Shape-Aware Weakly/Semi-Supervised Optic Disc and Cup Segmentation with Regional/Marginal Consistency

Yanda Meng[1], Xu Chen[2], Hongrun Zhang[1], Yitian Zhao[3], Dongxu Gao[4], Barbra Hamill[5], Godhuli Patri[6], Tunde Peto[5], Savita Madhusudhan[6], and Yalin Zheng[1(✉)]

[1] Department of Eye and Vision Science, University of Liverpool, Liverpool, United Kingdom
yalin.zheng@liverpool.ac.uk
[2] Department of Medicine, University of Cambridge, Cambridge, United Kingdom
[3] Cixi Institute of Biomedical Engineering, Ningbo Institute of Industrial Technology, Chinese Academy of Sciences, Ningbo, China
[4] School of Computing, University of Portsmouth, Portsmouth, United Kingdom.
[5] St Paul's Eye Unit, Liverpool University Hospitals NHS Foundation Trust, Liverpool, United Kingdom
[6] School of Medicine, Dentistry and Biomedical Sciences, Queen's University Belfast, Belfast, United Kingdom

Abstract. Glaucoma is a chronic eye disease that permanently impairs vision. Vertical cup to disc ratio (*vCDR*) is essential for glaucoma screening. Thus, accurately segmenting the optic disc (*OD*) and optic cup (*OC*) from colour fundus images is essential. Previous fully-supervised methods achieved accurate segmentation results; then, they calculated the *vCDR* with offline post-processing step. However, a large set of labeled segmentation images are required for the training, which is costly and time-consuming. To solve this, we propose a weakly/semi-supervised framework with the benefits of geometric associations and specific domain knowledge between pixel-wise segmentation probability map (*PM*), geometry-aware modified signed distance function representations (*mSDF*), and local boundary region of interest characteristics (*B-ROI*). Firstly, we propose a dual consistency regularisation based semi-supervised paradigm, where the regional and marginal consistency benefits the proposed model from the objects' inherent region and boundary coherence of a large amount of unlabeled data. Secondly, for the first time, we exploit the domain-specific knowledge between the boundary and region in terms of the perimeter and area of an oval shape of *OD* & *OC*, where a differentiable *vCDR* estimating module is proposed for the end-to-end training. Thus, our model does not need any offline post-process to generate *vCDR*. Furthermore, without requiring any additional laborious annotations, the supervision on *vCDR* can serve as a weakly-supervision for *OD* & *OC* region and boundary segmentation. Experiments on six large-scale datasets demonstrate that our method outperforms state-of-the-art semi-supervised approaches for segmentation of the optic disc and optic cup, and estimation of *vCDR*

for glaucoma assessment in colour fundus images, respectively. The implementation code is made available https://github.com/smallmax00/Shape-Aware-Weakly-Semi-Supervised-Optic-Disc-and-Cup-Segmentation.

Keywords: Optic Disc and Cup Segmentation · Weakly/Semi-supervised Learning.

