Cystoid Fluid Color Map Generation in Optical Coherence Tomography Images Using a Densely Connected Convolutional Neural Network

Plácido L. Vidal
CITIC Research center of Information and Communication Technologies
Universidade da Coruña
A Coruña, Spain
placido.francisco.lizancos.vidal@udc.es

Joaquim de Moura
CITIC Research center of Information and Communication Technologies
Universidade da Coruña
A Coruña, Spain
joaquim.demoura@udc.es

Jorge Novo
CITIC Research center of Information and Communication Technologies
Universidade da Coruña
A Coruña, Spain
jnovo@udc.es

Marcos Ortega
CITIC Research center of Information and Communication Technologies
Universidade da Coruña
mortega@udc.es

***Abstract*—Optical Coherence Tomography (OCT) is a medical imaging modality that is currently the focus of many advancements in the field of ophthalmology. It is widely used to diagnose relevant diseases like Diabetic Macular Edema (DME) or Age-related Macular Degeneration (AMD), both among the principal causes of blindness. These diseases have in common the presence of pathological cystoid fluid accumulations inside the retinal layers that tear its tissues, hindering the correct vision of the patient.**

In the last years, several works proposed a variety of methodologies to obtain a precise segmentation of these fluid regions. However, many cystoid patterns present several difficulties that harden significantly the process. In particular, some of these cystoid bodies present diffuse limits, others are deformed by shadows, appear mixed with other tissues and other complex situations. To overcome these drawbacks, a regional analysis has been proven to be reliable in these problematic regions.

In this work, we propose the use of the DenseNet architecture to perform this regional analysis instead of the classical machine learning approaches, and use it to represent the pathological identifications with an intuitive color map. We trained, validated and tested the DenseNet neural network with a dataset composed of 3247 samples labeled by an expert. They were extracted from 156 images taken with two of the principal OCT devices of the domain. Then, this network was used to generate the color map representations of the cystoid areas in the OCT images. Our proposal achieved robust results in these regions, with a satisfactory 97.48% $\pm$ 0.7611 mean test accuracy as well as a mean AUC of 0.9961 $\pm$ 0.0029.

This work is supported by the Instituto de Salud Carlos III, Government of Spain and FEDER funds of the European Union through the DTS18/00136 research projects and by the Ministerio de Economía y Competitividad, Government of Spain through the DPI2015-69948-R research project. Also, this work has received financial support from the European Union (European Regional Development Fund - ERDF) and the Xunta de Galicia, Centro singular de investigación de Galicia accreditation 2016-2019, Ref. ED431G/01; and Grupos de Referencia Competitiva, Ref. ED431C 2016-047.

I. Introduction

Optical Coherence Tomography (OCT) is a medical imaging technique that is used to study the retina, offering a cross-sectional visualization of its tissues and other relevant structures. The study of the retina is of great importance, as it allows to study and detect pathologies not only specific of the vision, but also others of general impact as diabetes [1], [2], abnormalities of the vascular system [3], [4] and also of the nervous system [5], [6]. In this work, we focus on the detection of cystoid fluid accumulations in these retinal tissues using OCT images, as they are one of the main symptoms of two of the main causes of blindness in developed countries: the Age-related Macular Degeneration (AMD) and Diabetic Macular Edema (DME) pathologies.

In both diseases, an early diagnosis is critical, as if their presence is not early detected and treated, invasive procedures may be the only possible solution. Thus, being OCT the principal and wide-available medical imaging technique and the relevance of the aforementioned pathologies, several works have arisen proposing approaches for the detection of these cystoid bodies. These methodologies intend to help the clinician to better assess the pathology, trying to isolate the subjectivity of the human expert from the analysis and diagnosis.

Most of the proposed automatic approaches followed the same main strategy: preprocessing step, initial candidate segmentation and posterior false positive (FP) filtering. As reference, *Wilkins et al.* [7] presented a segmentation process based on a binarization using an empirically set threshold and a filtering of FPs based on properties as the size and the intensity profiles of the candidates. The same way, *González et al.* [8] proposed a methodology based on the watershed transform

to generate the initial candidate set, filtering these candidates with texture-based features. And, similarly, *Sahoo et al.* [9] proposed a methodology based on the k-means clustering to generate the initial detection, but directly filtering the detected candidates with a complex set of rules.

Being OCT an imaging technique that creates consecutive sectional cuts of the retina (B-Scans), several works tried also to take advantage of this feature to create the precise segmentation of these cystoid bodies. An example of works that use these 3D capabilities are the ones of *Chen et al.* [10] or *Montuoro et al.* [11], both based on a 3D voxel classification and posterior contour refinement process to generate a precise 3D segmentation of these fluid regions.

Finally, recent approaches have applied deep learning techniques to create the precise segmentation of these cystoid bodies. Some examples of that can be found in the work of *Lee et al.* [12], using a Convolutional Neural Network (CNN); *Schlegl et al.* [13] with a decoder-encoder architecture; *Gopinath and Sivaswamy* [14] using a CNN with a posterior clustering step to filter FPs; *Venhuizen et al.* [15], *Tennakoon et al.* [16] and *Girish et al.* [17] with approaches based on derivations of the U-Net model, a network architecture designed in the context of medical imaging; and *Roy et al.* [18], with a proposal of a new architecture called "ReLayNet", to also solve this retinal cystoid fluid segmentation issue.

