

Automatic identification of the temporal retinal nerve fiber raphe from macular cube data

PHILLIP BEDGGOOD,^{1,2} FUMI TANABE,³ ALLISON M. MCKENDRICK,¹ AND ANDREW TURPIN^{2,*}

¹Department of Optometry & Vision Sciences, The University of Melbourne, 3010 Australia

²Department of Computing and Information Systems, The University of Melbourne, 3010 Australia

³Department of Ophthalmology, Kinki University Faculty of Medicine, Osaka, Japan

*aturpin@unimelb.edu.au

Abstract: We evaluated several approaches for automatic location of the temporal nerve fiber raphe from standard macular cubes acquired on a Heidelberg Spectralis OCT. Macular cubes with B-scan separation of 96–122 μm were acquired from 15 healthy participants, and “high density” cubes with scan separation of 11 μm were acquired from the same eyes. These latter scans were assigned to experienced graders for subjective location of the raphe, providing the ground truth by which to compare methods operating on the lower density data. A variety of OCT scan parameters and image processing strategies were trialed. Vertically oriented scans, purposeful misalignment of the pupil to avoid reflective artifacts, and the use of intensity as opposed to thickness of the nerve fiber layer were all critical to minimize error. The best performing approach “cFan” involved projection of a fan of lines from each of several locations across the foveal pit; in each fan the line of least average intensity was identified. The centroid of the crossing points of these lines provided the raphe orientation with an average error of 1.5° (max = 4.1°) relative to the human graders. The disc-fovea-rape angle was $172.4 \pm 2.3^\circ$ (range = $168.5\text{--}176.2^\circ$), which agrees well with other published estimates.

© 2016 Optical Society of America

OCIS codes: (170.4500) Optical coherence tomography; (100.0100) Image processing; (100.3008) Image recognition, algorithms and filters; (170.4580) Optical diagnostics for medicine.

References and links

1. F. Vrabec, “The temporal raphe of the human retina,” *Am. J. Ophthalmol.* **62**(5), 926–938 (1966).
2. D. C. Hood, A. S. Raza, C. G. V. de Moraes, J. M. Liebmann, and R. Ritch, “Glaucomatous damage of the macula,” *Prog. Retin. Eye Res.* **32**, 1–21 (2013).
3. N. M. Jansonius, J. Nevalainen, B. Selig, L. M. Zangwill, P. A. Sample, W. M. Budde, J. B. Jonas, W. A. Lagrèze, P. J. Airaksinen, R. Vonthein, L. A. Levin, J. Paetzold, and U. Schiefer, “A mathematical description of nerve fiber bundle trajectories and their variability in the human retina,” *Vision Res.* **49**(17), 2157–2163 (2009).
4. D. F. Garway-Heath, D. Poinosawmy, F. W. Fitzke, and R. A. Hitchings, “Mapping the visual field to the optic disc in normal tension glaucoma eyes,” *Ophthalmology* **107**(10), 1809–1815 (2000).
5. A. Turpin, G. P. Sampson, and A. M. McKendrick, “Combining ganglion cell topology and data of patients with glaucoma to determine a structure-function map,” *Invest. Ophthalmol. Vis. Sci.* **50**(7), 3249–3256 (2009).
6. A. M. McKendrick, J. Denniss, and A. Turpin, “Structure-function mapping for individuals - why individuals not populations are needed to determine the utility of customising structure-function mapping,” in *Association for Research in Vision and Ophthalmology Annual Meeting* (Denver Colorado, May 05–07, 2015).
7. J. Denniss, A. M. McKendrick, and A. Turpin, “An Anatomically Customizable Computational Model Relating the Visual Field to the Optic Nerve Head in Individual Eyes,” *Invest. Ophthalmol. Vis. Sci.* **53**(11), 6981–6990 (2012).
8. P. V. Le, O. Tan, V. Chopra, B. A. Francis, O. Ragab, R. Varma, and D. Huang, “Regional Correlation Among Ganglion Cell Complex, Nerve Fiber Layer, and Visual Field Loss in Glaucoma,” *Invest. Ophthalmol. Vis. Sci.* **54**(6), 4287–4295 (2013).
9. N. Amini, S. Nowroozizadeh, N. Cirineo, S. Henry, T. Chang, T. Chou, A. L. Coleman, J. Caprioli, and K. Nouri-Mahdavi, “Influence of the Disc-Fovea Angle on Limits of RNFL Variability and Glaucoma Discrimination,” *Invest. Ophthalmol. Vis. Sci.* **55**(11), 7332–7342 (2014).
10. Y. K. Kim, B. W. Yoo, H. C. Kim, and K. H. Park, “Automated Detection of Hemifield Difference across Horizontal Raphe on Ganglion Cell–Inner Plexiform Layer Thickness Map,” *Ophthalmology* **122**(11), 2252–2260 (2015).

