

# A SELF-TRAINING FRAMEWORK FOR GLAUCOMA GRADING IN OCT B-SCANS

Gabriel García\*, Adrián Colomer\*, Rafael Verdú-Monedero<sup>†</sup>, Jose Dolz<sup>‡</sup> and Valery Naranjo\*

\**Instituto de Investigación e Innovación en Bioingeniería, I3B Universitat Politècnica de València, Valencia, Spain.*

Email: {jogarpa7, adcogra, vnaranjo}@i3b.upv.es

<sup>†</sup>*Departamento de TICs, Universidad Politécnica de Cartagena, 30202, Cartagena, Spain.*

<sup>‡</sup>*Software Engineering, École de technologie supérieure, Montreal, QC H3C 1K3, Canada.*

**Abstract**—In this paper, we present a self-training-based framework for glaucoma grading using OCT B-scans under the presence of domain shift. Particularly, the proposed two-step learning methodology resorts to pseudo-labels generated during the first step to augment the training dataset on the target domain, which is then used to train the final target model. This allows transferring knowledge-domain from the unlabeled data. Additionally, we propose a novel glaucoma-specific backbone which introduces residual and attention modules via skip-connections to refine the embedding features of the latent space. By doing this, our model is capable of improving state-of-the-art from a quantitative and interpretability perspective. The reported results demonstrate that the proposed learning strategy can boost the performance of the model on the target dataset without incurring in additional annotation steps, by using only labels from the source examples. Our model consistently outperforms the baseline by 1-3% across different metrics and bridges the gap with respect to training the model on the labeled target data.

**Index Terms**—Glaucoma grading, Self-training, OCT

## I. INTRODUCTION

Glaucoma is a degenerative optic neuropathy that causes functional damage and visual field defects by altering several structures in the optic nerve head (ONH) of the retina [1]. Currently, the diagnostic procedure for detecting glaucoma requires several time-consuming tests, besides a visual examination of medical images, whose interpretation is often subjective for expert ophthalmologists at the grading time [2]. So, many state-of-the-art studies propose image-processing techniques to help experts via machine-learning solutions.

The optical coherence tomography (OCT) is the quintessential imaging modality for glaucomatous damage evaluation since it allows evidencing the deterioration of the cell layers of the optic nerve, which is intimately linked to the glaucoma severity (see Fig. 1). Specifically, the retinal nerve fibre layer (RNFL) has been reported in the clinical literature as the most important structure for glaucoma progression [3].

Some of the previous studies combined hand-driven learning algorithms with conventional machine-learning classifiers to discern between healthy and glaucomatous OCT samples [4], [5]. The aforementioned studies demand a previous segmentation of the retinal layers of interest to conduct the hand-

crafted feature extraction from specific regions of the B-scans, e.g. the RNFL thickness. However, approaches based on prior segmentation knowledge could transfer remaining errors for downstream classification tasks, according to the recent study [6]. To avoid this shortcoming, deep learning arises as an appealing alternative to derive high-performing computer-aided diagnosis systems in the ophthalmology field.

Nevertheless, despite the rise of these models in many computer vision and medical problems, their application on OCT images for glaucoma assessment still presents several limitations. First, most of the literature focuses on discerning between healthy and glaucoma classes via OCT B-scans [7], [8], SD-OCT volumes [9]–[11] or probability RNFL maps by combining fundus images and OCT samples [12], [13]. Indeed, to the best of our knowledge, only [2] proposes a glaucoma-based scenario beyond the healthy-glaucoma classification by including the suspect label. Furthermore, the limited size of available data sets may hamper the generalization capabilities of learned models, as they could easily lead to overfitting. This is particularly important in the presence of domain shift between training (*labeled*) and testing (*unlabeled*) data distribution. A naive solution would be to collect and labeled data that follows a similar distribution than testing images, which can be then employed to fine-tune the model on the new test data. However, this requires to label large amounts of testing images, involving a time-consuming human process, which is unrealistic in clinical practice. To alleviate the problem of domain shifts, unsupervised domain adaptation has recently appeared as an interesting learning strategy. These methods typically include an adversarial learning framework [14], [15], which can lead to unstable training and high sensitiveness to hyperparameters.

