




Deep learning identify retinal nerve fibre and choroid layers as markers of age-related macular degeneration in the classification of macular spectral-domain optical coherence tomography volumes

Arnt-Ole Tvenning,¹  Stian Rikstad Hanssen,² Dordi Austeng^{1,3}  and Tora Sund Morken^{1,3} 

¹Department of Ophthalmology, St. Olav Hospital, Trondheim University Hospital, Trondheim, Norway

²Department of Computer Science, Norwegian University of Science and Technology (NTNU), Trondheim, Norway

³Department of Neuromedicine and Movement Science, NTNU, Trondheim, Norway

ABSTRACT.

Purpose: Deep learning models excel in classifying medical image data but give little insight into the areas identified as pathology. Visualization of a deep learning model's point of interest (POI) may reveal unexpected areas associated with diseases such as age-related macular degeneration (AMD). In this study, a deep learning model coined OptiNet was trained to identify AMD in spectral-domain optical coherence tomography (SD-OCT) macular scans and the anatomical distribution of POIs was studied.

Methods: The deep learning model OptiNet was trained and validated on two data sets. Data set no. 1 consisted of 269 AMD cases and 115 controls with one scan per person. Data set no. 2 consisted of 337 scans from 40 AMD cases (62 eyes) and 46 from both eyes of 23 controls. POIs were visualized by calculating feature dependencies across the layer hierarchy in the deep learning architecture.

Results: The retinal nerve fibre and choroid layers were identified as POIs in 82 and 70% of cases classified as AMD by OptiNet respectively. Retinal pigment epithelium (98%) and drusen (97%) were the areas applied most frequently. OptiNet obtained area under the receiver operator curves of $\geq 99.7\%$.

Conclusion: POIs applied by the deep learning model OptiNet indicates alterations in the SD-OCT imaging regions that correspond to the retinal nerve fibre and choroid layers. If this finding represents a tissue change in macular tissue with AMD remains to be investigated, and future studies should investigate the role of the neuroretina and choroid in AMD development.

Key words: age-related macular degeneration – convolutional neural network – deep learning – explainable artificial intelligence – spectral-domain optical coherence tomography – visualization

†Arnt-Ole Tvenning and Stian Hanssen co-authored this paper.

The project has been made possible by the Dam Foundation (2016/FO80635 to A-O. Tvenning). Professor Keith Downing is acknowledged for his invaluable and constructive suggestions during this research work. The authors report no conflicts of interest.

Acta Ophthalmol. 2022; 100: 937–945

© 2022 The Authors. Acta Ophthalmologica published by John Wiley & Sons Ltd on behalf of Acta Ophthalmologica Scandinavica Foundation.

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

doi: 10.1111/aos.15126

Introduction

High-resolution spectral-domain optical coherence tomography (SD-OCT) has become an essential tool in the evaluation of age-related macular

degeneration (AMD) (Chen et al. 2005), but the ever-increasing number of SD-OCT images, and the information within, supersedes what ophthalmologists effectively can analyse. Therefore, the development of artificial intelligence tools is wanted to improve

diagnosis and research on AMD. Deep learning is a branch of artificial intelligence inspired by the brain that utilizes artificial neural networks structured with multiple layers of neurons to learn complex tasks. Deep learning models learn without *a priori*

knowledge their own features of a given disease, and these features may then be used to classify medical images such as SD-OCT as either pathological or normal.

Given this naivety of deep learning, it is conceivable that unexpected features or anatomical areas are identified, which may be exploited to increase insight into diseases such as AMD. However, the models do not give insight into the areas or features that are identified as pathological. Therefore, even though deep learning models perform equal to or better than experts in the field in the classification of retinal disease in SD-OCT images (Apostol Poulos et al. 2016; Lee et al. 2017; De Fauw et al. 2018; Kermany et al. 2018; Lu et al. 2018; Motozawa et al. 2019; Perdomo et al. 2019), their decisions may be considered a ‘black box’. Insight into which features a deep learning model apply to classify an image is highly needed to integrate artificial intelligence into clinical practice. *Visualization* of areas or points of interest (POIs) that deep learning utilizes to classify SD-OCT scans may be obtained via a method called convolutional neural network (CNN) fixations. CNN fixations display multiple POIs in SD-OCT images based on unbiased information from a deep learning model (Mopuri et al. 2019). The POIs are visualized by calculating feature dependencies across the layer hierarchy in the deep learning architecture. CNN fixations display the POIs with maximum activation of neurons used for classification (Bany & Yeasin 2021). Visualization with CNN fixations has to the best of our knowledge, not earlier been performed on SD-OCT images of AMD.

In this study, we developed and applied a deep learning model, OptiNet, to classify AMD and controls in SD-OCT scans and utilized a novel visualization method to display which anatomical parts of the SD-OCT scan OptiNet used to identify AMD.

Material and Methods

The study was conducted on two data sets. The first is an open-source data set from the Age-Related Eye Disease 2 (AREDS2) Ancillary SD-OCT (A2A SD-OCT) study (http://people.duke.edu/~sf59/RPEDC_Ophth_2013_dataset.htm) referred to as the A2A data set (Farsiu et al. 2014). The

second data set was obtained from the Norwegian Pigment Epithelial Detachment (NORPED) study (Tvenning et al. 2020).

