Guided Adversarial Adaptation Network for Retinal and Choroidal Layer Segmentation

Jingyu Zhao[1,2], Jiong Zhang[2], Bin Deng[1], Yalin Zheng[3], Jiang Liu[4], Ran Song[5], Yitian Zhao[2,*]

[1]School of Mechanical Engineering, Southwest Jiaotong University
[2]Cixi Institute of Biomedical Engineering, Ningbo Institute of Industrial Technology, Chinese Academy of Sciences
[3]University of Liverpool
[4]Department of Computer Science and Engineering, Southern University of Science and Technology
[5]School of Control Science and Engineering, Shandong University

Abstract. Morphological changes, e.g. thickness of retinal or choroidal layers in Optical coherence tomography (OCT), is of great importance in clinic applications as they reveal some specific eye diseases and other systemic conditions. However, there are many challenges in the accurate segmentation of retinal and choroidal layers, such as low contrast between different tissue layers and variations between images acquired from multiple devices. There is a strong demand on accurate and robust segmentation models with high generalization ability to deal with images from different devices. This paper proposes a new unsupervised guided adversarial adaptation (GAA) network to segment both retinal layers and the choroid in OCT images. To our best knowledge, this is the first work to extract retinal and choroidal layers in a unified manner. It first introduces a dual encoder structure to ensure that the encoding path of the source domain image is independent of that of the target domain image. By integrating the dual encoder into an adversarial framework, the holistic GAA network significantly alleviates the performance degradation of the source domain image segmentation caused by parameter entanglement with the encoder of the target domain and also improves the segmentation performance of the target domain images. Experimental results show that the proposed network outperforms other state-of-the-art methods in retinal and choroidal layer segmentation.

Keywords: OCT · domain adaptation · retinal and choroidal layer.

1 Introduction

Optical coherence tomography (OCT) is an indispensable ocular imaging tool and has been extensively used in clinics. Anatomically, the retina can be divided into nine cellular layers with varying thickness [1,2,3]. The choroid is a densely vascularized layer lying between the retina and the sclera of the eye. Fig. 1 (a) illustrates the boundaries of different retinal and choroidal layers in a B-scan

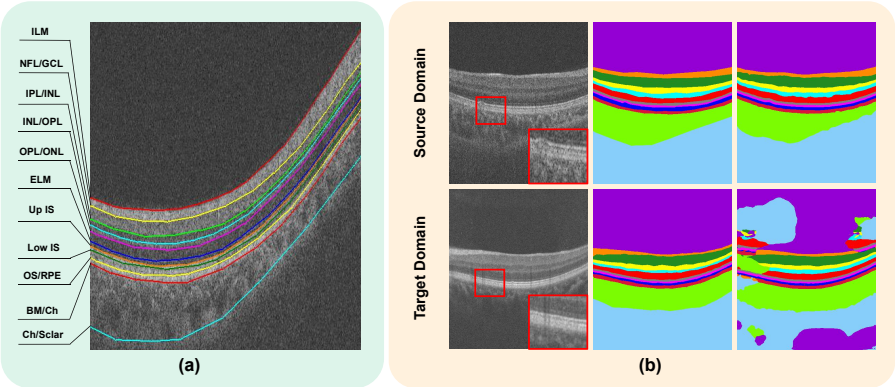

Fig. 1. Illustration of full layer segmentation of OCT images. (a) Boundaries of different layers manually annotated by an experienced clinician. ILM: internal limiting membrane, NFL: nerve fiber layer, GCL: ganglion cells layer, IPL: inner plexiform layer, INL:inner nuclear layer, OPL: outer plexiform layer, ONL: outer plexiform layer, ELM: external limiting membrane, IS: inner segment, OS: outer segment, RPE: retinal pigment epithelium, Ch: choroid. (b) Segmentation results by a pre-trained U-Net on images from the source and target domains. From left to right: example OCT B-scans, ground truth, and the segmentations by U-Net, which was trained on the source domain dataset.

