Published in final edited form as:
IEEE Trans Med Imaging. 2008 October ; 27(10): 1495–1505. doi:10.1109/TMI.2008.923966.

Intraretinal Layer Segmentation of Macular Optical Coherence Tomography Images Using Optimal 3-D Graph Search

Mona K. Garvin [Member, IEEE],
Department of Electrical and Computer Engineering and the Department of Biomedical Engineering, The University of Iowa, Iowa City, IA 52242 USA (e-mail: mona-garvin@uiowa.edu).

Michael D. Abràmoff [Member, IEEE],
Department of Ophthalmology and Visual Sciences and the Department of Electrical and Computer Engineering, The University of Iowa, Iowa City, IA 52242 USA, and also with the VA Medical Center, Iowa City, IA 52246 USA (e-mail: michael-abramoff@uiowa.edu).

Randy Kardon,
Department of Ophthalmology and Visual Sciences, The University of Iowa, Iowa City, IA52242 USA, and also with the VA Medical Center, Iowa City, IA 52246 USA (e-mail: randy-kardon@uiowa.edu).

Stephen R. Russell,
Department of Ophthalmology and Visual Sciences, The University of Iowa, Iowa City, IA 52242 USA (e-mail: steve-russell@uiowa.edu).

Xiaodong Wu [Senior Member, IEEE], and
Department of Electrical and Computer Engineering and the Department of Radiation Oncology, The University of Iowa, Iowa City, IA 52242 USA (e-mail: xiaodong-wu@uiowa.edu).

Milan Sonka [Fellow, IEEE]
Department of Electrical and Computer Engineering, the Department of Ophthalmology and Visual Sciences, and also with the Department of Radiation Oncology, The University of Iowa, Iowa City, IA 52242 USA (e-mail: milan-sonka@uiowa.edu).

Abstract

Current techniques for segmenting macular optical coherence tomography (OCT) images have been 2-D in nature. Furthermore, commercially available OCT systems have only focused on segmenting a single layer of the retina, even though each intraretinal layer may be affected differently by disease. We report an automated approach for segmenting (anisotropic) 3-D macular OCT scans into five layers. Each macular OCT dataset consisted of six linear radial scans centered at the fovea. The six surfaces defining the five layers were identified on each 3-D composite image by transforming the segmentation task into that of finding a minimum-cost closed set in a geometric graph constructed from edge/regional information and *a priori* determined surface smoothness and interaction constraints. The method was applied to the macular OCT scans of 12 patients (24 3-D composite image datasets) with unilateral anterior ischemic optic neuropathy (AION). Using the average of three experts' tracings as a reference standard resulted in an overall mean unsigned border positioning error of 6.1 ± 2.9 μm, a result comparable to the interobserver variability (6.9 ± 3.3 μm). Our quantitative analysis of the automated segmentation results from AION subject data revealed that the inner retinal layer thickness for the affected eye was 24.1 μm

Correspondence to: Mona K. Garvin.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

(21%) smaller on average than for the unaffected eye ($P < 0.001$), supporting the need for segmenting the layers separately.

Index Terms

Ophthalmology; optical coherence tomography; retina; segmentation; 3-D graph search

I. INTRODUCTION

Optical coherence tomography, first described in 1991 by Huang *et al.* [1], is increasingly being used in the diagnosis and management of a variety of ocular diseases such as glaucoma, diabetic macular edema, and age-related macular degeneration. As illustrated in Fig. 1, retinal time-domain optical coherence tomography (OCT) images are commonly acquired in the macular region or the peripapillary region (region near the optic disk). One common scanning protocol for acquiring scans in the macular region involves the acquisition of six linear radial scans in a "spoke pattern" centered at the fovea (e.g., the fast macular protocol on the commercially available Stratus OCT-3, Carl Zeiss Meditec, Inc., Dublin, CA). When acquiring scans surrounding the optic disk, it is common to use a number of circular scans. An example set of six radial images from a macular OCT series can be found in Fig. 2. Note that these protocols are reflective of the time-domain systems in current clinical usage; however, newer generation systems working in the Fourier-domain have just been released to the market, allowing the acquisition of substantially more data.

In the presence of ocular disease, each intraretinal layer may be affected differently. However, even though multiple layers of the retina are identifiable on OCT images, commercially available systems currently only segment and provide thickness measurements for one layer of the retina (i.e., the total retina on macular scans and the retinal nerve fiber layer on peripapillary scans). Although we do not know the proprietary details of the segmentation approach used, it is most likely a 2-D approach. Similarly, to the best of our knowledge, the reported approaches by others [2]–[9] for the segmentation of OCT images have been 2-D in nature (i.e., if multiple 2-D slices are available in a particular scanning sequence they are segmented independently). Thus, applying these approaches to the segmentation of 3-D images (e.g., by repeatedly applying the 2-D approach to each slice) does not take advantage of any 3-D contextual information. In fact, many of the reported approaches do not even take full advantage of 2-D contextual information as they first rely on finding intensity peaks from each preprocessed A-scan (column) separately [2]–[5], [7], [8]. While variations to each of the prior approaches exist for the segmentation of retinal boundaries, a typical approach proceeds as follows.

- Preprocess the image (e.g., with a median filter as in [2]–[5] or anisotropic diffusion filter as in [7]).
- Perform a 1-D peak detection algorithm on each A-scan (column) of the processed image to find points on each border of interest.
- For only a few approaches process the points further to correct for possible discontinuities in the 1-D border detection approaches (e.g., use Markov modeling to connect smaller segments to the largest continuous segment followed by spline-fit as in [2] and [3]).

Other prior approaches include the use of 2-D dynamic programming by Baroni *et al.* [9] and manually initialized deformable models for the segmentation of fluid-filled regions by Cabrera Fernández [6].

The prior segmentation approaches by others have attempted to find different numbers of boundaries of the retina. In particular, Koozekanani *et al.* [2], [3] found two, Baroni *et al.* [9] found three, Shahidi *et al.* found four [8], Ishikawa *et al.* [4], [5] found five, and Cabrera Fernández *et al.* found seven [7]. Because many authors reported limited quantitative validation (e.g., subjective failure rates), it is difficult to assess the robustness of the approaches. Nevertheless, it is clear that better segmentation approaches are needed for greater accuracy.

The specific purpose of this work was to develop an automated 3-D segmentation approach for the division of the retina on macular optical coherence tomography (OCT) scans into five layers and compare its performance on OCT scans of patients with unilateral anterior ischemic optic neuropathy to that of human experts. An early attempt at 3-D segmentation of the total retina (two bounding surfaces only) on peripapillary scans was presented in [10]. A preliminary version of this work outlining the method and giving a limited validation appeared in [11]. The work reported here provides additional studies in a set of subjects with unilateral anterior ischemic optic neuropathy (AION) and consists of three parts: 1) description of the method, 2) its validation in OCT slices for which manual tracings were available, and 3) employment of this validated segmentation approach to assess layer-specific differences in unilateral AION disease.

Fig. 3(b) shows an example of the six surfaces (labeled 1–6) we desired to find on each 3-D composite image. Based on histology and higher-resolution OCT images (from research scanners) published in the literature [12], we assumed the surfaces roughly had the following anatomical correspondence: surface 1 corresponded to the vitreo-retinal interface (VRI), surface 2 corresponded to the separation of the retinal nerve fiber layer (NFL) above from the ganglion cell layer (GCL) below, surface 3 corresponded to the separation of the inner plexiform layer (IPL) above from the inner nuclear layer (INL) below, surface 4 corresponded to separation of the outer plexiform layer (OPL) above from the outer nuclear layer (ONL) below, surface 5 corresponded to the junction between the photoreceptor inner and outer segments (IS/OS), and surface 6 corresponded to the separation of the photoreceptor outer segment (OS) from the retinal pigment epithelium (RPE). The corresponding five layers [labeled A–E in Fig. 3(b)] may well be associated with the following anatomical layers: A) NFL, B) GCL+IPL, C) INL+OPL, D) ONL+IS, E) OS. It is important to note that the actual segmentation was performed in 3-D. For example, Fig. 3(c) shows a 3-D visualization of surfaces 1, 3, and 4.

Our method found each surface (or set of surfaces) by transforming the 3-D segmentation problem into finding a minimum-cost closed set in a corresponding vertex-weighted geometric graph constructed from edge/regional image information and *a priori* surface smoothness and interaction constraints. This type of transformation for general 3-D multiple surface segmentation problems has been previously reported by Li *et al.* [13]. It extends a previously reported method for detecting a single optimal surface by Wu and Chen [14] by adding additional edges to model interactions between surfaces. One important advantage of using this surface detection method [13], [14] when compared to other previously reported 3-D based surface segmentation methods [15]–[17] is that it guarantees to find the three-dimensionally optimal solution with respect to the cost function.