All these works focused on the extraction of a precise segmentation for these cystoid bodies, but not all these fluid structures can be directly and clearly segmented. These fluid leakages tend to be mixed with the normal tissues of the retina (presenting diffuse borders). They can also be easily obscured in the OCT images by vessel shadows or from other dense structures in superior layers. Sometimes, they can even appear in grouped colonies with interlaced, diffuse and ripped tissue walls. An example of these complications can be seen in Figure 1, where two OCT images present clear cystoid regions but their exact segmentation is undefined due to several of the abnormalities mentioned before.

The works of *de Moura et al.* [19], [20], with a classical learning approach, were the first ones to tackle the issue of these irregular patterns. Instead of trying to find a precise segmentation like the rest of works, they proposed the study of classifying squared samples of a given size depending if they contain cystoid fluid or not. For this, they studied and extracted a significative number of intensity and texture descriptors from the samples, performed a comprehensive feature analysis and trained several classifiers.

This way of handling the images allowed the identification of the cystoid regions even if they presented the mentioned complications. Following this idea and using the classical learning approach of feature extraction, selection and classification, the methodology was evolved into a complementary paradigm to the classical segmentation approach by *Vidal et al.* [21], where it was improved and expanded to create intuitive fluid region maps. These maps proved to be able to create a robust representation of the cystoid regions even in the aforementioned zones of the OCT images, as they represent an abstraction of the fluid regions with different values depending on the underlying image. Moreover, these maps successfully create a representation of this information in a way that is easy for the expert clinician to analyze, and that could also be used to extract image features for the clinical characterization of the pathology.

In this work, we present an alternative powerful and simplified way of generating these maps. Instead of the mentioned classical approach that needs a manual definition of features that may be incomplete or suboptimal, a feature study combined with selection algorithms and a battery of trained models to find the proper classifier that is able to fit the data, we propose the direct use of a CNN to the problem. In particular, given its characteristics, we adapted to this issue a densely connected convolutional network (DenseNet) [22]. Hence, the chosen architecture is robust to overfitting thanks to its densely connected layers, as it implicitly creates its own features that may not even been considered by the classical approach, offering an extra of self-supervision thanks to the shorter paths between individual layers. Moreover, all these improvements could also result in a better final accuracy of the system.

This paper is divided in three sections. In Section 2, *Methodology*, we further explain the details of the procedure followed in this work. Section 3, *Results and Discussion*, presents the results of the experiments, as well as details the precise configuration used for each of the steps. Finally, in Section 4 *Conclusions*, we state the final thoughts and future work lines.

II. Methodology

The methodology firstly reduces the region of interest (ROI) of the input OCT image to the area between the innermost and outermost retinal layers: the Inner Limiting Membrane (ILM) and the Retinal Pigment Epithelium (RPE), respectively. From this ROI, several samples were extracted by an expert, belonging to both healthy and cystoid fluid samples. These samples were used to train a recent densely connected convolutional network (DenseNet) proposed by *Huang et al.* [22]. Using this trained model, with the design of an image sampling procedure, the target images are decomposed into overlapping samples of a given size and classified into two of the considered categories. Finally, using the results of this classification step, an intuitive color map is created indicating the certainty of the region of belonging to the cystoid fluid class. In the subsequent sections we will further explain each one of the steps.

A. Delimitation of the region of interest

As the cystoid fluid leakages can only appear inside the retinal layers, we firstly extract the region of interest (ROI) delimited by the ILM and RPE layers. By doing so, we are also improving the results in two different ways: first of all, we reduce as much as possible the area that the sampling algorithm has to evaluate. Secondly, we reduce the number of cases that the network needs to handle as above the ILM layer

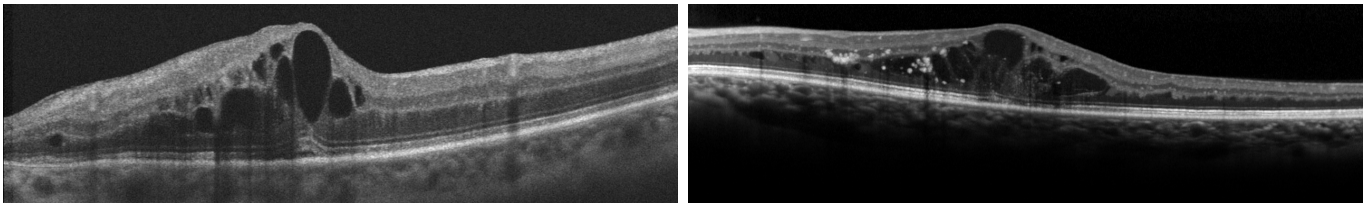

Fig. 1. OCT images from the two different OCT devices presenting diffuse cystoid limits, shadows and other complications.

we find the vitreous humor, a fluid inside the eyeballs that presents a very similar homogeneous pattern to the intraretinal cystoid fluid in the OCT images. Also, under the RPE layer, we can find the choroid region, a vascular layer that has circular patterns and borders that are very similar to the ones present in the target cystoid fluid regions.

To extract the ILM and RPE layers we used an approach based in the proposal of *Chiu et al.* [23], [24]. It uses the algorithm of Dijkstra [24] to find the minimum paths between both sides of the OCT image, using the gradients as weights paths that identify the target layers. In Figure 2, an example of OCT image with these ROI-marking layers that were using the algorithm is presented.

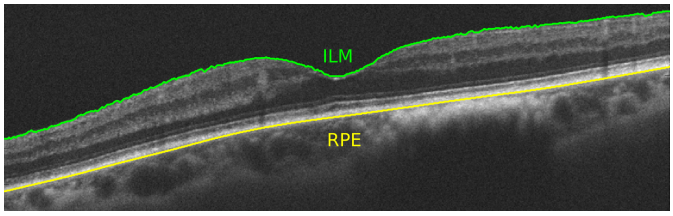

Fig. 2. Example of image with the relevant ILM and RPE layers extracted from the Cirrus OCT capture device.