Study populations

Participants in the A2A SD-OCT study were ≥50 years of age with intermediate AMD, drusen ≥125 µm in both eyes or one eligible eye and advanced AMD in the fellow eye. The patients had not previously undergone vitreoretinal surgery and were otherwise without ophthalmological diseases. The control group was age-matched and did not have any signs of AMD in either eye (Farsiu et al. 2014).

Participants in the NORPED study were ≥50 years of age with a drusenoid pigment epithelial detachment diameter ≥1000 µm in one eye and best-corrected visual acuity ≥20/400. The NORPED study is a longitudinal study, with imaging performed every 6 months, and the scans included in this study were from participants with up to 3 years of follow-up. Exclusion criteria were previous vitrectomy, corticosteroid therapy, intravitreal injections and photodynamic therapy. Eyes with subfoveal fibrosis, central geographic atrophy, choroidal neovascularization on fluorescein or indocyanine green angiography, glaucoma with central visual field defects, diabetic macular oedema, proliferative diabetic retinopathy and uveitis were also excluded. The fellow eyes of these patients were included regardless of the stage of AMD. Controls were hospital staff recruited from the Eye department at St. Olavs University Hospital, Trondheim, Norway.

Image acquisitions

In the A2A study, SD-OCT systems from Bioptigen, Inc. (Morrisville, NC,

USA) were used at four clinic sites to acquire volumetric rectangular scans. Both 0° and 90° volumes were captured centred on the fovea with 1000 A-scans per B-scan and 100 B-scans per volume (Farsiu et al. 2014). The NORPED study obtained macular cube scans with the Spectralis HRA-OCT (Heidelberg Engineering GmbH, Heidelberg, Germany). A dense scan pattern of 49 scans with eye-tracking and follow-up mode was enabled. The macular cube scans consisted of 18–30 averaged images recorded with the automatic real-time function.

Data set processing

The A2A data set (Table 1) contains SD-OCT scans from 269 AMD cases and 115 control cases. Each scan is of the shape 1 × 100 × 512 × 1000 pixels (channels × depth × height × width). The SD-OCT scans were cropped down to the shape 1 × 32 × 512 × 512 pixels. Centre crop was used for all dimensions except for height, where only the bottom was cropped off, to focus on the fovea. Furthermore, the SD-OCT scans were resized to the final shape 1 × 32 × 256 × 256 pixels. All cross sections in the SD-OCT scans had their mean set to 0 and variance to 1 to speed up training time (LeCun & Bottou 2012). The data set was split into two partitions: one for training and one for validation, with a ratio of 81% ($n = 312$) being training cases and 19% ($n = 72$) being validation cases. The SD-OCT scans from the same patient could only be in one of the partitions.

The NORPED data set (Table 2) contains 337 and 46 SD-OCT scans from 40 AMD patients (62 eyes) and 23 controls (46 eyes) respectively. Images were obtained every 6 months, with up to 3 years of follow-up from each study participant. Images in the NORPED

Table 1. The A2A data set.

	Training	Validation	Total
AMD cases, n	221	48	269
Control cases, n	91	24	115
Total cases, n	312	72	384
AMD cases, %	71%	67%	70%
Control cases, %	29%	33%	30%

Characteristics of an open-source data set from the Age-Related Eye Disease 2 (AREDS2) Ancillary SD-OCT (A2A SD-OCT) study referred to as the A2A data set (Farsiu et al. 2014). Abbreviations: AMD = age-related macular degeneration, n = number of individuals with one spectral-domain optical coherence tomography scan.

Table 2. The NORPED data set.

	Training partition	Validation partition	Total
AMD cases, <i>n</i>	235	102	337
Control cases, <i>n</i>	34	12	46
Total cases, <i>n</i>	269	114	383
AMD cases, %	87%	89%	88%
Control cases, %	13%	11%	12%

Characteristics of the Norwegian pigment epithelial detachment study (NORPED) data set. Abbreviations: AMD = age-related macular degeneration, *n* = number of longitudinally obtained spectral-domain optical coherence tomography macular scans from 40 patients (62 eyes) with AMD and 23 controls (46 eyes).

data set were centre cropped to the size of $1 \times 32 \times 384 \times 384$ pixels and then resized to $1 \times 32 \times 256 \times 256$ pixels. Their mean and variance were set to 0 and 1, respectively. The data set was split into two partitions: one for training and one for validation, with a ratio of 70% ($n = 269$) being training cases and 30% ($n = 114$) being validation cases. We ensured that all SD-OCT scans from one patient were placed in only one partition when dividing the data set.

Both data sets were augmented with horizontal flipping and stored in batches of 24 in HDF5 files to reduce file size and decrease loading time.

OptiNet model

OptiNet is a 3D deep learning model (Lecun et al. 1998) using (2 + 1)D convolution (Tran et al. 2018), trained

without *a priori* knowledge of AMD, or pretrained parts of other deep learning architectures (transfer learning). The model learns its feature extractors and then uses them to make a classification (Fig. 1). Each convolutional layer gradually learns a set of filters that extract specific information from the scan. In the end, this is given to classification layers where all neurons in one layer are connected to all neurons in the next. The first convolutional layers will learn simple features such as lines and edges. The consecutive layers build upon simple features, finding more complex features for each layer going deeper into the CNN. The CNN reduces the information it needs to handle with max-pooling layers (Scherer 2010), which can be viewed as reducing the size of the filtered SD-OCT scans. After the convolutional layers, a few final layers perform the

classification task using the deep learning model’s refined features. The final output of the CNN is put through a sigmoid function that makes the output a probability between 0 and 1. A value of 1 indicates a 100% certainty of AMD, while 0 indicates 0% certainty of AMD. When performing predictions to calculate accuracy, the output value is thresholded such that 0.5 and higher means an AMD classification, and a value lower than 0.5 is a control.