II. METHODS

A. Overview

As was indicated in Fig. 1(a)–(c) and Fig. 2, one macular OCT image series (using the fast macular Stratus OCT-3 protocol) consisted of six radial linear cross-sectional scans centered at the fovea. For each eye, repeated series were acquired (six, if possible), so that up to six

raw scans existed at each angular location. The overall goal of the segmentation method was to determine the six surfaces defining the five retinal layers on a composite 3-D image derived from the repeated raw scans. There were two stages to the overall approach: 1) the creation of a composite 3-D macular image from the raw scans and 2) the determination of the six surfaces on the 3-D composite image. An overview of the data flow in the segmentation process can be found in Fig. 4.

1) Overview of Stage I: Creation of Each 3-D Composite Image—The 3-D composite image associated with each eye was created in two major steps. In the first step (Fig. 5), raw scans for a particular angular location (e.g., all the vertical scans) were individually aligned so that that boundary 6 (the retinal pigment epithelium) appeared approximately straight in the aligned image. The purpose of the alignment was twofold: to aid in the final 3-D segmentation and to allow for better visualization. The alignment was based on the assumption that the retinal pigment epithelium was undisturbed, and that the retinal profile underwent a low spatial-frequency distortion caused by the curvature of the eye and patient motion during acquisition. Each scan was aligned by first finding boundaries 1, 5, and 6 simultaneously using an optimal graph search approach similar to that used during stage II (described in more detail in later sections), but performed in 2-D. To ensure smoothness, a least-squares spline was fit to boundary 6. The columns were then translated so that this spline would be a straight line in the aligned image.

In the second step of this stage, each aligned image was registered to the first image in its location set by exhaustively searching for the best whole-pixel translation (according to the mutual information registration metric [18], but based on our experiments, other metrics such as normalized cross correlation would work as well) to align each of its columns to the corresponding target image column. The position of boundary 6 determined during the first step was used as a guide to determine the range of translations to be tested for each column. The registered images in each location set were averaged together to form the composite image for that particular angular location. The purpose of averaging the images was to obtain a representative scan of that location that had a higher signal-to-noise ratio than any of the raw scans. An example of an individual scan and the corresponding 2-D composite scan is shown in Fig. 6. The set of 2-D composite images (one for each angular location) formed the 3-D composite image used in the next stage.

2) Overview of Stage II: Segmentation of Each 3-D Composite Image—In the second stage, the six surfaces were found on the 3-D composite image. As a preprocessing step, a speckle-reducing anisotropic diffusion method [19] was applied (Fig. 7). Surfaces 1, 5, and 6 were then simultaneously found using an optimal graph search approach (transforming the segmentation problem into finding a minimum-cost closed set in a geometric 3-D graph [13]). After the determination of surfaces 1, 5, and 6, the remaining surfaces were found sequentially (allowing the utilization of other surface locations in the cost functions) in the following order: surface 4, surface 3, and finally, surface 2. The graph search approach guaranteed that the optimal feasible (satisfied smoothness and interaction constraints) surfaces would be found with respect to the designed cost functions.

As the focus of this paper is on this 3-D segmentation stage, more details of the graph search approach and cost functions used will be described in the next sections. In particular, Section II-B will provide a more precise definition of the surface segmentation problem (the optimization problem to solve), Section II-C will briefly describe how the graph search was used to solve such an optimization problem, and Section II-D will describe the used cost functions in more detail.

B. Surface Segmentation Problem

The nature of the macular scans [Fig. 1(a)–(c)] made it natural to use a discrete cylindrical coordinate system when working with each 3-D composite image (the z axis coincided with the intersection of the six 2-D composite scans). The coordinates of each voxel could thus be described with the triple (r, θ, z), where r reflected the distance of the voxel from the z axis, θ reflected the angular location of the voxel (0, 30, 60, 90, 120, 150, 180, 210, 240, 270, 300, or 330), and z reflected the row of the voxel in the corresponding 2-D image. Note that with this coordinate system, voxels in the left half of each 2-D image had a different θ value than those in the right half (for example, for the vertical 2-D scan shown in red in Fig. 1(a)–(c), voxels in the right half of the image had a θ value of 90 while those in the left half had θ value of 270).

Each surface could be defined with a function $f(r, \theta)$, mapping (r, θ) pairs to z values. Furthermore, a surface was considered feasible if it satisfied certain smoothness and surface interaction constraints. In particular, a surface was considered feasible if:

- θ smoothness constraint: $|f(r,\theta + 30) - f(r,\theta)|$ was less than or equal to Δ_θ for all (r,θ), $\theta \leq 300$;
- circularity constraint: $|f(r, 0) - f(r, 330)|$ was less than or equal to Δ_θ for all r;
- r smoothness constraint: $|f(r + \mathbf{1},\theta) - f(r,\theta)|$ was less than or equal to Δ_r for all (r,θ);
- constraint to connect the left and right halves of the 2-D scans together: $|f(0,\theta) - f(0,\theta + 180)|$ was less than or equal to Δ_r for (r,θ) pairs in which $0 \leq \theta \leq 150$;
- surface interaction constraint for each pair of surfaces f_1 and f_2: The distance between the two surfaces was at least δ^l voxels and at most δ^u voxels [i.e., $\delta^l \leq f_1(r,\theta) - f_2(r,\theta) \leq \delta^u$ for all (r,θ)].

In essence, the smoothness constraints required the z values of neighboring surface points (see Fig. 8) on a particular surface to be within a specified range (given by $\pm\Delta_\theta$ or $\pm\Delta_r$) and the surface interaction constraints required the surface z values for a particular surface to be within a specified range of the corresponding points on the other surfaces.

Given a cost function $c(r, \theta, z)$ that measures the unlikeliness that each voxel belongs on a particular surface, the cost of a surface was defined as the summation of all voxel costs on the surface. Similarly, the cost of a set of surfaces was defined as summation of all the surface costs in the set. Consequently, the goal of the single surface detection problem (as used for finding each of surfaces 4, 3, and 2 sequentially) was to find the *feasible* surface with the lowest cost. The goal of the multiple surface detection problem (as used for finding surfaces 1, 5, and 6 simultaneously) was to find the set of *feasible* surfaces with the lowest cost.

C. Solving the Surface Segmentation Problem Using 3-D Graph Search

Each single and multiple surface problem was transformed into that of finding a minimum-cost (nonempty) closed set in a corresponding vertex-weighted geometric graph as described in [13] and [14]. Briefly, this involved first constructing a 3-D geometric graph for each surface to be found (each node in this graph corresponded to a voxel in the original image). Edges were added to the graph to enforce the feasibility constraints by ensuring that there was a one-to-one correspondence between each feasible surface and a closed set in the graph. Note that a closed set was a subset of the vertices of the graph such that no directed edges left the set. In the case of detecting multiple surfaces simultaneously, additional edges were added between the 3-D geometric graphs to model the surface interaction constraints. These edges ensured that there was a one-to-one correspondence between each feasible

surface *set* and a closed set in the overall graph. Next, the vertex costs were assigned so that the cost of each surface set corresponded to the cost of the corresponding closed set (within a constant). This ensured that the global minimum-cost surface set could be found by finding the minimum-cost closed set in this graph. The minimum-cost closed set was then found by computing a minimum *s-t* cut in a closely related graph.

D. Cost Functions

Clearly the defined cost functions were an important component in determining the desired surfaces. In this work, the cost function for each surface was constructed from a linear combination of base "intuitive" cost function terms so as to satisfy expected properties of the surface. For example, it was expected that the first surface could be characterized by a combination of the following two properties: 1) the presence of an edge with a dark-to-light transition and 2) the lack of bright voxels above the surface. Correspondingly, the cost function for the first surface was defined as a normalized combination of a signed edge image (to favor the dark-to-light transition) and a cumulative image (created starting at the top of the image so as to discourage the detection of surfaces for which there were many bright pixels above the surface).

The cost functions for all of the surfaces followed the general pattern of having an edge-based term (to either favor a dark-tolight transition or a light-to-dark transition) and one or more regional- based terms (such as the cumulative image used in the cost function for surface 1). Depending on the prior knowledge of the locations of other surfaces, regional information used in this work generally was acquired from the locations illustrated in Fig. 9. Because both surrounding surfaces of each surface were often not known (surface 2 was the only surface for which the two surrounding surfaces were known) before designing its cost function, it was common to only use regional information from a limited region surrounding the surface [e.g., as in Fig. 9(a)–(b)].

More specifically, each of the surface cost functions was constructed from a normalized combination of a set of the following terms.