B. Dataset creation

Once the ROI is delimited between the two limiting retinal layers, an expert extracted and labeled a battery of samples from all the images of both studied devices. Some of these extracted samples fell partially outside the ROI, as they were taken close to the retinal limiting layers. To prevent the inclusion of these extraretinal regions and patterns in the dataset, all these samples were cropped. This way, only the maximum rectangular region inside each sample that contained exclusively ROI region was used.

C. Training of the densely connected convolutional network

Thanks to the increased connectivity between layers of the used densely connected network, it offers a reduced risk of overfitting to problems with smaller training test sizes. Thus, we considered that it constitutes a suitable architecture to support our proposal. In particular, this densely connected network is arranged in dense blocks, where each layer is directly connected to every other layer in a feed-forward fashion through a composite function of a batch normalization, a rectified linear unit and a 3×3 convolution.

In this work, we use a modification on the DenseNet-161 configuration, whose main characteristics are depicted in Table I, where we adapted the classification layer to the problem we are dealing with. In particular, the original layer was designed to support up to 1000 classes from the ImageNet dataset, so we reinitialized that same linear layer to output only the two considered target classes.

To virtually increase the dataset size, the samples are randomly flipped horizontally (as the same patterns may appear naturally in any horizontal orientation depending on the afflicted eye and position). Additionally, as the network shape is squared and with a fixed size, we reshaped every sample to match the input size of the network of 224×224. This way, even the rectangular-shaped samples that are product of the ROI subsample extraction can be classified by the network. The dataset mean was also shifted to 0 and the standard deviation to 1.

Finally, to train the network, we based our procedure in the original work, using a Stochastic Gradient Descent (SGD) approach with Nesterov's accelerated gradient [25], [26]. Finally, as loss function, we used the Cross Entropy Loss, defined in Equation 1:

$$loss(x, class) = -log\left(\frac{exp(x[class])}{\sum_j exp(x[j])}\right) \quad (1)$$

D. Image sampling and color map generation

To create the target color maps, we used a sampling algorithm [21] that divides the biggest rectangular area that contains the ROI into overlapping windows of a given size. As some of the samples may partially fall outside the ROI, they are cropped the same way we did with the original training dataset (an example of this case can be seen in Figure 3). This way, we ensure that only the ROI patterns are classified by the network. Moreover, only samples centered in the retina are considered valid as, when cropped, smaller sample sizes may deform too much the information when they are resized to fit the input layer of the network.

As mentioned, after this sampling step, the samples are resized to fit the network input, normalized and classified using the previously trained DenseNet model. Each classified sample is used to vote over the region it overlapped in the original image. This way, the more overlapping samples consider a given region to be cystoid, the more confidence the color map will show. As some limiting regions may be overlapped by less samples than others, each pixel is normalized between 0

TABLE I
STRUCTURE OF THE DENSENET161-BASED NETWORK CONFIGURATION USED IN THIS WORK.

Layers	**Convolution**	**Pooling**	**Dense block**	**Transition layer**		**Dense block**	**Transition layer**		**Dense block**	**Transition layer**		**Dense block**	**Classification layer**	
Output Size	112×112	56×56	56×56	56×56	28×28	28×28	28×28	14×14	14×14	14×14	7×7	7×7	1×1	-
Structure	7×7 Convolution Stride 2	3×3 Maxpool Stride 2	$\begin{bmatrix} 1\times1 & conv \\ 3\times3 & conv \end{bmatrix} \times 6$	1×1 Convolution	2×2 Average pool Stride 2	$\begin{bmatrix} 1\times1 & conv \\ 3\times3 & conv \end{bmatrix} \times 12$	1×1 Convolution	2×2 Average pool Stride 2	$\begin{bmatrix} 1\times1 & conv \\ 3\times3 & conv \end{bmatrix} \times 36$	1×1 Convolution	2×2 Average pool Stride 2	$\begin{bmatrix} 1\times1 & conv \\ 3\times3 & conv \end{bmatrix} \times 24$	7×7 Global average pool	Fully-connected softmax 2 classes

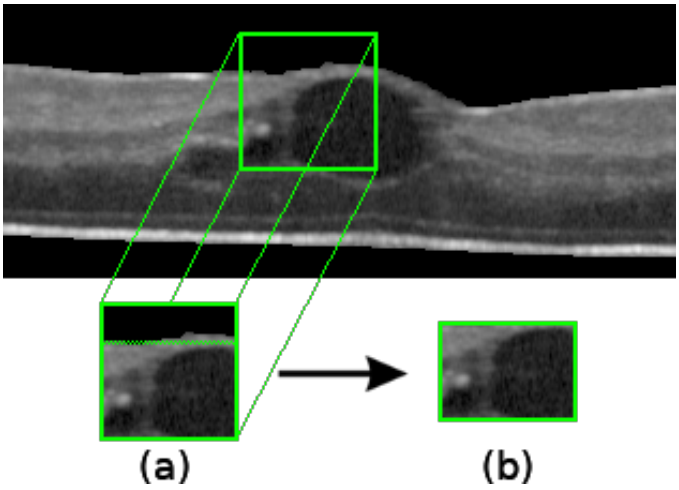

Fig. 3. Example of an extracted sample from an OCT image that falls outside the ROI (a) and the resulting cropped sample that will be used in the dataset (b).

and 1 depending on the total of samples that overlapped that given area. Finally, an intuitive color map is applied to the resulting confidence map to facilitate a better understanding of the results.

