



Published in final edited form as:

Proc SPIE Int Soc Opt Eng. 2014 ; 9034: . doi:10.1117/12.2043040.

An adaptive grid for graph-based segmentation in retinal OCT

Andrew Lang^a, Aaron Carass^a, Peter A. Calabresi^b, Howard S. Ying^c, and Jerry L. Prince^a

^aDepartment of Electrical and Computer Engineering, The Johns Hopkins University

^bDepartment of Neurology, The Johns Hopkins University School of Medicine

^cWilmer Eye Institute, The Johns Hopkins University School of Medicine

Abstract

Graph-based methods for retinal layer segmentation have proven to be popular due to their efficiency and accuracy. These methods build a graph with nodes at each voxel location and use edges connecting nodes to encode the hard constraints of each layer's thickness and smoothness. In this work, we explore deforming the regular voxel grid to allow adjacent vertices in the graph to more closely follow the natural curvature of the retina. This deformed grid is constructed by fixing node locations based on a regression model of each layer's thickness relative to the overall retina thickness, thus we generate a subject specific grid. Graph vertices are not at voxel locations, which allows for control over the resolution that the graph represents. By incorporating soft constraints between adjacent nodes, segmentation on this grid will favor smoothly varying surfaces consistent with the shape of the retina. Our final segmentation method then follows our previous work. Boundary probabilities are estimated using a random forest classifier followed by an optimal graph search algorithm on the new adaptive grid to produce a final segmentation. Our method is shown to produce a more consistent segmentation with an overall accuracy of 3.38 μm across all boundaries.

Keywords

OCT; retina; layer segmentation; adaptive grid; classification

1. INTRODUCTION

Retinal imaging of the macula using optical coherence tomography (OCT) is important for many applications. Although particularly well suited to examining more traditional ocular diseases such as glaucoma¹ and age-related macular degeneration,² it is becoming more popular for exploration of neurological diseases³ including Parkinson's disease⁴ and multiple sclerosis (MS).^{5, 6} In many cases, the morphological changes of the retina in these neurological diseases are not detectable by a simple visual inspection. To quantify the more subtle retinal changes in these diseases, it is necessary to delineate specific layers within the retina. In MS, for example, there is evidence that the RNFL and GCL layers undergo

thinning as the disease progresses.^{7,8} Since manual segmentation of the layers is time consuming, automatic segmentation methods are preferable.

Recent developments in retinal layer segmentation have focused on graph-based algorithms which have shown promise in being more robust to different scanners, image quality, and finding more layers.⁹⁻¹³ In Garvin et al.,⁹ segmentation is carried out on a 4D graph, constructed following the methods described in Li et al.¹⁴ The final boundaries are found by solving a minimum closure problem on the graph using a max-flow, min-cut algorithm. Hard constraints are used to enforce smoothness and to limit the thickness of each layer. This work was extended by adding soft smoothness penalties generated by a prior information model,^{10, 12} which are enforced by including additional edges in the graph, increasing its complexity. Additionally, in both of these works, the size of the retina is not well accounted for by the prior information model since it is constructed based on the location of a single specified reference surface. It will therefore not perform as well for excessively large, small, and misaligned retinas.

In this paper, we enhance an existing retinal segmentation algorithm¹³ by deforming, or adapting, the graph that the segmentation is carried out on to account for the size, orientation, and curvature of each individual's retina. Traditionally, a graph is constructed by placing vertices, or nodes, at the location of each voxel, with edges used to guide the segmentation. In this work, we instead place vertices at locations relative to estimates of the top and the bottom surfaces of the retina, the inner limiting membrane (ILM) and Bruch's membrane (BM), respectively. Careful placement of the nodes in the graph is done by using a streamlining algorithm in conjunction with a regression model. As a result, vertices are placed at anatomically consistent locations with horizontal edges following the natural curvature of the retina. The final segmentation closely follows the state-of-the-art,^{9, 12, 13} utilizing a pixel classification approach to estimate boundary locations and a max-flow, min-cut algorithm to find the resulting segmentation. Overall, there are three main benefits to constructing the graph in this new way. Firstly, we control the scale of the graph, allowing for subvoxel resolution where necessary (for example, the fovea). Secondly, by allowing the graph to take the shape of the retina, graph construction is simplified enabling a more natural inclusion of smoothness constraints. Finally, we are able to smoothly fill in areas that have a lack of boundary evidence (for example, in the shadow of a blood vessel). A flowchart of our algorithm is shown in Fig. 1.

