Self-Loop Uncertainty: A Novel Pseudo-Label for Semi-supervised Medical Image Segmentation

Yuexiang Li[(✉)], Jiawei Chen, Xinpeng Xie, Kai Ma, and Yefeng Zheng

Tencent Jarvis Lab, Shenzhen, China
vicyxli@tencent.com

Abstract. Witnessing the success of deep learning neural networks in natural image processing, an increasing number of studies have been proposed to develop deep-learning-based frameworks for medical image segmentation. However, since the pixel-wise annotation of medical images is laborious and expensive, the amount of annotated data is usually deficient to well-train a neural network. In this paper, we propose a semi-supervised approach to train neural networks with limited labeled data and a large quantity of unlabeled images for medical image segmentation. A novel pseudo-label (namely self-loop uncertainty), generated by recurrently optimizing the neural network with a self-supervised task, is adopted as the ground-truth for the unlabeled images to augment the training set and boost the segmentation accuracy. The proposed self-loop uncertainty can be seen as an approximation of the uncertainty estimation yielded by ensembling multiple models with a significant reduction of inference time. Experimental results on two publicly available datasets demonstrate the effectiveness of our semi-supervised approach.

Keywords: Semi-supervised learning · Pseudo-label · Jigsaw puzzles

1 Introduction

Deep neural networks often require large quantity of labeled images to achieve satisfactory performance. However, since annotating medical images requires experienced physicians to spend hours or days to investigate, which is laborious and expensive, the labeled medical images are often very deficient, especially for the tasks requiring pixel-wise annotations (e.g., segmentation). To tackle this problem, many researches [1,2,17,19] have been proposed to improve the segmentation performance of deep neural networks through exploiting the information from unlabeled data. Using pseudo-labels of unlabeled data (generated automatically by a segmentation algorithm via uncertainty estimation) is one of the potential solutions, which has been extensively studied. The most popular approaches are: 1) softmax probability map [1], 2) Monte Carlo (MC) dropout [17,19], and 3) uncertainty estimation via network ensemble [8]. Specifically,

Bai et al. [1] proposed a semi-supervised approach for the cardiac magnetic resonance volume segmentation. The proposed approach first used a limited number of labeled data to train the neural network and then utilized the softmax probability maps predicted by the neural network as the pseudo-label for the unlabeled volumes to augment the training set. In a more recent study, Sedai et al. [17] proposed an uncertainty guided semi-supervised learning framework for the segmentation of retinal layers in optical coherence tomography images. The pseudo-label for semi-supervised learning was generated using the Monte Carlo (MC) dropout [4], which can be viewed as an approximation of Bayesian uncertainty. Uncertainty estimation via model ensemble [8] is another form of approximation of Bayesian uncertainty, which separately trained K networks and combined the softmax probability map of each network k by averaging as the ensemble uncertainty (i.e., $\frac{1}{K} \sum_{k=1}^{K} p_k$, where p is the probability map).

Due to the variety of existing uncertainty estimation methods, Jungo et al. [7] conducted experiments to evaluate the reliability and limitation of existing approaches and concluded several observations. Two of them cause our interests: 1) the widely-used MC-dropout-based approaches are heavily dependent on the influence of dropout on the segmentation performance; 2) the computational-expensive ensemble method yields the most reliable results and is typically a good choice if the resources allow. To this end, an efficient way to yield the reliable ensemble uncertainty is worthwhile to investigate.

In this paper, we propose a novel pseudo-label, namely self-loop uncertainty, for the semi-supervised medical image segmentation. The proposed self-loop uncertainty is generated by recurrently optimizing the encoder of a fully convolutional network (FCN) with a self-supervised sub-task (e.g., Jigsaw puzzles). The benefits of integrating self-supervised learning into our framework can be summarized in two folds: 1) the self-supervised learning sub-task encourages the neural network to deeply mine the information from raw data and benefits the image segmentation task; 2) the same network at different stages during the self-supervised sub-task optimization can be seen as different models, which leads our self-loop uncertainty to an approximation of ensemble uncertainty with much lower computational cost. We evaluate the proposed semi-supervised learning approach on two medical image segmentation tasks—nuclei segmentation and skin lesion segmentation. Experimental results show that our self-loop uncertainty can significantly improve the segmentation accuracy of the neural network, which outperforms the currently widely-used pseudo-label (e.g., softmax probability map and MC dropout).