III. RESULTS AND DISCUSSION

The training of the system was performed with 61×61 samples that were extracted from two representative ophthalmological capture devices: a CIRRUS HD-OCT Carl Zeiss Meditec confocal scanning laser ophthalmoscope and a Spectralis OCT confocal scanning laser ophthalmoscope from Heidelberg Engineering. All the samples were extracted from OCT images centered in the macula by an expert, trying to include representative samples from both pathological and non-pathological regions. These samples were extracted from different patients and also from both left and right eyes. From the Cirrus device, a total of 1613 samples from 83 OCT histological images were considered, 806 belonging to fluid regions and 807 from non-fluid ones. On the other hand, from the Spectralis device, 1634 samples from 73 OCT histological images were considered, 778 fluid-containing samples and 856 from non-fluid regions.

The dataset was divided into 3 subsets: 50% of the samples were randomly selected from the total to train the network, 30% of the samples were randomly selected for the validation step and the remaining 20% of samples were reserved to test the final trained model with a subset of unbiased samples.

For the training of the DenseNet model, an initial learning rate of 0.01 and a momentum of 0.9 were empirically defined as the starting conditions. Moreover, as previous tests showed that this learning rate may overfit the network in long runs, the learning rate during the training step was decreased by a factor of 3 if the validation loss did not improve over 10 epochs.

As stopping criterion, instead of defining a fixed number of epochs, we established an early-stopping rule. This rule automatically stops the training of the network after 15 epochs if no further improvement in the validation loss was achieved. After the training step, the model that obtained the best validation accuracy is the one that is afterwards used to generate the intuitive color maps.

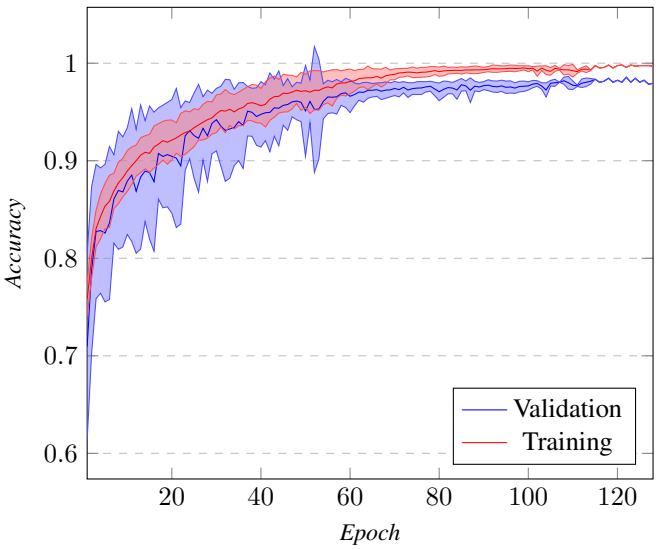

Fig. 4. Mean training and validation accuracy after 50 independent randomly-distributed experiments.

To correctly assess the feasibility of the problem, the experiments were repeated, following the previous criteria, 50 times. In Figure 4, we can see the result of these experiments over all the epochs. As seen in this figure, around epochs 60 - 70 the training reaches stability and the difference between models is not significant, as even the standard deviation of the accuracy from the epoch 70 to the last epoch is only of 0.23 in training and 0.34 in validation. This way, these experiments reach a final mean test accuracy of 97.48% $\pm$ 0.7611 as well as an AUC of 0.9961 $\pm$ 0.0029. This stagnation can be further seen in Figure 5, where the first epochs show a significant variance compared to the more stable section after the aforementioned epochs. Moreover, it outperforms the

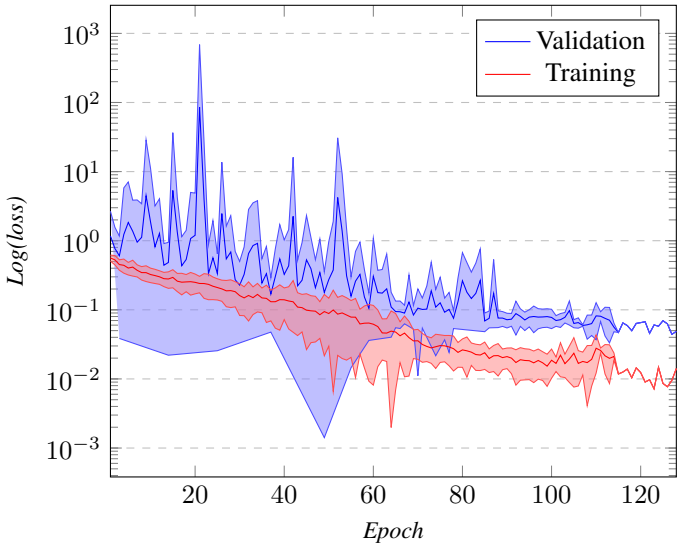

Fig. 5. Mean training and validation loss after 50 independent randomly-distributed experiments. A logarithmic scale has been set to correctly display the loss values for a better understanding of the results.

original approach of *Vidal et al.* [21], that used the same OCT image dataset and reached a 94.01% of accuracy with classical machine learning techniques.

TABLE II
RESULTING CONFUSION MATRIX FROM THE TEST SUBSET OF SAMPLES AND THE FINAL TRAINED NETWORK.

