

Segmentation of the Surfaces of the Retinal Layer from OCT Images

Mona Haeker^{1,2}, Michael Abramoff^{1,3}, Randy Kardon³, and Milan Sonka^{1,3}

¹ Department of Electrical and Computer Engineering,

² Department of Biomedical Engineering,

³ Department of Ophthalmology and Visual Sciences,

The University of Iowa, Iowa City, IA 52242, USA

{mona-haeker, michael-abramoff, randy-kardon, milan-sonka}@uiowa.edu

Abstract. We have developed a method for the automated segmentation of the internal limiting membrane and the pigment epithelium in 3-D OCT retinal images. Each surface was found as a minimum $s-t$ cut from a geometric graph constructed from edge/regional information and *a priori*-determined surface constraints. Our approach was tested on 18 3-D data sets (9 from patients with normal optic discs and 9 from patients with papilledema) obtained using a Stratus OCT-3 scanner. Qualitative analysis of surface detection correctness indicates that our method consistently found the correct surfaces and outperformed the proprietary algorithm used in the Stratus OCT-3 scanner. For example, for the internal limiting membrane, 4% of the 2-D scans had minor failures with no major failures using our approach, but 19% of the 2-D scans using the Stratus OCT-3 scanner had minor or complete failures.

1 Introduction

Retinal thickness measurements obtained from optical coherence tomography (OCT) provides clinically useful information for the diagnosis of diseases of the optic nerve head, such as glaucoma [1]. However, to date, very little has been reported on the automated analysis of such images [2]. Furthermore, although “volumetric” data is often obtained by performing multiple circular scans (typically six), current methods rely on detecting the retinal boundaries in each 2-D cross-sectional image individually [1, 2]. To our knowledge, nothing has been reported about the automated detection of the *surfaces* of the retinal layers. This work presents a method for automatically detecting the surfaces of the retinal layer in 3-D OCT images: the internal limiting membrane and the pigment epithelium, thereby allowing quantitative measurements of retinal thickness. In particular, our method finds each surface by finding a minimum $s-t$ cut in a 3-D geometric graph constructed from edge/regional information and *a priori*-determined surface constraints. The internal limiting membrane is found first and the pigment epithelium second, thereby allowing additional information to be used in the construction of the graph for detecting the pigment epithelial surface.

The process of transforming a multiple surface segmentation problem into that of finding a minimum s - t cut in a geometric graph has been previously reported by Li *et al.* [3]. It extends a previously reported method for detecting a single optimal surface by Wu and Chen [4] by adding additional edges to model interactions between the surfaces. While our method does not find the surfaces simultaneously as in [3], it imposes similar surface interaction constraints. One important advantage of using the minimum cost s - t cut surface detection method [3, 4] when compared to other previously-reported 3-D based surface segmentation methods [5–7] is that it guarantees to find the optimal solution with respect to the cost function.

2 Methods

2.1 The Sequential Multiple Surface Segmentation Problem

As OCT images were acquired as a sequence of circular scans (Fig. 1), the images naturally used a cylindrical coordinate system. Let $I(r, \theta, z)$ represent a 3-D OCT image, with corresponding image sizes R , T , and Z . R reflects the number of circular scans (typically 6), T reflects the number of angles used in each circular scan (typically 128), and Z reflects the number of points sampled at each (r, θ) point (typically 1024).

When using such a coordinate system, the surfaces of the retinal layer appear terrain-like. In particular, for each column of z -values corresponding to a specific (r, θ) pair, each surface intersects the column only once. Thus, each surface could be defined with a function $f(r, \theta)$, mapping (r, θ) pairs to z -values. Let $f_{ilm}(r, \theta)$ define the internal limiting membrane and $f_{rpe}(r, \theta)$ define the pigment epithelial surface of the retina. It then became the goal of our multiple surface segmentation method to determine $f_{ilm}(r, \theta)$ and $f_{rpe}(r, \theta)$.

Given a cost function $c(r, \theta, z)$ that specifies a surface unlikelihood measurement for each voxel, the cost of each surface was defined as the total cost of all voxels on the surface. An optimal surface was defined as the feasible surface for




Fig. 1. Each 3-D OCT image is composed of a number of 2-D circular scans, as schematically shown on the left. One example circular scan is shown on the right. The internal limiting membrane is the “upper” surface of the retinal layer (smaller z -values) and the pigment epithelium is the “lower” surface of the retinal layer (larger z -values).

which the surface cost was minimal. With a careful transformation of the costs of the voxels as described in [3] or [4] into the weights of a directed 3-D graph, an optimal surface directly corresponded to the minimum-cost (nonempty) closed set in the graph. Note that a closed set was a subset of the graph such that no directed edges left the set. The minimum-cost closed set in the transformed graph could be determined by finding the minimum s - t cut of the graph.