11. F. Tanabe, C. Matsumoto, S. Okuyama, S. Takada, T. Numata, T. Kayazawa, M. Eura, S. Hashimoto, E. Koike, and Y. Shimomura, "Imaging of temporal retinal nerve fiber trajectory with Transverse Section Analysis," in *Association for Research in Vision and Ophthalmology Annual Meeting* (Orlando Florida, May 04–08, 2014).
12. B. C. Chauhan, G. P. Sharpe, and D. M. Hutchison, "Imaging of the temporal raphe with optical coherence tomography," *Ophthalmology* **121**(11), 2287–2288 (2014).
13. G. Huang, T. J. Gast, and S. A. Burns, "In vivo adaptive optics imaging of the temporal raphe and its relationship to the optic disc and fovea in the human retina," *Invest. Ophthalmol. Vis. Sci.* **55**(9), 5952–5961 (2014).
14. G. Huang, T. Luo, T. J. Gast, S. A. Burns, V. E. Malinovsky, and W. H. Swanson, "Imaging Glaucomatous Damage Across the Temporal Raphe," *Invest. Ophthalmol. Vis. Sci.* **56**(6), 3496–3504 (2015).
15. J. D'Errico, "Inpaint NaNs," (2005), MATLAB Central File Exchange: <http://www.mathworks.com/matlabcentral/fileexchange/4551-inpaint-nans>, Accessed Mar 24, 2016.
16. A. S. Reis, G. P. Sharpe, H. Yang, M. T. Nicolela, C. F. Burgoyne, and B. C. Chauhan, "Optic disc margin anatomy in patients with glaucoma and normal controls with spectral domain optical coherence tomography," *Ophthalmology* **119**(4), 738–747 (2012).

1. Introduction

Glaucoma is an optic neuropathy that results in progressive loss of retinal ganglion cells and their axons, leading to visual field loss. Consequently, contemporary clinical assessment of eyes with glaucoma typically includes both anatomical assessment (with optical coherence tomography imaging, or OCT) in addition to the assessment of peripheral visual function (measured subjectively with perimetry). In order to combine OCT and perimetric data, it is necessary to map locations in visual space to the anatomical images. Since axons originating in the temporal retina follow a path to the optic nerve head that avoids the fovea, such maps need to incorporate the location of the *temporal raphe*: the boundary between axons that follow a superior course and those that follow an inferior course [1]. The location of the temporal raphe is particularly important for glaucoma because characteristic visual field losses, such as nasal steps and arcuate shaped patterns, occur in this region [2].

The geometry of the temporal raphe can be usefully described by the relative angles made between the fovea and the disc ("FoDi"), and the fovea and the raphe ("FoRaph") – these angles can be combined to define the disc-fovea-raphe angle. The FoDi and FoRaph angles are depicted in Fig. 1(a). Previous tools for analysis of anatomical images and visual field data have assumed either that the temporal raphe lies along a horizontal line [3, 4], i.e. the line R1 in Fig. 1; or that the temporal raphe is a continuation of the line connecting the center of the disc to the center of the fovea [5], i.e. the line R2 in Fig. 1.

Small deviations in the orientation of the temporal raphe from the assumed orientation can result in large deviations in the structure-function map, since key points in the visual field will then map to the opposite side of the optic nerve head. For example in Fig. 1(b), if the temporal raphe is at R1 or R2 then all points in the lower half of the visual field will follow a superior course (due to inversion of the visual scene imaged by the eye) and so map to the superior side of the optic disc. However, if the raphe is actually at R3, two or three of those points now lie above the raphe and so, by definition, will connect with the disc on the inferior side. Modelling work incorporating measured population statistics for various parameters of ocular biometry predicts that 14% of the patient population should undergo such a "flipping" in their structure-function map, with 50% likelihood of a lesser shift of up to 30° [6]. Since visual sensitivity and thickness can change substantially over such distances in glaucoma, the utility of applying any structure-function map based on population norms becomes severely reduced in many individual patients.

Rather than assuming a particular standard orientation of the temporal raphe, it has been suggested that its orientation be measured on a per-patient basis, and a customized mapping of visual field and OCT data created [7]. Some efforts have been made towards measuring or inferring the orientation of the raphe with currently available clinical tools. Hood et al. located the axis of minimum RNFL thickness from 128 healthy eyes, reporting approximately horizontal orientation of the raphe despite a large physiological variation in the relative positions of the optic disc and fovea [2]. In glaucomatous eyes, Le et al., Amini et al. and

Kim et al. have exploited characteristic patterns of glaucomatous damage coupled with measured visual field loss to infer the axis between superior and inferior fields, again relying on changes in thickness of particular layers measured with OCT [8–10]. Each of the above studies concluded that the temporal raphe is oriented roughly along a horizontal line.