To fill these gaps in the literature, we propose an alternative learning strategy for glaucoma grading that differs from current literature in several ways. First, unlike existing methods, our approach addresses the problem of glaucoma grading, according to the clinical annotation criteria [16]. Second, it demonstrates to improve the testing performance of a model trained in the presence of domain shift, approaching the results obtained by full-supervision. Third, the simplicity of our model facilitates the training convergence, contrary to complex adversarial learning-based methods. And last, we propose architectural changes that result in enhanced useful

This work has been partially funded by the following projects: [H2020-ICT-2016-2017, 732613], [DPI2016-77869-C2-1-R, AES2017-PI17/00771, AES2017-PI17/00821 and 20901/PI18]. The work of Gabriel García has been supported by the State Research Spanish Agency PTA2017-14610-I.

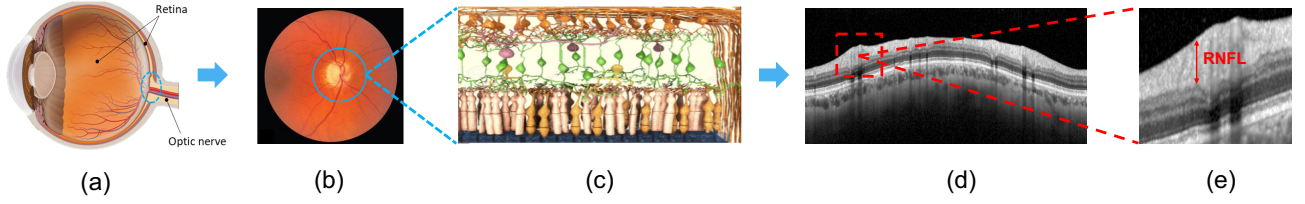


Fig. 1. (a) Eyeball showing the regions of interest. (b) Fundus image highlighting the optic nerve head (ONH). (c) Arrangement of the cell fibre layers of the retina. (d) Typical OCT B-scan evidencing the retinal fibre layers by different grey-intensity levels. (e) Cropping of the RNFL structure.

representations from the OCT B-scans, leading to a better performance and more meaningful prediction interpretations compared to conventional architectures.

## II. METHODS

### A. Self-training strategy

Self-training, or self-supervised learning, aims at automatically generating a supervisory signal for a given task, which can be then used to enhance the representation learning of features or to label an independent dataset. The former typically involves integrating an unsupervised pretext task [17], relational reasoning [18] or contrastive learning [19]. Nevertheless, in our glaucoma grading scenario, we advocate for the use of a sequential step, where the model is first trained on a labeled source dataset. Then, this model performs an inference on the unlabeled dataset to generate the target pseudo-labels, which are later used to train the model, mimicking full supervision on the test data (see Fig. 2). This learning strategy has demonstrated a high classification performance on different imaging modalities, such as histopathology [20] or natural images from ImageNet [21]. Formally, we denote  $D = \{X, Y\}$  as the independent training set, where  $x_i \in X$  refers to the  $i$ -sample of the database with its corresponding ground truth  $y_i \in Y$ , being  $i = 1, 2, \dots, \mathcal{N}$  the number of training image pairs. Furthermore, we use  $T = \{Y, \hat{p}(Y|X)\}$  to represent a given task, where  $\hat{p}(Y|X)$  is usually learnt by a neural network. In the current work, we leverage the scenario where the source and target domains are related ( $D^S \sim D^T$ ) since both correspond to OCT samples, but acquired from different hardware settings. Furthermore, tasks across models are the same ( $T^S = T^T$ ), as we focus on multi-class B-scans classification to discern between healthy, early and advanced glaucoma classes.