Our method used a sequential approach in detecting both surfaces. In particular, after finding the first surface using the single-surface detection method, we used the information from the location of the first surface to aid in creating the geometric graph used for finding the second surface (e.g., to define a more appropriate cost function and to limit the search range of the second surface).

2.2 Surface Feasibility

It was the structure of each graph (vertices and edges) that guaranteed that a feasible surface would be found. Two major types of surface constraints were imposed: (1) smoothness constraints for each individual surface (i.e., how much the radial and theta values were allowed to vary from column to column) and (2) surface interaction constraints (i.e., the minimum and maximum allowed distances between the surfaces). In detecting the internal limiting membrane, only smoothness constraints were imposed since this surface was determined first. Once the internal limiting membrane was found, the interaction constraints could be added for detection of the pigment epithelial surface by limiting the search range.

More specifically, both surfaces were defined to be feasible if they satisfied the following properties:

- Smoothness constraint: The difference in the z -value between the points on the surface in adjacent columns did not exceed a constant value. In particular, $|f(r, \theta + 1) - f(r, \theta)|$ was required to be less than or equal to $\Delta\theta$ for all (r, θ) and $|f(r + 1, \theta) - f(r, \theta)|$ was required to be less than or equal to Δr for all (r, θ) . Furthermore, since a larger physical distance was associated with each unit change of r , Δr was set to be greater than $\Delta\theta$. Note that because we were using cylindrical coordinates, columns with $\theta = T - 1$ were defined as being adjacent to columns with $\theta = 0$ (circularity constraint).
- Interaction constraint: The distance between the two surfaces was at least δ^l voxels and at most δ^u voxels (i.e., $\delta^l \leq f_{rpe}(r, \theta) - f_{ilm}(r, \theta) \leq \delta^u$ for all (r, θ)).

2.3 Cost Function

Whereas the structure of the graph defined the set of feasible surfaces, it was the cost function itself that was integral in defining which surfaces were finally identified. Intuitively, the cost functions in this work were based on the following *a priori* observable properties about the retinal layer: a strong dark-to-light transition occurred at the internal limiting membrane, a dark-to-light transition

occurred at the pigment epithelial surface, very few voxels above the internal limiting membrane were bright, and the voxels immediately following (larger z -value) each surface had a higher intensity, on average, than the voxels preceding each surface. Reflecting these properties, the cost functions $c(r, \theta, z)$ used in this work had two major components: an edge-based term $e(r, \theta, z)$ and a region-based term $v(r, \theta, z)$:

$$c(r, \theta, z) = \alpha e(r, \theta, z) + \beta v(r, \theta, z), \quad (1)$$

where α and β were constant normalizing terms. Furthermore, as the two surfaces were expected to have somewhat different properties, a different cost function was used for each.

Image pre-processing. Like ultrasound images, OCT images contain speckle noise. Thus, as a pre-processing step, a speckle-reducing anisotropic diffusion (SRAD) method [8] was first applied to each circular scan image before computing each cost function. This was in an attempt to reduce the effects of the speckle noise, while maintaining edge information.

Edge information. A 3x3 Sobel kernel was convolved with each pre-processed image, resulting in an estimate the strength of the edge in the z -direction for each voxel (giving positive values for light-to-dark transitions and negative values for dark-to-light transitions). Because we expected a dark-to-light transition to occur at both surfaces, we maintained the signed edge values in the cost function, thus favoring dark-to-light transitions and discouraging light-to-dark transitions.

Regional information. The regional cost terms were added to help aid in the detection of the correct surface even when edge-information was lacking. For the internal limiting membrane, we used a normalized cumulative image as our regional cost term. This was based on the observation that very few bright voxels exist above (have a lower z -value than) the internal limiting membrane in each column. Incorporating a cumulative image into the cost also had the effect of discouraging the selection of voxels in the pigment epithelial surface, a necessity since we were finding the internal limiting membrane first.