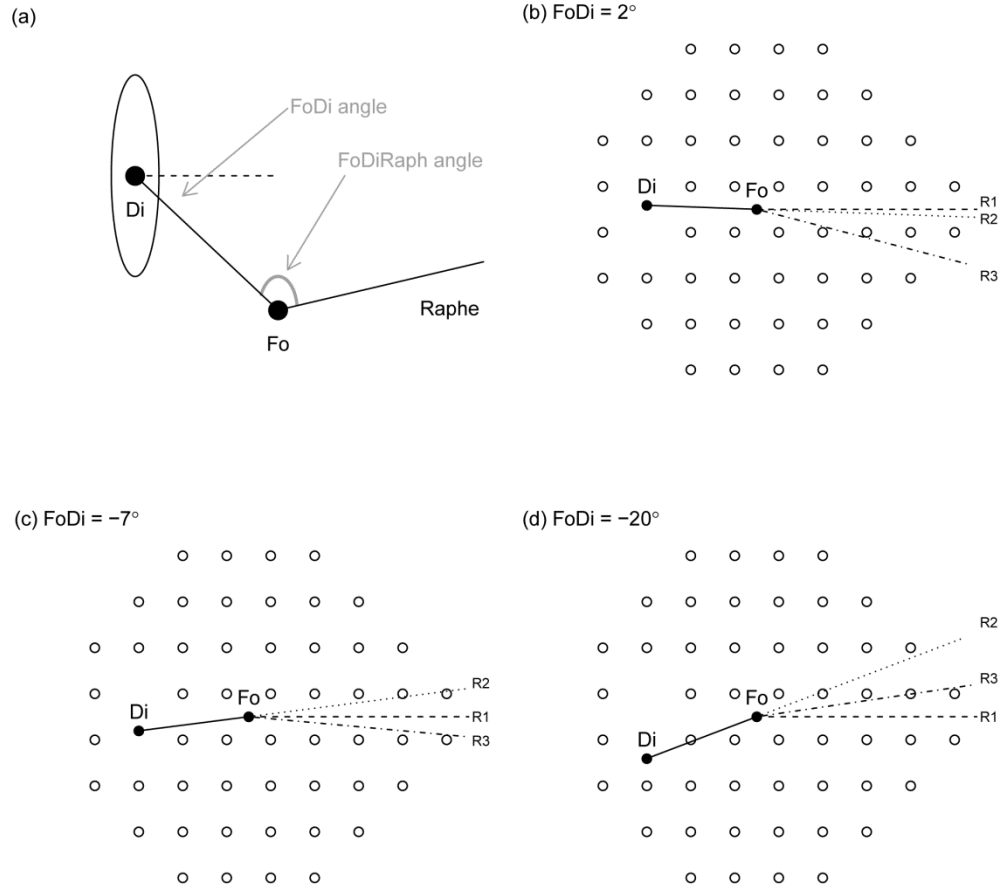


Fig. 1. Small differences in the raphe, or axis of symmetry between superior and inferior visual fields, can have a large impact on the mapping of certain visual field test locations to anatomical tissue. a) Shows the definition of angles described in this paper, where Fo indicates the fovea, Di the disc (or more properly, the blind spot), FoDi the angle between fovea and disc, and FoDiRaph the angle between fovea, disc and temporal raphe. Panels (b), (c) and (d) show the commonly used 24-2 visual field pattern with 3 different FoDi angles, each with 3 definitions of the FoRaph angle: R1 always has the raphe horizontal ($FoRaph = 180^\circ + FoDi$); R2 continues the line from Di to Fo ($FoRaph = 180^\circ$); and R3 has $FoRaph = 172^\circ$, which is close to the population average.

Recently, the feasibility of obtaining more detailed visualizations of the temporal raphe has been demonstrated with state of the art imaging techniques that feature high axial and/or transverse resolution, such as optical coherence tomography (OCT) or adaptive optics scanning laser ophthalmoscopy (AOSLO) [11–14]. These methods allow more precise characterization of the raphe, and have reported a deviation of temporal raphe from horizontal, in healthy eyes, by up to $5\text{--}9^\circ$ depending on the individual. Unfortunately, these methods require long acquisition and analysis times, and so are not currently well suited for broad application to clinical populations. Nevertheless, such high resolution data demonstrates that there is substantial variation in the position of the temporal raphe in different eyes.

The aim of this study was to develop improved analysis methods for evaluating the orientation of the temporal raphe from OCT, using lower resolution macular cube scan protocols which are suitable for routine clinical use. We utilize high density spectral domain OCT scans to localize the temporal raphe, and compare the results to those obtained from lower density or “standard” macular cube data from the same eyes. This approach allows us to evaluate the precision with which the raphe can be automatically located from standard clinical scans of the macula. We explore a variety of approaches including the use of thickness versus intensity of the retinal nerve fiber layer (RNFL), horizontal versus vertical scan orientation, and evaluate several algorithms for automatic localization of the raphe.