- Signed edge term (using Sobel kernel) favoring a dark-tolight transition (used for surfaces 1, 5, and 6).
- Signed edge term (using Sobel kernel) favoring a light-todark transition (used for surfaces 2, 3, and 4).
- Summation of pixel intensities in a limited region [Fig. 9(a)] above each potential surface voxel to encourage favoring surfaces with dark regions above surface (used for surfaces 5 and 6).
- Negated summation of pixel intensities in a limited region [Fig. 9(a)] above each potential surface voxel to encourage favoring surfaces with bright regions above surface (used for surface 2).
- Summation of pixel intensities in a limited region [Fig. 9(b)] below each potential surface voxel to encourage favoring surfaces with dark regions below (used for surface 3).
- Negated summation of pixel intensities in a limited region [Fig. 9(b)] below each potential surface voxel to encourage favoring surfaces with bright regions below (used for surfaces 5 and 6).
- Cumulative term acquired starting at the top of the image and accumulating downwards [Fig. 9(c)] to discourage finding surfaces with bright pixels above the surface (used for surface 1).

- Cumulative term acquired starting from the known boundary below and accumulating upwards [Fig. 9(d)] to discourage finding surfaces with bright pixels below the surface (used for surface 4).
- Chan-Vese [20] inspired term that attempted to minimize the intensity variances surrounding the surface. *A priori* estimated means of the two regions separated by the surface were computed from a region surrounding each known surface (as shown in Fig. 9(e) with the lighter intensity region indicated by a dashed line). Because the best use of this term required the prior location of the two surrounding surfaces, only surface 2 used this term.

III. EXPERIMENTAL METHODS

The algorithm was tested on fast macular scans from 12 subjects with unilateral chronic anterior ischemic optic neuropathy. Note that the unilateral nature of the disease meant that we had data for 24 eyes, 12 of which were affected by optic neuropathy, 12 of which were not. In almost all cases (21/24 eyes), six repeated series were used to create the 3-D composite image for each eye. (Each of the remaining three eyes used fewer than six repeated series to create the 3-D composite image.) The resulting 24 3-D composite images were each comprised of six composite 2-D scans (144 total composite 2-D scans) of size 128 × 1024 pixels. The corresponding reported physical width and height of the 2-D raw scans (and thus also the composite scans) was 6 mm × 2 mm, resulting in a pixel size of approximately 50 μm (horizontally) × 2 μm (vertically).

Based on our prior expectations of maximum change in surface shape, the smoothness constraints (in pixels) for use in the graph search were set as follows: $\Delta_r = 10$ for surface 1, $\Delta_r = 5$ for surface 2, $\Delta_r = 2$ for surface 3, $\Delta_r = 5$ for surface 4, $\Delta_r = 5$ for surface 5, $\Delta_r = 5$ for surface 6, and $\Delta_\theta = 10$ for all surfaces. The surface interaction constraints used in the simultaneous segmentation of surfaces 1, 5, and 6 were set as follows: $\delta^l = 50/\delta^u = 800$ for the allowed minimum/maximum distances between surfaces 1 and 5 and $\delta^l = 10/\delta^u = 50$ for the allowed minimum/maximum distances between surfaces 5 and 6. Surface interaction constraints for the sequential detection of each interior surface were incorporated by modifying the cost functions to have a high value outside a feasible region defined based on the location of the previously determined surfaces (e.g., the previously determined surfaces for the detection of surface 4 were surfaces 1 and 5).

For validation purposes, one raw scan from each eye was independently traced by three human experts with the average of the three tracings being used as the reference standard. Experts were masked as to the clinical status of the scanned eye. The experts did not attempt to trace borders that they did not consider visible. The algorithmic result on the corresponding composite 2-D scan was converted into the coordinate system of the raw scan (undoing alignment/registration) and the mean signed and unsigned border positioning errors for each border were computed (the middle 30 pixels were not included to exclude the fovea because not all layers were visible in this region). The signed and unsigned border positioning errors were also computed between the observers in the following ways:

- observer 1 versus observer 2;
- observer 2 versus observer 3;
- observer 1 versus observer 3;
- observer 1 versus the average of observers 2 and 3;
- observer 2 versus the average of observers 1 and 3;
- observer 3 versus the average of observers 1 and 2.

Standard deviations (over images) were also computed.

In addition, in order to compute an example clinically meaningful measure, the mean thickness (again, not including the middle 30 pixels to exclude the fovea) of layers defined from the first border to each of the remaining borders was computed using the algorithm and each observer. The average thicknesses computed from the three observers was used as a reference standard and the absolute differences between the algorithmic thicknesses and reference thicknesses were computed. Furthermore, absolute thickness differences were computed between the observers using the same cases as was done for computing the border positioning errors.

Finally, thickness differences between the affected and unaffected eyes were computed based on the 3-D algorithmic results. In particular, we computed the thickness of the inner retinal layer assumed to contain the ganglion cells (from surface 1 to surface 3 in Fig. 3: NFL+GCL+IPL) and the thickness of the remaining outer layers (from surface 3 to surface 6 in Fig. 3). We also computed the total retinal thickness. (In all cases, the middle 30 pixels were not included to exclude the fovea.) Paired *t*-tests were used to test for significant thickness differences (p values < 0.05 were considered significant).

IV. RESULTS

Our approach successfully segmented all six intraretinal surfaces in the set of 24 AION OCT images (there were no segmentation failures). While local inaccuracies existed and their quantitative assessment is given below, they were minor and no manual editing of the segmentation results was performed prior to the subsequent quantitative analysis. Table I and Table II summarize the computed border position errors. Table III and Fig. 10 summarize the thickness difference results. The border positioning errors and thickness differences between the algorithm and the reference standard were very similar to those computed between the observers. For example, the algorithm's overall unsigned border positioning error was 6.1 ± 2.9 μm, while the overall observer error averaged 6.9 ± 3.3 μm (ranging from 5.5 ± 3.0 μm for observer 2 versus the average of observers 1 and 3 to 8.2 ± 3.8 μm for observer 1 versus observer 3). In terms of thickness differences, the smallest errors for both the algorithm and the observers were for the layer defined by surfaces 1 and 5 (2.2 ± 1.8 μm for the algorithm and an average of 2.2 ± 1.9 μm for the observers), while the largest errors for the algorithm were for the layer defined by surfaces 1 and 6 (6.2 ± 3.9 μm). (However, note that surface six was also the surface displaying the largest bias based on the signed border positioning errors.) The largest thickness differences between the observers were for the layer defined by surfaces 1 and 4 (overall average error of 8.2 ± 4.5 μm, which is larger than the algorithm's error of 4.8 ± 4.7 μm for this layer). Three example results (reflecting the best case, the median case, and the worst case according to the overall unsigned border positioning error) are shown in Fig. 11. The mean intraretinal layer segmentation time (after alignment/registration) was 4.1 ± 0.9 min (using a Windows XP workstation with a 3.2-GHz Intel Xeon CPU).

A summary of the resulting thickness values for the affected and unaffected eyes of each subject are shown in Fig. 12. The inner retinal layer thickness for the affected eye was 24.1 μm smaller on average than for the unaffected eye (or about 21%, statistically significant, $p < 0.001$), while the outer retinal layer thickness was 3.7 μm larger on average (not statistically significant, $p = 0.14$). The total retinal thickness for the affected eye was 20.4 μm smaller on average than for the unaffected eye (or about 7%, statistically significant, $p < 0.001$). Thus, it was the thickness of the inner layers that showed the largest difference between the affected and unaffected eyes.

V. DISCUSSION AND CONCLUSION

As indicated by the reported quantitative results, our method performed very well overall, based on two different and independent error metrics. First, comparing the performance of the algorithm to the border tracings by human experts showed that the method's mean unsigned border error of 6.1 ± 2.9 μm compared favorably to that between all three paired comparisons of human observers (ranging between 7.0 ± 3.5 and 8.2 ± 3.8 μm), and performed better than two out of three comparisons of a single human expert against the two others (6.5 ± 3.2, 5.5 ± 3.0, and 7.0 ± 3.0 μm). Second, the estimates of the inner retinal thickness by our method in these patients with unilateral AION was able to correctly identify the affected eye in 12/12 cases. In addition, as also suggested by others, our results further support the clinical utility of segmenting individual retinal layers in the macula. This is in contrast to the standard clinical practice of only determining the total retinal thickness. In the studied AION patients, the total retinal thickness showed an average difference of 20.4 μm (7%). Notably, the thickness of the inner retina showed a larger average difference of 24.1 μm (21%) and may be a better assessment factor for early disease detection and quantification.

It is also important to recognize some of the limitations of using this approach. One example is that in some cases, a specific surface may not visible on the images (e.g., surface 2 in some of our data). In such cases, the multisurface detection algorithm attempts to find all surfaces even if some of them are not present. A human observer, on the other hand, is able to indicate that the surface is not visible. Of course, having an undetectable surface does not necessarily imply that the corresponding layer is missing, as the ophthalmic community is still in the process of learning the precise anatomical correspondence of the visible OCT layers. In addition, the images may be too noisy to resolve the layer boundaries in some cases. Nevertheless, in such cases of an undetectable surface, it would be important for a human to review the segmentation data and not use the identified surface corresponding to missing image features.

