Automatic segmentation of retinal and choroidal thickness in OCT images using convolutional neural networks

David Alonso-Caneiro, Scott A. Read, Jared Hamwood, Stephen J. Vincent, Michael J. Collins

Contact Lens and Visual Optics Laboratory, School of Optometry and Vision Science

Queensland University of Technology, Brisbane, Australia

PURPOSE

To evaluate the performance of a fully automatic method based on a deep learning approach to segment retinal and choroidal boundaries in OCT images, and derive retinal thickness (RT) and choroidal thickness (ChT) using data obtained from a healthy pediatric cohort[1].

METHODS

Custom designed convolutional neural networks (CNN) were trained to classify three boundaries; the inner limiting membrane (ILM), the retinal pigment epithelium (RPE) and the chorio-scleral interface (CSI) (Fig 1).

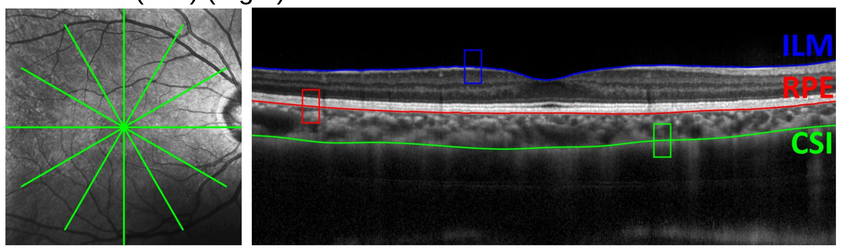

Figure 1. Example of the spectral domain OCT B-scan captured using the instrument's high resolution scanning protocol including the fundus image and B-scan. The B-scan is shown with the three boundaries of interest, including the retinal pigment epithelium (RPE, red), the inner limiting membrane (ILM, blue) and the chorio-scleral interface (CSI, green). The boxes provides an example of the patch size used for training.

The CNN predicts the boundary location and provides a per-layer probability map, which is then used to trace the boundary using graph-search methods.[2] The method runs a patch (window) of fixed size through the image to provide the probability of a particular boundary being present in the center of that window.

The CNN uses a rectangular patch size of 61x31 (VxH) pixels to train the network and three network-input options were tested during training; (i) standard intensity, (ii) attenuation coefficient[3] equivalent and (iii) a combination of both (dual). For each option, the network was trained on the same 137 randomly selected B-scans (70 subjects) and validated on 28 images to ensure adequate training. The network was then tested on 30 different B-scans from 30 different subjects.

To test repeatability, consecutive images from the same subject/location were used. The CNN outputs a probability map for each boundary position that was traced with a graph-search technique. The results from the automatic method were compared to data from manual segmentation by an experienced observer.

Table 1. The difference in boundary position and thickness error for each network-input and the manual observer for the entire dataset. The results are reported as the mean value (standard deviation) in pixel units (1 pixel = 3.9 µm).

	Network input		
	Intensity	Attenuation Coefficient	Dual
Boundary absolute error – mean (SD) pixels			
ILM	0.64 (2.40)	0.64 (2.37)	0.62 (0.47)
RPE	0.54 (0.44)	0.60 (0.51)	0.56 (0.47)
CSI	3.84 (5.64)	3.64 (5.28)	3.64 (5.25)
Thickness error – mean (SD) pixels			
RT	-0.17 (1.08)	-0.12 (1.85)	-0.13 (1.28)
ChT	-1.59 (7.32)	-1.97 (7.55)	-1.41 (7.28)

RESULTS

The well-defined ILM and RPE boundaries showed small errors (<1 pixel) in comparison to the CSI which exhibited slightly larger errors (4 pixels) across all tested options (Table 1, Fig 2). The different CNN inputs had a small effect on the boundary error, with the dual input yielding a slightly smaller mean error and SD.

Analysis of the ChT and RT, revealed errors of -0.1 and -2 pixels respectively. The mean repeatability difference results (in pixels) for the RT [1.10;1.09;1.08] and ChT [3.40;3.43;3.38] across all input options, were comparable with the repeatability from manual segmentation [RT 1.24,ChT 2.51].