1 Introduction

The relative size of the *OD* and the *OC* in fundus images can be used to assess glaucomatous damage to the optic nerve head [3]. As a general rule, a greater *vCDR* indicates a higher risk of developing glaucoma and vice versa [23]. Thus, accurate *OD* & *OC* segmentation is critical for glaucoma assessment via *vCDR* measurement. Recently, numerous deep learning-based segmentation models [3,26,23,17,18,21,19] have been proposed, significantly improving the *OD* & *OC* segmentation accuracy. However, they still use a fully supervised paradigm, which requires large annotations which is time-consuming, laborious and costly. Semi-supervised learning frameworks [15,16,10] can obtain high-quality segmentation results by directly learning from a small set of labeled data and a large set of unlabeled data. Numerous of them have been developed to investigate unsupervised consistency regularisation. For instance, they introduced noises at the data-level [25,28] into unlabeled samples and required consistency between model predictions on the original and perturbed data. Furthermore, the feature-level of perturbations are incorporated into multiple output branches [16,10], to ensure the consistency of model predictions across output branches. On the contrary, the consistency regularisation at task-level in semi-supervised learning has received little attention until recently in a variety of computer vision tasks, including crowd counting [20], 3D object detection [13], and 3D medical image segmentation [15]. For example, if we can map the predictions of different tasks into the same predefined space and then evaluate them using the same criterion, the results will undoubtedly be less than optimal, as there are prediction perturbations between tasks. To this end, we learned a dual-task level of geometric consistency via *PM* segmentation and *mSDF* regress. Additionally, we investigated the boundary quality of consistency regularisation at the task-level as the second consistency learning. Specifically, we derived the *B-ROI* masks from the *PM* segmentation and *mSDF* regress branches, respectively. Then the supervised and unsupervised losses are applied to learn more accurate boundary segmentation results with the help of labeled and unlabeled data. Note that a high-quality object boundary is more critical than that of the regional pixel-wise coverage in medical image segmentation tasks [21]. On the other hand, previous weakly supervised learning methods [8,9,11] segmented images using bounding boxes [9], scribbles [11], or image-level tags [8] rather than pixel-by-pixel annotation, which alleviates the burden of annotations. Differently, in this work, we investigated the task-specific domain knowledge of oval shape for the *OD* & *OC* segmentation task. Along with the estimated region and boundary predictions

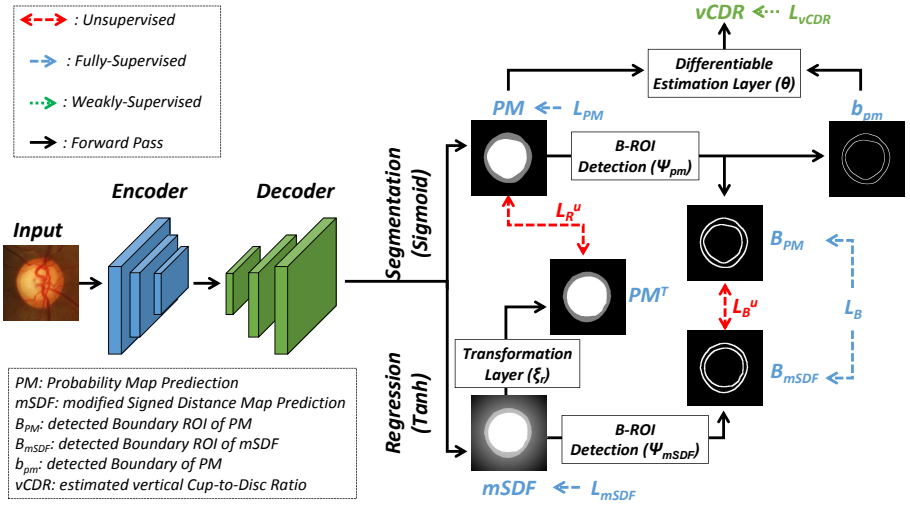

Fig. 1: Overview of the proposed weakly/semi-supervised learning pipeline. *PM* and *mSDF* both have two channels to represent the output of *OC* and *OD*. We overlapped them for better visualisation.

of *OD* & *OC*, we proposed a novel differentiable *vCDR* estimation layer. As a result, our model is capable of estimating the *vCDR* end-to-end on the basis of *OD* & *OC* segmentation. Simultaneously, the information gain from *vCDR* ground truth can weakly-supervise the segmentation process for both region and boundary of *OD* & *OC*.

Despite human graders' instinctive use of both domains, previous methods approach to segment biomedical images frequently overlooked the underlying relationships between the region and boundary characteristics. This article demonstrates how to rationally leverage geometric associations between *OD* & *OC* in terms of region and boundary on semi-supervised consistency learning and differentiable weakly-supervised *vCDR* estimation.

2 Methods

Fig. 1 depicts an overview of the proposed learning pipeline, which consists of two tasks of *PM* segmentation and *mSDF* regression. The geometric associations of two tasks in terms of the region and boundary are exploited in the proposed weakly/semi-supervised learning manner, respectively. The details are elaborated as follows.