Thus, let  $x_i \in \mathbb{R}^{M \times N}$  be a raw B-scan of dimensions  $M \times N = 248 \times 384$ , a first model is defined by training a base encoder network  $f_\phi(\cdot)$  on the labeled source dataset  $D^S$  composed of  $\mathcal{N}^S$  samples, where each training instance is denoted by  $(x_i, y_i)$ , as observed in Fig. 2. Then, the embedding representations  $z_i = f_\phi(x_i)$  are fed into a classification layer  $g_\theta$  to extract the logit scores  $s_i$  that are transformed via softmax function to obtain a class-probability  $\hat{y}_i$  (see Algorithm 1). The coefficients  $\phi$  and  $\theta$  of the first model are updated during the back-propagation step at every epoch  $e = 1, 2, \dots, \epsilon$ . Once the training is finished on this model, it is used to predict the class of each sample  $x_j \in X^T$  from the unlabeled target

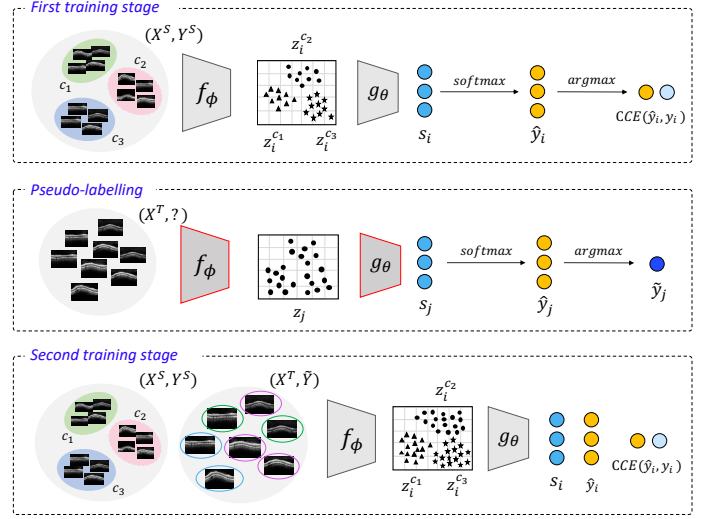


Fig. 2. Illustration of the proposed learning strategy broken down by stages.

dataset  $D^T$ , with  $j = 1, 2, \dots, \mathcal{N}^T$ , leading to the corresponding pseudo-labels  $\hat{Y}^T$  (Algorithm 2). Last, the pseudo-labels are used to augment the training dataset, which results in  $D^{S,T} = \{X^{S,T}, Y^{S,T}\}$ , where  $Y^{S,T} = Y^S \cup \hat{Y}^T$ . In this way, the model in the last step is trained on the augmented pseudo-labeled dataset, following Algorithm 1, under the hypothesis that substantial improvements could be reported at the test time because two reasons: i) the increase of labeled training samples and ii) the knowledge distilled from the unlabeled target dataset.

---

### Algorithm 1: Model Training

---

**Data:** Training set  $D^S = \{(x_1, y_1), \dots, (x_{\mathcal{N}^S}, y_{\mathcal{N}^S})\}$ .

**Results:** Trained coefficients  $\phi, \theta$  ;

**Algorithm:**

$\phi, \theta \leftarrow \text{random};$

**for**  $e \leftarrow 1$  **to**  $\epsilon$  **do**

**for**  $i \leftarrow 1$  **to**  $\mathcal{N}^S$  **do**

$z_i \leftarrow f_\phi(x_i)$  ;

$s_i \leftarrow g_\theta(z_i)$  ;

$\hat{y}_i \leftarrow \frac{\exp(s_i)}{\sum_{c'} \exp(-s_{c'})}$  ;

$\mathcal{L}(y, \hat{y}) \leftarrow - \sum_i y_i \log(\hat{y}_i)$  ;