2. METHODS

Our layer segmentation algorithm works by generating a patient specific adaptive grid to carry out an optimal graph-based segmentation algorithm. To do this, we assume that there is an initial segmentation for the inner and outer boundaries of the retina, which we find using a previously described algorithm.¹³ Briefly, the ILM is estimated from the largest positive gradient along each A-scan (vertical scan line in an OCT image), while the BM is estimated using the largest negative gradient. Final boundary positions were constrained by distance from an approximate guess of the IS-OS boundary (also a large negative gradient), median filtered to remove outliers, and finally smoothed using a Gaussian kernel.

2.1 Adaptive grid construction

There are three steps for generating a patient-specific adaptive grid, as illustrated in Figs. 2 and 4. Due to differences in the orientation and curvature between different subject's retinal OCT scans, we do *not* place vertices on the grid at voxel locations along each A-scan. Instead, we place vertices along streamlines generated from the inner to the outer retina surface (Fig. 2(a)). Such a streamlining approach has previously been used to aid in the computation of cortical thickness of brain MRI scans,¹⁵ as well as in the construction of a grid for annular tissues.¹⁶ In the second step, we use a regression model to produce an estimate of the location of each layer boundary between the two outer boundaries (Fig. 2(b)). A *separate* regression model is learned at each point on the surface of the ILM, thereby incorporating spatial information into the estimated positions. Finally, we fill in the graph with nodes placed along the streamlines based on the regression (Fig. 4(a)).

2.1.1 Streamline generation—We generate streamlines from the ILM to the BM by solving Laplace's equation between the two surfaces.¹⁵ This approach has the property that the streamlines are guaranteed not to cross each other; a property that lends itself nicely to construction of a grid. Laplace's equation takes the form

$$\nabla^2 u = \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} = 0, \quad (1)$$

subject to the boundary conditions $u(S_{\text{ILM}}) = 0$ and $u(S_{\text{BM}}) = 1$, where u represents the harmonic function to be solved for, and S_{ILM} and S_{BM} are the locations of the inner and outer retina surfaces, respectively. Since the retina does not form a closed object, we additionally enforce boundary conditions along the sides of each image to maintain their initial values. Also note that since the out-of-plane resolution of our data is disproportionately larger than the in-plane resolution, we only solve the equation in 2D, on each B-scan independently. Were the data closer to isotropic, then solving the equation would make more sense in 3D (in particular because we will later interpolate the volume along the streamlines).

We solve Eq. 1 for u using a red-black Gauss-Seidel method. The interior values of u (between S_{ILM} and S_{BM}) are initialized to the proportional distance of each pixel from the top surface to the bottom, along each A-scan. This initialization provides a good approximation leading to fast convergence. We then construct streamlines between the two surfaces by integrating over the tangent field of u .¹⁵ A tangent vector field, \vec{N} between the two surfaces is constructed from

$$\vec{N} = \frac{\nabla u}{\|\nabla u\|}. \quad (2)$$

By integrating \vec{N} from every point on S_{ILM} to some point on S_{BM} , we can generate correspondence trajectories, or streamlines, between the two surfaces. We carry out the integration using a Runge-Kutta 4th order scheme. Finally, to account for not having

isotropic pixel size, we modify both of the previous equations to incorporate the resolution of the data.¹⁷ An example of the streamlines generated between the retinal boundaries is shown in Fig. 2(a). Notice how the streamlines flow orthogonally from each surface.