2. Methods

2.1 Overview

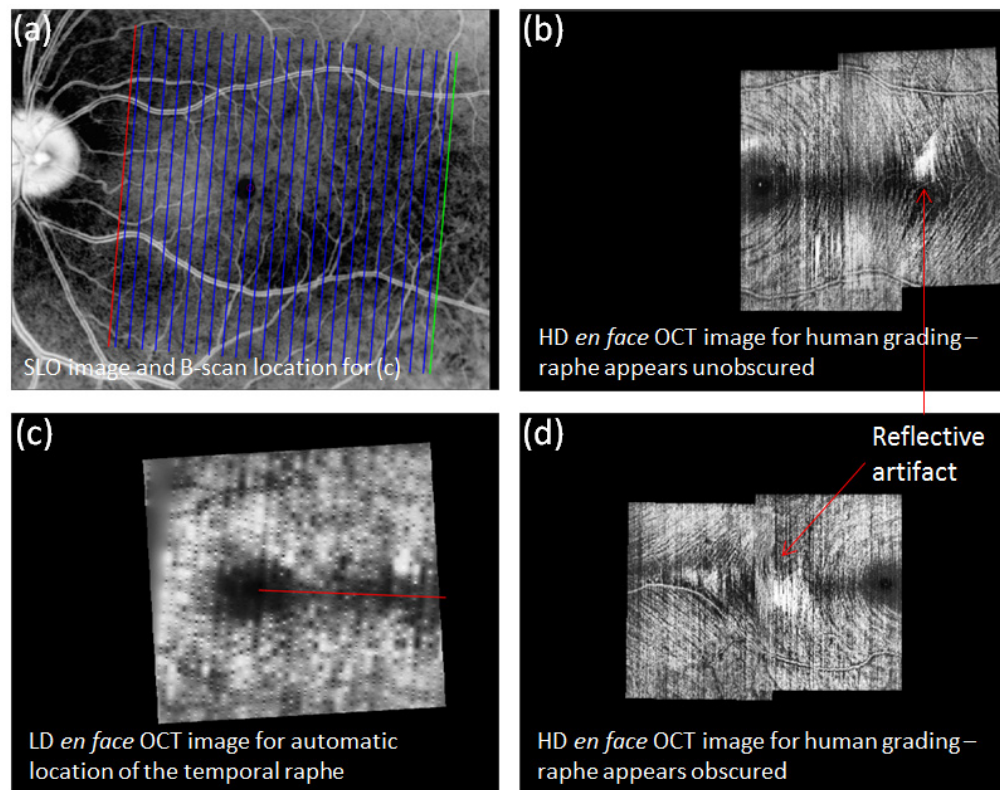


Fig. 2. Example of data acquisition and automated localization of the temporal raphe in two subjects. a) IR SLO image of the fundus for the subject of a, b and c. Vertical lines: location of B-scans used to generate image shown in c. b) High resolution *en face* OCT image of B-scan intensity at the RNFL, used for human grading. Image generated from two sets of high-resolution, overlapping scans. An area of dark (and thinned) nerve fiber layer extends temporally from the fovea, which we take to be the temporal raphe. A much lower resolution version of this area is evident in c. c) Similar to b, but with data obtained from a standard macular cube which has more widely spaced scans. Red line: example of an automated attempt to define the fovea-raphe angle. d) Similar to b but obtained from a different subject with a particularly reflective fundus. Large reflective patches (arrow) obscure much of the temporal raphe.

Using the Heidelberg Spectralis OCT (Heidelberg Engineering GmbH, Germany), a series of macular scans were acquired from each subject with inter-scan separation either $11\ \mu\text{m}$, or $96 - 122\ \mu\text{m}$. The latter is commensurate with a typical “macular cube” acquired routinely in clinical examinations and takes only seconds to acquire, while the former produces very

detailed images but requires several minutes to acquire, and so is unsuited for routine clinical use. We refer to these two scan profiles as “high density” and “low density”, respectively. We assigned the high density scans to three trained, independent graders to label the position of the raphe by visual inspection, and compared these labeled positions to the output of several automated approaches operating on the low density scans only. An overview of our data acquisition pipeline is shown in Fig. 2.

2.2 Subjects

We recruited 15 healthy participants from staff and students at the Department of Optometry and Vision Sciences from the University of Melbourne, ranging in age from 19 to 45 years. Subjects with refractive error above 5 DS or 2 DC were excluded, as were those with any notable ocular pathology. Subject pupils were not dilated.

2.3 Generation of *en face* images

Both thickness [2, 8, 9] and intensity [11–14] of the RNFL have been used by other groups to estimate the orientation of the temporal raphe, and both are reasonable indicators since the apparent “missing” fibers at the raphe should reduce both thickness and intensity recorded by OCT at the level of the RNFL.

The location of each B-scan from the Spectralis is specified by coordinates on the SLO image acquired at the beginning of the scan, as shown by the vertical lines in Fig. 2(a). To generate an *en face* image from the B-scans, we placed 1D information extracted from each B-scan into an image array at the specified coordinates.

In the case of the *en face* thickness map, the information used was the output of the automatic segmentation feature of the Spectralis software, which reports the thickness of the nerve fiber layer. No errors in the automated segmentation were apparent in this young, healthy data set, and so no manual correction of segmentation was required.