OptiNet consists of a repeating pattern called a base block (Figure S1). The base block is made of (2 + 1)D convolutional layers modified from Tran et al. (2018), each with a batch normalization layer and a rectified linear activation unit function (Nair & Hinton 2010; Ioffe & Szegedy 2015). The last layer in the base block performs 3D max pooling. (2 + 1)D convolution is a variation of 3D convolution, where the CNN first identifies features in 2D and subsequently performs 1D convolution to define relations between cross sections. Finally, the model ends with a sigmoid activation function.

OptiNet was trained on an NVIDIA Tesla V100 PCI-E 16 GB GPU. The hyper-parameters and architectural design decisions (Table 3) have been applied in deep learning models such as VGG and ResNet (Simonyan & Zisserman 2015; He et al. 2016).

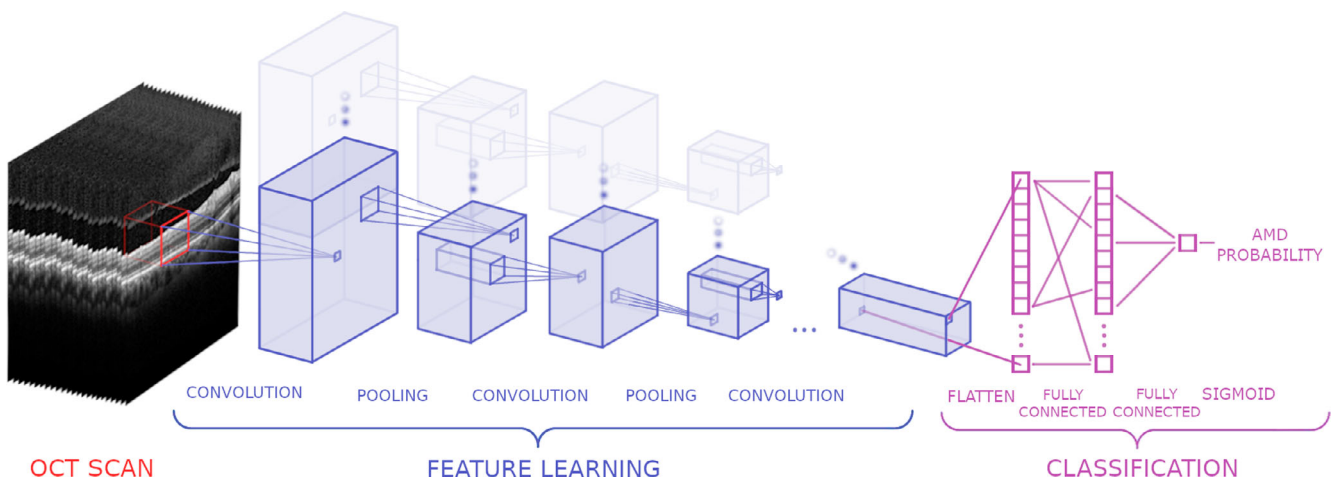


Fig. 1. Simplified illustration of the 3D deep learning model OptiNet. The spectral-domain optical coherence tomography (SD-OCT) scan seen to the left is the input to the deep learning model OptiNet. OptiNet learns to extract specific features in the SD-OCT scan through convolutional and pooling layers during feature learning. The blue boxes represent filtered versions of the SD-OCT scan. The faded blue boxes deeper into the image represent all the filtered versions of the SD-OCT scan created from the learned filters of each convolutional layer. The compact set of features coming out of the feature learning section is flattened out and given to a final neural network performing the classification. The classification section of the deep learning model can be viewed as clusters of neurons where all neurons in one layer are connected to all neurons in the next. The network’s final output is put through a sigmoid function that makes the output a probability between 0 and 1.

Table 3. Hyper-parameters and design decisions for the deep learning model OptiNet.

Hyper-parameter/Decision	Choice
Training mini-batch size	1
Validation mini-batch size	1
Learning rate	0.00001
Loss function	Binary Cross-Entropy
Bias initialization	All values set to 0
Weight initialization	Xavier (Glorot & Bengio 2010)
Optimiser	Adam (Kingma & Ba 2015)

The deep learning models were saved each time that a new best validation accuracy was achieved during training. The training was stopped when a model overfitted the data. The training of each model took approximately 4 hr.

In the A2A data set, 10 independent OptiNet models were trained, where each model had its data set shuffled before training started. Each model




Fig. 2. Visualization of points of interest in three cases with age-related macular degeneration. The figure shows how the deep learning model, OptiNet, displayed results to the retina specialist on three cases with age-related macular degeneration. Three images are given for each case, displaying the regular spectral-domain optical coherence tomography scan, points of interest and a localization heat map in a row. Convolutional neural network fixations show the points of interest (red dots) that were important for the classification. A high value on the colour bar, displayed in red, implies high interest from the model.




