Prediction of Treatment Response from Retinal OCT in Patients with Exudative Age-Related Macular Degeneration

Hrvoje Bogunović[1], Michael D. Abràmoff[2,1,3,4], Li Zhang[1], and Milan Sonka[1,2]

[1]Department of Electrical and Computer Engineering
[2]Department of Ophthalmology and Visual Sciences
[3]Department of Biomedical Engineering
The University of Iowa, Iowa City, Iowa, USA
[4]Department of Veterans Affairs, Iowa City, USA
{hrvoje-bogunovic,michael-abramoff,li-zhang-1,milan-sonka}@uiowa.edu

Abstract. Age related macular degeneration is a major cause of blindness and visual impairment in older adults. Its exudative form, where fluids leak into the macula, is especially damaging. The standard treatment involves injections of anti-VEGF (vascular endothelial growth factor) agents into the eye, which prevent further vascular growth and leakage, and can restore vision. These intravitreal injections have a risk of devastating complications including blindness from infection and are expensive. Optimizing the interval between injections in a patient specific manner is of great interest, as the retinal response is partially patient-specific. In this paper we propose a machine learning approach to predict the retinal response at the end of a standardized 12-week induction phase of the treatment. From a longitudinal series of optical coherence tomography (OCT) images, a number of quantitative measurements are extracted, describing the underlying retinal structure and pathology and its response to initial treatment. After initial feature selection, the selected set of features is used to predict the treatment response status at the end of the induction phase using the support vector machine classifier. On a population of 30 patients, leave-one-out cross-validation showed the classification success rate of 87% of predicting whether the subject will show a response to the treatment at the next visit. The proposed methodology is a promising step towards the much needed image-guided prediction of patient-specific treatment response.

1 Introduction

The primary cause of blindness in older adults is age-related macular degeneration (AMD) [1]. It is a disease which affects the macula of the retina and destroys the sharp, central vision. The most damaging form of AMD is exudative or wet AMD, also known as choroidal neovascularization (CNV), which is caused by the growth of abnormal blood vessels from the choroid vasculature. The abnormal vessels leak fluid into the macula, leading to its structural damage. Such fluid-filled regions consist of intraretinal and subretinal fluid as well as

pigment epithelial detachments (PED) that we jointly call symptomatic exudate associated derangements (SEAD).

A standard and effective treatment for CNV is to inject into the eye anti-vascular endothelial growth factor (anti-VEGF) agents, which suppress further blood vessel growth [2]. Such treatment enables restoring the visual acuity close to normal but requires frequent retreatments, which carries several risks. Besides the risk of devastating endophthalmitis that can result in blindness and the very high cost of the drug ($2,000 per injection), the risks of frequent treatments may lead to outer retinal atrophy and nerve fiber layer loss due to the drugs' anti-angiogenic effect.

Patients respond differently to anti-VEGF treatments. Thus, the optimal treatment regime is patient-specific and would consist of the smallest number of anti-VEGF injections that are still effective. In current clinical practice the treatment frequency is primarily guided by SEAD presence, which is subjectively assessed from images noninvasively acquired with optical coherence tomography (OCT). This calls for designing objective, patient-specific treatment dosing regimen.

Most of the prior work on quantitative analysis of OCT images was devoted to the difficult tasks of retinal layer and SEAD segmentation [3–6]. The availability of those segmentation methods now allows for quantitative and objective analysis of retinal structure and pathology from OCT images. Currently it is not clear which quantitative indices are associated with desirable anti-VEGF treatment response. In this paper, we attempt to learn the patient-specific treatment response from the standardized 12-week long induction phase. Using objectively-defined layer and SEAD morphologic properties extracted from the longitudinal series of OCT images, we determine associations and predict the outcome of the induction phase.

2   Materials and Methods

2.1   The Induction Phase

The AMD patients underwent an initial 12-week induction phase during which treatment and imaging was uniform for all patients. The induction phase was defined as the process of administration of three intravitreal injections at 4 week intervals (weeks 0, 4, 8) supplemented by OCT imaging sessions every two weeks (Fig. 1). After such standardized period, the treatment was patient-specific. In this work we focus on the treatment's induction phase only.