In the case of the *en face* intensity image (e.g. Fig. 2(b), 2(c), 2(d)), the information used was the intensity of each B-scan at a desired depth. Depths assessed ranged from the Inner Limiting Membrane (ILM) to the RNFL in steps of 20% of the local RNFL thickness for each A-scan (6 depth locations). Each algorithm was evaluated at the posterior edge of the RNFL, which appeared subjectively to maximize overall contrast, but graders were supplied with data from the full range of depths.

For each type of image, we interpolated to smoothly fill the missing data between adjacent scans. 2D image interpolation is a non-trivial task, especially considering the relative sparsity of the acquired scans in a typical macular cube. The method that we adopted [15] propagates intensity information from known to unknown areas in a manner similar to the solution of heat flow problems in thermodynamics (that is, a linear system of partial differential equations is established over the imaged area, and solved for the unknown regions). Other interpolation approaches may also be suitable but were not explored.

2.4 Avoidance of reflective artifacts

Especially in younger patients, reflective artifacts (e.g. Fig. 2(b) and 2(d)) at the level of the RNFL can manifest in the B-scans, depending on patient alignment. In general, we have found that these can be avoided entirely if purposeful misalignment of the eye is introduced. However, since our analysis was intended to apply to images as they would typically be acquired in a clinical setting, in several patients reflections remained and interfered with the ability of automated approaches to accurately determine the orientation of the temporal raphe. Two main strategies were trialed to minimize the influence of these reflections.

- Threshold the image to remove pixels above a certain fraction of the maximum image intensity (chosen empirically as 70% of the maximum image intensity).

- Along each line projected from the fovea (explained further below), find the optimal 100-pixel window that minimizes the criterion of interest (as opposed to calculating the criterion along the entire line).

2.5 Automated detection of reflective artifacts

In addition to the above strategies, we manually rejected and re-acquired high- and low-density scans deemed to suffer from an excess of reflective artifacts in the region of the raphe. We were able to develop a criterion that can be used to automatically identify such images, which we term “DeltaRef”. The criterion is based on the observation that reflections transition more rapidly with depth than does the appearance of the anatomical tissue itself. This observation held true at the level of the nerve fiber layer in each of the eyes for which reflective artifacts were apparent. Therefore, we calculated the difference in *en face* image intensity between the nerve fiber layer and the internal limiting membrane. The image histogram of this difference image, in the region of the raphe, was markedly skewed towards higher intensities in images suffering from marked reflective artifacts. The performance of this approach agrees well with manual image classification as shown in Fig. 3.

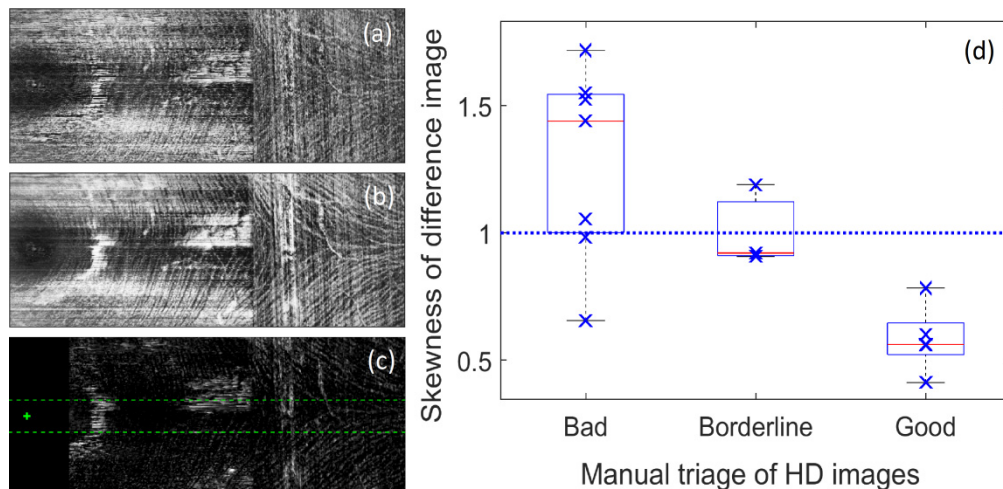


Fig. 3. “DeltaRef” approach for automated identification of reflective artifacts at the temporal nerve fiber raphe. a) Shows an *en face* intensity image at the level of the ILM; b) shows an *en face* intensity image averaged over the NFL; c) shows the difference between these images. The center of the fovea is marked with a cross – tissue within 0.75 mm of this point was excluded (black rectangle). The influence of reflections was calculated in a region of interest indicated by the green dashed lines; d) DeltaRef statistic, which shows skewness of the region of interest in the difference image, is plotted against manual classification of image reflections. A skewness of ~ 1.0 is a common rule of thumb which appears to be a reasonable criterion for triage of image quality to minimize reflective artifacts.

Although the DeltaRef approach robustly detects reflective artifacts at the level of the NFL, such artifacts may not always corrupt the measured raphe angle. However, since it is difficult to ascertain whether any given artifact will influence the raphe angle or not, we recommend a conservative approach whereby images showing high values for DeltaRef are rejected and the scans re-acquired using purposeful misalignment as described above.