OCT image annotated manually by a senior ophthalmologist. In clinics, layer thickness is an important biomarkers for the diagnosis of many different types of eye diseases. For instance, glaucoma leads to the thinning of the nerve fiber layer (NFL) [4,5]. Age-related macular degeneration (AMD) causes a thinner choroid [6] whilst central serous chorioretinopathy [7] and polypoidal choroidal vasculopathy [6] may lead to choroidal thickening. In consequence, the accurate measurement of thickness of retinal and choroidal layers is vital for diagnosing and monitoring disease progression. However, manual annotation of a large number of images is an exhausting task for clinicians and vulnerable to human errors. Current proprietary segmentation programs of clinical OCT devices still lack accuracy and robustness.

With the rapid development of deep learning, many segmentation networks, such as FCN [8], U-Net [9], CS-Net [10,11] and CE-Net [3], have been employed for retinal layer segmentation tasks. However, to the best of our knowledge, no existing method is dedicated to the segmentation of retinal and choroidal layers in a unified model. In addition, although the retina and choroid of the human eye share similarity, different imaging devices could produce large domain discrepancy even of the same eye due to different noise distributions, i.e., domain gap between the training (source) and test (target) images. This often causes low generalization of a pre-trained model - high performance in the source domain and low performance in target domain, as demonstrated in Fig. 1 (b). Hence, supervised model often requires re-annotating pixel-level ground truth and thus require high labour costs. To this end, it is essential to establish a model trained

on an existing dataset with manual annotations and can be generalized to new test data from another domain (e.g. different device or with varying protocols).

In order to overcome these shortcomings, several unsupervised domain adaptation techniques [12] based on Generative Adversarial Network (GAN) [13,14,15], have been proposed to close the gap between the source and target domains, where manual labels are not available in the target domain. Although some typical approaches such as Adversarial Discriminative Domain Adaptation [16] achieved promising results, the input images of the source and target domains are encoded using the same path, which means that the source and the target domain segmentation networks share the same parameters. As a result, the parameters of the two networks will be entangled with each other and affect the overall performance of the model.

In this paper, we develop a Guided Adversarial Adaptive (GAA) framework for full layer segmentation in OCT images. We use the source domain encoder to guide the target domain encoder for learning segmentation network parameters. The dual encoder structure makes the encoding path of the source domain independent of that of the target domain, and thus does not produce parameter entanglement. Simultaneously, we carry out adversarial adaptation both in the feature and output space of the two domain images, to minimize the feature discrepancy between the source and target domains after encoding. Consequently, the target domain encoder can make continuous progress.

The contributions of our work can be summarized in three-fold: **1)** This is the first attempt to segment full layers (both retinal and choroidal layers) in OCT imagery by a single segmentation model, and it also demonstrates the ability of data adaptation for different imaging devices. **2)** We propose a guided dual-encoder joint structure to guarantee the mutual independence between the encoding paths of the source and target domains for parameter entanglement. **3)** We show that without the need of any manual annotations on the target domain, our method outperforms supervised learning using annotations in the target domain by a large margin.

2 Proposed Method

In this section, we first provide an overview to the proposed method, and then elaborate its two main components, i.e., the guided dual-encoding and the adversarial adaption, respectively.

2.1 Overview

As shown in the Fig. 2, our framework consists of five basic modules: a source domain encoder E_s, a target domain encoder E_t, a sharing decoder D_{sh}, a encoding discriminator Dis_{en} and a decoding discriminator Dis_{de}. Thus, E_s and D_{sh} constitute the source domain segmentation network (SDSN), and E_t and D_{sh} constitute the target domain segmentation network (TDSN). The input OCT image from the source domain is denoted as $X_s \in R^{C\times H\times W}$ with its corresponding annotation L_s while the one from the target domain denoted as $X_t \in R^{C\times H\times W}$ has no annotation.