    Update  $\phi, \theta$  using  $\nabla_{\phi, \theta} \mathcal{L}$  ;

---

---

**Algorithm 2:** Pseudo-labeling

---

**Data:** Target set  $X^T = \{x_1, \dots, x_{N^T}\}$ .

**Results:** Target pseudo-labels  $\tilde{y}$ ;

**Algorithm:**

$\phi, \theta \leftarrow \text{frozen};$

**for**  $j \leftarrow 1$  **to**  $N^T$  **do**

$z_j \leftarrow f_\phi(x_j);$   
     $s_j \leftarrow g_\theta(z_j);$   
     $\tilde{y}_j \leftarrow \frac{\exp(s_j)}{\sum_{c'} \exp(s_{c'})};$

---

### B. Proposed architecture

In addition to the presented learning strategy, we propose several architectural changes that improve over existing ones both quantitatively and from an interpretability perspective. In our recent works focused on glaucoma detection from raw OCT samples [7]–[9], we conducted an in-depth experimental analysis of different deep-learning architectures, both trained from scratch and fine-tuning the most popular models in the literature. Specifically, the VGG family of networks reported the best results in [7]. So, in [9], we employed them as a strong backbone to develop a new feature extractor able to learn useful representations from the slides of a spectral-domain (SD)-OCT volume. Following this findings, in this paper, we adopt our previous work [9] as a benchmark to infer the encoder architecture by introducing slight nuances that allow reinforcing the learning of the intrinsic knowledge of OCT samples for glaucoma grading.

As observed in Fig. 3, we freeze the three first convolutional blocks of the VGG architecture by applying a deep fine-tuning strategy in order to leverage the knowledge acquired by the network when it was pre-trained on the *ImageNet* dataset. Following [9], we include a residual block via convolutional-skip connections and an attention module by means of an identity shortcut to give rise to the *RAGNet\_v2* architecture. As a novelty, we refine the filters of the residual convolutions to optimize the glaucoma learning process by leveraging the domain-specific knowledge of the OCT samples. In particular, we introduce a tailored kernel size of  $3 \times 1$  (yellow box) to enforce the network to focus on critical glaucoma-specific regions which underlie contrast changes along the vertical axis of the B-scans. A concatenation aggregation function is used to combine the outputs from the residual block and VGG architecture. Then, a convolution  $1 \times 1$  is applied to reduce the filters' dimension without affecting the properties of the feature maps. This structure is introduced via skip-connections to refine the embedded space throughout a convolutional autoencoder  $1 \times 1$  with a sigmoid function aimed to recalibrate the feature learning. Again, a concatenation operation is defined to combine the information from the attention block to the feature map of the main branch. An additional  $1 \times 1$  convolutional layer is included to provide an embedding volume of features  $F$  of dimensions  $H \times W \times C = 7 \times 12 \times 60$ , where  $C = 60$  was empirically calculated as a multiple of the number of classes to encourage a better

learning convergence. Finally, a global average-pooling (GAP) layer is applied to compute a spatial squeeze from the feature volume  $F$  such that  $F \in \mathbb{R}^{H \times W \times C} \rightarrow Z \in \mathbb{R}^C$ . In this way, given an input OCT image  $x_i \in \mathbb{R}^{M \times N}$ , an embedding representation map  $z_i \in \mathbb{R}^C$  is achieved by the backbone network  $f_\phi : X \rightarrow Z$ . Regarding the classification stage  $g_\theta$ , a single output layer with three neurons, corresponding to the number of classes, is implemented through a softmax activation function to determine the probability that  $x_i$  belong to the class  $c$  (see Fig. 3).