2.5 Avoidance of the fovea

Since our metrics of thickness and intensity are also diminished at the fovea, the fovea should be excluded from the image for the purposes of locating the raphe. We assumed a foveal radius of 0.75 mm (2.5°) and excluded image pixels that lay closer than this to the fovea. Whilst this may not entirely exclude the fovea in each eye, this is not critical – we must only

exclude a sufficient portion of fovea such that the metric of interest (thickness or intensity) is minimized at the raphe, instead of at the fovea.

2.5 Reduction of image non-uniformities

Due to small movements of the head and eyes, overall brightness of each B-scan can change throughout the course of acquisition. To somewhat mitigate the resulting non-uniformities in the image, we equalized the histogram of each image using the default CLAHE (Contrast Limited Adaptive Histogram Equalization) algorithm included in the Matlab Image Processing Toolbox R2015b.

2.6 Location of anatomical landmarks

The fovea was located by finding the minimum total retinal thickness within a 6° area centered on the expected foveal location. This position was calculated independently for each scan rather than basing the location on any particular scan, as we found that small image distortions due to intra-frame eye movements can significantly shift the location of the fovea in the image.

The Spectralis is able to automatically determine the location of the optic disc by finding the center of the opening in Bruch's membrane, which may be more accurate than labeling based on the apparent position of the optic disc from a fundus image [16]. In our protocol, for each patient a series of radial scans was made through the optic nerve head, from which the Spectralis automatically finds the axis connecting the fovea and center of Bruch's membrane opening ("FoDi" axis as in Fig. 1(a)). The subsequent macular scans were aligned either parallel ("horizontal" scans) or perpendicular ("vertical" scans) to this axis by reference to a SLO image acquired just before the OCT scan commenced; this allowed correction for differences in head tilt between scans which would otherwise make the raphe appear to change direction.

2.7 Automated determination of raphe angle

Several approaches to identify the temporal raphe were developed and implemented, and evaluated by comparing the returned raphe orientation to the "ground truth" angles obtained by human graders. These approaches are specified below.

Each algorithm was evaluated for both thickness and intensity images, and for both thresholded images and windowed projections, as described above.

1. Lowest averaging ray projected from the foveal center – "fov"

Lines were traced from the fovea in a fan spanning $\pm 20^\circ$ in 0.25° increments. The line that minimized either thickness or intensity was chosen as the raphe (e.g. "fov" in Fig. 4).

2. Lowest averaging ray projected from the foveal pit – "opt"

Rays were traced not only from the foveal center but from points spanning ± 0.5 mm (in 0.05 mm increments) from the foveal center in a vertical line. The line with minimal average thickness or intensity, from the set of all lines, was chosen as the raphe ("opt" in Fig. 4). This method avoids the assumption that the raphe trajectory should intersect the foveal center, and can potentially avoid small reflective artifacts.

3. Median of lowest averaging rays projected from the foveal pit – "med"

This approach uses the same data from the "opt" method, however, instead of identifying the minimum of all lines considered, the minimal line from each individual point is determined and the median of these line angles is then taken as the raphe.




Fig. 4. Illustration of methods evaluated for automatic determination of temporal raphe angle. In each method, fans of lines (blue) are traced from a point in the foveal pit (f) and a “minimal line” is determined; that is, the line in each fan which encounters the minimum average intensity (examples in green) or thickness. For the “fov” algorithm, the raphe is the minimal line stemming from the fovea. For “opt”, the raphe is the minimum of all minimal lines. For “med” (not illustrated), the median of the minimal lines is determined. For “cFan”, the intersections of each minimal line with each other minimal line are found, and the centroid of these crossing points (yellow circle) is taken to lie along the raphe.

4. Centroid of intersections of lines projected from the foveal pit – “cFan”

This method also uses the same data from the “opt” method. As depicted in Fig. 4, the intersections of the minimal line from each point with those from every other point are determined. This creates a cluster of line intersections, the centroid of which is taken to lie along the raphe (yellow circle in Fig. 4). The angle that this centroid makes with respect to the fovea is taken as the raphe angle.

2.8 Human grader assessment of raphe angle

For each eye, high density *en face* intensity images were generated for 6 equally spaced depths spanning the full depth of the NFL, as well as an average across the entire NFL. Images were cropped to share a common area with the corresponding low density scan, and blinded to de-identify subject information before being passed to the graders. Graders were asked to independently label the apparent trajectory of the raphe by labeling between 1 and 5 points across the image. The angle made by each point with the center of the fovea was calculated, and an average taken to determine the raphe angle assigned by that grader. The average of the 3 grader raphe angles was taken as the ground truth for comparison of algorithm performance.

3. Results

The differences in raphe angle between each algorithm and the human graders are shown in Table 1 (average error) and Table 2 (maximum error). The fovea-raphe lines for each grader were averaged to provide the gold standard against which performance of our automated approaches was compared. Errors are expressed as the difference from this average.