While others have presented automated 2-D approaches for the segmentation of intraretinal layers in OCT images [4], [7], to our knowledge, this the first reported approach for the automated 3-D segmentation of intraretinal layers. Having such a 3-D approach that takes advantage of 3-D contextual information will be especially important when segmenting more densely acquired OCT images, such as those resulting from truly 3-D spectral-domain systems. For example, when three-dimensionally segmenting sets of 2-D macular time-domain OCT images, the smoothness constraint between radial slices (i.e., Δ_θ) must be set to a large enough value to allow for the increasing distances between slices as you move towards the periphery. Because these increasing distances do not occur in spectral-domain OCT images, the interslice constraints can serve a more prominent role. In applying this approach to spectral-domain images, one can also more appropriately use 3-D methods for computing cost functions. For example, in this work, we applied a 2-D spectral reducing anisotropic diffusion filter as a preprocessing step because of the sparse sampling of the data. When segmenting spectral-domain OCT images, a 3-D speckle-reducing anisotropic diffusion (SRAD) approach could more appropriately be applied.

In addition, this work reports performance and application of a novel method in which a 3-D graph search approach was employed to find so many interrelated layers for the first time. Theoretically, the graph search would be capable of simultaneously finding as many layers as desired. However, when identifying multiple layers simultaneously, the cost functions for all surfaces must be specified upfront. We have found that a divide-and-conquer approach is preferred, whereby some cost functions are designed by allowing the use of previously found surfaces. Thus our strategy led to first simultaneously finding three "easier" surfaces,

followed by sequentially finding the remaining surfaces in an effort to best utilize *a priori* information in the cost function design. We have also recently explored extending the graph search approach itself so that true regional information can be included in the cost functions, thus allowing the simultaneous detection of surfaces optimized over regional layer properties [21]. This extension has also allowed us to automatically detect an additional layer (seven surfaces instead of six).We have also explored extending the graph search to allow for varying constraints [22], so that more precise surface feasibility constraints (learned from examples) can be specified.

In any optimization approach, the cost function design influences the method's behavior and performance. While the best way to design cost functions is to perform its optimization in an independent set of images for which ground truth is available, cost functions in this work were designed experimentally and intuitively based on a visual inspection of OCT images from a variety of diseases. As such, the cost functions were not AION-specific and were not optimized in any way with respect to the AION datasets analyzed in this work.

In summary, we have presented an automated 3-D approach for the segmentation of intraretinal layers on macular OCT images, thus enabling the separate computation of individual layer properties, such as thickness. Having separate layer properties will be especially important in cases in which a retinal disease affects the individual layers differently (e.g., one layer may thin due to neuron loss, while another one may thicken due to the presence of fluid). In addition, it will aid in clinical studies designed to pinpoint which layers are actually altered during disease processes.

Acknowledgments

The authors would like to thank Dr. V. A. Shah for his involvement in image acquisition.

This work was supported in part by the National Institutes of Health–National Institute of Biomedical Imaging and BioEngineering (NIH-NIBIB) under Grant R01-EB004640-01, in part by the National Institutes of Health–National Eye Institute (NIH-NEI) under Grant R01-EY017066, in part by a grant from the Veterans Administration (Merit Review and Rehabilitation Division), and in part by an unrestricted grant from Research to Prevent Blindness, New York, NY.

REFERENCES

1. Huang D, Swanson EA, Lin CP, Schuman JS, Stinson WG, Chang W, Hee MR, Flotte T, Gregory K, Puliafito CA. Optical coherence tomography. Science. 1991 Nov.vol. 254(no. 5035):1178–1181. [PubMed: 1957169]
2. Koozekanani, D.; Boyer, K.; Roberts, C. Retinal thickness measurements in optical coherence tomography using a Markov boundary model; Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR); 2000 Jun.. p. 363-370.
3. Koozekanani D, Boyer K, Roberts C. Retinal thickness measurements from optical coherence tomography using a Markov boundary model. IEEE Trans. Med. Imag. 2001 Sep.vol. 20(no. 9): 900–916.
4. Ishikawa H, Stein DM, Wollstein G, Beaton S, Fujimoto JG, Schuman JS. Macular segmentation with optical coherence tomography. Investigative Ophthalmol. Vis. Sci. 2005 Jun.vol. 46(no. 6): 2012–2017.
5. Chan A, Duker JS, Ishikawa H, Ko TH, Schuman JS, Fujimoto JG. Quantification of photoreceptor layer thickness in normal eyes using optical coherence tomography. Retina. 2006; vol. 26(no. 6): 655–660. [PubMed: 16829808]
6. Cabrera Fernández D. Delineating fluid-filled region boundaries in optical coherence tomography images of the retina. IEEE Trans. Med. Imag. 2005 Aug.vol. 24(no. 8):929–945.
7. Cabrera Fernández D, Salinas HM, Puliafito CA. Automated detection of retinal layer structures on optical coherence tomography images. Opt. Express. 2005; vol. 13(no. 25):10 200–10 216.

8. Shahidi M, Wang Z, Zelkha R. Quantitative thickness measurement of retinal layers imaged by optical coherence tomography. Am J. Ophthalmol. 2005 Jun.vol. 139(no. 6):1056–1061. [PubMed: 15953436]
9. Baroni M, Fortunato P, Torre AL. Towards quantitative analysis of retinal features in optical coherence tomography. Med. Eng. Phys. 2007 May; vol. 29(no. 4):432–441. [PubMed: 16860587]
10. Haeker, M.; Abràmoff, MD.; Kardon, R.; Sonka, M. Segmentation of the surfaces of the retinal layer from OCT images. Proceedings of the 9th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI 2006), Part I; New York. Springer; 2006. p. 800-807.Lecture Notes Computer Science
11. Haeker M, Sonka M, Kardon R, Shah VA, Wu X, Abràmoff MD. Automated segmentation of intraretinal layers from macular optical coherence tomography images. Proc. SPIE Med. Imag. 2007: Image Process. 2007; vol. 6512:651214.
12. Wollstein G, Paunescu LA, Ko TH, Fujimoto JG, Kowalevicz A, Hartl I, Beaton S, Ishikawa H, Mattox C, Singh O, Duker J, Drexler W, Schuman JS. Ultrahigh-resolution optical coherence tomography in glaucoma. Ophthalmology. 2005; vol. 112(no. 2):229–237. [PubMed: 15691556]
13. Li K, Wu X, Chen DZ, Sonka M. Optimal surface segmentation in volumetric images—A graph-theoretic approach. IEEE Trans. Pattern Anal. Machine Intell. 2006; vol. 28(no. 1):119–134.
14. Wu, X.; Chen, DZ. Optimal net surface problems with applications. Proceedings 29th International Colloquium on Automata, Languages, and Programming (ICALP); New York. Springer; 2002. p. 1029-1042.Lecture Notes Computer Science
15. Thedens DR, Skorton DJ, Fleagle SR. Methods of graph searching for border detection in image sequences with applications to cardiac magnetic resonance imaging. IEEE Trans. Med. Imag. 1995 Jan.vol. 14(no. 1):42–55.
16. Suri JS, Liu K, Singh S, Laxinarayan SN, 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. 2002 Jan.vol. 6(no. 1):8–28.
17. McInerney T, Terzoppoulos D. Deformable models in medical image analysis: A survey. Med. Image Anal. 1996; vol. 1(no. 2):91–108. [PubMed: 9873923]
18. Maes F, Collignon A, Vandermeulen D, Marchal G, Suetens P. Multimodality image registration by maximization of mutual information. IEEE Trans. Med. Imag. 1997 Apr.vol. 16(no. 2):187–198.
19. Yu Y, Acton ST. Speckle reducing anisotropic diffusion. IEEE Trans. Image Process. 2002 Nov.vol. 11(no. 11):1260–1270. [PubMed: 18249696]
20. Chan TF, Vese LA. Active contours without edges. IEEE Trans. Image Process. 2001 Feb.vol. 10(no. 2):266–277. [PubMed: 18249617]
21. Haeker, M.; Wu, X.; Abràmoff, MD.; Kardon, R.; Sonka, M. Information Processing in Medical Imaging (IPMI). Vol. vol. 4584. New York: Springer; 2007. Incorporation of regional information in optimal 3-D graph search with application for intraretinal layer segmentation of optical coherence tomography images; p. 607-618.Lecture Notes in Computer Science
22. Haeker, M.; Abràmoff, MD.; Wu, X.; Kardon, R.; Sonka, M. Proceedings of the 10th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI 2007). Vol. vol. 4791. New York: Springer; 2007. Incorporation of shape into optimal 3-D graph search for segmentation of macular optical coherence tomography images; p. 244-251.Lecture Notes Computer Science