## III. ABLATION EXPERIMENTS

### A. Data sets

To evaluate the proposed learning methodology, we resort to two independent databases containing circumpapillary B-scans around the optic nerve head (ONH) of the retina. Note that the OCT samples from both source  $D^S$  and target  $D^T$  data sets were acquired from different Hospitals using the Heidelberg Spectralis OCT system under distinct settings conditions, e.g. illumination, noisy, contrast, etc. A different senior ophthalmologist (with more than 25 years of clinical experience) annotated each B-scan according to the European Guideline for Glaucoma Diagnosis. We considered  $D^T$  as an unlabeled data set during the entire learning process to conduct the proposed methodology. We only used the target labels at the test time to evaluate the models' performance. Information about the data sets distribution per patient and per sample is detailed in Table I. Note that the study was conducted according to the guidelines of the Declaration of Helsinki and approved by the Institutional Review Board of each implicated Hospital. Informed consent was obtained from all subjects involved in the study.

TABLE I  
PATIENTS (PAT.) AND SAMPLES (SAMP.) PER DATA SET GROUPED BY CATEGORIES, ACCORDING TO THE EXPERTS' ANNOTATION

	Healthy (pat./samp.)	Early (pat./samp.)	Advanced (pat./samp.)	TOTAL (pat./samp.)
Source $D^S$	32 / 41	28 / 35	25 / 31	85 / 107
Target $D^T$	26 / 49	24 / 37	21 / 26	71 / 112
TOTAL	58 / 90	52 / 72	46 / 57	156 / 219

**Data partitioning.** At the first stage, we performed a patient-level data partitioning to divide the source data set  $D^S$  into five different subsets. A 5-fold cross-validation strategy was addressed to provide robust models and reliable results. In each of the five iterations,  $\frac{4}{5}$  of the data were used to train the first model, whereas the remaining samples were employed as a validation subset to prevent overfitting. Otherwise, we randomly selected  $\frac{2}{3}$  from the  $D^T$  data set to generate the pseudo-labels from which training the model at the second stage. The rest of the target data was used as a test set.

### B. Validation of the backbone architecture

In this stage, we conduct a comparison between the proposed model and the state-of-the-art architectures focusing on

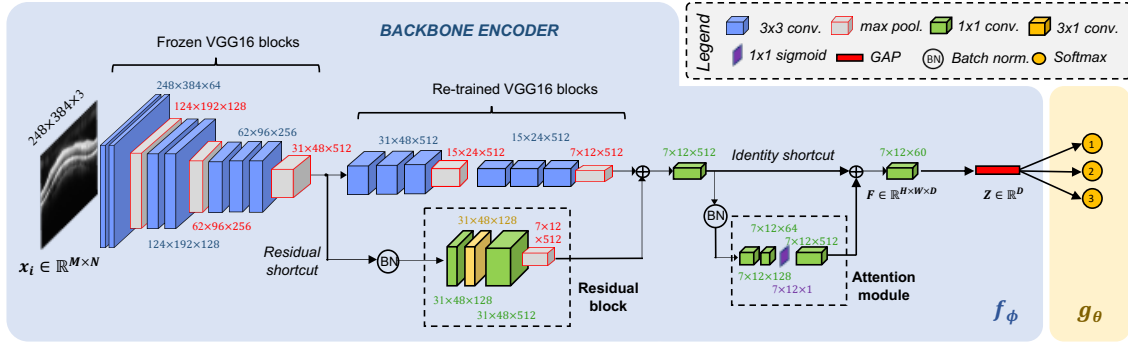


Fig. 3. Backbone architecture used to train the model during the first and second stages. The blue region corresponds to the encoder network structure, whereas the yellow frame denotes the classification layer to discern between healthy, early and advanced glaucomatous cases.

OCT-based glaucoma identification. Following the experimental setup carried out in [9], we contrast here the canonical VGG family of networks and the proposed *RAGNet\_v2* architecture using as a backbone both VGG16 and VGG19 architectures. In Table II, we show the performance of the aforementioned networks during the training of the model at the first stage in a multi-class scenario for glaucoma grading. To this end, different figures of merit, e.g. sensitivity (SN), specificity (SP), F-score (FS), accuracy (ACC) and area under the ROC curve (AUC), are considered. Note that results correspond to the average and standard deviation from the cross-validation stage in terms of micro-average per class.