2.1 Modified Signed Distance Function (*mSDF*)

Previous works [15,27] adopted *SDF* to represent the target mask in segmentation tasks because it enables the network to learn a distance-aware representation

w.r.t the object boundary, emphasising the spatial perception of the input images. Inspired by [6], we propose a modified Signed Distance Function (*mSDF*), which is defined as:

$$mSDF(x) = \begin{cases} 1, & x \in B_{in} \\ 0, & x \in \Delta B \\ \underset{y \in \Delta B}{-inf} ||x - y||_2, & x \in B_{out} \end{cases} \tag{1}$$

where $||x - y||_2$ is the Euclidean distance between pixel x and y. Besides, B_{out}, B_{in} and ΔB, denote the outside, inside and boundary of the object, respectively. According to the above definition, outside each object, *mSDF* takes negative values, proportional to the distance from the boundary, while it is simply 1 inside of the object and 0 on the boundary. In this way, dual tasks can acquire the coherent semantic features, meanwhile the *mSDF* regression task benefits from the distance-aware spatial information supervision.

2.2 Dual Consistency Regularisation of Semi-Supervision

Under semi-supervised conditions, the dual consistency regularisation imposes regional and marginal consistency at the task level. As for region-wise consistency, similar to [15,20,27], we propose a transformation layer to convert the *mSDF* to *PM* in a differentiable way. To be precise, the region-wise transformation layer ξ_r is defined as:

$$\xi_r(z) = 2 * Sigmoid(K \cdot ReLu(z)) - 1, \tag{2}$$

where z denotes the *mSDF* value at pixel x; K is a very large value; *Sigmoid* and *ReLu* are the non-linear activation functions. The larger K value indicates a closer approximation, and it is adopted as 5000 in this work. With Eq. 2, we can obtain the transformed segmentation maps PM^t, for example, $PM^t = \xi_r(mSDF)$. For all of the unlabeled input, we apply a *Dice* loss (L_{R^u}) between *PM* and PM^t to enforce the unsupervised regional consistency regularisation.

Concerning the boundary-wise consistency, we derive the spatial gradient of *PM* and *mSDF*, as the respectively estimated contours. Previous studies [21,2] have proven that such narrow contours with a width of one pixel are challenging to optimise due to the extremely unbalanced foreground and background, resulting in weakened consistency regularisations. Rather than focusing exclusively on the thin contour locations, we consider the *ROI* within a certain distance (boundary width) of the corresponding estimated contours. A simple yet efficient *B-ROI* detection layer (ψ) is proposed for *PM* and *mSDF*. For example, ψ_{PM} and ψ_{mSDF} are defined as :

$$\psi_{PM} = PM + Maxpooling2D(-PM), \tag{3}$$

$$\psi_{mSDF} = \xi_r(mSDF) + Maxpooling2D\big(-\xi_r(mSDF)\big), \tag{4}$$

It is worth noting that the output width of ψ can be determined by varying the kernel size, stride, and padding value of the Maxpooling2D operation. We empirically set the output boundary width of ψ_{PM} and ψ_{mSDF} to 4 pixels in this work. After ψ_{PM} and ψ_{mSDF}, we refer to such *B-ROI* of PM and $mSDF$ as B_{pm} and B_{mSDF}, respectively. Ideally, B_{PM} and B_{mSDF} should be close enough to one another. Thus, a *Dice* loss (L_{B^u}) between B_{PM} and B_{mSDD} is applied to enforce the unsupervised marginal consistency regularisation of unlabeled data. Meanwhile, we apply a *Dice* loss (L_B) on both B_{PM} and B_{mSDF} to supervise the dual boundary predictions of labeled data.