Regardless of the particular algorithm used, the results show substantially less error when using intensity as opposed to thickness data, and when using vertical as opposed to horizontal scans. Attempts at minimizing the impact of reflective artifacts by applying a threshold or window were largely ineffective, with angles derived by all algorithms being at least as close, and usually closer, to the human graders when using raw images.

Table 1. Mean absolute difference between the angle estimated by each algorithm (column headings) and the mean of 3 human graders (degrees). The smallest value for each scan type is shown in bold.

Image Type	Processing	Macular cube (vertical scans)				Macular cube (horizontal scans)			
		fov	cFan	med	opt	fov	cFan	med	opt
Intensity	Raw	1.7	1.5	1.6	5.3	5.4	3.7	3.5	4.9
	Threshold	1.9	1.7	1.8	4.2	5.7	4.3	4.0	5.7
	Window	2.7	2.8	2.9	5.3	6.4	5.0	4.3	5.0
Thickness	Raw	11.2	7.5	7.5	6.1	11.3	7.5	9.1	9.6
	Thresh	11.2	7.5	7.5	6.4	11.3	7.5	9.1	9.6
	Window	8.3	4.6	5.3	8.8	9.7	6.5	7.4	7.1

Table 2. Maximum absolute difference between the angle estimated by each algorithm (column headings) and the mean of 3 human graders (degrees). The smallest value for each scan type is shown in bold.

Image Type	Processing	Macular cube (vertical scans)				Macular cube (horizontal scans)			
		fov	cFan	med	opt	fov	cFan	med	opt
Intensity	Raw	4.8	4.1	4.8	12.3	13.7	12.2	7.9	9.8
	Threshold	8.5	6.6	8.5	11.0	13.7	12.2	7.9	11.4
	Window	6.1	6.2	6.1	17.7	15.2	11.1	9.6	10.1
Thickness	Raw	23.5	15.1	16.0	17.5	20.9	14.3	17.3	21.1
	Thresh	23.5	14.2	15.6	17.5	20.9	14.3	17.3	21.1
	Window	18.0	12.4	12.1	20.3	19.2	15.1	15.0	13.8

Of the algorithms considered, “cFan” gave the lowest error when vertical scans were used, with “med” being slightly better on horizontal scans. These differences were not statistically significant ($p > 0.28$), however, the sample size is small ($n = 15$). The “opt” algorithm gave substantially higher error compared to human grader angles as expressed relative to the fovea; this suggests that the raphe may not point precisely at the fovea (discussed further below).




Fig. 5. Algorithm performance for HD (left) and LD image (right) from subject with highest error. Red lines show algorithm, yellow spots show one of our grader’s labels. Black bar in left image resulted from non-overlapping scans. The LD image has some error but it is hard to fault it based on this image alone; intensity artifacts due to the low spatial sampling appear to be the limiting factor.

Given errors as large as 4° , it is instructive to look at individual examples to determine whether the algorithms perform with as much precision as should be expected given the low density of the scans. Figure 5 shows the high density (left) and low density (right) scans for the eye that yielded maximum error. The labels made by one grader on the high density image only are indicated on both images with yellow crosses, and the output of the cFan algorithm for each image is shown with a red line. It can be seen that the algorithm and grader agreed well for the HD image, but the angle found by the algorithm differs somewhat in the low density image (by 4.1°). Nonetheless it is not obvious how to make a better/different

approximation to the raphe in this low density image alone, and so we deem the algorithm successful in this case.

To assess whether the cFan algorithm conferred any systematic bias, we plotted the error in the FoRaph angles returned by the algorithm against the ground truth grader angle (Fig. 6(a)) and the FoDi angle (Fig. 6(b)). No significant correlation was found with either parameter, indicating a lack of systematic bias.

Another check to confirm the validity of our approach is to compare the apparent population spread in raphe orientation with that determined from previous studies. The absolute values of both the FoDi and FoRaph angles defined in Fig. 1 are subject to rotations of the head/eyes, but the angle formed by all 3 structures (disc-fovea-raphe) is not. In our 15 subjects this angle was $172.4 \pm 2.3^\circ$ (ranging from 168.5 to 176.2°), which agrees well with prior estimates [9, 12, 13].

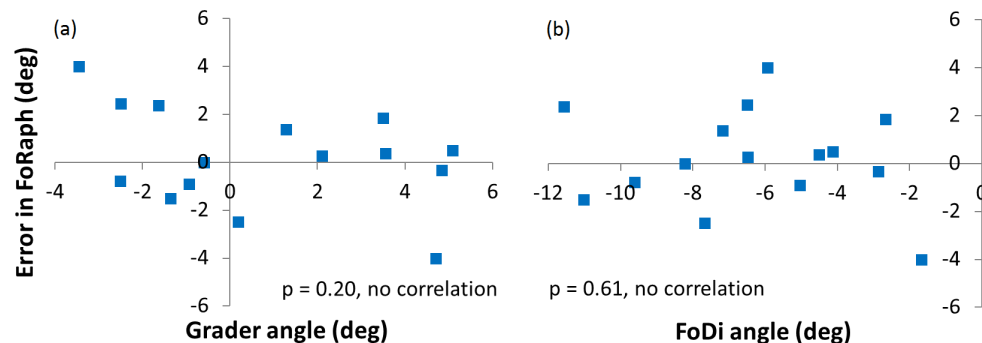


Fig. 6. Assessment of systematic bias of cFan algorithm. a) Plotted as a function of ground truth FoDiRaph angle determined from average of 3 grader estimates. b) Plotted as a function of FoDi angle as returned from the Spectralis device. No significant bias is apparent in either case.