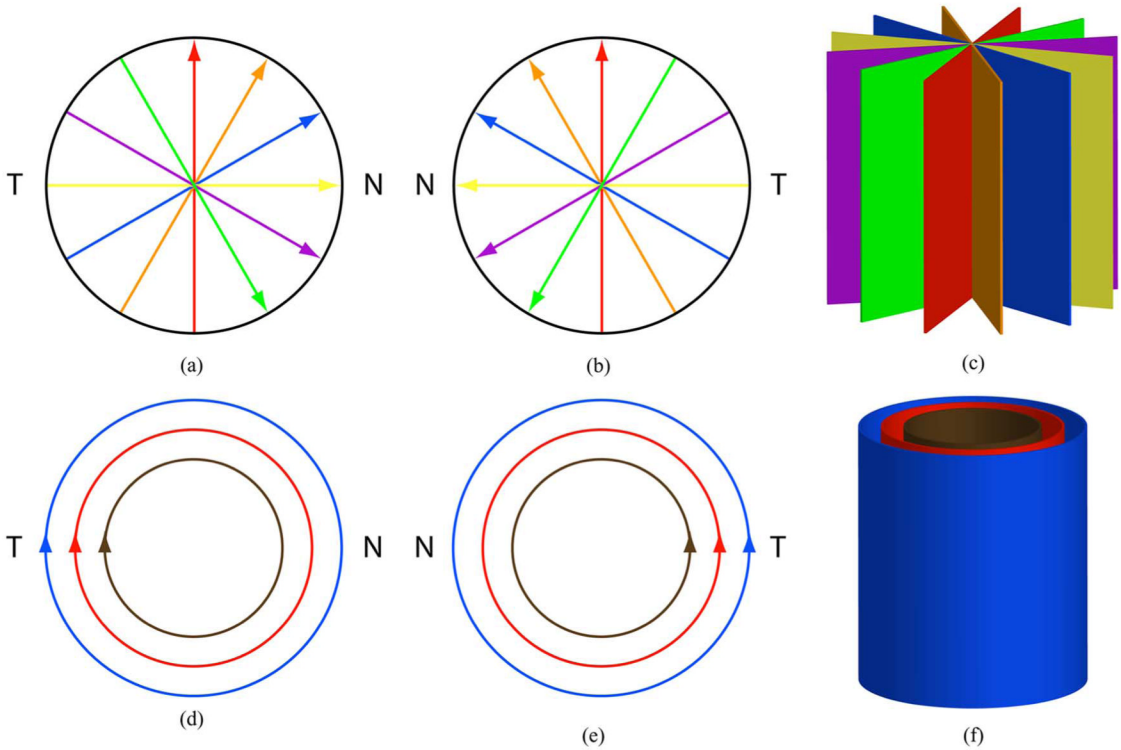

Fig 1.

Schematic view of the macular (a–c) and circular (d–f) scanning protocols on time-domain OCT systems. (a) Scans in macular series on the right eye. (N = nasal; T = temporal.) (b) Scans in macular series on the left eye. (c) Visualization of acquired macular scans for one eye in 3-D. Each color represents a different 2-D scan. (d) Scans in peripapillary circular series on the right eye. (e) Scans in peripapillary circular series on the left eye. (f) Visualization of acquired circular scans for one eye in 3-D.

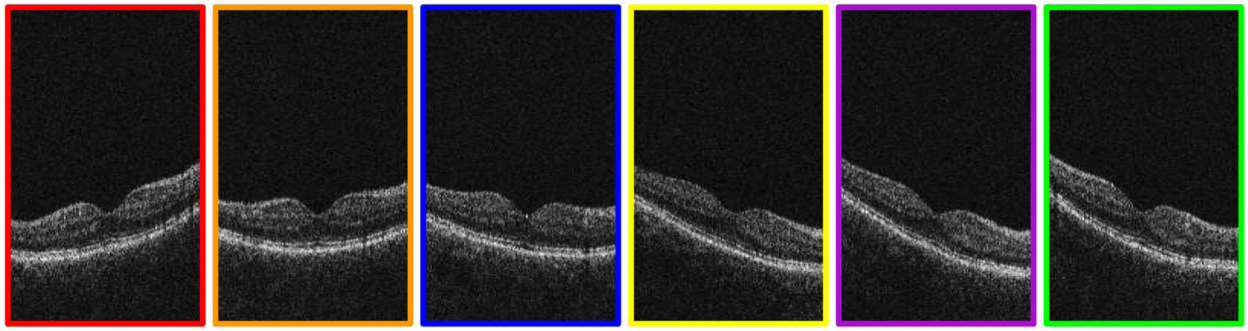

Fig 2.
Example six raw scans in a macular scan series. Note that the colored borders correspond to those found in Fig. 1(a)–(c)

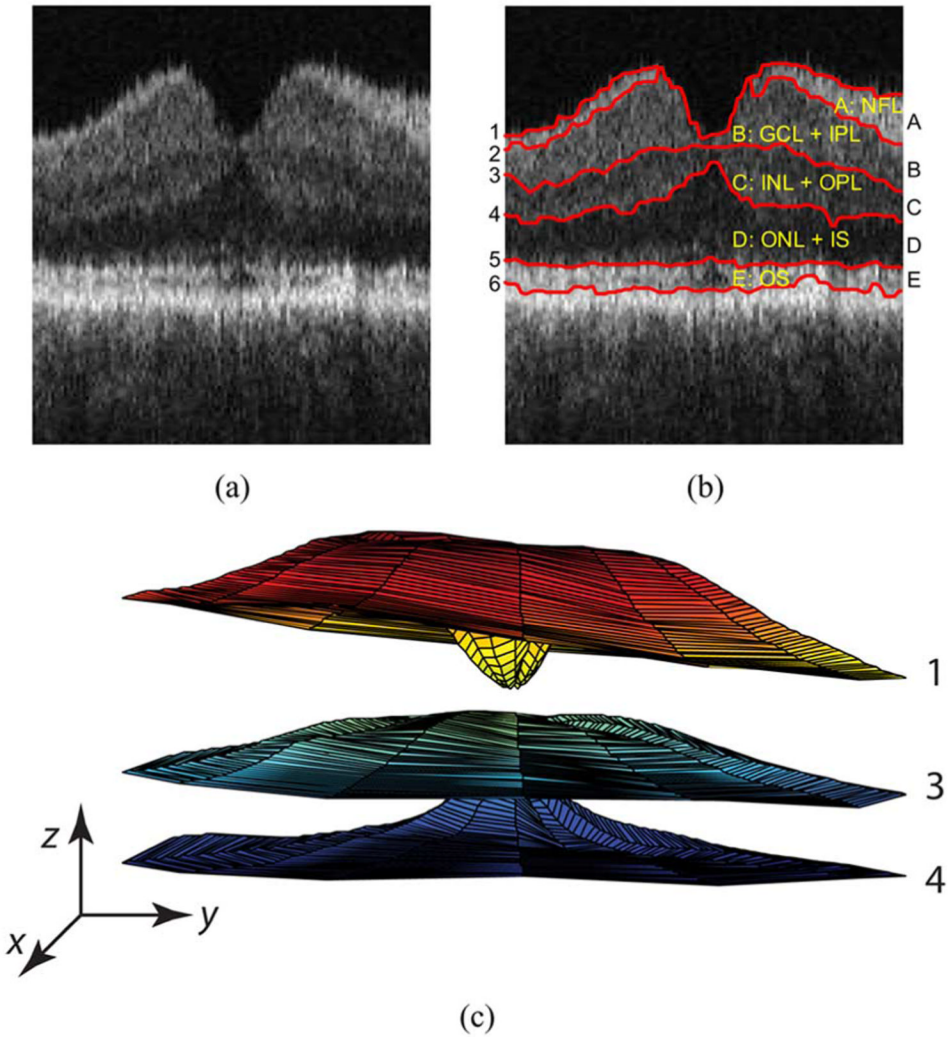

Fig 3.

Example composite image with labeled intralayer segmentation and 3-D visualization of three surfaces (top and bottom of images have been cropped to aid in visualization). (a) Composite image. (b) Six surfaces (labeled 1–6) and five corresponding intralayers (labeled A–E). The anatomical correspondence is our current presumption based on histology and example images from higher-resolution research OCT scanners [12]: (A) NFL (nerve fiber layer), (B) GCL + IPL (ganglion cell layer and inner plexiform layer), (C) INL+OPL (inner nuclear layer and outer plexiform layer), (D) ONL + IS (outer nuclear layer and photoreceptor inner segments), (E) OS (photoreceptor outer segments). (c) Example 3-D visualization of surfaces 1, 3, and 4.