2.3 Differentiable *vCDR* estimation of Weakly Supervision

Because the shape of *OD* & *OC* are oval-like [23], previous methods adopt to offline post-process the segmentation predictions with ellipse fitting to improve the segmentation accuracy [3], or to calculate the *vCDR* using the approximated diameters of the *OD* & *OC* in long axis [18,17,21]. However, they overlooked the underlying supervision value of it in *OD* & *OC* segmentation task. To address this issue, we take advantage of the specific domain knowledge between the boundary and region in terms of the perimeter and area of an oval-like shape to approximate the *vCDR* in a differentiable way. To be precise, the *vCDR* is defined as the ratio of dividing the measured diameters of the cup by disc in the long axis. While, such ratio can also be estimated given the size of perimeter and the area of *OD* and *OC*. According to the *Euler's Method* [12], the area (A_o) and perimeter (P_o) of the oval shape are defined as:

$$A_o = \pi \cdot a \cdot b, \tag{5}$$

$$P_o = \pi \cdot \sqrt{2(a^2 + b^2)}. \tag{6}$$

where a and b denote the semi-axis of the long and short axis of oval shape, respectively. We approximate A_o with the summed pixel value of PM, which can be regarded as the area of oval shape in pixel level. Furthermore, we derive the spatial gradient of PM via the *B-ROI* detection layer (ψ_{PM}), to detect the boundary (b_{pm}) with width = 1. Then the summed pixel values of b_{pm} is approximately regarded as P_o. With Eq. 5 and Eq. 6, we can approximate a with A_o and P_o, such as:

$$a = \sqrt{\frac{(P_o)^2 + \sqrt{(4\pi A_o + (P_o)^2) \cdot |(4\pi A_o - (P_o)^2))|}}{4\pi^2}}), \tag{7}$$

where $|\cdot|$ is used to prevent sqrt from returning a negative value during the initial learning period. Given Eq. 7, we can calculate the *OD* long semi-axis (a^{OD}) and the *OC* long semi-axis (a^{OC}) with the respective P_o and A_o. Then, the *vCDR* estimation layer θ can be defined as:

$$\theta(vCDR) = \frac{a^{OC} + e^{-6}}{a^{OD} + e^{-6}}, \tag{8}$$

where, e^{-6} is added to avoid dividing by zero errors. Given the prediction of *vCDR*, we apply a *MSE* loss (L_{vCDR}) between the prediction and ground truth to fully-supervise the *vCDR* estimation and weakly-supervise the *OD* & *OC* segmentation.

3 Experiments

3.1 Datasets and Implementation Details

SEG dataset: following the previous methods [21,19], we pooled 2,068 images from five public available datasets (Refuge [23], Drishti-GS [24], ORIGA [29], RIGA [1], RIM-ONE [4]). These five datasets provide the fundus images and the ground truth masks, then we generate the corresponding ground truth of *mSDF*, B_{PM}, B_{mSDF} and *vCDR* with Eq. 1, 3, 4 and 8. Following the previous methods [21,19], 613 fundus images were randomly selected as the test dataset, leaving the other 1,315 images for training and 140 images for validation.
UKBB dataset: The UK Biobank [1] is a large-scale biomedical database and research resource, that contains detailed health information on half a million participants from the United Kingdom. Participants were scanned using the TOPCON 3D OCT 1000 Mk2 camera (Topcon Inc, Japan). There are 117,832 fundus images with *vCDR* scalars are available, of which 38,421 are randomly selected as the weakly/semi-supervised training dataset, and the rest 79,411 are used as test datasets.
Implementation: We cropped the image of 256 $\times$ 256 pixels with the same way of [17,21,19], then randomly rotated and flipped the training dataset with a probability of 0.5. The rotation ranges from -20 to 20 degree. The stochastic gradient descent with a momentum of 0.9 is used to optimise the overall parameters. We trained the model around 10,000 iterations for all the experiments, with a learning rate of 1e-2 and a step decay rate of 0.999 every 100 iterations. The batch size was set as 56, consisting of 28 labeled and 28 unlabeled images. A backbone network [5] is used for ours and all the compared methods. All the training processes were performed on a server with four *GEFORCE RTX 3090 24GiB GPUs*, and all the test experiments were conducted on a workstation with *Intel(R) Xeon(R) W-2104 CPU* and *Geforce RTX 2080Ti GPU* with 11GB memory. We use the output of the *PM* as the segmentation result, and a fixed threshold 0.5 is employed to get a binary mask. For a fair comparison, we do not use any post-processing or ensemble methods. Given the previously discussed loss function terms, we defined the overall loss function as:

$$L_{overall} = L_{PM} + L_{mSDF} + L_B + \lambda * (L_{R^u} + L_{B^u} + L_{vCDR}) \tag{9}$$

where L_{PM} is *Dice* loss for supervised segmentation, L_{mSDF} is *MSE* loss for supervised regression. λ is adopted from [7] as the time-dependent Gaussian