2.1.2 Boundary regression model—With the streamlines computed, we estimate the positions of the seven interior retinal boundaries along each streamline by using a regression model trained on manual segmentation data. This regression takes the form

$$p_{i,j} = \alpha_{i,j} + \beta_{i,j} d_i, \quad (3)$$

where $p_{i,j}$ is the relative distance along streamline $i \in I$ from the ILM to surface $j \in \{1, \dots, 7\}$, d_i is the total length of the streamline, and I is the set of all streamlines across the retina. Intuitively, this regression is trying to estimate the position of each boundary based on the total thickness of the retina (represented by the total streamline length). Since the actual boundary location along the streamline is computed as $p_{i,j} \cdot d_i$, the equation is actually quadratic in the total thickness, allowing us to capture more variability. Additionally, to get a better spatial alignment of each subject for the regression model, the thicknesses of each subject were aligned at the center of the fovea, which was estimated as the location of the shortest streamline within the central 40% (2.4 mm) region of the macula.

Fig. 3 shows an example of the regression on the difference $(p_{i,j} - p_{i,j-1})$ at three locations on the retina. We thought it would be more informative to show $(p_{i,j} - p_{i,j-1})$ instead of $p_{i,j}$ since $p_{i,j}$ represents the total distance of that boundary from the ILM (and therefore exhibits a cumulative effect over layer thicknesses). Looking at $(p_{i,j} - p_{i,j-1})$, we see the direct effect of the regression on each individual layer thickness. Note that in Fig. 3, looking across the columns, RNFL, GCIP, ..., RPE are the clinical acronyms of the layers in order from the top of the retina to the bottom.

From Fig. 3, we can see that the linear fit works well. In fact, it is quite a bit better than simply using the average relative position of each boundary, which would be a simple horizontal line in each plot. Intuitively, a horizontal line means that as the total retina thickness increases, the thickness of that particular layer is linearly increasing at the same rate, i.e. always estimate the OPL thickness as 10% of the distance from the ILM to BM. A line with a strong positive slope means that as the total retina thickness increases, the layer thickness increases at a greater rate than the total thickness (a quadratic increase). For example, we estimate the GCIP thickness as 10% of the distance from the ILM to BM if the total distance is 250 μm and 20% of the distance if the total thickness is 330 μm . One effect of this regression is to help normalize the thickness values due to discrepancies in scale between subjects. For example, if the foveal area is ‘zoomed in’ in one subject relative to another (perhaps due to a difference in the focal point), the foveal cup diameter would appear artificially larger, making the thickness at a corresponding location near the foveal rim appear thinner in the zoomed in subject. The GCIP plot in the middle row of Fig. 3 shows this phenomenon, whereas the retina gets thinner, the GCIP is getting smaller at a faster rate since it eventually disappears at the fovea.

2.1.3 Final graph construction—Given a set of streamlines between the retinal boundaries and a learned regression model for each streamline at each boundary, we place “base nodes” in the graph according to Eq. 3. In other words, for boundary i we place a vertex at a distance p_{ij} along streamline i . We denote these vertices as base nodes since they anchor the graph at locations consistent with the location of each boundary, also allowing us to fill in the rest of the graph with an arbitrary number of vertices between each base node. All of the base nodes for a respective boundary are connected with horizontal edges between adjacent A-scans, allowing them to follow the natural shape of the retina (see Fig. 2(b)).

We fill in the rest of the graph, between the base nodes, by adding N_j vertices along the streamlines between the base-nodes surrounding layer j . Nodes are also added above and below the top and bottom boundaries to allow for small changes to the fit of these surfaces. An example of a final constructed graph overlaid on a B-scan is shown in Fig. 4(a). Note that, for display purposes, this is not the full density of the final graph we use for our algorithm. We choose N_j such that the average distance between nodes within a layer is equal to s (i.e., if the average distance between base nodes within a layer is $40\text{ }\mu\text{m}$, and we would like an average of $s = 4\text{ }\mu\text{m}$ between nodes, $N_j = 40/4 - 1 = 9$ equally spaced nodes would be added). Since the distance between base nodes is variable over the area of the retina, the actual distance between inserted nodes is different across the retina, providing better resolution in areas where it is more useful, like where the layers come together at the fovea.

