel tree. The orientation is also a variable to be optimized in the objective function, in contrast to being fixed in existing methods such as those reported [1][3]. The approach has two phases for each iteration: in the first phase, the estimated orientation is regularized by fixing the segmentation mask; in the second phase, the segmentation mask is updated by a min-cut/max-flow algorithm.

Future plans include applications of our method for nodule and abnormality detection. Also we want to examine quantitative validation of orientation and connectivity, and compare them with existing methods in a large scale test. Regarding computation efficiency, current unoptimized implementation took about 40 minutes on the $256\times256\times200$ volumes for one iteration on a Intel Xeon 3GHz CPU. This can be improved, since the method, especially the regularization step, is highly parallelizable. And about 60% of the time is spent sampling the Gaussian PDF for D^S (see [8]), which can be reduced by a better sampling function.

Acknowledgment

This research was supported in part by CONACYT Mexico by a postdoctorate scholarship to Alonso Ramirez-Manzanares.

References

1. Shikata, H., Hoffman, E.A., Sonka, M.: Automated segmentation of pulmonary vascular tree from 3d ct images. Volume 5369., SPIE (2004) 107–116

2. Tozaki, T., Kawata, Y., Niki, N., Ohmatsu, H., Moriyama, N.: 3-d visualization of blood vessels and tumor using thin slice ct. In: IEEE Nuclear Science Symposium and Medical Imaging Conference. Volume 3. (1995) 1470–1474
3. Wu, C., Agam, G., Roy, A.S., Armato, III, S.G.: Regulated morphology approach to fuzzy shape analysis with application to blood vessel extraction in thoracic CT scans. In: Medical Imaging 2004: Image Processing. Volume 5370. (2004) 1262–1270
4. Lorigo, L.M., Faugeras, O.D., Grimson, W.E.L., Keriven, R., Kikinis, R., Nabavi, A., Westin, C.F.: Curves: Curve evolution for vessel segmentation. Medical Image Analysis **5** (2001) 195–206
5. Nain, D., Yezzi, A.J., Turk, G.: Vessel segmentation using a shape driven flow. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI). Volume 1. (2004) 51–59
6. Boykov, Y., Funka-Lea, G.: Graph cuts and efficient n-d image segmentation. International Journal of Computer Vision **70**(2) (2006) 109–131
7. Sinop, A.K., Grady, L.: Accurate banded graph cut segmentation of thin structures laplacian pyramids. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI). Volume 2. (2006) 896–903
8. Song, Z., Awate, S.P., Licht, D., Gee, J.C.: Clinical neonatal brain MRI segmentation using adaptive nonparametric data models and intensity-based Markov priors. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI). Volume 1. (2007) 883–90
9. Nishino, M., Boiselle, P., Copeland, J., Raptopoulos, V., Hatabu, H.: Value of volumetric data acquisition in expiratory high-resolution computed tomography of the lung. J Comput Assist Tomogr. **28**(2) (2004) 209–214
10. Morrone, M.C., Owens, R.A.: Feature detection from local energy. Pattern Recognition Letters **6**(5) (1987) 303–313
11. Leung, T.K., Malik, J.: Contour continuity in region based image segmentation. In: ECCV '98: Proceedings of the 5th European Conference on Computer Vision-Volume I, London, UK, Springer-Verlag (1998) 544–559
12. Ramirez-Manzanares, A., Rivera, M., Kornprobst, P., Lauze, F.: A variational approach for multi-valued velocity field estimation in transparent sequences. In: Proceedings of International Conference on Scale Space and Variational Methods in Computer Vision. LNCS 4485, Ischia, Italy (2007) 227–238
13. Yushkevich, P.A., Piven, J., Hazlett, H.C., Smith, R.G., Ho, S., Gee, J.C., Gerig, G.: User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. Neuroimage **31**(3) (2006) 1116–1128