[1] https://www.ukbiobank.ac.uk/

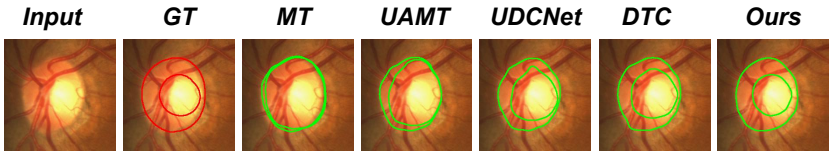

Fig. 2: Qualitative results of *OD* & *OC* segmentation in *SEG* test dataset. We compare our model with *MT* [25], *UAMT* [28], *UDCNet* [10] and *DTC* [15].

Methods	SEG (OC)		SEG (OD)		*SEG (vCDR)*		*UKBB (vCDR)*	
	Dice (%)↑	*BIoU(%)*↑	*Dice (%)*↑	*BIoU(%)*↑	*MAE* ↓	*Corr*↑	*MAE* ↓	*Corr*↑
MT [25]	84.1 (81.8, 85.7)	78.2 (77.0, 79.6)	94.3 (94.0, 94.7)	86.5 (85.0, 87.3)	0.091 (0.080, 0.099)	0.683 (0.641, 0.701)	0.145 (0.139, 0.150)	0.307 (0.276, 0.340)
UAMT [28]	85.3 (82.8, 86.9)	80.2 (79.0, 81.7)	95.2 (94.7, 95.6)	86.4 (85.1, 87.7)	0.075 (0.063, 0.081)	0.692 (0.642, 0.723)	0.134 (0.127, 0.139)	0.339 (0.301, 0.361)
URPC [16]	86.1 (83.1, 87.2)	81.2 (79.6, 82.0)	96.0 (95.4, 96.3)	87.3 (85.0, 87.9)	0.067 (0.059, 0.073)	0.701 (0.659, 0.742)	0.126 (0.121, 0.135)	0.361 (0.337, 0.382)
DTC [15]	86.1 (83.0, 87.4)	81.1 (79.5, 82.8)	96.1 (95.3, 96.4)	87.0 (85.2, 87.8)	0.065 (0.060, 0.072)	0.703 (0.661, 0.739)	0.126 (0.120, 0.137)	0.364 (0.339, 0.389)
UDCNet [10]	86.2 (83.3, 87.1)	81.4 (79.6, 83.0)	96.2 (95.7, 96.5)	87.1 (85.6, 87.9)	0.067 (0.059, 0.071)	0.714 (0.663, 0.742)	0.127 (0.119, 0.135)	0.389 (0.365, 0.412)
Ours (Semi)	**87.1** (86.4, 87.8)	**83.4** (81.0, 85.5)	**97.2** (97.1, 97.3)	**89.3** (88.2, 89.9)	**0.052** (0.049, 0.056)	**0.817** (0.777, 0.852)	**0.102** (0.099, 0.104)	**0.453** (0.439, 0.477)

Table 1: Quantitative segmentation results of *OD* & *OC* and glaucoma assessment on *SEG* and *UKBB* test datasets. The performance is reported as *Dice* (%), *BIoU* (%), *Corr* and *MAE*. 95% confidence intervals are presented in the brackets, respectively. The implementation of the compared semi-supervised state-of-the-art works is mainly based on an open-source codebase [14].

ramp-up weighting coefficient to account for the trade-off between the supervised, unsupervised, and weakly-supervised losses. This avoids the network getting stuck in a degenerate solution during the initial training period. Because no meaningful prediction of the unlabeled data, as well as *vCDR*, are obtained.