		Expert label		
		Positive	Negative	Total
Predicted	Positive	313	8	321
	Negative	9	320	329
	Total	322	328	650

To further demonstrate the capabilities of the methodology, a model has been chosen randomly from these experiments. This model achieved a test accuracy of 97.38% with an AUC of 0.9935, and its confusion matrix with the test dataset is shown in Table II. As seen in this matrix, the sensitivity and specificity of the system is well balanced, not favoring neither the pathological nor the healthy classes.

By examining the unbiased test samples of this model, we can further understand the behavior of the net, in particular the patterns that where learned by the network. First of all, Figure 6 presents representative samples with different cystoid bodies that were correctly detected by the network. We can see how the network is able to successfully detect regions with diffuse borders, cystoid bodies with similar size to the used window, fluid bodies that are present in poor quality OCT images, fluid regions of small size closely agglomerated (even with diffuse borders themselves) and mixed with other pathological structures like the hyperreflective exudates that tend to neighbor the fluid structures faced in this work.

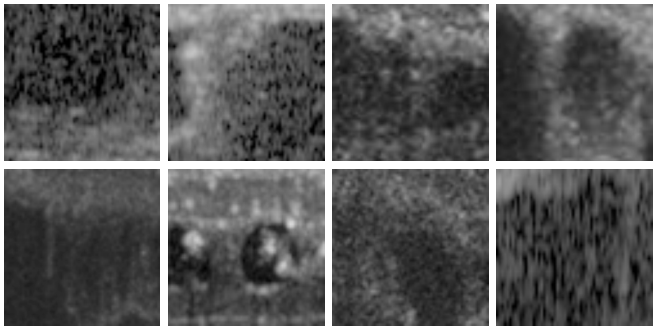

Fig. 6. Examples of correctly classified positive test samples extracted from images of both devices.

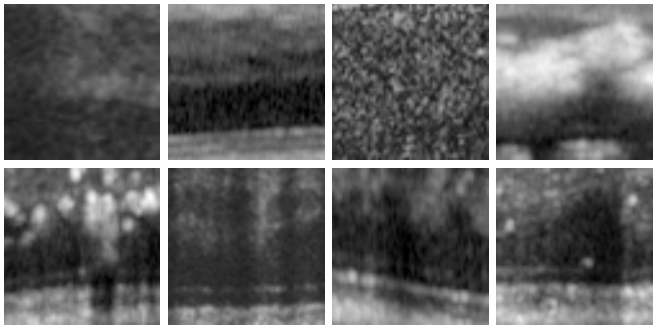

Fig. 7. Examples of correctly classified negative test samples extracted from images of both devices.

On the other hand, the significant robustness of the system is emphasized with scenarios as the ones of Figure 7. In this figure, we can check how the trained network was clearly able to distinguish between exudates mixed with pathological fluid (as shown in Figure 6) and exudates that cast a shadow due to the nature of the imaging technique and their higher density. Moreover, other dark structures like darkened (without any fluid presence) retinal layers, shadows with the source outside the sample window and dark-to-light contrasts that tend to confuse other approaches are successfully classified into the correct class by the proposed method.

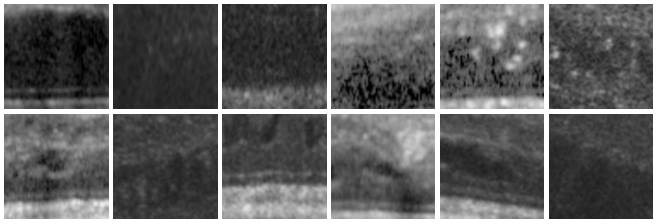

Fig. 8. Examples of test samples incorrectly classified as positives (1st row) and test samples incorrectly classified as negatives (2nd row).

Finally, Figure 8 presents some representative examples where the system failed to correctly assess the class of the test sample. In the FPs presented in Figure 8, 1st row, we can distinguish two types of behaviors. The system deemed as negative some cases where the internal fluid covered the

entire sample, as well as samples where the tissues presented a pattern very similar to the ones in the borders of the problematic cysts. The dark homogeneous pattern is commonly found in the foveal region of the images taken with the Cirrus device (an example of this foveal pattern can be seen in Figure 2, where the foveal region presents a very similar texture to the eyeball internal fluid: the mentioned vitreous humor) so, without further information about the surrounding structures, the network considered the pathological case as the most probable (which, considering we are talking about the medical field, is the most logical option in case of doubt). Also, this represents a reasonable decision, as even the expert opinion could reach this conclusion with the information contained in the analyzed sample.

Finally, regarding the false negatives (FN) of Figure 8, 2nd row, the system presents some confusion when cysts presents smaller sizes (called microcysts or pseudocysts) and are mixed with darker retinal tissues. Probably, due to the size of these cystoid structures and the rarity of the case, the dataset includes an smaller set of representative samples to satisfy the network training.

With all the information above, it is shown that the main disadvantage of the system is suffering from its lack of information from adjacent windows. Nonetheless, using the methodology for the color map creation, this problem is palliated by a voting system between overlapping adjacent samples. In Figures 9 and 10, some examples of color maps with different levels of complexity are shown from both Cirrus and Spectralis capture devices, respectively. In the 1st row of both figures, an OCT image without any kind of cystoid fluid is also represented with satisfactory results. On the other hand, the 2nd row presents a case with different complications mixed with the cystoid bodies. Both images present clear cystoid bodies where part of their structure is mixed with normal retinal tissues and shadows cast by other structures. Additionally, the Spectralis image (Figure 9, 2nd row) presents other dense bodies that could affect the detection, but the system correctly shows the cystoid regions even in these unclear zones and correctly skips the other dense pathological bodies.