TABLE II  
MICRO-AVERAGE CROSS-VALIDATION RESULTS ACHIEVED DURING THE TRAINING OF THE FIRST MODEL USING THE SOURCE DATABASE  $D^S$

	VGG16	VGG19	RAGNet_v2 (with VGG16)	RAGNet_v2 (with VGG19)
SN	0.67 ± 0.06	0.75 ± 0.11	0.76 ± 0.11	<b>0.77 ± 0.06</b>
SP	0.84 ± 0.03	0.87 ± 0.05	0.88 ± 0.06	<b>0.89 ± 0.03</b>
FS	0.67 ± 0.06	0.75 ± 0.11	0.76 ± 0.11	<b>0.77 ± 0.06</b>
ACC	0.78 ± 0.04	0.83 ± 0.07	0.84 ± 0.08	<b>0.85 ± 0.04</b>
AUC	0.76 ± 0.06	0.82 ± 0.08	0.82 ± 0.09	<b>0.83 ± 0.05</b>

Based on the results from Table II, we selected the network relative to the *RAGNet\_v2* (with VGG19) as a baseline to address the pseudo-labeling stage since it outperformed conventional architectures for both VGG16 and VGG19 approaches. In a non-realistic setting, we calculate the performance of the selected backbone at the pseudo-labeling time to determine the usefulness of the proposed approach. To this end, the baseline trained on  $(X^S, Y^S)$  was tested on  $X^T$ , which resulted in accuracy  $ACC = 0.8376$ . Besides, the qualitative class activation maps (CAMs) computed in Fig. 4 further strengthen our confidence in the proposed backbone encoder, since it is evident how the heat maps provided by the attention module (Fig. 4 (b)) focus on more localized and glaucoma-specific regions than conventional VGG networks (Fig. 4 (a)). Also, the findings from the CAMs are directly in line with the clinicians' opinion, since the generated heat maps keep an evident relationship between the RNFL thickness and the predicted class, according to the clinical statements [3].

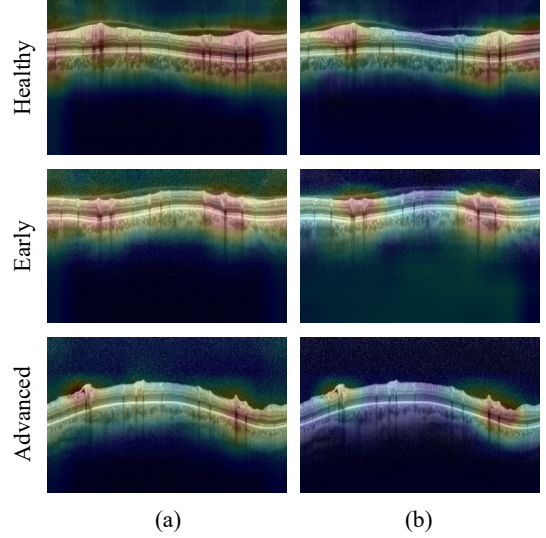


Fig. 4. Class activation maps (CAMs). (a) Heat maps extracted from the VGG19 architecture. (b) Heat maps achieved from the RAGNet\_v2 (with VGG19 as a backbone) at the output from the attention module.

#### IV. PREDICTION RESULTS

Once the pseudo-labels  $\tilde{y}$  were generated during the first stage, we trained the model making use of the same *RAGNet\_v2* (with VGG19) architecture. All the experiments were conducted under the same conditions in order to provide a reliable comparison between the different approaches. In particular, all models were trained during 100 epochs using 16 B-scans per batch and the Adadelata optimizer to minimize the categorical cross-entropy (CEE) loss function. At this point, it is important to note that there are no public databases to make a comparison with the literature. In addition, no state-of-the-art studies have been performed to grade the glaucoma severity, so replicating previous glaucoma-based methods would lead to an unreliable and non-objective comparison.