4 Results

In this section, we show qualitative (Fig. 2) and quantitative results (Tab. 1) of the *OD* & *OC* segmentation and glaucoma assessment tasks. More qualitative results are shown in the Supplementary. Dice similarity score (*Dice*) and boundary intersection-over-union (*BIoU*) [2] are used as the segmentation accuracy metrics; Mean Absolute Error (*MAE*) and Pearson's correlation coefficients [22] (*Corr*) are used as the *vCDR* estimation metrics. The best result in each category is highlighted in bold. 95% confidence intervals were generated by using 2000 sample bootstrapping.

Optic Disc & Cup Segmentation In Tab. 1, we present the results that are trained with 5 % of *SEG* training dataset and all of *UKBB* training dataset. *Ours (Semi)* obtains an average 87.1 % and 97.2 % *Dice* on *OC* and *OD* segmentation, respectively, outperforms data-level consistency regularisation based

Methods	SEG (OC)		SEG (OD)		UKBB (vCDR)	
	Dice (%)↑	*BIoU (%)*↑	*Dice (%)*↑	*BIoU (%)*↑	*MAE* ↓	*Corr* ↑
w/o L_{R^u}	85.7	80.6	95.9	86.6	0.151	0.319
w/o L_{B^u}	86.1	81.1	96.1	87.0	0.133	0.338
w/ Both	86.5	82.1	96.5	88.0	0.127	0.341
w/ L_{vCDR}	86.7	82.6	96.7	88.5	0.112	0.410
w/ L_{B^u} +L_{vCDR}	86.9	82.9	96.7	88.6	0.109	0.427
w/ L_{R^u} +L_{vCDR}	86.9	82.8	96.6	88.7	0.105	0.439
Ours (Label-only)	80.3	70.5	91.2	75.5	0.631	0.112
Ours (Semi)	**87.1**	**83.4**	**97.2**	**89.3**	**0.102**	**0.453**

Table 2: Ablation study on weakly/semi-supervision components. The performance is reported as *Dice* (%), *BIoU* (%), *MAE*, and *Corr*.

methods *MT* [25], *UAMT* [28] by 3.3 % and 2.1 %, outperforms feature-level regularisation based methods *URPC* [16],*UDCNet* [10] by 1.2 % and 1.0 %.

Clinical Evaluation Along with assessing computer vision evaluation metrics, we evaluated our method's performance via *vCDR* in glaucoma assessment. Tab. 1 illustrates the *vCDR* evaluation results on *SEG* and *UKBB* test dataset respectively. The *UKBB (vCDR)* has 79,411 test images, which is significantly larger than the *SEG (vCDR)* (619 images). Therefore, the performance on *UKBB (vCDR)* may more accurately reflect the situation in real-world. Specifically, *Ours (semi)* achieved the best performance of 0.102 *MAE*, which outperforms *UDCNet* [10] by 19.7 %. Please note that, the segmentation-unlabeled 38,421 images of *UKBB* training dataset also serve as the fully supervision for *vCDR* estimation.

Ablation Study. We conducted extensive ablation studies, and all the results demonstrate our model's effectiveness. To illustrate, the ablation results for weakly/semi-supervision components are shown in Tab. 2. More ablation experiment results (e.g., Data Utilisation Efficiency) are given in the Supplementary. Specifically, we conduct experiments to evaluate the effectiveness of the proposed dual consistency regularisation in a semi-supervised manner and the propose differentiable *vCDR* estimation module in a weakly-supervised manner. We represent our model that is trained with only 5 % *SEG* training data as *Ours (Label-only)*. Firstly, we retain the same model structure and eliminate the *vCDR* estimation loss to focus on the dual consistency regularisation losses (*w/ Both*). Following that, we remove the region-wise unsupervised loss (*w/o* L_{R^u}), boundary-wise unsupervised loss (*w/o* L_{B^u}) respectively. Secondly, we remove both of the consistency losses and only apply the weakly-supervised *vCDR* estimation loss (*w/* L_{vCDR}). Then we add the other two unsupervised consistency losses individually (*w/* L_{B^u} +L_{vCDR} and *w/* L_{R^u} +L_{vCDR}) to see if the performance are boosted. Tab. 2 demonstrates that the proposed unsupervised dual consistency losses and weakly supervised loss can improve the model by 6.7 % and 6.9 % *Dice* respectively for *OD* & *OC* segmentation. Particularly, the boundary-wide unsupervised loss can increase the model by 15.1 % *BIoU*,