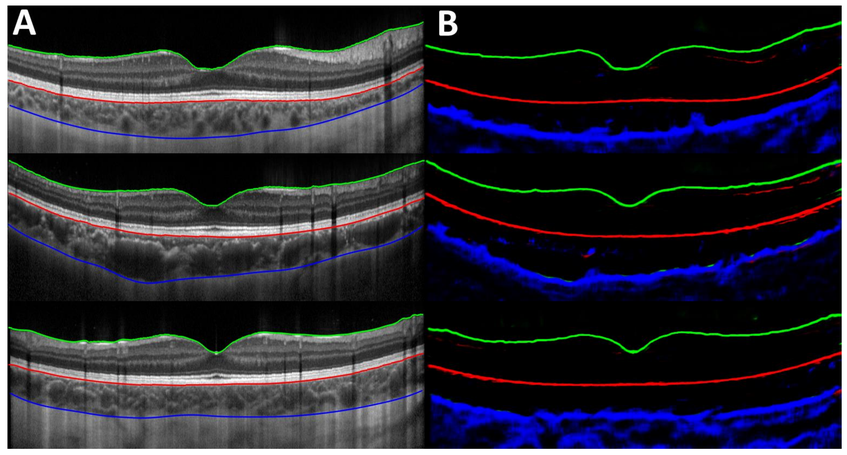

Figure 2. Example B-scans with manual segmentation of the layers of interest (A) and the corresponding probability maps (B) for the three layers (ILM in green, RPE in red and CSI in blue). Each colour indicates a high probability of a boundary being present in that location. The rectangular 31x61 CNN with standard intensity image was used to generate these results.

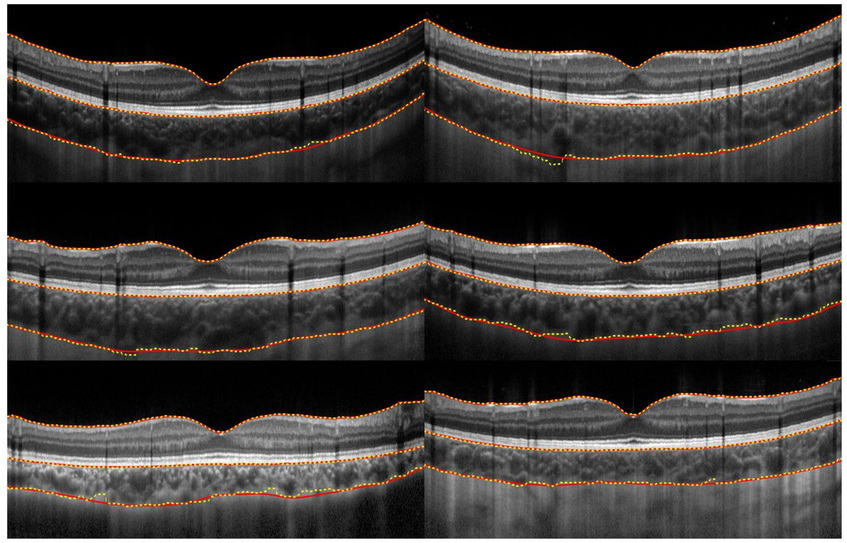

Figure 3. Six B-scans representing different examples from different participants, with typical variations in overall thickness and contrast. Manual segmentation is shown in red-solid and automatic is shown in yellow-dotted. The rectangular 31x61 CNN with standard intensity image was used to generate these results.

CONCLUSIONS

The automatic and manual segmentation methods showed very close agreement (Fig 3). Thus suggesting, that the rectangular CNN architecture used in this work, which utilizes the rich feature details along the axial direction of the scan, provides a robust detection of the retinal and choroidal boundaries of interest.

REFERENCES

1. Read, S.A., Collins, M.J., Vincent, S.J. and Alonso-Caneiro, D., 2013. Choroidal thickness in myopic and nonmyopic children assessed with enhanced depth imaging optical coherence tomography. Investigative Ophthalmology & Visual Science, 54(12), pp.7578-7586.

2. Fang, L., Cunefare, D., Wang, C., Guymer, R.H., Li, S. and Farsiu, S., 2017. Automatic segmentation of nine retinal layer boundaries in OCT images of non-exudative AMD patients using deep learning and graph search. Biomedical Optics Express, 8(5), pp.2732-2744.

3. Girard, M.J., Strouthidis, N.G., Ethier, C.R. and Mari, J.M., 2011. Shadow removal and contrast enhancement in optical coherence tomography images of the human optic nerve head. Investigative Ophthalmology & Visual Science, 52(10), pp.7738-7748.

Disclosures : N (none)

Email : d.alonsocaneiro@qut.edu.au

Web : https://research.qut.edu.au/clvol/