In Table III, we report the results achieved by the trained model both in the first stage (*baseline*) and the second stage (*proposed*). Also, as a reference point, we show the performance of the approaches relative to the *upper bound* (model

TABLE III  
TEST RESULTS ACHIEVED DURING THE PREDICTION OF THE TARGET SET

	$X^S$	$X^T$	$Y^S$	$Y^T$	$\tilde{Y}^T$	SN	SP	FS	ACC	AUC
Baseline	✓	–	✓	–	–	0.7059	0.8529	0.7059	0.8039	0.7613
Proposed	✓	✓	✓	–	✓	<b>0.7353</b>	<b>0.8676</b>	<b>0.7353</b>	<b>0.8235</b>	<b>0.7853</b>
Lower bound	–	✓	–	–	✓	0.6765	0.8382	0.6765	0.7843	0.7463
Upper bound	–	✓	–	✓	–	0.7647	0.8824	0.7647	0.8431	0.8050

trained with target labels  $Y^T$ ) and the *lower bound* (model just trained with target pseudo-labels  $\tilde{Y}^T$ ). We can observe that the *proposed* learning strategy, which does not require additional target labeled data, consistently outperforms the *baseline* methodology across all the metrics, with improvements of 1-3%. Note that the *upper bound* scenario is considered to evidence how large the gap in performance is between the fully and semi-supervised approaches. In this case, reported values reveal compelling results, as we observe small differences (2-3%) between the *upper bound* and the *proposed* strategy. Furthermore, the model just trained on the target dataset making use of the pseudo-labels (*lower bound*) results in a poor-performance with respect to the rest of approaches, as expected, with differences ranging from 3% to 6%. This evidences that an increase of the training set via the proposed pseudo-labeling strategy improves the prediction performance for glaucoma grading, as a result of a knowledge transfer between the source  $D^S$  and target  $D^T$  domains.

## V. CONCLUSION

The proposed self-training learning strategy has been successfully applied to grade the glaucoma severity from OCT B-scans in the presence of domain shift. Results have demonstrated that including pseudo-labels in the training-loop can enhance the performance over a model trained only on labeled source data, without incurring on extra annotation steps. In addition, the results achieved by the proposed model surpass those reached by the conventional architectures for glaucoma grading, leading to better predictions from both quantitative and interpretability perspective. These findings are evident in the provided heat maps, which highlight more localized glaucoma-specific areas, which are clinically relevant. As a future work, we intend to evaluate our learning strategy across more datasets that might contain larger domain shifts.

## ACKNOWLEDGMENT

We gratefully acknowledge the support of the Generalitat Valenciana (GVA) for the donation of the DGX A100 used for this work, action co-financed by the European Union through the Programa Operativo del Fondo Europeo de Desarrollo Regional (FEDER) de la Comunitat Valenciana 2014-2020 (IDIFEDER/2020/030).

## REFERENCES

- [1] R. N. Weinreb and P. T. Khaw, "Primary open-angle glaucoma," *The Lancet*, vol. 363, no. 9422, pp. 1711–1720, 2004.
- [2] H. Raja, T. Hassan, M. U. Akram, and N. Werghi, "Clinically verified hybrid deep learning system for retinal ganglion cells aware grading of glaucomatous progression," *IEEE TBME*, 2020.