which leads to a better boundary segmentation quality. The weakly supervised loss can bring a large improvement of 82.2 % *MAE* of *vCDR* estimation, which is the ultimate goal for *OD* & *OC* segmentation task *w.r.t* clinic application.

5 Conclusion

We propose a novel weakly/semi-supervised segmentation framework. The geometric associations and specific domain knowledge between the modified signed distance function representations, object boundary characteristics, and pixel-wise probability map features are exploited in the proposed semi-supervised consistency regularisations, and weakly-supervised guidance. Our experiments have demonstrated that the proposed model can effectively leverage semantic region features and spatial boundary features for segmentation of optic disc & optic cup and *vCDR* estimation of glaucoma assessment from retinal images.

References

1. Almazroa, A., Alodhayb, S., Osman, E., Ramadan, E., Hummadi, M., Dlaim, M., Alkatee, M., Raahemifar, K., Lakshminarayanan, V.: Retinal fundus images for glaucoma analysis: the RIGA dataset. In: Medical Imaging 2018: Imaging Informatics for Healthcare, Research, and Applications. vol. 10579, p. 105790B. International Society for Optics and Photonics (2018)
2. Cheng, B., Girshick, R., Dollár, P., Berg, A.C., Kirillov, A.: Boundary IoU: Improving object-centric image segmentation evaluation. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2021)
3. Fu, H., Cheng, J., Xu, Y., Wong, D.W.K., Liu, J., Cao, X.: Joint optic disc and cup segmentation based on multi-label deep network and polar transformation. IEEE Transactions on Medical Imaging **37**(7), 1597–1605 (2018)
4. Fumero, F., Alayón, S., Sanchez, J.L., Sigut, J., Gonzalez-Hernandez, M.: RIM-ONE: An open retinal image database for optic nerve evaluation. In: 24th International Symposium on Computer-based Medical Systems (CBMS). pp. 1–6. IEEE (2011)
5. Gao, S., Cheng, M.M., Zhao, K., Zhang, X.Y., Yang, M.H., Torr, P.H.: Res2net: A new multi-scale backbone architecture. IEEE Transactions on Pattern Analysis and Machine Intelligence (2019)
6. Jiang, W., Kolotouros, N., Pavlakos, G., Zhou, X., Daniilidis, K.: Coherent reconstruction of multiple humans from a single image. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5579–5588 (2020)
7. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. In: International Conference on Learning Representations (ICLR) (2017)
8. Lee, J., Kim, E., Lee, S., Lee, J., Yoon, S.: Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5267–5276 (2019)
9. Lee, J., Yi, J., Shin, C., Yoon, S.: Bbam: Bounding box attribution map for weakly supervised semantic and instance segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2643–2652 (2021)