Fig. 3. The high-resolution visualization of points of interest in a patient with age-related macular degeneration. (A) The standard presentation of points of interest (red dots). (B-C) Increasing magnification of the spectral-domain optical coherence tomography scan shows the high-resolution points of interest in specific retinal layers and structures.

was trained on the same number of cases, but the selection of cases in each partition differed. In the same manner,

OptiNet was tested on the NORPED data set by training 10 independent models. However, in this data set, partitions were not shuffled. This was due to the dependencies in the NORPED data set, which made it challenging to perform shuffling that needs to be random while maintaining the same number of cases in each partition. The trained models were validated on the validation partition of their data set.

Visualization

The visualization was performed through computing POIs on the validation partitions of the data sets. The technique called CNN fixations reveals what the model identifies as AMD in the SD-OCT scan when making a classification. (2 + 1)D convolution was handled by expanding the general concept from CNN fixations (Mopuri et al. 2019). POIs were calculated for 1D convolution and then for 2D

convolution while transforming the locations’ dimensionality according to (2 + 1)D convolution implementation. Similar to Mopuri et al. (2019), a 3D heat map was created with a gaussian blur of the CNN fixations, which illustrates the densities of POIs in the SD-OCT scan. Outlier removal was used, and the POIs are only shown if 5% of the total CNN fixations are within a sphere with a radius equal to 10% of the diagonal of the SD-OCT scan. During the evaluation of POIs, each SD-OCT scan with AMD was evaluated by a retina specialist, and the POIs in each retinal layer and features of AMD selected by OptiNet were registered. OptiNet was trained to output a probability of AMD, and the POIs in the control group are therefore limited to visualize features of AMD.

Statistical analyses

In both test cases, the following metrics were collected: area under the receiver operating characteristics curve (AUC), accuracy, sensitivity, specificity, negative predictive value and positive predictive value. The total error rate was calculated as 1—accuracy. The scientific computing package (SciPy), in the programming language Python 3 and Stata 15.1 (StataCorp LLC, College Station, TX, USA), was used for statistical analyses.

Ethics

Four institutional review boards at Devers Eye Institute, Duke Eye Center, Emory Eye Center and National Eye Institute have approved using the SD-OCT images from the A2A SD-OCT study (Clinical trials identifier: NCT00734487), and informed consent had been obtained from all subjects in that study. The NORPED study was approved by the Regional Committee for Medical and Health Research Ethics Central Norway (2012/1743) and followed the tenets of the Declaration of Helsinki. Written informed consent was obtained from study participants.

Results

POI distribution was presented to the evaluating retina specialist as localization heat maps (Fig. 2), and the high

Table 4. The distribution of points of interest in retinal layers, zones and features of age-related macular degeneration in the validation partitions on optical coherence tomography images of age-related macular degeneration.

	NORPED (n = 102*)		A2A (n = 48*)		Combined (n = 150*)	
	n	(%)	n	(%)	n	(%)
Retinal pigment epithelium	99	(97)	48	(100)	147	(98)
Interdigitation zone	87	(85)	39	(81)	126	(84)
Ellipsoid zone	86	(85)	46	(96)	132	(88)
Outer segment	83	(81)	34	(73)	117	(78)
Retinal nerve fibre	81	(79)	42	(88)	123	(82)
Choroid	59	(58)	46	(96)	105	(70)
Inner plexiform	45	(44)	30	(63)	75	(50)
Outer plexiform	43	(42)	29	(61)	72	(48)
Ganglion cell	41	(40)	33	(69)	74	(49)
Henle nerve fibre	37	(36)	24	(50)	61	(41)
Vitreous	36	(35)	22	(46)	58	(39)
Inner nuclear	29	(28)	21	(44)	50	(33)
Outer nuclear	16	(16)	22	(45)	38	(25)
External limiting membrane	15	(15)	7	(15)	22	(15)
Bruch’s membrane	8	(8)	7	(15)	15	(10)
Myoid zone	7	(7)	9	(19)	16	(11)
Drusen	99	(97)	47	(98)	146	(97)
Subretinal drusenoid deposits	42	(41)	11	(23)	53	(35)
Intraretinal hyperreflective foci	32	(31)	4	(8)	36	(24)
Acquired vitelliform lesion	24	(24)	-	-	24	(16)
Hypertransmission	13	(13)	5	(10)	18	(12)
Subretinal fluid	3	(3)	1	(2)	4	(3)
Intraretinal cysts	1	(1)	-	-	1	(1)

Notes: The numbers (*) represent the number of spectral-domain optical coherence tomography scans from the validation partitions, the parts of the data set that had not been used for training. (%) = the percentage of SD-OCT scans to have at least one point of interest and their distribution in that respective retinal layers, zones and structures.

Abbreviations: A2A = Age-Related Eye Disease 2 Ancillary spectral-domain optical coherence tomography study, n = number, NORPED = Norwegian pigment epithelial detachment study.




Fig. 4. Visualization of points of interest in the retinal nerve fibre and choroid layers in spectral-domain optical coherence tomography scans of age-related macular degeneration. The figure shows points of interest (red dots) that the deep learning model OptiNet identified as markers of age-related macular degeneration (AMD) during classification. The images (A)–(L) represent AMD cases from the validation partition that consisted of spectral-domain optical coherence tomography scans that the deep learning model had not seen during training. Angled arrowheads show points of interest in the retinal nerve fibre layer, and vertical arrows show points of interest in the choroid layer.

resolution of the POIs could be visualized with increasing magnification (Fig. 3). The distribution of POIs in retinal layers, zones and features of AMD on SD-OCT images revealed that OptiNet applied the retinal nerve fibre layer (RNFL) and choroid in 82 and 70% of cases classified as AMD (Table 4). The retinal pigment epithelium (98%) and drusen (97%) were the areas applied most frequently. Fig. 4 show the presentations of POIs in the RNFL and choroid layers in SD-OCT scans of cases with AMD.