Figure 8, 3rd row, presents an image example from each device including nearby groups of cystoid bodies. The Cirrus map is specially interesting, as these cystoid bodies are very dim and difficult to analyze for both the expert and the network (as can be seen in the FNs of Figure 8). This image shows the benefits of the voting system, compensating the information deficiency of a single sample, as even the two dim small cystoid structures are detected with a sensible confidence by the methodology.

IV. Conclusions

The OCT imaging modality has proven to be one of the most reliable means of analyzing and helping to the early diagnosis of pathologies like the Age-related Macular Degeneration and Diabetic Macular Edema, both diseases among the main causes of blindness in developed countries. The intraretinal fluid accumulation that characterizes these diseases was mainly faced by different proposals with direct and specific segmentation approaches. Despite that many fluid regions can be correctly segmented, many others present significant drawbacks for these proposals. In particular, several of the fluid bodies in the retina present diffuse contours, appear mixed with other pathological structures or are corrupted in the OCT image due to shadows and other darkened structures by the nature of the technique. To overcome these complications, an alternative paradigm to the segmentation approach was also recently proposed based on a regional analysis and a color map representation of these cystoid fluid bodies. Moreover, this way of representing the cystoid regions allowed the clinicians to analyze the problematic cases of the precise segmentation approach too.

This alternative paradigm mitigates these limitations by generating a confidence map instead of a segmentation by means of an intensive sampling of the image and a voting system. Additionally, this paradigm could be easily adapted to other pathologies and medical imaging modalities, as the ground truth is represented by a reduced set of representative samples and there is no need for a complete precise segmentation (if was even possible in the first place).

This alternative paradigm demonstrated its viability and suitability being previously faced by a classical learning approach, which implied the necessity of a manually defined large set of features that had to be filtered using feature selection techniques. After this feature ranking, several models had to be trained and tested to find the optimal number of features and parameters that better represented the defined problem. Finally, a valid model would be obtained and used in the map creation step.

In contrast, by using Deep Learning techniques, there is no need for this analysis of relevant features, feature selection and model training/testing to find the optimal subset of features whatsoever. Instead, the own network training process and weight adjusting performs all these processes implicitly, simplifying the entire approach. Moreover, the used network architecture offers extra robustness to overfitting with small datasets, being able to create its own features that may not been considered by the manually defined set. These advantages not only reduce the complexity of the training and map generation processes, but also sensibly improve the final results.

To train and validate the methodology, we used a dataset composed by 156 OCT images from two representative ophthalmological devices: a CIRRUS HD-OCT Carl Zeiss Meditec confocal scanning laser ophthalmoscope and a Spectralis OCT confocal scanning laser ophthalmoscope from Heidelberg Engineering. From this dataset, 3247 squared windows were marked and extracted by an expert into fluid or non-fluid samples. To further study the capabilities of the proposed approach, we trained 50 DenseNet networks randomly dividing the dataset into 50% for training, 30% for validation and the remaining 20% for testing. These experiments proved that the network was suitable for this issue and that was correctly adapted to this task, achieving a mean test accuracy of $97.48\% \pm 0.7611$ and a mean area under the ROC curve of 0.9961 ± 0.0029. Furthermore, a coherent analysis on the

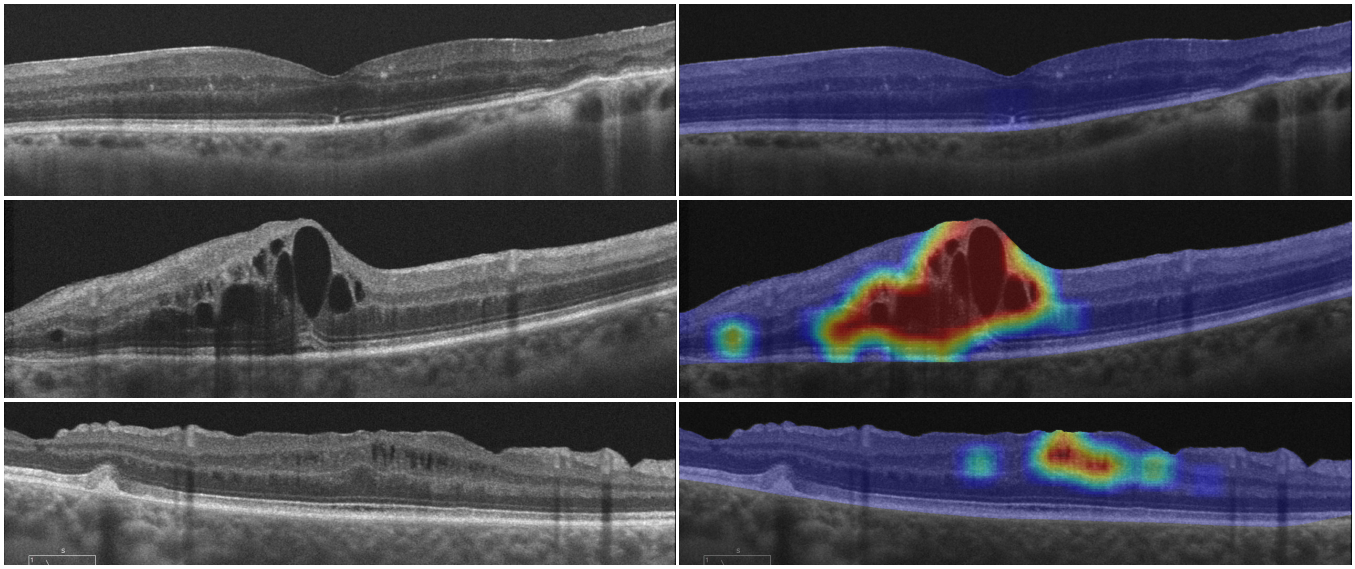

(a) (b)

Fig. 9. Examples of color maps that were generated with images coming from the CIRRUS device (b) with their respective original image (a).