10. Li, Y., Luo, L., Lin, H., Chen, H., Heng, P.A.: Dual-consistency semi-supervised learning with uncertainty quantification for covid-19 lesion segmentation from ct images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 199–209. Springer (2021)
11. Liu, X., Yuan, Q., Gao, Y., He, K., Wang, S., Tang, X., Tang, J., Shen, D.: Weakly supervised segmentation of covid19 infection with scribble annotation on ct images. Pattern recognition **122**, 108341 (2022)
12. Lockwood, E.: Length of ellipse. The Mathematical Gazette **16**(220), 269–270 (1932)
13. Lu, Y., Pirk, S., Dlabal, J., Brohan, A., Pasad, A., Chen, Z., Casser, V., Angelova, A., Gordon, A.: Taskology: Utilizing task relations at scale. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8700–8709 (2021)
14. Luo, X.: Ssl4mis. https://github.com/hilab-git/ssl4mis (2020)
15. Luo, X., Chen, J., Song, T., Wang, G.: Semi-supervised medical image segmentation through dual-task consistency. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 8801–8809 (2021)
16. Luo, X., Liao, W., Chen, J., Song, T., Chen, Y., Zhang, S., Chen, N., Wang, G., Zhang, S.: Efficient semi-supervised gross target volume of nasopharyngeal carcinoma segmentation via uncertainty rectified pyramid consistency. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 318–329. Springer (2021)
17. Meng, Y., Meng, W., Gao, D., Zhao, Y., Yang, X., Huang, X., Zheng, Y.: Regression of instance boundary by aggregated cnn and gcn. In: European Conference on Computer Vision. pp. 190–207. Springer (2020)
18. Meng, Y., Wei, M., Gao, D., Zhao, Y., Yang, X., Huang, X., Zheng, Y.: CNN-GCN aggregation enabled boundary regression for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 352–362. Springer (2020)
19. Meng, Y., Zhang, H., Gao, D., Zhao, Y., Yang, X., Qian, X., Huang, X., Zheng, Y.: Bi-gcn: Boundary-aware input-dependent graph convolution network for biomedical image segmentation. In: 32nd British Machine Vision Conference: BMVC 2021. British Machine Vision Association (2021)
20. Meng, Y., Zhang, H., Zhao, Y., Yang, X., Qian, X., Huang, X., Zheng, Y.: Spatial uncertainty-aware semi-supervised crowd counting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15549–15559 (2021)
21. Meng, Y., Zhang, H., Zhao, Y., Yang, X., Qiao, Y., MacCormick, I.J., Huang, X., Zheng, Y.: Graph-based region and boundary aggregation for biomedical image segmentation. IEEE Transactions on Medical Imaging (2021)
22. Mukaka, M.M.: A guide to appropriate use of correlation coefficient in medical research. Malawi medical journal **24**(3), 69–71 (2012)
23. Orlando, J.I., Fu, H., Breda, J.B., van Keer, K., Bathula, D.R., Diaz-Pinto, A., Fang, R., Heng, P.A., Kim, J., Lee, J., et al.: REFUGE challenge: A unified framework for evaluating automated methods for glaucoma assessment from fundus photographs. Medical Image Analysis **59**, 101570 (2020)
24. Sivaswamy, J., Krishnadas, S., Joshi, G.D., Jain, M., Tabish, A.U.S.: Drishti-GS: Retinal image dataset for optic nerve head (ONH) segmentation. In: 2014 IEEE 11th International Symposium on Biomedical Imaging (ISBI). pp. 53–56. IEEE (2014)

25. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: Advances in Neural Information Processing Systems. pp. 1195–1204 (2017)
26. Wu, J., Wang, K., Shang, Z., Xu, J., Ding, D., Li, X., Yang, G.: Oval shape constraint based optic disc and cup segmentation in fundus photographs. In: BMVC. p. 265 (2019)
27. Xue, Y., Tang, H., Qiao, Z., Gong, G., Yin, Y., Qian, Z., Huang, C., Fan, W., Huang, X.: Shape-aware organ segmentation by predicting signed distance maps. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 12565–12572 (2020)
28. Yu, L., Wang, S., Li, X., Fu, C.W., Heng, P.A.: Uncertainty-aware self-ensembling model for semi-supervised 3d left atrium segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 605–613. Springer (2019)
29. Zhang, Z., Yin, F.S., Liu, J., Wong, W.K., Tan, N.M., Lee, B.H., Cheng, J., Wong, T.Y.: ORIGA-light: An online retinal fundus image database for glaucoma analysis and research. In: 2010 Annual International Conference of the IEEE Engineering in Medicine and Biology. pp. 3065–3068. IEEE (2010)