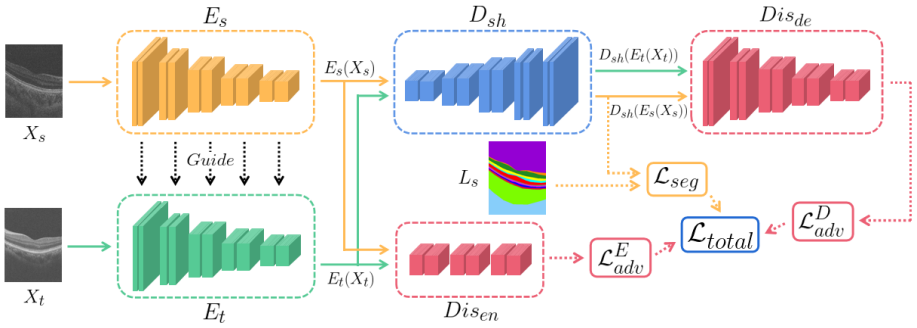

Fig. 2. Overview of the Guided Adversarial Adaptation (GAA) Network. The yellow and the green arrows indicate the source and target domain paths, respectively. The black dashed arrows denote the parameter guidance and the red dashed arrows denote the adversarial learning.

Adversarial methods can reduce the domain discrepancy [17] and thus make the output feature space of TDSN consistent with that of SDSN through training. However, for most adversarial methods, since the parameters of the two networks are shared which trigger the parameter entanglement, the SDSN couldn't gain the optimal solution, so that TDSN will often end with a compromise performance, which is better than that of the model trained only with the source domain data but worse than that of the model trained with the target domain data (assuming that the annotations are available). To alleviate this problem, our idea is to lift the performance of TDSN by allowing it more independence while still keeping its training guided by SDSN for domain adaption. Therefore, we propose a guided dual-encoding architecture where the two encoders of SDSN and TDSN are not shared and the domain adaption for transferring the segmentation knowledge from SDSN to TDSN is delivered through a parameter guidance process and an architecture of adversarial learning.

2.2 Guided Dual-Encoding

In this work, we use two individual encoders for X_s and X_t for SDSN and TDSN, respectively so that there is no parameter sharing in their encoding paths during the training. We build a teacher-student structure which aims to use E_s to guide E_t for encoding the input images of the same modality but acquired by different OCT devices into the same feature space.

Then, two levels of adversarial learning is adopted to promote the continuous progress of E_t and achieve the same encoding effect as E_s. Here, we apply the Exponential Moving Average (EMA) in order to guide E_t to learn parameters from E_s:

$$\begin{cases} \phi_t^n = \gamma \hat{\phi}_t^{n-1} + (1-\gamma)\phi_s^n & (n \geq 2) \\ \phi_t^n = \phi_s^n & (n = 1) \end{cases} \tag{1}$$

$$\hat{\phi}_t^{n-1} = \phi_t^{n-1} - \alpha \nabla J(\phi_t^{n-1}) \quad (n \geq 2) \tag{2}$$

where ϕ_s^n and ϕ_t^n denote the parameters of E_s and E_t before the adversarial training, respectively. $\hat{\phi}_t^n$ denote the parameters of E_t after the adversarial training where n is the iteration index. γ as a hyperparameter is a smoothing coefficient. $J(\phi)$ represents the loss function of E_t and α denotes the learning rate.

2.3 Adversarial Adaptation

We regard the encoder and decoder of the segmentation network as two levels of generators, which conduct adversarial learning with different discriminators in the intermediary feature space and the output space of the whole model, respectively. In the encoding stage, we adopt the adversarial process between E_t and the Dis_{en} to reduce the gap between the feature spaces of $E_s(X_s)$ and $E_t(X_t)$, which aims to encode the X_s and X_t from E_s and E_t respectively to an identical feature space. The Dis_{en} loss $\mathcal{L}_d^E$ and the adversarial loss $\mathcal{L}_{adv}^E$ for E_t can be expressed as follows:

$$\begin{aligned}\mathcal{L}_d^E(X_s, X_t) = -\sum\, & z\, log(Dis_{en}(E_s(X_s))) \\ & + (1-z)(1 - log(Dis_{en}(E_t(X_t))))\end{aligned} \tag{3}$$

$$\mathcal{L}_{adv}^E(X_t) = -\sum log(Dis_{en}(D_{sh}(E_t(X_t)))) \tag{4}$$

where $z = 1$ if the encoding prediction is from S, and $z = 0$ if from T.