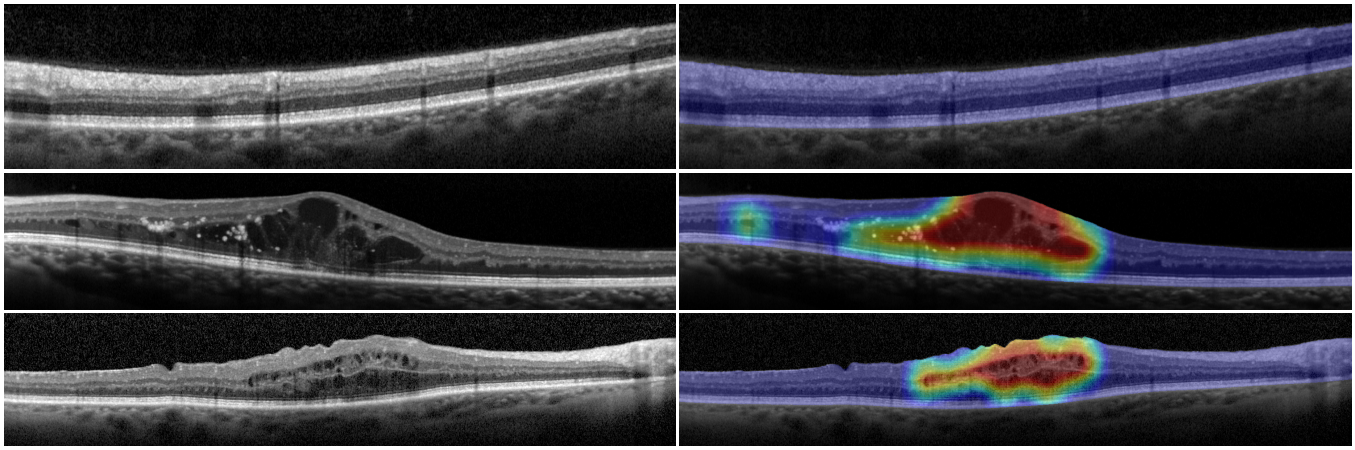

(a) (b)

Fig. 10. Examples of color maps that weregenerated with images coming from the SPECTRALIS device (b) with their respective original image (a).

different cases that the network may encounter was conducted, using a network from the previous experiments that was chosen at random. The test samples and color maps showed that the system was able to satisfactorily classify the pathological cystoid regions, independently of the presence of another pathologies or complications caused by the OCT technique itself.

Additionally, the intuitive color map generation process implicitly integrates information about the neighboring patterns to each classification through the voting system. Thus, each detected zone is generated with a compendium of information coming from the neighboring regions. This means that, even if some misclassifications occur, they present a negligible confidence value, being barely represented in the final map. So, this methodology not only offers robustness in areas where the cystoid bodies are mixed with other retinal patterns and tissues, but also to its own misclassifications. For this reason, we consider the proposed system as an appropriate approach to be used in the medical field, helping the clinicians in the early detection and assessment of the severity of the considered pathologies.

As future work, we plan to expand this methodology to different types of pathological structures that are also present in the retina, as well as the design of several biomarkers that can be extracted from these maps that could help in the assessment of the studied pathologies. Finally, an study on possible kernels for the voting process and window sizes is interesting to further diminish the influence of the limited “field of view” of the network when classifying the samples.

REFERENCES

[1] D. J. Browning, C. M. Fraser, and S. Clark, "The relationship of macular thickness to clinically graded diabetic retinopathy severity in eyes without clinically detected diabetic macular edema," *Ophthalmology*, vol. 115, no. 3, pp. 533 – 539.e2, 2008.

[2] S. Vujosevic, A. Muraca, V. Gatti, L. Masocro, M. Brambilla, B. Cannillo, E. Villani, P. Nucci, and S. De Cilla, "Peripapillary Microvascular and Neural Changes in Diabetes Mellitus: An OCT-Angiography Study," *Investigate Ophthalmology & visual science*, vol. 59, no. 12, pp. 5074–5081, OCT 2018.

[3] N. M. Taha, H. T. Asklany, A. H. Mahmoud, L. Hammoda, H. R. Attallah, A. M. Kamel, and M. A. AbdelWahab, "Retinal fluorescein angiography: A sensitive and specific tool to predict coronary slow flow," *The Egyptian Heart Journal*, vol. 70, no. 3, pp. 167 – 171, 2018.

[4] J. de Moura, J. Novo, J. Rouco, N. Barreira, M. Penedo, and M. Ortega, "Retinal vasculature identification and characterization using oct imaging," in *OCT*, M. Lanza, Ed. Rijeka: IntechOpen, 2018, ch. 2.

[5] S. M. Silverstein, D. Paterno, L. Cherneski, and S. Green, "Optical coherence tomography indices of structural retinal pathology in schizophrenia," *Psychological Medicine*, vol. 48, no. 12, p. 2023–2033, 2018.