2.2 Segmentation algorithm

Our final segmentation method follows closely from our previous work.¹³ The algorithm uses a boundary classification approach whereby at every pixel, the probability of belonging to each boundary (or background) is calculated using a random forest classifier.¹⁸ In total, 27 features consisting of spatial positions, intensities, and gradients are used for the final classification.¹³ The resulting probabilities are then input into an optimal graph search algorithm¹⁴ for computing the final set of boundary surfaces which maximize the probability distributed over each surface. The difference between the current method and previous method is in the use of the adaptive graph, described above, and the edge weights within the graph. Since the graph nodes do not generally fall on the regular grid of the OCT volume, we use linear interpolation to extract the probability values corresponding to each node. Also note that, as shown in Fig 1, we compute the features and do the boundary classification after flattening the data to the BM. This is done to improve the response of the horizontal gradient features used in the algorithm, improving the quality of the classifier. The resulting probabilities are then “unflattened” when interpolating the values to the deformed grid locations.

If we think of the deformed grid on a regular lattice grid (Fig. 4(b)), we see that the overall effect is to flatten the data to each boundary. In doing so, setting up the final graph-search segmentation is much simpler. Following the method described in Li et al.,¹⁴ we use hard constraints to limit the the minimum and maximum distance between adjacent surfaces. These constraints take the form of directed edges connecting vertices with an infinite weight. Since our graph contains a fixed number of vertices between base nodes, we simply set the

maximum distance to be $2(N_j + 1)$ voxels for layer j . The minimum distance is set to one voxel to allow for the possibility of surfaces coming together near the fovea. Additionally, we limit the maximum distance a surface may change between adjacent pixels. This distance is set to one voxel in each direction since, as a result of the implicit boundary flattening, the surfaces should not change much between A- and B-scans.

Finally, we add soft constraints to our graph to penalize surfaces which deviate from the natural curvature provided by our adaptive grid. This type of soft constraint has been used before in such a segmentation problem by simply assigning a non-zero edge weight to horizontally connected edges.^{10, 12} We assign a constant weight of w to every horizontal edge connecting A-scans. In fact, setting w to infinity would limit the set of feasible surfaces to those connected by horizontal edges (see Fig. 4(a), ignoring the vertical streamlines).

3. RESULTS

To test our algorithm, we used macular OCT scans ($20^\circ \times 20^\circ$) from the right eye of 38 subjects. All scans were acquired using a Spectralis OCT scanner (Heidelberg Engineering, Heidelberg, Germany). Of the 38 subjects, 15 were healthy controls and 23 were diagnosed with MS. Each volume contained 49 B-scans, 1024 A-scans, 496 pixels per A-scan, although we downsampled the lateral direction to 512 A-scans. The total imaged area of the macula was approximately 6×6 mm. All nine layer boundaries were manually segmented on all B-scans of all subjects by a single trained rater. The boundaries were found by clicking on approximately 20–50 points along each layer border with cubic B-spline interpolation used to fill in the boundary between points.

The regression model from Sec. 2.1.2 was trained using all of the manual segmentation data. A cross-validation strategy could have been employed here, but we found the regression parameters to be robust to the data used. The random forest classifier used for Sec. 2.2 was trained using data from seven randomly selected subjects. The training used eight B-scans from each of these seven subjects for a total of 56 B-scans (exact details of the training can be found in Ref. 13). For the random forest, we used 60 trees, selecting 10 random features to examine at each split node of the trees.