Image acquisitions were performed using the commercially available Topcon spectral-domain OCT scanner (Topcon Inc., Pyramus, NJ), centered on the macula of the retina. The device acquires anisotropic 3D OCT images having $512 \times 128 \times 885$ voxels with $11.72 \times 46.88 \times 2.6 \ \mu m^3$ spacing, and covering the volume of $6 \times 6 \times 2.3 \ mm^3$.
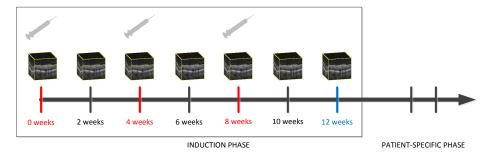
Fig. 1: Illustration of the induction phase of patient treatment, with the denoted imaging and treatment (in red) time points.

2.2   Retinal Layer and SEAD Segmentation

To quantify the retinal anatomical structure, the layer segmentations are performed with a graph-search based method [3, 4]. The method is able to efficiently find the globally optimal solution of multiple interacting surfaces, given surface cost functions and a set of geometric constraints defining intra-surface smoothness and inter-surface distance variability. Sets of multiple surfaces are hierarchically detected starting from the most evident ones and ending with the most subtle layer interfaces. The method segments ten retinal layers defined by eleven surfaces from the inner limiting membrane (ILM) to the retinal pigment epithelium (RPE) (Fig. 2). Even though under the presence of SEAD the intraretinal layers are not always accurately segmented, general layer properties are well captured.

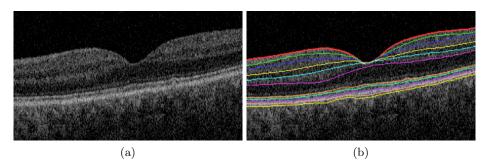

(a)   (b)

Fig. 2: Example of ten layer (eleven surface) segmentation.

The layer segmentation is followed by automated intraretinal and subretinal SEAD segmentation method [7]. Although the shape, size, and location of fluid regions varies, SEADs exhibit layer dependent properties. For example, intraretinal fluid normally appears above the outer plexiform layer and has a cystoid shape while the subretinal fluid lies beneath the outer segment layer but above the RPE. These properties are used by the segmentation method, which consists of a supervised voxel classification approach. During the training phase,

for each image voxel, a set of features composed of textural, structural, and layer location information is calculated. During the testing phase, for each voxel a k-nearest-neighbor classifier is employed to assign its probability of belonging to a SEAD. Lastly, the probability map is thresholded at 50% probability level to obtain binary SEAD segmentations (Fig. 3). No further post-processing of the resulting binary segmentations was performed.
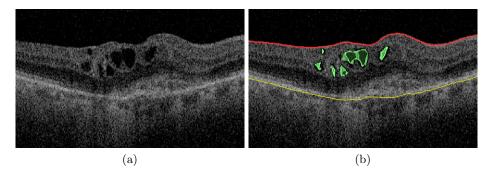

(a)   (b)

Fig. 3: Example of intraretinal SEAD segmentation (in green).

To segment the remaining sub-RPE regions of SEADs, a recently proposed method was used [6]. The method extends the previously segmented layers by adding outer retinal—subretinal layer (ORSR). It is based on the graph-search layer segmentation method where by using the image properties describing the local structural abnormalities it is able to locally increase the inter-surface distance constraints accordingly. While the method does not detect fluid, it focuses on segmenting the RPE and the Bruch's membrane in the presence of PED. The sub-RPE fluid is subsequently segmented by analyzing the obtained ORSR layer thickness.

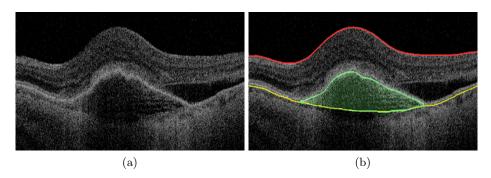(a)   (b)

Fig. 4: Example of a correctly identified Bruch's membrane surface (yellow) in the presence of PED, and the resulting sub-RPE fluid segmentation (green).