4. Discussion

Based on the range of approaches trialed here, automated estimation of the temporal raphe angle from macular cubes is most accurate when using intensity images derived from vertically oriented B-scans. Generally, it was better to avoid reflective artifacts by purposeful misalignment of the pupil with respect to the device, rather than attempt to use thresholding or windowing to remove the reflections *post hoc*. We presented the DeltaRef method which can be used to objectively identify images that suffer from excessive reflective artifacts in the region of the temporal raphe, which could be added to existing OCT software to assist with the collection of reliable data.

For estimation of the trajectory of the temporal raphe, the best performing algorithm (cFan) made independent estimations of the raphe angle from a variety of positions along a vertical line spanning the foveal pit; the resulting raphe lines all tend to cross each other near the center of the raphe, giving a robust estimate of the raphe angle. Using this approach in 15 healthy eyes the raphe angle was on average 1.5° away from the average value assigned by the human graders, and at most 4.1° from these values.

The success of the intensity compared with thickness data is likely a result of the inherent difficulty in using automated methods to reliably determine the thickness of the nerve fiber layer close to the raphe, where the nerve fiber layer is so thin. Methods based on thickness [2, 9] could possibly be improved by manual exclusion or editing of the automated segmentation of retinal layers, however, we advocate an alternate approach based on intensity.

The success of vertical compared with horizontal scans is likely a product of the inherent difficulty in appropriately scaling the intensity of each acquired B-scan in order to compensate for head and eye movements. In the vertical scans, the diversity of anatomical landmarks makes this an easier task; that is, each scan contains both the dark raphe and the

bright nerve fiber bundles. In contrast, some horizontal scans through the raphe contain only the dark raphe area, making it more difficult for the device to determine the appropriate intensity scaling for those B-scans and resulting in more visible horizontally oriented artifacts in the *en face* image. As Fig. 5 (right) showed, even the vertical scans do contain some degree of acquisition artifact that can displace the apparent trajectory of the raphe.

Our “opt” algorithm produced substantially higher error values than our “fov” algorithm. Both methods select the dimmest line projected from the foveal pit; in the latter the only point considered is the center of the fovea itself, whereas in the former the full gamut of points across the vertical extent of the fovea are considered. The fact that this algorithm was consistent in returning an angle markedly different (4-5°) from the human grader angles, which were referenced to the foveal center, implies that the raphe may not be a straight line that extends from the fovea as is commonly assumed. The offset producing the dimmest of all projected lines ranged in value across our images from -0.5 to + 0.5 mm from the foveal center, with an average close to zero. This suggests that if the raphe does not point straight at the foveal center, there is no clear population bias towards the superior or inferior direction.

High density scans were required here as “ground truth”, enabling validation of measurements made from more clinically-suited scans in a young healthy population. It is possible in principle to repeat this exercise in a glaucoma patient population, however, we have found the high density scans very difficult to acquire in older, untrained patients due to large fluctuations in fixation and tear film quality. This substantially increases the time required for validation scans, since the device must continually re-acquire data until image quality is sufficient for registration to occur. This underscores the benefit of the present work which suggests that high density scans, impractical to acquire in older patients, are not actually necessary for reliable measurement of raphe orientation. Further confirmation that our method is as applicable to the patient population as it is to healthy individuals could be achieved by using low density macular cube data to compare the distributions in raphe orientation between the two groups. However, even this may produce unexpected results due to changes in raphe shape that apparently occur as glaucoma progresses [2, 8–10,14]. The most instructive approach to improve our understanding in the future will be measurement of the raphe orientation when the retina appears healthy, and tracking the raphe shape and/or orientation in the same patient as the disease progresses. The method published here makes this approach far more clinically feasible, and we hope that the work will facilitate more widespread adoption of customized structure-function mapping in the field.

5. Conclusion

We present an algorithm to automatically detect the trajectory of the temporal nerve fiber layer from standard OCT macular cubes with an average error of ~1.5°. Based on our visual inspection of the data, it would probably be difficult to improve on this accuracy without increasing scan density. To obtain this level of accuracy it is recommended to avoid reflective artifacts through purposeful misalignment of the retina during image acquisition, use vertical as opposed to horizontal scans, and assess the intensity of the nerve fiber layer as opposed to its thickness.

Funding

Australian Research Council (ARC) Linkage Project LP13100055.

Commercial disclosure: Research support from Heidelberg Engineering GmbH, Heidelberg, Germany.