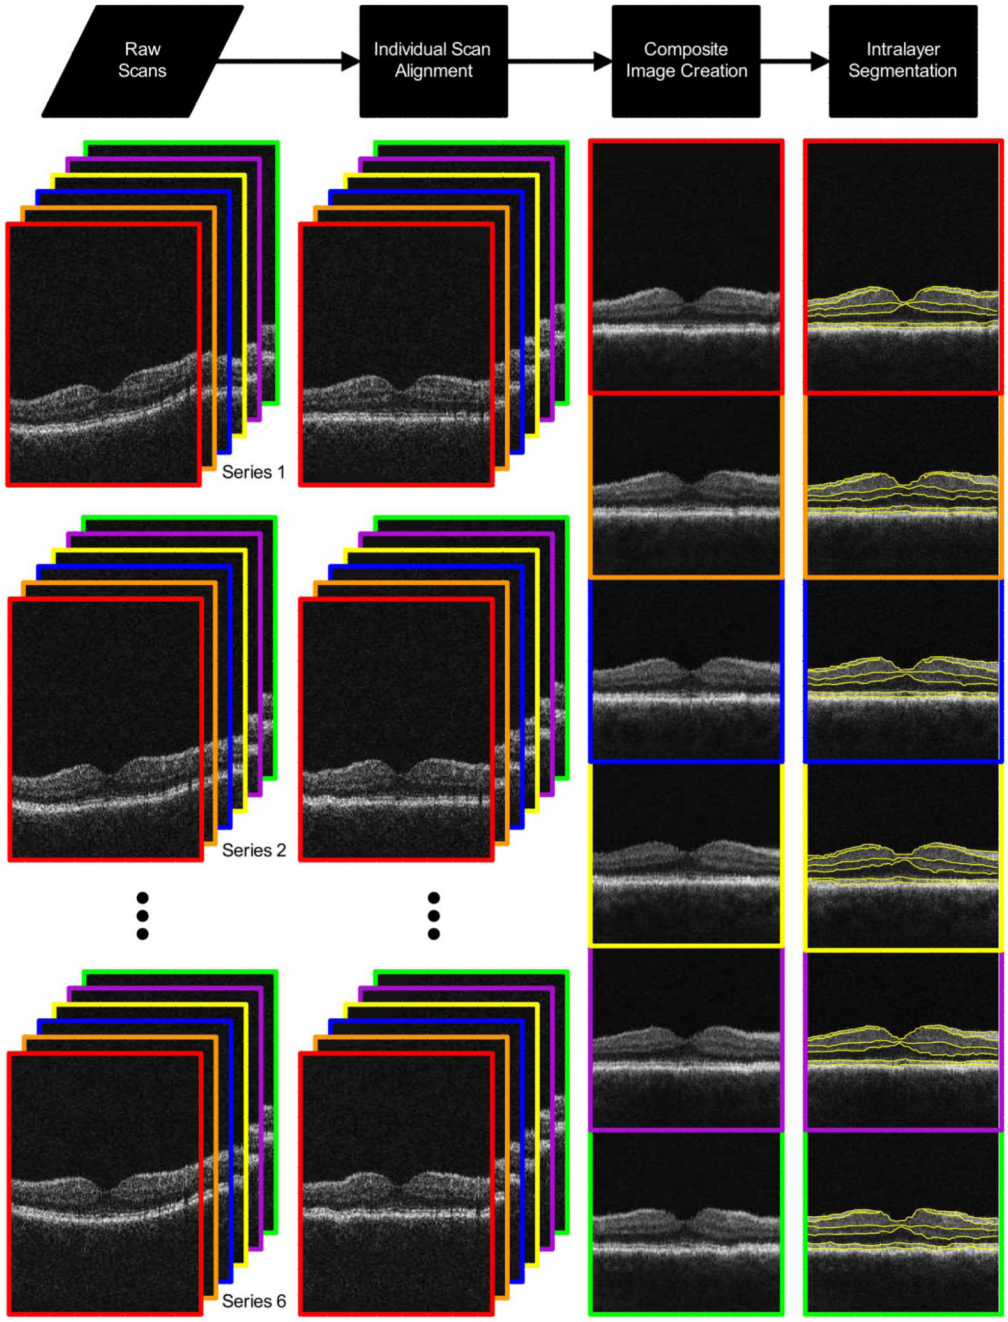

Fig 4.

Overview of segmentation steps for the data associated with one eye. First, each individual scan was aligned so that the RPE (boundary 6) was approximately horizontal in the image. Second, images from each location were registered and averaged to form a composite image. Finally, the intralayer surfaces were determined using a 3-D graph-search approach. All steps were performed automatically.

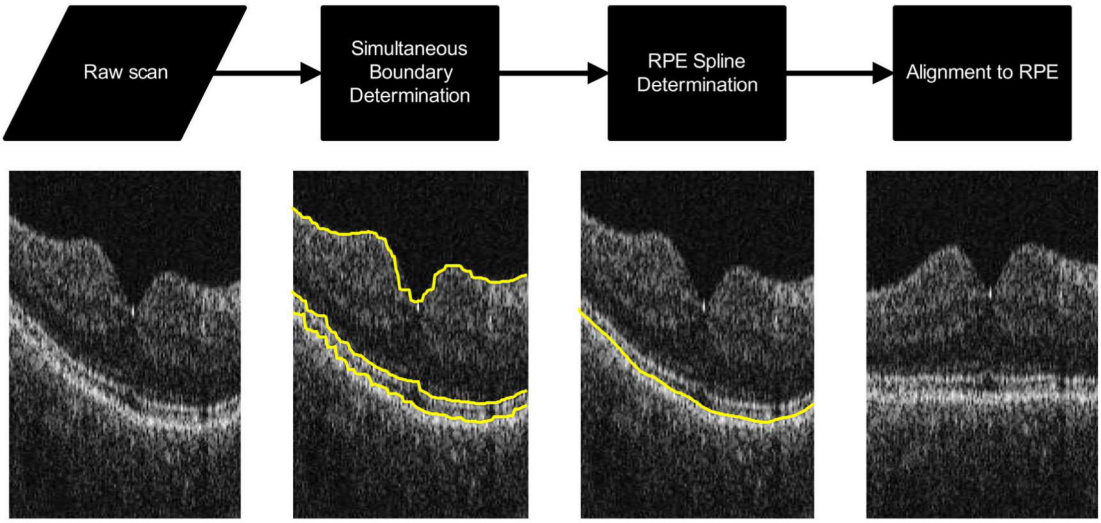

Fig 5.
Individual scan alignment (top and bottom of images have been cropped to aid in visualization).

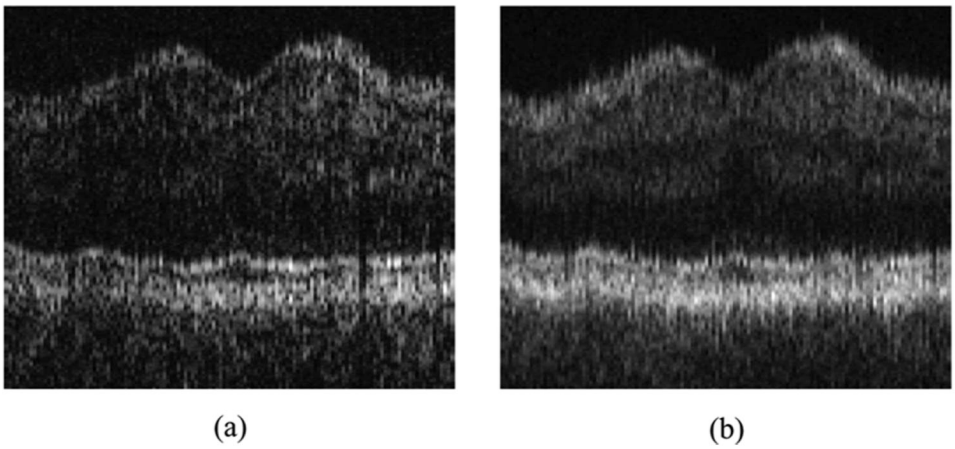

Fig 6.
Comparison between an individual scan and a 2-D composite scan (top and bottom of images have been cropped to aid in visualization). (a) Individual scan. (b) Composite scan.