In the decoding stage, D_{sh} starts with the encoded features $E_s(X_s)$ and $E_t(X_t)$, and fuse the multi-scale features outputs from different levels of the two encoders concurrently through skip connections. Although such a popular network architecture is well know for improving the segmentation mainly due to the preservation of low-level features, it hinders the restoration of high-level features after the adversarial encoding. Therefore, we use Dis_{de} in the output space of D_{sh} to eliminate the potential impact of the skip connections. Dis_{de} can further enhance the effect of domain adaptation and make the output $D_{sh}(E_t(X_t))$ more similar to $D_{sh}(E_s(X_s))$. The Dis_{en} loss $\mathcal{L}_d^D$ and the adversarial loss $\mathcal{L}_{adv}^D$ for E_t are expressed as follows:

$$\begin{aligned}\mathcal{L}_d^D(X_s, X_t) = -\sum\, & z\, log(Dis_{de}(D_{sh}(E_s(X_s)))) \\ & + (1-z)(1 - log(Dis_{de}(D_{sh}(E_t(X_t))))))\end{aligned} \tag{5}$$

$$\mathcal{L}_{adv}^D(X_t) = -\sum log(Dis_{de}(D_{sh}(E_t(X_t))))). \tag{6}$$

We adopt the mean square error (MSE) loss function to train SDSN with supervised learning.

$$\mathcal{L}_{seg}(X_s) = \frac{1}{2m}\sum_{i=1}^{m}\left(l_s^{(i)} - D_{sh}(E_s(x_s^{(i)}))\right)^2 \tag{7}$$

where $l_s^{(i)}$ and $x_s^{(i)} \in X_s$ denote the i^{th} ground truth and the input image in the source domain, respectively. m is the total number of the source domain OCT images. The overall training objective for our framework is:

$$\mathcal{L}_{total}(X_s, X_t) = \mathcal{L}_{seg}(X_s) + \mathcal{L}_{adv}^E(X_t) + \mathcal{L}_{adv}^D(X_t). \tag{8}$$

Based on Eq. (8), we optimize the following min-max criterion:

$$\min_{G} \max_{D} = \mathcal{L}_{total}(X_s, X_t). \tag{9}$$

Where G denotes the generator and D denotes the discriminator. The ultimate objective is to minimize the segmentation loss for source image, while fooling the discriminators Dis_{en} and Dis_{de} by maximizing the probability of $E_t(X_t)$ and $D_{sh}(E_t(X_t))$ being considered as $E_s(X_s)$ and $D_{sh}(E_s(X_s))$ in the feature and the output spaces, respectively.

3 Experimental Results

3.1 Datasets

Our experiments are performed on two OCT image datasets acquired by two different devices. The first dataset provides layer segmentation annotations and thus is considered as the *source* dataset to train SDSN in a supervised manner. The images in the second dataset are used as *target* domain images and will be used to evaluate the model generalization ability.

Source: **Topcon** dataset consists of 1,280 OCT B-scans with the resolution of 992 $\times$ 512 pixels. All the images were captured by a Topcon DRI-OCT-1 system from 20 subjects. Each image has a corresponding pixel-level manual annotation of the retinal and choroidal layers provided by experts. We make use of 640 images for training and 640 images for testing.

Target: **Optovue** dataset comprises 670 OCT B-scans in total taken by an Optovue RTVue-XR device, with the resolution of 640 $\times$ 400 pixels. In particular, 640 images (without layer manual annotations) were used for training, and 30 images (with manual annotations) were used for testing. All the images were acquired with regulatory approvals and patient consents as appropriate.

3.2 Implementation details

In this experiment, DCGAN [13] and ResNet [18] were employed as the encoding and decoding discriminator, respectively. Both the source and target domain images were cropped to 512 $\times$ 400 pixels automatically, where all the cropped images contain retinal and choroidal layers in either source or target domain. During the training, batch size was set to 4 and we adopt the Adam optimizer with a weight decay of $5e^{-4}$ to train the entire network end-to-end. The smoothing coefficient γ of EMA was set to 0.8.