As the classifier was trained using seven subjects, we examined the results of our algorithm on the remaining 31 subjects. We first explored the effect of changing two parameters in our algorithm, s and w , the average distance between vertices and the amount of smoothness, respectively. We changed the values of s from $2 \mu\text{m}$ to $6 \mu\text{m}$ in increments of 1, while we changed the value of w from 0 to 1 in increments of 0.2. For each pair of values, we looked at the average absolute error across all subjects and all boundaries. Fig. 5 shows the average and standard deviation of the errors as heatmaps. The minimum error was found for values of $s = 3$ and $w = 0.2$. These are the values used to produce the results in Table 1. It makes sense that a smaller value of s produced the best accuracy since lower values of s lead to higher resolution of the grid. Since our manual segmentation protocol leads to smoother segmentations, it makes sense that a small amount of smoothing also improves the results.

In Table 1 we compare the results of running the segmentation on the adaptive grid for each subject to the results of not using the adaptive grid — i.e., using the voxel grid as done by our previous segmentation method.¹³ The values represent the mean absolute error over all subjects, with distances computed along A-scans. Looking only at the overall errors, we see that the OS-RPE boundary has the largest error. This is likely due to the ambiguity which sometimes arises as there are multiple dark bands in this region. Overall, we see that the most of the rest of the boundary errors are within the axial resolution of our data ($\approx 3.9 \mu\text{m}$). In comparison to the voxel grid method, the new method has smaller errors for 8 out the 9 boundaries. Fig. 6 shows the results of both algorithms on two B-scan images. The manual and both automatic segmentation results are displayed for comparison. In particular, the results in the right column of Fig. 6 show the new algorithm performing better in two areas where the voxel grid algorithm performed poorly due to errors in the pixel classification. Overall, we see similarities in agreement with the numerical results given above. Note some of the discretization effects of our adaptive graph on the final segmentation that arise due to the graph construction not lying on the pixel boundaries, which could be smoothed out using smaller/larger values of s/w or with additional post-processing such as smoothing or a deformable model.

4. CONCLUSIONS

A new layer segmentation method was developed by using a subject adapted graph. It has many advantages over previously developed methods including added smoothness without additional complexity, the ability to smoothly fill in areas with low boundary probabilities by smoothly following the retina shape, and finally, the possibility for subvoxel resolution in areas where the layers converge together. We showed accuracy that is as good or better than our previous algorithm, which was shown to be competitive with the current state-of-the-art. One of the main drawbacks of the method is the additional computational burden from streamlining. Fortunately, every streamline can be generated independently, which would allow for a large speedup with parallel processing. Alternatively, fast algorithms for computing correspondence trajectories using a hybrid Euler-Lagrange approach could be used.¹⁶ Finally, since training was required for several steps of the algorithm (streamline regression and random forest classification), it is not likely to work with data containing gross pathological issues. This atypical data could be included in the training process, but other possibilities for improvement include removing the spatial component from the classifier training and using a nonlinear regression model for the graph construction.

Acknowledgments

This project was supported by NIH/NEI grant 5R21EY022150 and NIH/NINDS grant 1R01NS082347 (PAC). The authors would also like to thank Matthew Hauser for help with manual segmentation.

References

1. Guedes V, Schuman JS, Hertzmark E, Wollstein G, Correnti A, Mancini R, Lederer D, Voskanyan S, Velazquez L, Pakter HM, Pedut-Kloizman T, Fujimoto JG, Mattox C. Optical coherence tomography measurement of macular and nerve fiber layer thickness in normal and glaucomatous human eyes. *Ophthalmology*. 2003; 110(1):177–189. [PubMed: 12511364]