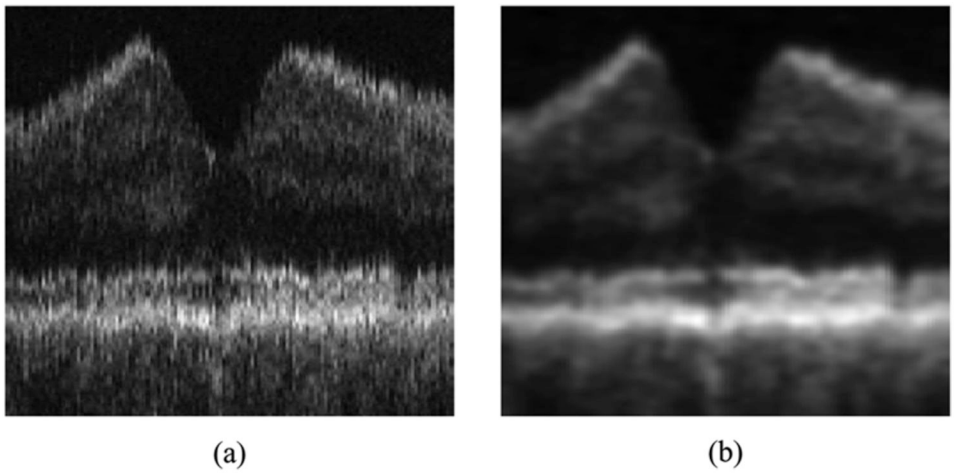

Fig 7.
Example of using a SRAD method as a preprocessing step (top and bottom of images have been cropped to aid in visualization). (a) Composite scan. (b) Composite scan after application of the SRAD method.

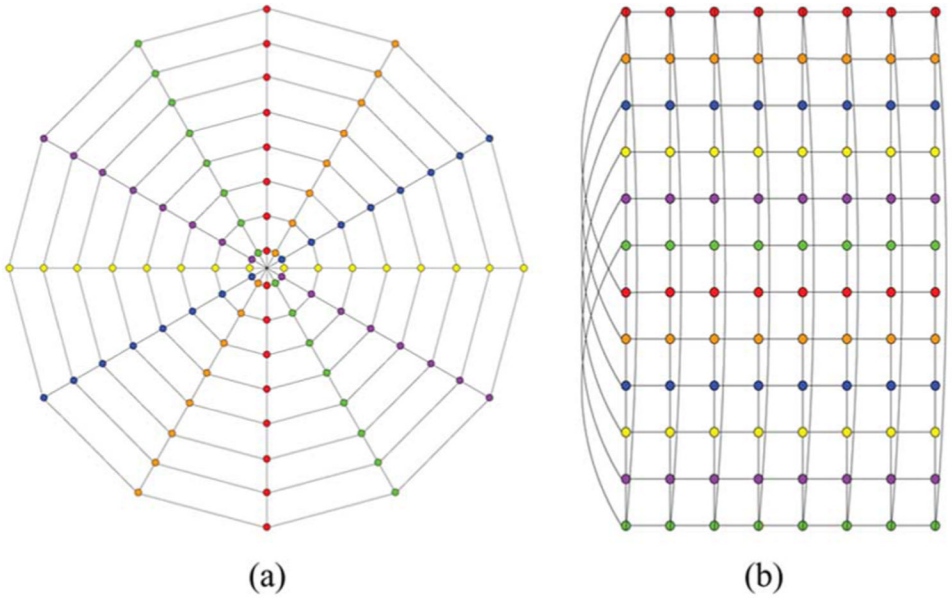

Fig 8.

Schematic view of neighbor relationship for 3-D macular OCT segmentation. The edges indicate neighborhood connectivity of one "column" of z values at a $(r; \theta)$ pair to another. For each edge shown, smoothness constraints existed between corresponding voxel z columns for the two $(r; \theta)$ pairs connected to the edge. (a) Base graph using cylindrical coordinates. (b) Base graph using unwrapped coordinate system (as might be stored in the computer).

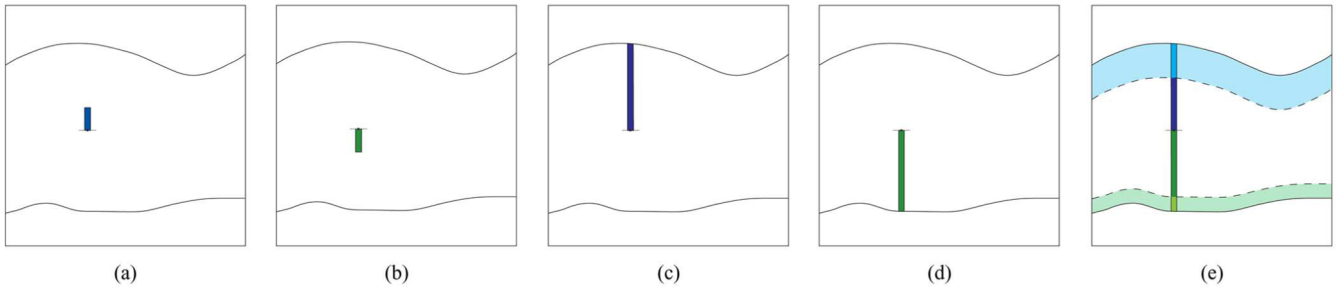

Fig 9.

Some examples for where the image information comes from in a regional cost function term. Dark borders represent surrounding surfaces (may not be known) of the surface for which the cost function term is being defined. In cases for which an upper or lower surrounding surface does not exist (i.e., the first and last surfaces), the corresponding dark border represents the boundary of the image.

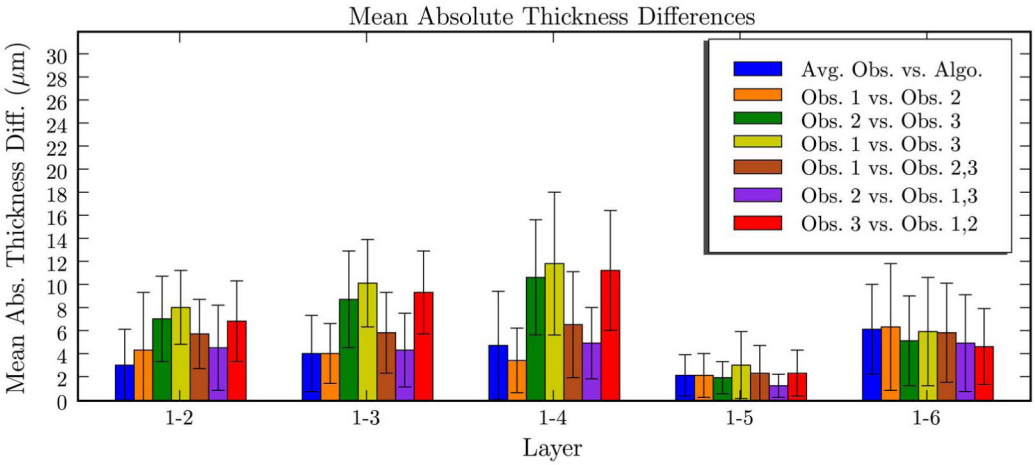

Fig 10.
Bar chart of mean thickness differences (error bars reflect standard deviations).

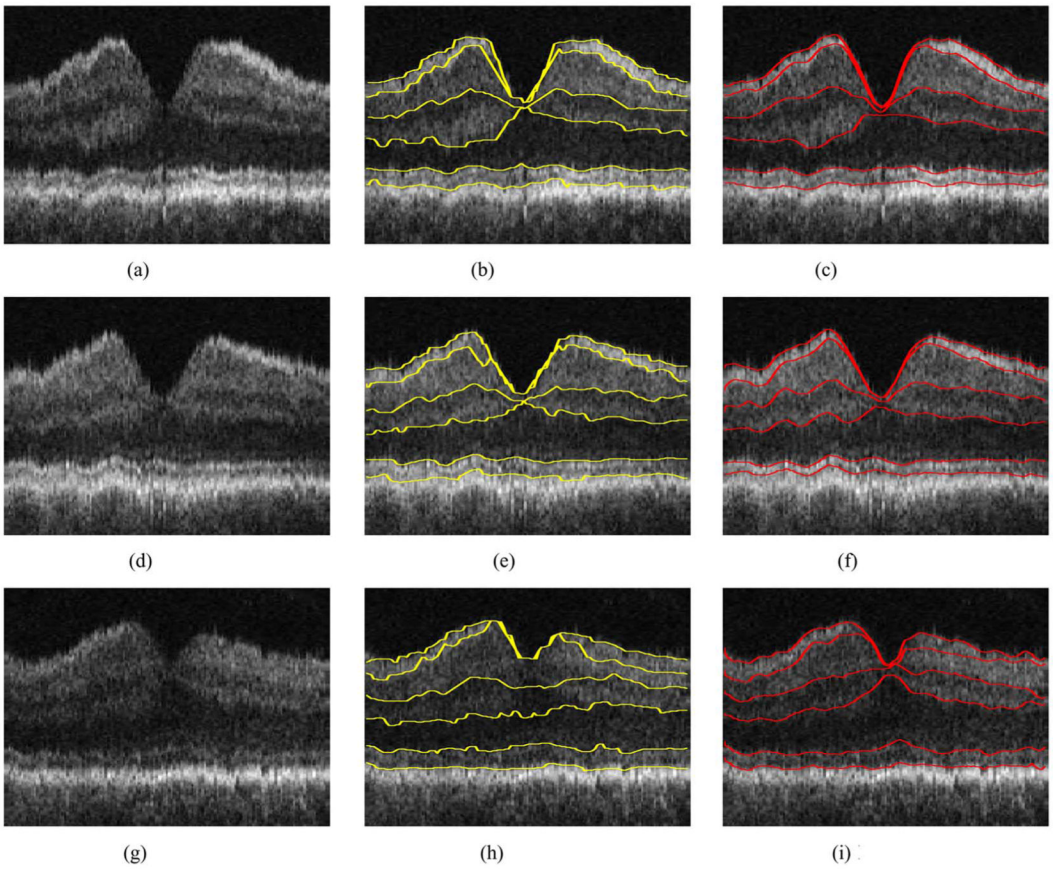

Fig 11.