3.3 Evaluation Metrics

In order to quantitatively evaluate the performance of our framework, the following metrics were calculated: the Dice coefficient (Dice) and the Intersection over Union (IoU). In addition, we introduce the mean absolute error of the boundaries (MAE (pixels)) to evaluate boundary segmentation performance. It is defined as the mean error of retinal and choroidal interfaces:

$$MAE_{R/C} = \frac{1}{M \times N} \sum_{i=1}^{N} \sum_{j=1}^{M} \left| l_j^{(i)} - y_j^{(i)} \right| \tag{10}$$

where $l_j^{(i)}$ and $y_j^{(i)}$ denote the j^{th} boundary mean coordinates of the ground truth label and the prediction of the i^{th} testing OCT image, respectively. We choose $M = 10$ when computing the MAE of the retina layers, and M=2 of the choroid layer.

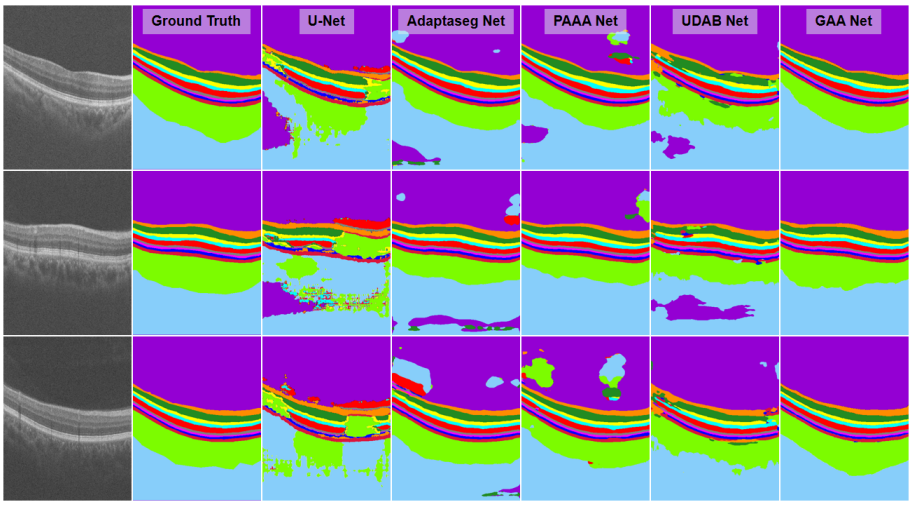

Fig. 3. Visual results of different segmentation networks with domain adaptation.

3.4 Results

In the following sections, we report the segmentation performance under different scenarios, i.e., different segmentation models with and without our adaptation module. For comparisons, we use the well-known network architectures such as U-Net [9], CE-Net [3] and CS-Net [10,11] as SDSN and TDSN.

Domain Adaptation To justify the superiority of the proposed method in domain adaptation, we compared our GAA Net with state-of-the-art domain adaptation methods: Adapt Structured Output Space for Semantic Segmentation (AdaptSeg Net) [19], Perceptual-assisted Adversarial Adaptation(PAAA Net) [20], Unsupervised domain adaptation by backpropagation (UDAB Net) [17], with the U-Net applied as backbone for the segmentation of the target domain images.

Table 1 shows the evaluation results of the proposed GAA Net against the state-of-the-art methods. Compared to utilizing U-Net to segment the target domain images directly, the Dice and the IoU of domain adaptation methods have great improvements in the retina and choroid segmentation. The GAA Net outperforms all competing methods consistently in terms of all evaluation metrics. In particular, we calculate the MAE of all domain adaptation methods after post-processing, our method achieves a much lower MAE in both retina and choroid segmentation, which indicates that it is an accurate and reliable method for measuring the retinal layers and choroid thickness. Fig. 3 shows some visual results for retina and choroid segmentation. Moreover, we also set CE-Net and CS-Net as the segmentation backbone to verify the versatility of our method. The results in Table 1 show that compared to the pre-trained CE-Net and CS-Net, our method leads to a roughly 18% improvement in terms of Dice score for both segmentation backbones.