- [3] A. El-Naby *et al.*, "Correlation of retinal nerve fiber layer thickness and perimetric changes in primary open-angle glaucoma," *Journal of the Egyptian Ophthalmological Society*, vol. 111, no. 1, 2018.
- [4] E. Gao, B. Chen, J. Yang, F. Shi, W. Zhu, D. Xiang, and *et al.*, "Comparison of retinal thickness measurements between the topcon algorithm and a graph-based algorithm in normal and glaucoma eyes," *PloS one*, vol. 10, no. 6, pp. 1–13, 2015.
- [5] S. J. Kim and K. J. e. a. Cho, "Development of machine learning models for diagnosis of glaucoma," *PLoS One*, vol. 12, no. 5, pp. 1–16, 2017.
- [6] A. C. Thompson, A. A. Jammal, and F. A. Medeiros, "A review of deep learning for screening, diagnosis, and detection of glaucoma progression," *Translational Vision Science & Technology*, vol. 9, no. 2, pp. 42–42, 2020.
- [7] G. García, R. d. Amor, A. Colomer, and V. Naranjo, "Glaucoma detection from raw circumpapillary oct images using fully convolutional neural networks," in *IEEE ICIP*, 2020, pp. 2526–2530.
- [8] G. García, A. Colomer, and V. Naranjo, "Analysis of hand-crafted and automatic-learned features for glaucoma detection through raw circumpapillary oct images," in *International Conference on Intelligent Data Engineering and Automated Learning*, 2020, pp. 156–164.
- [9] G. García, A. Colomer, and V. Naranjo, "Glaucoma detection from raw sd-oct volumes: a novel approach focused on spatial dependencies," *Computer Methods and Programs in Biomedicine*, p. 105855, 2020.
- [10] S. Maetschke, B. Antony, H. Ishikawa, G. Wollstein, J. Schuman, and R. Garnavi, "A feature agnostic approach for glaucoma detection in oct volumes," *PloS one*, vol. 14, no. 7, 2019.
- [11] A. R. Ran *et al.*, "Detection of glaucomatous optic neuropathy with spectral-domain optical coherence tomography: a retrospective training and validation deep-learning analysis," *The Lancet Digital Health*, vol. 1, no. 4, pp. e172–e182, 2019.
- [12] K. A. Thakoor, X. Li, E. Tsamis, P. Sajda, and D. C. Hood, "Enhancing the accuracy of glaucoma detection from oct probability maps using convolutional neural networks," in *IEEE EMBC*, 2019, pp. 2036–2040.
- [13] T. Shehryar, M. U. Akram, S. Khalid, S. Nasreen, A. Tariq, A. Perwaiz, and A. Shaukat, "Improved automated detection of glaucoma by correlating fundus and sd-oct image analysis," *International Journal of Imaging Systems and Technology*, 2020.
- [14] S. Yang, X. Zhou, J. Wang, G. Xie, C. Lv, P. Gao, and B. Lv, "Unsupervised domain adaptation for cross-device oct lesion detection via learning adaptive features," in *IEEE ISBI*, 2020, pp. 1570–1573.
- [15] J. Wang *et al.*, "Domain adaptation model for retinopathy detection from cross-domain oct images," in *Medical Imaging with Deep Learning*, 2020, pp. 795–810.
- [16] J. Flammer, "The concept of visual field indices," *Graefes archive for clinical and experimental ophthalmology*, vol. 224, no. 5, pp. 389–392, 1986.
- [17] S. Gidaris, P. Singh, and N. Komodakis, "Unsupervised representation learning by predicting image rotations," in *ICLR*, 2018.
- [18] M. Patacchiola and A. Storkey, "Self-supervised relational reasoning for representation learning," *arXiv preprint arXiv:2006.05849*, 2020.
- [19] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *ICML*, 2020, pp. 1597–1607.
- [20] J. Silva-Rodríguez, A. Colomer, J. Dolz, and V. Naranjo, "Self-learning for weakly supervised gleason grading of local patterns," *IEEE journal of biomedical and health informatics*, 2021.
- [21] Q. Xie, M.-T. Luong, E. Hovy, and Q. V. Le, "Self-training with noisy student improves imagenet classification," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 10 687–10 698.