2.3   Feature Extraction and Tretment Response Prediction

For each patient, from its longitudinal series of images and derived segmentations, we extract a set of quantitative features characterizing the underlying anatomy. To limit the dimensionality of the resulting feature vector, we divide the retina into nine subvolumes (Fig. 5) inspired by the partitioning used in Early Treatment Diabetic Retinopathy Study (ETDRS). The simplified ETDRS grid is centered on the fovea, which was detected automatically as a local dip in the ILM surface. However as heavily diseased patients can have large retinal distortions, in this pilot study the fovea was not always easily identifiable and its locations have been visually inspected and corrected if needed. Once the subvolumes are defined, a number of quantitative values is extracted for each subvolume, and the entire parafovea and perifovea. The features are based on the properties of the image intensity distribution and the segmented layers and SEADs (Table 1). In order for the features to correspond across subjects, all left eye scans were mirrored to conform to the scans of the right eye.
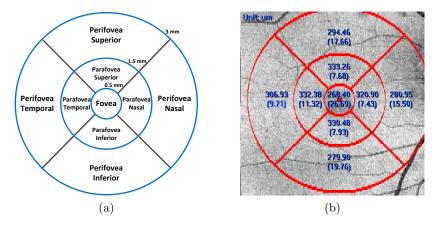

Fig. 5: Subvolume analysis. (a) Nine regions of interest defined by ETDRS. (b) The corresponding grid with total retinal thickness mean (std) values overlayed on the OCT projection image.

The set of extracted features are employed for the task of predicting whether the subjects will respond to the anti-VEGF treatment. The subjects were assumed to be in the non-responsive group if the total segmented SEAD volume present at the start (week 0) increases, stays the same or drops by less than 20% at the end of the induction phase (week 12). Since there are seven imaging sessions (week 0 − week 12), the features are extracted from the first six imaging sessions of each patient producing a total of 756 features. To take into account the relation between the features measuring the same characteristic at different time-points, the features in week 2 - week 10 are made relative to their value at the start of the treatment (week 0). Furthermore since such a large number

Table 1: Set of features computed per image per subvolume of interest.

Feature	Description
Intensity	Mean, std, skewness and kurtosis
Layer	Mean and std of the total retinal thickness
SEAD (sub-RPE separately)	Total volume, 2D en-face area Per entire image only: number of connected components, volume of the biggest component, mean and std distance to fovea, mean and std distance to BMO

of features is much larger than the number of subjects in our dataset, a feature selection is required. For each feature we evaluate how well it distinguishes between the two categories (responders, non-responders) by performing a t-test. Features that lead to statistically significant difference with a p-value < 0.01 were selected. Finally, the selected features were used to train a support vector machine to predict the treatment responses.

3   Results

The performance of the method was evaluated on longitudinal images forming the induction phase of 30 patients suffering from wet AMD. Thus, 210 OCT images were processed and analyzed. In our studied population, there were 70% (21/30) responders and 30% (9/30) non-responders, which is representative of a general population.

For descriptive analysis, we look at the change during the induction phase of the total retinal thickness (TRT) at the fovea (Fig. 6). TRT is often used as a measure of overall retinal state as the presence of SEAD causes an increase of TRT. From the figure it can be observed that already after the first treatment (week 2) the TRT becomes well correlated with its value at the end of the induction phase. This shows the potential to infer the retinal response after only one injection.

As a result of feature selection, the extracted features that had the highest association with the treatment outcome are shown in Table 2. As expected, the majority of important features are based on 3D and 2D SEAD measurements. The measurements of SEAD volume and its 2D en-face area, in particular in the perifoveal area, was found to be highly associated with the response status. The top three features represent measurements at week 4, which confirms that observing the retinal response and its status after only one injection can be indicative of the final treatment response.