Parameter entanglement We also evaluate the effect of parameter entanglement (i.e. SDSN and TDSN share the same encoder). In order to reasonably verify the performance degradation caused by parameter entanglement in the source domain, the

Table 1. Layer segmentation performances over **target domain** image by different domain adaptation methods with different segmentation networks.

Method	Retinal layer			Choroidal layer		
	IoU ↑	Dice ↑	MAE ↓	IoU ↑	Dice ↑	MAE ↓
U-Net [9]	64.09%	77.96%	—	71.32%	83.26%	—
Adaptseg Net (U) [19]	82.22%	90.15%	2.918	77.93%	87.40%	8.928
PAAA Net (U) [20]	83.56%	91.00%	2.141	70.21%	82.34%	10.689
UDAB Net (U) [17]	82.55%	90.40%	3.334	74.12%	85.04%	10.614
GAA Net(U)	**92.41%**	**96.05%**	**1.099**	**85.01%**	**91.84%**	**5.374**
CE-Net [3]	65.52%	79.17%	—	57.09%	72.69%	—
GAA Net(CE)	**93.74%**	**96.76%**	**0.939**	**87.01%**	**93.06%**	**4.124**
CS-Net [10]	62.06%	76.53%	—	76.51%	86.01%	—
GAA Net(CS)	**90.12%**	**94.79%**	**1.026**	**83.37%**	**90.84%**	**5.610**

* Adaptseg Net (U) denotes the Adaptseg Net uses U-Net as the segmentation backbone, and so on.

Table 2. Performance degradation between U-Net and domain adaptation methods over **source domain** image.

Method	Retinal layer			Choroidal layer		
	E_{IoU} ↓	E_{Dice} ↓	E_{MAE} ↓	E_{IoU} ↓	E_{Dice} ↓	E_{MAE} ↓
U-Net [9]	—	—	—	—	—	—
Adaptseg Net [19]	4.37%	2.44%	1.285	0.55%	0.32%	1.699
PAAA Net [20]	5.06%	2.80%	1.145	0.96%	0.56%	2.474
UDAB Net [17]	6.17%	3.44%	1.145	2.93%	1.69%	2.592
GAA Net	**0.68%**	**0.70%**	**0.70**	**0.10%**	**0.05%**	**0.178**

SDSN and TDSN of the five competing methods listed in Table 2 all adopt U-Net with the same setting. The parameters of the SDSN and TDSN are shared in AdaptSeg Net, PAAA Net and UDAB Net but not in the GAA Net that we propose. We take the performance mertics of U-Net as the baseline, and computed the error of IoU (E_{IoU}), Dice (E_{Dice}) and MAE (E_{MAE}) between U-Net and the other methods. It can be seen from Table 2 that the metrics of retinal layer segmentation of all methods whose parameters are shared, have significant margin when compared to U-Net. By contrast, our GAA Net achieves a comparable performance with the original U-Net. This indicates that the SDSN module in our method is capable of retaining the segmentation performance of the source domain images, while the TDSN module can alleviate the performance degradation caused by the parameter sharing. For the choroidal layer segmentation, similarly, GAA Net performs better than all other competing methods in terms of all metrics. The results demonstrate that the parameter entanglement generally leads to the performance degradation of the SDSN, such that the TDSN cannot gain the best performance when combined with SDSN together.

4 Conclusion

This paper have proposed a guided adversarial adaptation (GAA) framework for the segmentation of retinal and choroidal layers in OCT images acquired from different devices. By using a dual-encoder structure, the source domain encoder guides the learning of the target domain encoder. This helps to avoid the degradation of source domain segmentation caused by parameter entanglement. In addition, through an adversarial scheme, the target domain segmentations are also enhacned with good performance as the source domain segmentations. In the future work, we will focus on applying the GAA framework to the diagnosis of various ophthalmic diseases.