For the pigment epithelial surface, we incorporated a region-based term into the cost function that attempted to minimize the regional variances above and below the surface. As an example, suppose there are only two relatively-homogeneous regions in a 3-D image separated by a surface S with average intensities a_1 (from voxel locations (x, y, z) inside S) and a_2 (from voxel locations (x, y, z) outside S). Then a regional cost function term as presented by Chan and Vese [9] can be defined as

$$F(S, a_1, a_2) = \int_{inside(S)} |I(x, y, z) - a_1|^2 dx dy dz \quad (2) \\ + \int_{outside(S)} |I(x, y, z) - a_2|^2 dx dy dz.$$

In our work, we maintained the idea of minimizing the variance on either side of the surface, but we limited the computation to a region of limited size. This was because the assumption of having only two homogeneous regions in OCT images was only applicable to a limited area around each surface. Furthermore, because our cost function was voxel-based and not surface-based (the cost of a surface in this work was computed from a summation of voxel-based costs and thus could not include terms that depended on the overall surface properties), we added the variances of the voxels above and below each voxel for which we were computing the cost:

$$\widehat{v}(r, \theta, z) = \sum_{k=z-\gamma}^{z-1} (I(r, \theta, k) - \bar{I}_-)^2 + \sum_{k=z+1}^{z+\gamma} (I(r, \theta, k) - \bar{I}_+)^2, \quad (3)$$

where γ was a constant defining how far above and below the voxel we computed the variance, \bar{I}_- was the average intensity of voxels with smaller z -values than the voxel at position (r, θ, z) , and \bar{I}_+ was the average intensity of voxels with larger z -values than the voxel at position (r, θ, z) . However, since Eq. (3) would also favor voxels that were not of interest (e.g., background voxels), the actual regional cost term $v(r, \theta, z)$ would use the cost defined by Eq. (3) only if \bar{I}_+ were greater than \bar{I}_- :

$$v(r, \theta, z) = \begin{cases} \widehat{v}(r, \theta, z) & \text{if } \bar{I}_+ + \varepsilon > \bar{I}_- \\ c & \text{otherwise,} \end{cases} \quad (4)$$

where c was a relatively large constant, thereby giving a large cost to potential surface voxels that do not separate a high intensity region “below” (larger z -values) from a low intensity region “above” (smaller z -values) the voxel.

3 Experimental Methods

The method was tested on 18 3-D OCT data sets obtained from a Stratus OCT-3 scanner: 9 were from patients with normal optic discs and 9 were from patients with papilledema (optic nerve swelling). The dimensions of each data set were $6 \times 128 \times 1024$ ($R = 6$, $T = 128$, $Z = 1024$). The axial scan length was 2 mm, resulting in a voxel size of approximately $2 \mu\text{m}$ in the z -direction. The size of the voxels in the r and θ varied due to the radial nature of the scans. The following surface constraints were used:

- $\Delta\theta = 10$ pixels for both surfaces,
- $\Delta r = 150$ pixels for the internal limiting membrane, $\Delta r = 100$ pixels for the pigment epithelial surface,
- $\delta^l = 60$ pixels (minimum distance between surfaces), and
- $\delta^u = 600$ pixels (maximum distance between surfaces).

Each resulting surface was plotted on the six corresponding circular scans and its correctness assessed by an expert observer. In addition, the internal limiting

membrane resulting from our method was compared to that of the method used by Stratus OCT-3 system (Stratus OCT-3 results for the pigment epithelium were not available).

4 Results

Expert analysis of all the data sets indicated that our method found the correct surfaces. Fig. 2 displays an example of typical segmentation results. The corresponding retinal thickness plots are shown in Fig. 3.




Fig. 2. Typical results for one 3-D data set. The results are shown in increasing circular scan order such that the results for the inner circular scan are shown in (a), while the results for the outer circular scan are shown in (f). Note that for easier visualization of the borders, the top and bottom of the images have been cropped.

Table 1 shows a comparison of the “failures” of our method versus that of the Stratus OCT-3 for the internal limiting membrane. The counts given in the table refer to the number of 2-D circular scans on which some type of failure was found out of all of the circular scans for each scan group (54 normal scans, 54 papilledema scans, 108 total scans). A circular scan was considered to have a “minor failure” if the visible deviation of the result from the correct surface was small enough not to require correction in a clinical setting. A scan was considered to have a “complete failure” if the result had a large deviation that would require correction in order for the segmentation to be clinically useful. The Stratus OCT-3 system had two such failures, as shown in Fig. 4 (our results are also shown in the figure for comparison). Minor failures (4/108) of our method occurred in cases in which the image data did not satisfy our circularity assumption. Note that we consider our result in Fig. 4 (c) our worst “minor failure.” Each example is from an inner circular scan from a patient with papilledema.