Ten OptiNet models achieved a mean AUC of 99.7% and 99.9%, calculated from the mean ROC curves illustrated in Fig. 5(A) and 5(B) for the A2A and NORPED data sets. The mean total error rates were 3.5% and 1.1% for the A2A and NORPED data sets. Additional performance metrics of OptiNet are presented in Table 5. An example of a control case is shown in Fig. 6.

The mean age of patients was 75 (51 to 87) and 67 years (51 to 83) for controls in the A2A data set. In the

NORPED data set, the mean age was 72 (range 57 to 86) and 47 (range 28 to 66) years for patients and controls respectively.

Discussion

For the first time, the anatomical areas utilized by a deep learning model to identify AMD are shown with a novel visualization method. This method revealed that regions such as the RNFL and choroid might be altered in AMD and demonstrate the potential of deep learning as a method not only in identification but also in the exploration of retinal disease. As a result, the methods applied in this study may be expanded to include multiple retinal diseases.

The RNFL was a frequent POI in the classification of AMD in both data sets. Few studies have investigated the effect of AMD on the inner retinal layers, and most have investigated RNFL volume measurements and reported conflicting results. A prospective study did not find a difference in the change of RNFL volume between dry AMD compared with controls after 2 years of follow-up (Lamin et al. 2019), while a cross-sectional study reported that the RNFL volume was lower in treatment-naïve patients with choroidal neovascularization compared with controls (Zucchiatti et al. 2015). Another cross-sectional study of patients with incident early AMD showed that the RNFL was progressively thinner with increasing severity of AMD (Farinha et al. 2021). Other cross-sectional studies have not found a difference in RNFL thickness between early and intermediate AMD (Savastano et al. 2014; Zucchiatti et al. 2015), or those that progressed to geographic atrophy compared to controls (Zucchiatti et al. 2015). Choroidal degeneration has earlier been proposed to be involved in retinal diseases (Rivera et al. 2017). Indeed, in the present study, the choroidal layer was a frequent POI in both data sets, especially for the A2A data set with an age-matched control group. The choroid and its association with AMD are uncertain, maybe due to the relative inaccessibility of imaging this part of the retina. Ting et al. (2016) reported an increased choroidal layer in neovascular AMD. Choroidal thickness has also been decreased following

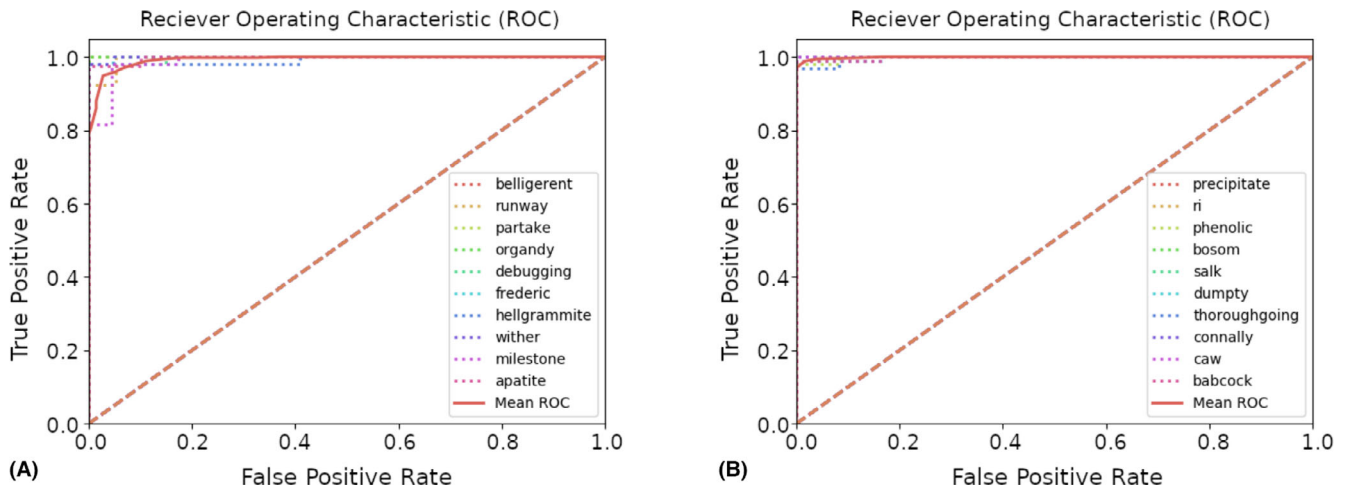


Fig. 5. ROC curves computed from the models of OptiNet trained on the A2A and NORPED data sets. (A) ROC curves calculated for all 10 models of OptiNet trained on the A2A data set. (B) ROC curves calculated for all 10 models of OptiNet trained on the NORPED data set. Each model has a randomly generated name, which is shown in the lower-right corner of both figures. The mean ROC curve is identified as the solid line in the graph. Abbreviations: A2A = Age-Related Eye Disease 2 Ancillary SD-OCT study, NORPED = Norwegian Pigment Epithelial Detachment Study, ROC = receiver operating characteristic.