2. Keane PA, Patel PJ, Liakopoulos S, Heussen FM, Sadda SR, Tufail A. Evaluation of age-related macular degeneration with optical coherence tomography. *Survey of Ophthalmology*. 2012; 57(5): 389–414. [PubMed: 22898648]
3. Jindahra P, Hedges TR, Mendoza-Santiesteban CE, Plant GT. Optical coherence tomography of the retina: applications in neurology. *Curr Opin Neurol*. 2010; 23(1):16–23. [PubMed: 20009925]
4. Hajee ME, March WF, Lazzaro DR, Wolintz AH, Shrier EM, Glazman S, Bodis-Wollner IG. Inner retinal layer thinning in Parkinson disease. *Arch Ophthalmol*. 2009; 127(6):737–741. [PubMed: 19506190]
5. Frohman EM, Fujimoto JG, Frohman TC, Calabresi PA, Cutter G, Balcer LJ. Optical coherence tomography: a window into the mechanisms of multiple sclerosis. *Nat Clin Pract Neuro*. 2008; 4(12):664–675.
6. Saidha S, Syc SB, Ibrahim MA, Eckstein C, Warner CV, Farrell SK, Oakley JD, Durbin MK, Meyer SA, Balcer LJ, Frohman EM, Rosenzweig JM, Newsome SD, Ratchford JN, Nguyen QD, Calabresi PA. Primary retinal pathology in multiple sclerosis as detected by optical coherence tomography. *Brain*. 2011; 134(2):518–533. [PubMed: 21252110]
7. Fisher JB, Jacobs DA, Markowitz CE, Galetta SL, Volpe NJ, Nano-Schiavi ML, Baier ML, Frohman EM, Winslow H, Frohman TC, Calabresi PA, Maguire MG, Cutter GR, Balcer LJ. Relation of visual function to retinal nerve fiber layer thickness in multiple sclerosis. *Ophthalmology*. 2006; 113(2):324–332. [PubMed: 16406539]
8. Saidha S, Syc SB, Durbin MK, Eckstein C, Oakley JD, Meyer SA, Conger A, Frohman TC, Newsome S, Ratchford JN, Frohman EM, Calabresi PA. Visual dysfunction in multiple sclerosis correlates better with optical coherence tomography derived estimates of macular ganglion cell layer thickness than peripapillary retinal nerve fiber layer thickness. *Multiple Sclerosis Journal*. 2011; 17(12):1449–1463. [PubMed: 21865411]
9. Garvin M, Abramoff M, Kardon R, Russell S, Wu X, Sonka M. Intraretinal layer segmentation of macular optical coherence tomography images using optimal 3-D graph search. *IEEE Trans Med Imag*. 2008; 27(10):1495–1505.
10. Song Q, Bai J, Garvin M, Sonka M, Buatti J, Wu X. Optimal multiple surface segmentation with shape and context priors. *IEEE Trans Med Imag*. 2013; 32(2):376–386.
11. Chiu SJ, Li XT, Nicholas P, Toth CA, Izatt JA, Farsiu S. Automatic segmentation of seven retinal layers in SDOCT images congruent with expert manual segmentation. *Opt Express*. 2010; 18(18): 19413–19428. [PubMed: 20940837]
12. Dufour PA, Ceklic L, Abdillahi H, Schroder S, De Zanet S, Wolf-Schnurrbusch U, Kowal J. Graph-based multi-surface segmentation of OCT data using trained hard and soft constraints. *IEEE Trans Med Imag*. 2013; 32(3):531–543.
13. Lang A, Carass A, Hauser M, Sotirchos ES, Calabresi PA, Ying HS, Prince JL. Retinal layer segmentation of macular OCT images using boundary classification. *Biomed Opt Express*. 2013; 4(7):1133–1152. [PubMed: 23847738]
14. Li K, Wu X, Chen D, Sonka M. Optimal surface segmentation in volumetric images - a graph-theoretic approach. *IEEE Trans Pattern Anal Mach Intell*. 2006; 28(1):119–134. [PubMed: 16402624]
15. Jones SE, Buchbinder BR, Aharon I. Three-dimensional mapping of cortical thickness using Laplace's equation. *Human Brain Mapping*. 2000; 11(1):12–32. [PubMed: 10997850]
16. Rocha K, Yezzi A, Prince J. A hybrid Eulerian-Lagrangian approach for thickness, correspondence, and gridding of annular tissues. *IEEE Trans Med Imag*. 2007; 16(3):636–648.
17. Acosta O, Bourgeat P, Zuluaga MA, Frapp J, Salvado O, Ourselin S. Automated voxel-based 3D cortical thickness measurement in a combined Lagrangian-Eulerian PDE approach using partial volume maps. *Med Image Anal*. 2009; 13(5):730–743. [PubMed: 19648050]
18. Breiman L. Random forests. *Machine Learning*. 2001; 45(1):5–32.