Fig. 3. Color-coded retinal thickness plots corresponding to results shown in Fig. 2. (a) 3-D view of the internal limiting membrane with colors reflecting retinal thickness (in pixels). (b) 2-D view of retinal thickness values (in pixels).

Table 1. Count of circular scan failures in our method and the Stratus OCT-3 method for the internal limiting membrane

Scan group	Method	Minor failures	Complete failures	Total failures
Normal scans	3-D Graph Search	0 (0.0%)	0 (0.0%)	0 (0.0%)
Normal scans	Stratus OCT-3	4 (7.4%)	0 (0.0%)	4 (7.4%)
Papilledema scans	3-D Graph Search	4 (7.4%)	0 (0.0%)	4 (7.4%)
Papilledema scans	Stratus OCT-3	15 (27.8%)	2 (3.7%)	17 (31.5%)
All scans	3-D Graph Search	4 (3.7%)	0 (0.0%)	4 (3.7%)
All scans	Stratus OCT-3	19 (17.6%)	2 (1.8%)	21 (19.4%)



(a) Our method. (b) Stratus method. (c) Our method. (d) Stratus method.

Fig. 4. Examples of “complete failures” (arrows) for the Stratus OCT-3 method and our corresponding results for the internal limiting membrane. (a)–(b) Our method versus the Stratus OCT-3 method for an inner circular scan on a patient with papilledema, (b) is considered a “complete failure.” (c)–(d) Our method versus the Stratus OCT-3 method for an inner circular scan on a different patient with papilledema, (c) is considered a “minor failure” and (d) is considered a “complete failure.”

5 Discussion and Conclusion

The presented qualitative analysis of the results from all 18 3-D OCT data sets (108 total circular scans) indicates that our method identified the correct

border for the internal limiting membrane in 96% of the circular scans (93% in papilledema scans). This offers a substantial improvement over the Stratus OCT-3 system which only found the correct border 81% of the time (69% in papilledema scans). The 4/108 minor failures of our method (e.g., see Fig. 4 (c)) can be attributed to the invalidity of the circularity constraint in these cases.

We have presented a fully automated 3-D method for detecting the internal limiting membrane and pigment epithelial surface of the retina from OCT images that we hope to evaluate on clinical datasets.

References

1. Jaffe, G.J., Caprioli, J.: Optical coherence tomography to detect and manage retinal disease and glaucoma. *Am. J. Ophthalmol.* **137**(1) (2004) 156–169
2. Koozekanani, D., Boyer, K., Roberts, C.: Retinal thickness measurements from optical coherence tomography using a Markov boundary model. *IEEE Trans. Med. Imag.* **20**(9) (2001) 900–916
3. Li, K., Wu, X., Chen, D.Z., Sonka, M.: Optimal surface segmentation in volumetric images – a graph-theoretic approach. *IEEE Trans. Pattern Anal. Machine Intell.* **28**(1) (2006) 119–134
4. Wu, X., Chen, D.Z.: Optimal net surface problems with applications. In: *Proc. of the 29th International Colloquium on Automata, Languages, and Programming (ICALP)*, LNCS 2380, Springer-Verlag (2002) 1029–1042
5. Thedens, D.R., Skorton, D.J., Fleagle, S.R.: Methods of graph searching for border detection in image sequences with applications to cardiac magnetic resonance imaging. *IEEE Trans. Med. Imag.* **14**(1) (1995) 42–55
6. Suri, J.S., Liu, K., Singh, S., Laxinarayan, S.N., Zeng, X., Reden, L.: Shape recovery algorithms using level sets in 2-D/3-D medical imagery: A state-of-the-art review. *IEEE Trans. Inform. Technol. Biomed.* **6**(1) (2002) 8–28
7. McInerney, T., Terzopoulos, D.: Deformable models in medical image analysis: A survey. *Medical Image Analysis* **1**(2) (1996) 91–108
8. Yu, Y., Acton, S.T.: Speckle reducing anisotropic diffusion. *IEEE Trans. Image Processing* **11**(11) (2002) 1260–1270
9. Chan, T.F., Vese, L.A.: Active contours without edges. *IEEE Trans. Image Processing* **10**(2) (2001) 266–277