Finally, we quantitatively evaluate the performance of the classification task, i.e., predicting treatment response at the end of the induction phase. Leave-one-out cross-validation was performed on support vector machine having the radial basis function kernel. The best classification accuracy obtained by accounting for the features from all the imaging sessions (week 0 − week 10) was 87% (Fig. 7).
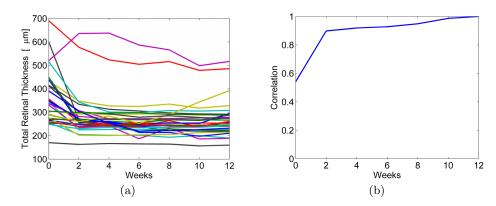
Fig. 6: Total retinal thickness (TRT) during the induction phase. (a) Population variability (each subject is a line). (b) Correlation of the TRT at a certain week with the TRT measured at the end of the induction phase (week 12).

Table 2: Feature Selection: First six features ordered by importance.

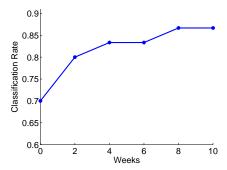| Feature Description	Measured at week
SEAD En-Face Area in Perifovea	4
SEAD Volume in Perifovea	4
Total SEAD Volume	4
Mean Total Retinal Thickness in Parafovea	8
Total sub-RPE SEAD Volume	4
SEAD Volume in Perifovea	10

Fig. 7: Classification rate as the selected features are becoming available with the progress of the induction phase.

4   Conclusion

In this paper, we presented a method to learn the retinal treatment responses from a longitudinal series of OCT scans acquired during the induction phase.

To the best of our knowledge, this is the first time such a task was attempted. The obtained classification rates of 83% after four weeks (one injection) and 87% after 8 weeks (two injections) are promising for achieving early prediction of the subject's response. Additional contribution is the identification of quantitative features which are most strongly associated with the treatment outcome. The measurements of SEAD volume and its 2D en-face area, in particular in the perifoveal area, were found to be highly associated with the response status.

The study has two main limitations. First, the population size (30 patients) is small. This is mainly due to the difficulty of recruiting patients that are able to follow a long and strict imaging and treatment protocol, which is even more involved than the already-demanding sequence of standard-of-care imaging and treatments. This further motivates our study and the need to differentiate between responders and non-responders as early as possible. Second, the quantitative measures extracted from OCT images of such diseased patients are still not sufficiently validated. Both the layer segmentation in the presence of a disease and the SEAD segmentation are difficult and still-unsolved tasks. As a consequence, our responder/non-responder categorization could have been affected by the segmentation errors.

Future work is to extend the study to the subject-specific treatment phase, to predict the time to retreatment or the number of required injections per year.

References

1. Jager, R.D., Mieler, W.F., Miller, J.W.: Age-related macular degeneration. N. Engl. J. Med. **358**(24) (2008) 2606–17
2. Folk, J.C., Stone, E.M.: Ranibizumab therapy for neovascular age-related macular degeneration. N. Engl. J. Med. **363**(17) (2010) 1648–1655
3. Garvin, M.K., Abramoff, M.D., Wu, X., Russell, S.R., Burns, T.L., Sonka, M.: Automated 3-D intraretinal layer segmentation of macular spectral-domain optical coherence tomography images. IEEE Trans. Med. Imag. **28**(9) (2009) 1436–47
4. Quellec, G., Lee, K., Dolejsi, M., Garvin, M.K., Abramoff, M.D., Sonka, M.: Three-dimensional analysis of retinal layer texture: identification of fluid-filled regions in SD-OCT of the macula. IEEE Trans. Med. Imag. **29**(6) (2010) 1321–30
5. Chen, X., Niemeijer, M., Zhang, L., Lee, K., Abramoff, M.D., Sonka, M.: Three-dimensional segmentation of fluid-associated abnormalities in retinal OCT: probability constrained graph-search-graph-cut. IEEE Trans. Med. Imag. **31**(8) (2012) 1521–31
6. Zhang, L., Sonka, M., Russell, S., Folk, J., Abràmoff, M.D.: Quantifying disrupted outer retina-subretinal layer in SD-OCT images in choroidal neovascularization. Invest. Ophthalmol. Vis. Sci. (5) (2014)
7. Xu, X., Zhang, L., Lee, K., Wahle, A., Chen, X., Wu, X., Abramoff, M.D., Sonka, M.: Automated choroidal neovascularization associated abnormality detection and quantitative analysis from clinical SD-OCT. In: ARVO. (2013) 54(6):5510.