Figure 1.

Flowchart of our layer segmentation algorithm. Note that feature computation and boundary classification are done on ‘flattened’ images, which is undone to build the adaptive grid.




Figure 2.

Images showing the construction of our adaptive grid. **(a)** Streamlines are generated between initial estimates of the outer retinal boundaries. **(b)** Estimates of each boundary as determined by using a regression model.

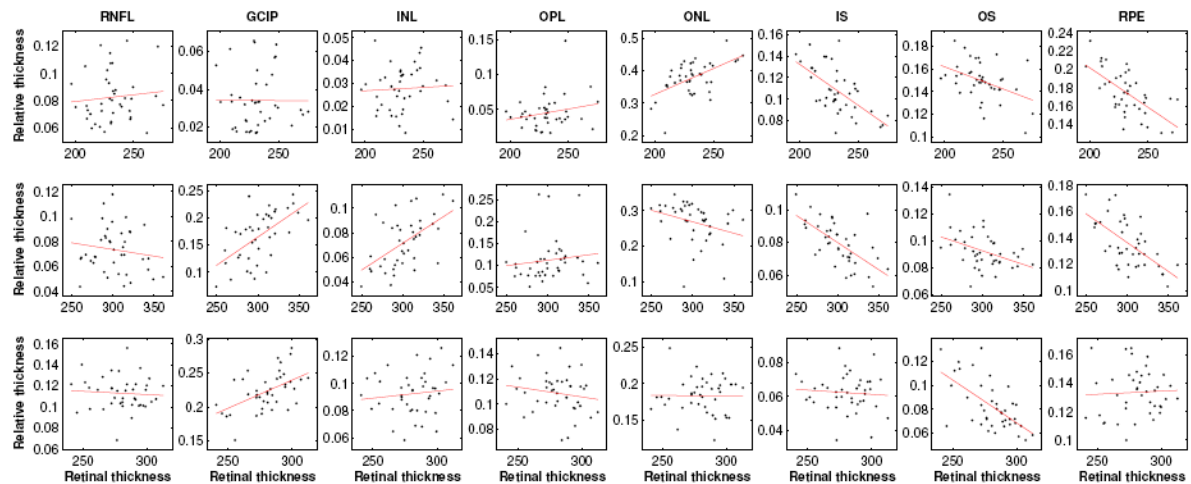


Figure 3.

Examples of fitting the regression model at three separate locations in the retina (from three separate streamlines). Results are shown for each layer (**top row**) at the fovea center, (**middle row**) 0.5 mm from the fovea in the nasal direction to the fovea, and (**bottom row**) 1.5 mm from the fovea in the temporal superior direction. Each black dot represents the measurement from a separate subject (by manual segmentation), with total retina thickness (i.e. the total length of the streamline, in μm) on the x-axis, and the thickness of the layer with respect to the total distance from the ILM to the BM on the y-axis.