References

1. Liu, X., Shen, M., Yuan, Y., Huang, S., Zhu, D., Ma, Q., Ye, X., Lu, F.: Macular thickness profiles of intraretinal layers in myopia evaluated by ultrahigh-resolution optical coherence tomography. American journal of ophthalmology **160**(1) (2015) 53–61
2. Cheng, J., Tao, D., Quan, Y., Wong, D.W.K., Cheung, G.C.M., Akiba, M., Liu, J.: Speckle reduction in 3d optical coherence tomography of retina by a-scan reconstruction. IEEE Transactions on Medical Imaging **35**(10) (2016) 2270–2279
3. Gu, Z., Cheng, J., Fu, H., Zhou, K., Hao, H., Zhao, Y., Zhang, T., Gao, S., Liu, J.: Ce-net: Context encoder network for 2d medical image segmentation. IEEE Transactions on Medical Imaging **38**(10) (2019) 2281–2292
4. Bowd, C., Weinreb, R.N., Williams, J.M., Zangwill, L.M.: The retinal nerve fiber layer thickness in ocular hypertensive, normal, and glaucomatous eyes with optical coherence tomography. Archives of ophthalmology **118**(1) (2000) 22–26
5. Charng, J., Simcoe, M., Sanfilippo, P.G., Allingham, R.R., Hewitt, A.W., Hammond, C.J., Mackey, D.A., Yazar, S.: Age-dependent regional retinal nerve fibre changes in six1/six6 polymorphism. Scientific Reports **10**(1) (2020) 1–7
6. Chung, S.E., Kang, S.W., Lee, J.H., Kim, Y.T.: Choroidal thickness in polypoidal choroidal vasculopathy and exudative age-related macular degeneration. Ophthalmology **118**(5) (2011) 840–845
7. Yang, L., Jonas, J.B., Wei, W.: Choroidal vessel diameter in central serous chorioretinopathy. Acta ophthalmologica **91**(5) (2013) e358–e362
8. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE 33rd Conference on Computer Vision and Pattern Recognition. (2015) 3431–3440
9. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Proceedings of the 18th International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer (2015) 234–241
10. Mou, L., Zhao, Y., Chen, L., Cheng, J., Gu, Z., Hao, H., Qi, H., Zheng, Y., Frangi, A., Liu, J.: Cs-net: channel and spatial attention network for curvilinear structure segmentation. In: Proceedings of the 22nd International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer (2019) 721–730
11. Mou, L., Zhao, Y., Fu, H., Liu, Y., Cheng, J., Zheng, Y., Su, P., Yang, J., Chen, L., Frangi, A.F., et al.: Cs2-net: Deep learning segmentation of curvilinear structures in medical imaging. Medical Image Analysis **67** (2021) 101874

12. Kouw, W.M., Loog, M.: A review of domain adaptation without target labels. IEEE Transactions on Pattern Analysis and Machine Intelligence **43**(3) (2021) 766–785
13. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE 35th Conference on Computer Vision and Pattern Recognition. (2017) 1125–1134
14. Yi, X., Walia, E., Babyn, P.: Generative adversarial network in medical imaging: A review. Medical Image Analysis **58** (2019) 101552
15. Ma, X., Niu, Y., Gu, L., Wang, Y., Zhao, Y., Bailey, J., Lu, F.: Understanding adversarial attacks on deep learning based medical image analysis systems. Pattern Recognition **110** (2021) 107332
16. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: Proceedings of the IEEE 35th Conference on Computer Vision and Pattern Recognition. (2017) 7167–7176
17. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: Proceedings of the 32nd International Conference on Machine Learning. (2015) 1180–1189
18. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE 34th Conference on Computer Vision and Pattern Recognition. (2016) 770–778
19. Tsai, Y.H., Hung, W.C., Schulter, S., Sohn, K., Yang, M.H., Chandraker, M.: Learning to adapt structured output space for semantic segmentation. In: Proceedings of the IEEE 36th Conference on Computer Vision and Pattern Recognition. (2018) 7472–7481
20. Chai, Z., Zhou, K., Yang, J., Ma, Y., Chen, Z., Gao, S., Liu, J.: Perceptual-assisted adversarial adaptation for choroid segmentation in optical coherence tomography. In: Proceedings of the 2020 IEEE 17th International Symposium on Biomedical Imaging. (2020) 1966–1970