Table 5. The performance of OptiNet on the validation partition of the A2A and NORPED data sets.

Metric	A2A data set	NORPED data set
Accuracy	96.5 ± 2.1%	98.9 ± 0.8%
Positive predictive value	97.6 ± 3.2%	99.5 ± 0.9%
Sensitivity	97.6 ± 2.3%	99.2 ± 0.6%
Negative predictive value	94.8 ± 5.1%	93.9 ± 4.6%
Specificity	93.8 ± 8.3%	95.8 ± 7.7%
AUC	99.7 ± 0.4%	99.1 ± 0.1%

The results are of the format mean ± standard deviation based on ten trained models from the OptiNet architecture for each of the two data sets. Abbreviations: A2A = Age-Related Eye Disease 2 Ancillary SD-OCT study, AUC = area under the receiver operator curve, NORPED = Norwegian Pigment Epithelial Detachment Study.

Table 6. Previous studies and performance on the A2A data set.

Study	Method	Results
(Apostolopoulos et al. 2016)	Convolutional neural network	AUC: 99.7%
(Santos et al. 2018)	Support vector machine	AUC: 98.9%
(Venhuizen et al. 2015)	Unsupervised feature learning, bag of words	AUC: 98.4%
(Farsiu et al. 2014)	Generalized linear regression	AUC: 99.2%

Abbreviations: AUC = area under the receiver operator curve, A2A data set = Age-Related Eye Disease 2 Ancillary spectral-domain optical coherence tomography study.

drusenoid pigment epithelium detachment regression and progression of geographic atrophy (Dolz-Marco et al. 2018). Age, refractive error and axial length affect the choroidal thickness, and most studies that adjust for these factors do not find associations to AMD (Jonas et al. 2014; Yiu et al. 2015; Ho et al. 2018). Since these studies mainly investigate volume changes in these layers, it is conceivable

that there might be other SD-OCT features than the volume of the RNFL and choroid layer that changes with the development of AMD and that is recognized by the deep learning model. Such features may be focal alterations, reflectivity or configuration that the deep learning model identifies during feature learning and emphasizes in its classification. Age may be a confounding factor since both RNFL and

choroid layer thickness are known to decrease with age (Budenz et al. 2007). However, the RNFL and choroid layers were marked as POIs in both data sets, of which the A2A data set was age-matched, indicating that alterations in these areas were not caused by age in this study.

Deep learning methods give little insight into their classification and are often considered as ‘black boxes’. Visualization methods reveal which parts of an image that were important for the classification and are an integral part of explainable artificial intelligence. The advantage of CNN fixations is that it does not alter the architecture of the deep learning model and can calculate the most important POIs for the classification in addition to conventional localization heat maps (Mopuri et al. 2019). Occlusion testing is another visualization method used by many studies (Lee et al. 2017; Kermay et al. 2018). This method occludes an area of the image, and if it negatively affects the classification, that area is important (Zeiler & Fergus 2014). However, this process must be repeated for every possible 20x20 pixel area in the image, requiring many calculations. The advantage of CNN fixations is the need for only one computation through the deep learning model with higher resolution and precision than occlusion testing. Perdomo et al. (2019) used class activation mapping with detailed visualization; however, it alters the deep learning architecture, potentially




Fig. 6. Visualization of points of interest in a control case. The points of interest (red dots) are mainly located in the layer of the retinal pigment epithelium and photoreceptors. The locations appear to overlap with the points of interest of age-related macular degeneration because the deep learning model was trained to identify the disease, and convolutional neural network fixations are therefore limited to their visualization. The output probability in the deep learning model for control cases is <0.5 .

limiting its performance (Mopuri et al. 2019). De Fauw et al. (2018) used automated segmentation to transform the SD-OCT images through learning, which can be viewed as visualization. However, the method only displays the segmentation results and analyses the segmentation outputs in contrast to CNN fixations that recognize the raw image features. Convolutional neural network (CNN) fixations provide unbiased high-resolution POIs and localization heat maps as is implemented in OptiNet, which can uncover new biomarkers learned by the model and can be expanded to visualize multiple retinal diseases.

OptiNet shows promising results, with AUC scores similar to other high performing deep learning models (Table 6) in identifying AMD (Apostolopoulos et al. 2016; Lee et al. 2017; De Fauw et al. 2018; Kermayn et al. 2018). Despite the relatively small data sets used in our study, the AUCs and visual feedback to the retina specialist indicate that OptiNet learned to identify AMD. Deep learning models can perform well on small data sets, probably when the differences between images are sufficiently large (Cho et al. 2015). In our study, the presence of drusen or no drusen or alterations in the retinal pigment epithelium and photoreceptor layers may have been such an easy task for a deep learning model to learn. (2 + 1)D convolution used by OptiNet enables the learning of more complex features with the same amount of training as 2D (Tran et al. 2018). Most studies use parts from well-known deep learning architectures previously

trained on 2D images in the ImageNet Large Scale Visual Recognition Challenge (Russakovsky et al. 2015), limiting them to 2D image classification. 2D models can train on all the SD-OCT scans, facilitating much larger data sets. However, this is also a limitation of 2D classification since the relations between cross sections in the retinal SD-OCT scan are not investigated and applied. The mean AUC score of OptiNet was identical to the highest performing study on the A2A data set by Apostolopoulos et al. (2016). Still, the method in that study involved two 2D deep learning models with transfer learning. OptiNet was trained without such *a priori* knowledge, and our results indicate that a two-step process might not be necessary. Indeed, 3D deep learning models outperform 2D models in lung nodule detection for lung cancer (Dou et al. 2017), and studies indicate that 3D deep learning models are superior when evaluating 3D data (Tran et al. 2018).