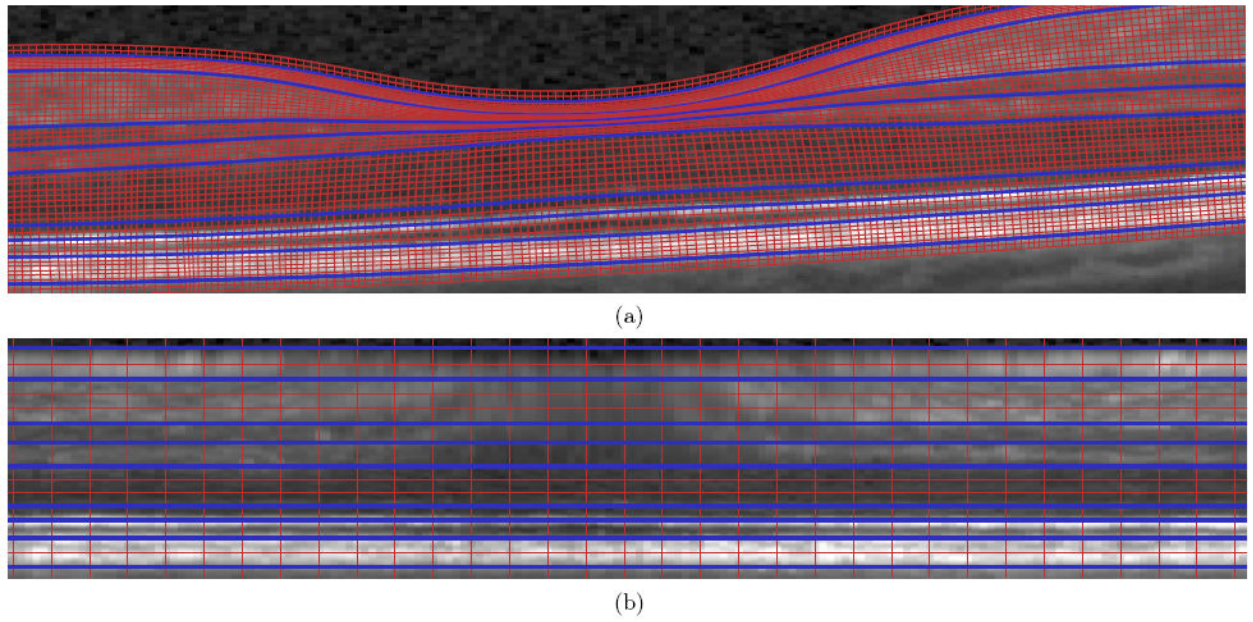


Figure 4.

(a) The final grid overlaid on the retina, constructed after filling in the graph along the streamlines between the regression estimates (shown in blue). The vertices of the graph are located at the intersection of the lines. We denote ‘base nodes’ as those at the intersection of the regression (blue) lines. (b) When looking at the deformed grid as a 4-connected rectangular lattice, we can think of the deformation as ‘flattening’ the data to each boundary (significantly downsampled for visualization).




Figure 5.

Images showing the value of the average absolute error (left) and standard deviation (right) as the parameters for smoothness (w) and for grid size (s) are changed.




Figure 6.

Two B-scan images with overlaid segmentations from (a) the manual ground truth, (b) the voxel grid algorithm, and (c) the deformed grid algorithm. The results from a zoomed in region of the fovea are shown in (d). Images are from separate subjects and have been scaled 3× in the vertical direction.

Table 1

Values of the mean absolute error (standard deviation) for each boundary for the adaptive grid algorithm in comparison to using the voxel grid (as in Lang et al.¹³). Values are in μm and standard deviations are computed across subjects.

	ILM	RNFL-GCL	IPL-INL	INL-OPL	OPL-ONL
Voxel Grid ¹³	3.89 (0.72)	4.06 (0.62)	3.85 (0.49)	3.68 (0.68)	3.52 (0.82)
Adaptive Grid	3.49 (0.65)	3.42 (0.72)	3.81 (0.55)	3.83 (0.64)	3.32 (0.94)
	ELM	IS-OS	OS-RPE	BM	Overall
Voxel Grid ¹³	3.05 (0.40)	2.45 (0.54)	4.60 (1.84)	3.35 (1.68)	3.60 (1.14)
Adaptive Grid	2.81 (0.39)	2.24 (0.61)	4.28 (1.89)	3.20 (1.63)	3.38 (1.15)