Three example results reflecting the best, median, and worst performances according to the overall unsigned border positioning error. (a) Best case composite image. (b) Best case composite image with segmented borders. (c) Best case composite image with average manual tracing. (d) Median case composite image. (e) Median case composite image with segmented borders. (f) Median case composite image with average manual tracing. (g) Worst case composite image. (h) Worst case composite image with segmented borders. (i) Worst case composite image with average manual tracing.

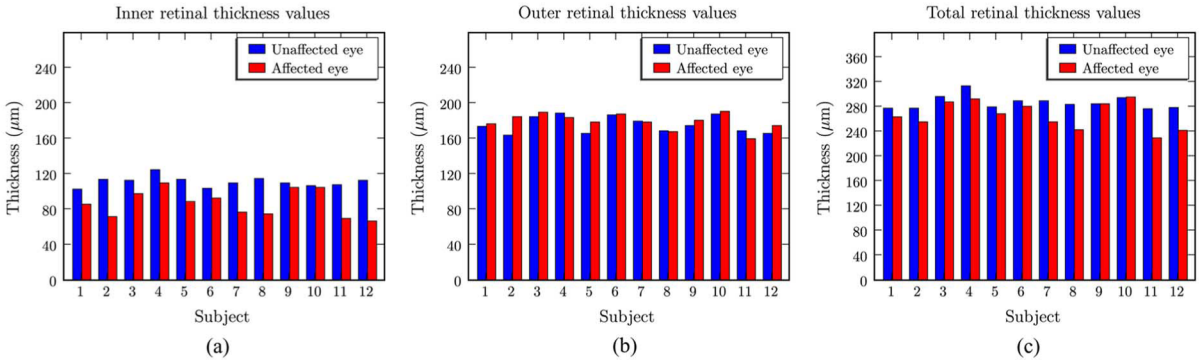

Fig 12.

Summary of thickness values based on our intraretinal layer segmentation approach. The thickness differences between the affected and unaffected eyes were largest on average for the inner retinal layer. Inner layer used in (a) contains the retinal ganglion cells and axons.

TABLE I

Summary of mean unsigned border positioning errors † for 24 scans in micrometers

Border	Avg. Obs. vs. Algo.	Obs.1 vs. Obs. 2	Obs. 2 vs. Obs. 3	Obs. 1 vs. Obs. 3	Obs. 1 vs. Obs. 2,3	Obs. 2 vs. Obs. 1,3	Obs. 3 vs. Obs. 1,2
1	4.4 ± 1.2	4.0 ± 1.1	4.8 ± 1.4	5.9 ± 1.3	4.3 ± 0.9	3.3 ± 1.0	5.0 ± 1.3
2*	7.6 ± 3.1	6.9 ± 4.4	7.5 ± 4.6	5.8 ± 1.2	5.3 ± 2.3	6.5 ± 4.4	5.9 ± 2.3
3*	6.9 ± 2.0	8.2 ± 2.1	8.6 ± 2.6	8.1 ± 3.0	6.9 ± 2.4	7.5 ± 1.7	7.3 ± 2.3
4	7.0 ± 2.9	7.5 ± 3.0	9.1 ± 3.6	9.4 ± 4.4	7.3 ± 3.2	6.8 ± 2.4	8.4 ± 3.9
5	3.0 ± 1.1	4.4 ± 1.6	5.5 ± 1.8	7.8 ± 2.8	5.6 ± 2.1	3.4 ± 1.0	6.4 ± 2.3
6	7.6 ± 2.9	8.5 ± 4.8	8.0 ± 3.0	11.5 ± 4.6	9.2 ± 4.5	6.2 ± 3.1	9.0 ± 3.0
overall	6.1 ± 2.9	7.0 ± 3.5	7.2 ± 3.3	8.2 ± 3.8	6.5 ± 3.2	5.5 ± 3.0	7.0 ± 3.0

† Mean ± SD. For each boundary differences were not computed for the middle 30 pixels (out of 128) to exclude the fovea.

* Errors were not computed for those scans in which boundary was determined to not be visible by at least one expert.

IEEE Trans Med Imaging. Author manuscript; available in PMC 2009 January 06.

TABLE II

Summary of mean signed border positioning errors † for 24 scans in micrometers

Border	Avg. Obs. vs. Algo.	Obs. 1 vs. Obs. 2	Obs. 2 vs. Obs. 3	Obs. 1 vs. Obs. 3	Obs. 1 vs. Obs. 2,3	Obs. 2 vs. Obs. 1,3	Obs. 3 vs. Obs. 1,2
1	0.1 ± 2.5	1.4 ± 1.6	3.8 ± 1.9	5.2 ± 1.7	3.3 ± 1.3	1.2 ± 1.6	−4.5 ± 1.6
2*	1.4 ± 3.9	−1.4 ± 5.7	−1.4 ± 5.8	−2.7 ± 2.5	−2.0 ± 3.3	0.0 ± 5.6	2.0 ± 3.5
3*	−3.9 ± 3.4	−0.1 ± 3.9	−5.0 ± 4.0	−5.0 ± 3.6	−2.5 ± 3.2	−2.5 ± 3.5	5.0 ± 3.3
4	−3.0 ± 4.9	0.2 ± 4.8	−6.9 ± 5.1	−6.7 ± 6.1	−3.3 ± 4.8	−3.5 ± 3.9	6.8 ± 5.1
5	0.9 ± 1.9	2.7 ± 2.1	4.8 ± 2.4	7.5 ± 3.1	5.1 ± 2.4	1.1 ± 1.7	−6.2 ± 2.6
6	−5.6 ± 4.0	6.1 ± 6.9	5.0 ± 6.1	11.1 ± 4.9	8.6 ± 5.2	−0.5 ± 6.0	−8.0 ± 4.3
overall	−1.8 ± 4.4	1.6 ± 5.1	0.2 ± 6.5	1.9 ± 7.8	1.7 ± 5.7	−0.7 ± 4.4	−1.0 ± 6.8

† Mean ± SD. For each boundary, differences were not computed for the middle 30 pixels (out of 128) to exclude the fovea.

* Errors were not computed for those scans in which boundary was determined to not be visible by at least one expert.

TABLE III

Summary of mean absolute thickness differences [†] for 24 scans in micrometers

Border	Avg. Obs. vs. Algo.	Obs. 1 vs. Obs. 2	Obs. 2 vs. Obs. 3	Obs. 1 vs. Obs. 3	Obs. 1 vs. Obs. 2,3	Obs. 2 vs. Obs.1,3	Obs. 3 vs. Obs.1,2
1–2 [*]	3.1 ± 3.1	4.4 ± 5.0	7.1 ± 3.7	8.1 ± 3.2	5.8 ± 3.0	4.6 ± 3.7	6.9 ± 3.5
1–3 [*]	4.1 ± 3.3	4.1 ± 2.6	8.8 ± 4.2	10.2 ± 3.8	5.9 ± 3.5	4.4 ± 3.2	9.4 ± 3.6
1–4	4.8 ± 4.7	3.5 ± 2.8	10.7 ± 5.0	11.9 ± 6.2	6.6 ± 4.6	5.0 ± 3.1	11.3 ± 5.2
1–5	2.2 ± 1.8	2.2 ± 1.9	2.0 ± 1.4	3.1 ± 2.9	2.4 ± 2.4	1.3 ± 1.0	2.4 ± 2.0
1–6	6.2 ± 3.9	6.4 ± 5.5	5.2 ± 3.9	6.0 ± 4.7	5.9 ± 4.3	5.0 ± 4.2	4.7 ± 3.3

[†] Mean ± SD. For each boundary, thickness differences were not computed for the middle 30 pixels (out of 128) to exclude the fovea.

[*] Errors were not computed for those scans inn which boundary was determined to not be visible by at least one expert.