[6] H. Liao, Z. Zhu, and Y. Peng, "Potential Utility of Retinal Imaging for Alzheimer's Disease: A Review," *Frontiers in aging neuroscience*, vol. 10, Jun 2018.

[7] G. Wilkins, O. Houghton, and A. Oldenburg, "Automated segmentation of intraretinal cystoid fluid in optical coherence tomography," *IEEE Transactions on Biomedical Engineering*, vol. 59, no. 4, pp. 1109–1114, 2012.

[8] A. González, B. Remeseiro, M. Ortega, M. Penedo, and P. Charlón, "Automatic cyst detection in OCT retinal images combining region flooding and texture analysis," *IEEE Int. Symp. on Computer-Based Medical Systems*, pp. 397–400, 2013.

[9] M. Sahoo, S. Pal, and M. Mitra, "Automatic segmentation of accumulated fluid inside the retinal layers from optical coherence tomography images," *Measurement*, vol. 101, pp. 138 – 144, 2017.

[10] X. Chen, M. Niemeijer, L. Zhang, K. Lee, M. D. Abramoff, and M. Sonka, "Three-dimensional segmentation of fluid-associated abnormalities in retinal oct: Probability constrained graph-search-graph-cut," *IEEE Transactions on Medical Imaging*, vol. 31, no. 8, pp. 1521–1531, Aug 2012.

[11] A. Montuoro, S. M. Waldstein, B. S. Gerendas, U. Schmidt-Erfurth, and H. Bogunović, "Joint retinal layer and fluid segmentation in oct scans of eyes with severe macular edema using unsupervised representation and auto-context," *Biomed. Opt. Express*, vol. 8, no. 3, pp. 1874–1888, Mar 2017.

[12] C. S. Lee, A. J. Tyring, N. P. Deruyter, Y. Wu, A. Rokem, and A. Y. Lee, "Deep-learning based, automated segmentation of macular edema in optical coherence tomography," *Biomed. Opt. Express*, vol. 8, no. 7, pp. 3440–3448, Jul 2017.

[13] T. Schlegl, S. M. Waldstein, H. Bogunovic, F. Endstraßer, A. Sadeghipour, A.-M. Philip, D. Podkowinski, B. S. Gerendas, G. Langs, and U. Schmidt-Erfurth, "Fully automated detection and quantification of macular fluid in oct using deep learning," *Ophthalmology*, vol. 125, no. 4, pp. 549 – 558, 2018.

[14] K. Gopinath and J. Sivaswamy, "Segmentation of retinal cysts from optical coherence tomography volumes via selective enhancement," *IEEE Journal of Biomedical and Health Informatics*, pp. 1–1, 1 2018.

[15] F. G. Venhuizen, B. van Ginneken, B. Liefers, F. van Asten, V. Schreur, S. Fauser, C. Hoyng, T. Theelen, and C. I. Sánchez, "Deep learning approach for the detection and quantification of intraretinal cystoid fluid in multivendor optical coherence tomography," *Biomed. Opt. Express*, vol. 9, no. 4, pp. 1545–1569, Apr 2018.

[16] R. Tennakoon, A. K. Gostar, R. Hoseinnezhad, and A. Bab-Hadiashar, "Retinal fluid segmentation in oct images using adversarial loss based convolutional neural networks," in *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, April 2018, pp. 1436–1440.

[17] G. Girish, B. Thakur, S. Roychowdhury, A. Kothari, and J. Rajan, "Segmentation of intra-retinal cysts from optical coherence tomography images using a fully convolutional neural network model," *IEEE Journal of Biomedical and Health Informatics*, vol. 99, pp. 1–1, 2 2018.

[18] A. G. Roy, S. Conjeti, S. P. K. Karri, D. Sheet, A. Katouzian, C. Wachinger, and N. Navab, "Relaynet: Retinal layer and fluid segmentation of macular optical coherence tomography using fully convolutional network," *CoRR*, vol. abs/1704.02161, 2017.

[19] J. Moura, J. Novo, J. Rouco, M. Penedo, and M. Ortega, "Automatic identification of intraretinal cystoid regions in optical coherence tomography," *Conference on Artificial Intelligence in Medicine in Europe - AIME'17*, pp. 305–315, 2017.

[20] J. Moura, P. L. Vidal, J. Novo, J. Rouco, and M. Ortega, "Feature definition, analysis and selection for cystoid region characterization in optical coherence tomography," in *Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 21st International Conference KES-2017, Marseille, France, 6-8 September 2017.*, 2017, pp. 1369–1377.

[21] P. L. Vidal, J. de Moura, J. Novo, M. G. Penedo, and M. Ortega, "Intraretinal fluid identification via enhanced maps using optical coherence tomography images," *Biomed. Opt. Express*, vol. 9, no. 10, pp. 4730–4754, Oct 2018.

[22] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks." in *CVPR*, vol. 1, no. 2, 2017, p. 3.

[23] S. Chiu, X. Li, P. Nicholas, C. Toth, J. Izatt, and S. Farsiu, "Automatic segmentation of seven retinal layers in SDOCT images congruent with expert manual segmentation," *Optics Express*, vol. 10, no. 10, pp. 19 413–19 428, 2010.

[24] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, no. 1, pp. 269–271, Dec 1959.

[25] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, "On the importance of initialization and momentum in deep learning," in *International conference on machine learning*, 2013, pp. 1139–1147.

[26] Y. E. Nesterov, "A method of solving a convex programming problem with convergence rate O(1ˆk)," *Dokl. Akad. Nauk SSSR*, vol. 269, pp. 543–547, 1983.