The limitations of this study are that the data sets used were relatively small, unbalanced and only included one macular disease. Although OptiNet's performance is promising, it has not been tested on a general population cohort where most subjects are normal. Our trained models would likely obtain a higher number of false positives when presented with other macular diseases, and the POIs used by OptiNet might not be unique for AMD. In addition, enhanced depth imaging was not used in this study, and the precise localization of POIs in the deeper choroid layers could not always be determined.

Conclusion

POIs applied by the deep learning model OptiNet indicate that the SD-OCT imaging regions that represent the RNFL and choroid layers may be altered in AMD. However, if these findings represent a tissue change in macular tissue with AMD remains to be investigated, and future studies should investigate the role of the neuroretina and choroid in AMD development. Furthermore, the application of CNN fixations as a visualization method should be expanded to multiple retinal diseases. It is conceivable that visual feedback to the clinician from a deep learning model may be useful as a diagnostic and research tool.

References

- Apostolopoulos S, Ciller C, De Zanet SI, Wolf S & Sznitman R (2016): RetiNet: Automatic AMD identification in OCT volumetric data. arXiv preprint: arXiv:1610.03628.
- Bany M & Yeasin M (2021): Eigen-CAM: visual explanations for deep convolutional neural networks. *SN Comput Sci* 2: 47.
- Budenz DL, Anderson DR, Varma R et al. (2007): Determinants of normal retinal nerve fiber layer thickness measured by Stratus OCT. *Ophthalmology* 114: 1046–1052.
- Chen TC, Cense B, Pierce MC et al. (2005): Spectral domain optical coherence tomography: ultra-high speed, ultra-high resolution ophthalmic imaging. *Arch Ophthalmol* 123: 1715–1720.
- Cho J, Lee K, Shin E, Choy G & Do S (2015): How much data is needed to train a medical image deep learning system to achieve necessary high accuracy?. arXiv preprint: arXiv:1511.06348.

- De Fauw J, Ledsam JR, Romera-Paredes B et al. (2018): Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nat Med* **24**: 1342–1350.
- Dolz-Marco R, Balaratnasingam C, Gattoussi S, Ahn S, Yannuzzi LA & Freund KB (2018): Long-term Choroidal Thickness Changes in Eyes With Drusenoid Pigment Epithelium Detachment. *Am J Ophthalmol* **191**: 23–33. Spectral.
- Dou Q, Chen H, Yu L, Qin J & Heng PA (2017): Multilevel contextual 3-D CNNs for false positive reduction in pulmonary nodule detection. *IEEE Trans Biomed Eng* **64**: 1558–1567.
- Farinha C, Silva AL, Coimbra R et al. (2021): Retinal layer thicknesses and neurodegeneration in early age-related macular degeneration: insights from the Coimbra Eye Study. *Graefes Arch Clin Exp Ophthalmol* **259**: 2545–2557.
- Farsiu S, Chiu SJ, O'Connell RV et al. (2014): Quantitative classification of eyes with and without intermediate age-related macular degeneration using optical coherence tomography. *Ophthalmology* **121**: 162–172.
- Glorot X & Bengio Y (2010): Understanding the difficulty of training deep feedforward neural networks. Proceedings of the thirteenth international conference on artificial intelligence and statistics. Sardinia, Italy **9**. Cambridge, MA: PMLR pp. 249–256.
- He K, Zhang X, Ren S & Sun J (2016): Deep residual learning for image recognition. In: *Proc IEEE Comput Soc Conf Comput Vis Pattern Recognit*. Silver Spring, MD: IEEE Computer Society Press pp. 770–778.
- Ho CY, Lek JJ, Aung KZ, McGuinness MB, Luu CD & Guymer RH (2018): Relationship between reticular pseudodrusen and choroidal thickness in intermediate age-related macular degeneration. *Clin Exp Ophthalmol* **46**: 485–494.
- Ioffe S & Szegedy C (2015): Batch normalisation: Accelerating deep network training by reducing internal covariate shift. 32nd International Conference on Machine Learning. Lille, France. **1**. Cambridge, MA: PMLR pp. 448–456.
- Jonas JB, Forster TM, Steinmetz P, Schlichtenbrede FC & Harder BC (2014): Choroidal thickness in age-related macular degeneration. *Retina* **34**: 1149–1155.
- Kermany DS, Goldbaum M, Cai W et al. (2018): Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell* **172**: 1122–1131. e1129.
- Kingma DP & Ba JL (2015): Adam: A method for stochastic optimisation. arXiv preprint: arXiv:1412.6980
- Lamin A, Oakley JD, Dubis AM, Russakoff DB & Sivaprasad S (2019): Changes in volume of various retinal layers over time in early and intermediate age-related macular degeneration. *Eye (Lond)* **33**: 428–434.
- Lecun Y, Bottou L, Bengio Y & Haffner P (1998): Gradient-based learning applied to document recognition. *Proc IEEE Inst Electr Electron Eng* **86**: 2278–2324.
- LeCun YA, Bottou L, Orr GB & Müller KR (2012): Efficient backProp. In: Montavon G, Orr GB & Müller K (eds). *Neural networks: tricks of the trade*. Berlin, Heidelberg: Springer pp. 9–48.
- Lee CS, Baughman DM & Lee AY (2017): Deep learning is effective for the classification of OCT images of normal versus age-related macular degeneration. *Ophthalmol Retina* **1**: 322–327.
- Lu W, Tong Y, Yu Y, Xing Y, Chen C & Shen Y (2018): Deep learning-based automated classification of multi-categorical abnormalities from optical coherence tomography images. *Transl Vis Sci Technol* **7**: 41.
- Mopuri KR, Garg U & Venkatesh R (2019): CNN fixations: an unraveling approach to visualise the discriminative image regions. *IEEE Trans Image Process* **28**: 2116–2125.
- Motozawa N, An G, Takagi S et al. (2019): Optical coherence tomography-based deep-learning models for classifying normal and age-related macular degeneration and exudative and non-exudative age-related macular degeneration changes. *Ophthalmol Ther* **8**: 527–539.
- Nair V & Hinton GE (2010): Rectified Linear Units Improve Restricted Boltzmann Machines. Proceedings of the 27th International Conference on Machine Learning. Haifa, Israel. Madison, WI: Omnipress pp. 807–814.
- Perdomo O, Rios H, Rodriguez FJ et al. (2019): Classification of diabetes-related retinal diseases using a deep learning approach in optical coherence tomography. *Comput Methods Programs Biomed* **178**: 181–189.
- Rivera JC, Holm M, Austeng D et al. (2017): Retinopathy of prematurity: inflammation, choroidal degeneration, and novel promising therapeutic strategies. *J Neuroinflammation* **14**: 165.
- Russakovsky O, Deng J, Su H et al. (2015): ImageNet large scale visual recognition challenge. *Int J Comput Vis* **115**: 211–252.
- Santos AM, Paiva AC, Santos APM et al. (2018): Semivariogram and Semimadogram functions as descriptors for AMD diagnosis on SD-OCT topographic maps using Support Vector Machine. *Biomed Eng Online* **17**: 160.
- Savastano MC, Minnella AM, Tamburrino A, Giovinco G, Ventre S & Falsini B (2014): Differential vulnerability of retinal layers to early age-related macular degeneration: evidence by SD-OCT segmentation analysis. *Invest Ophthalmol Vis Sci* **55**: 560–566.
- Scherer D, Müller A & Behnke S (2010): Evaluation of pooling operations in convolutional architectures for object recognition. In: Diamantaras K, Duch W & Iliadis LS (eds). *Artificial Neural Networks - ICANN 2010*. Berlin, Heidelberg: Springer pp. 92–101.
- Simonyan K & Zisserman A (2015): Very deep convolutional networks for large-scale image recognition. arXiv preprint: arXiv:1409.1556.
- Ting DS, Ng WY, Ng SR et al. (2016): Choroidal thickness changes in age-related macular degeneration and polypoidal choroidal vasculopathy: a 12-month prospective study. *Am J Ophthalmol* **164**: 128–136.
- Tran D, Wang H, Torresani L, Ray J, LeCun Y & Paluri M (2018): A Closer Look at Spatiotemporal Convolutions for Action Recognition. In: *Proc IEEE Comput Soc Conf Comput Vis Pattern Recognit*. Silver Spring, MD: IEEE Computer Society Press pp. 6450–6459.
- Tvenning AO, Krohn J, Forsaa V, Malmin A, Hedels C & Austeng D (2020): Drusenoid pigment epithelial detachment volume is associated with a decrease in best-corrected visual acuity and central retinal thickness: the Norwegian Pigment Epithelial Detachment Study (NORPED) report no. 1. *Acta Ophthalmol* **98**: 701–708.
- Venhuizen F, van Ginneken B, Bloemen B et al. (2015): Automated age-related macular degeneration classification in OCT using unsupervised feature learning. Proceedings SPIE Medical Imaging. Orlando, USA. Bellingham, WA: SPIE p. 9414.
- Yiu G, Chiu SJ, Petrou PA et al. (2015): Relationship of central choroidal thickness with age-related macular degeneration status. *Am J Ophthalmol* **159**: 617–626.
- Zeiler MD & Fergus R (2014): In: Fleet D, Pajdla T, Schiele B & Tuytelaars T (eds). *Computer Vision – ECCV 2014*. Cham: Springer pp. 818–833.
- Zucchiatti I, Parodi MB, Pierro L et al. (2015): Macular ganglion cell complex and retinal nerve fiber layer comparison in different stages of age-related macular degeneration. *Am J Ophthalmol* **160**: 602–607. e601.

Received on November 11th, 2021.
Accepted on February 18th, 2022.

Correspondence:

Arnt-Ole Tvenning, MD, PhD
Department of Ophthalmology
St. Olav Hospital
Trondheim University Hospital
Trondheim 7006
Norway
Tel: +47 72575012
Fax: +47 72574833
Email: arnt-ole.tvenning@stolav.no

Supporting Information

Additional Supporting Information may be found in the online version of this article:

Figure S1. The Architecture of Opti-Net. The top of the Figure displays the content of a base block. Each box represents a layer or a module in the architecture, with their kernel sizes, strides, padding, and activation functions given.