2   Method

The proposed semi-supervised segmentation framework is illustrated in Fig. 1. The training set for our semi-supervised framework consists of labeled data D_L and unlabeled data D_U. The proposed semi-supervised framework involves three losses (i.e., $\mathcal{L}_{SEG}$, $\mathcal{L}_{UG}$, and $\mathcal{L}_{SS}$) to supervise the network training with D_L and D_U, respectively. The colored arrows in Fig. 1 represent the information flows of
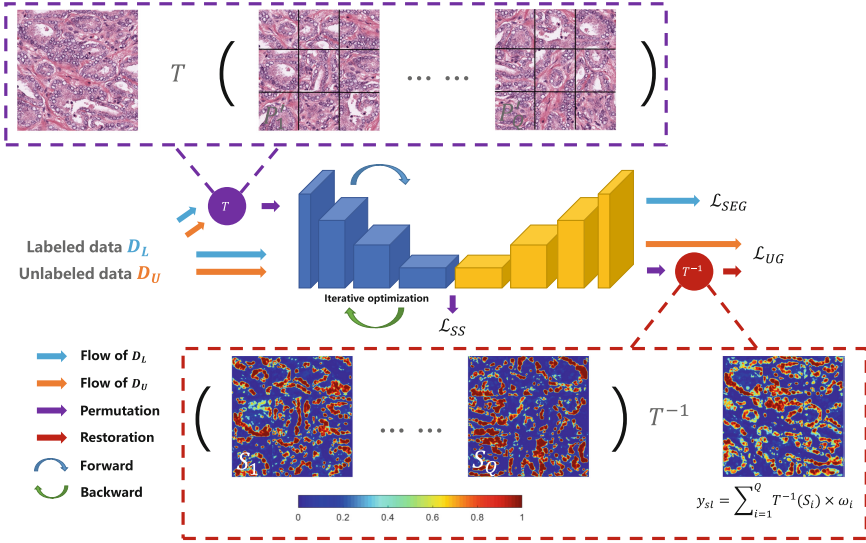
Fig. 1. The pipeline of our semi-supervised segmentation framework. The proposed framework recurrently optimizes the encoder part of FCN by addressing the self-supervised learning task (i.e., supervised by $\mathcal{L}_{SS}$) to generate the pseudo-label for the unlabeled data. There are two losses, i.e., segmentation loss $\mathcal{L}_{SEG}$ and uncertainty-guided loss $\mathcal{L}_{UG}$, adopted in our framework to supervise the segmentation of labeled and unlabeled data. Our framework generates Q permutations $(P_1', \ldots P_Q')$ for an image (either labeled or unlabeled) and yields corresponding Q segmentation predictions $(S_1, \ldots S_Q)$ for the estimation of self-loop uncertainty y_{sl} (as illustrated in Algorithm 1). (Color figure online)

D_U (orange) and D_L (cyan). For a batch containing images from D_L and D_U, we calculate the supervised segmentation loss $\mathcal{L}_{SEG}$ (i.e., binary cross-entropy loss in our experiment) for labeled data with pixel-wise annotation to ensure the FCN has the segmentation capacity, self-supervised loss $\mathcal{L}_{SS}$ for both D_L and D_U to exploit rich information from raw data and generate the self-loop uncertainty, and uncertainty-guided loss $\mathcal{L}_{UG}$ for the unlabeled images to boost the segmentation performance of FCN with unlabeled data.

2.1 Self-supervised Sub-task

As aforementioned, the self-supervised loss $\mathcal{L}_{SS}$ aims to exploit rich information contained in raw data and generate the self-loop uncertainty. Various pretext tasks, such as rotation prediction [5] and colorization [9], can be adopted to achieve this goal. In this study, we use Jigsaw puzzles [14] consisting of translation and rotation transformations as the self-supervised sub-task to recurrently optimize the encoder of an FCN and yield the self-loop uncertainty.

Similar to the standard Jigsaw puzzles, we partition the image into several tiles, e.g., nine tiles for 3×3 Jigsaw puzzles. To formulate the Jigsaw puzzles

---

Algorithm 1. Generation of self-loop uncertainty.

1: **Input:**
2: Network weights: θ_e of the encoder and θ_d of the decoder.
3: Unlabeled data: $x \in D_U$.
4: **Function:**
5: $f(x; \theta)$ neural network forward function.
6: $update(.)$ backpropagation to update the neural network weights.
7: $T(.)$ permuted transformation of Jigsaw puzzles.
8: $T^{-1}(.)$ inverse-permuted transformation.
9: $\mathcal{L}_{SS}(p, g)$ calculation of the self-supervised loss with prediction p and self-supervised signal g.
10: **Procedure†:**
11: Q permutations are randomly selected from $\mathbb{P}'$: $\{P'_1, ..., P'_Q \in \mathbb{P}'\}$.
12: **for** $i \in \{1, ..., Q\}$ **do**
13:     $p_i \leftarrow f(T_{P'_i}(x); \theta_e); \ S_i \leftarrow f(T_{P'_i}(x); \{\theta_e, \theta_d\});$
14:     $l_i \leftarrow \mathcal{L}_{SS}(p_i, g_i); \ \theta_e^i \leftarrow update(l_i);$
15:     $\theta_e \leftarrow \theta_e^i.$
16: **end for**
17: $y_{sl} = \sum_{i=1}^{Q} T_{P'_i}^{-1}(S_i) \times norm(\omega_i)$, where $norm(.) = \frac{\omega_i}{\sum_{i=1}^{Q} \omega_i}$ and $\omega_i = 1 - \frac{l_i}{\sum_{i=1}^{Q} l_i}$.
    † S is the segmentation prediction of FCN.
    † l is the calculated self-supervised loss.
18: **Output:** self-loop uncertainty y_{sl} of input x.

---

sub-task, we permute the tiles using the approach proposed by [14]—a small subset $\mathbb{P}'$ of the large permutation pool, i.e., $\mathbb{P} = (P_1, P_2, ..., P_{9!})$ is formed by selecting the K permutations with the largest Hamming distance between each other. In each training iteration, the input image is repeatedly disarranged (Q times in total where $Q \ll K$, $Q = 10$ and $K = 100$ in our experiments) by one randomly selected permutation from $\mathbb{P}'$. Meanwhile, the encoder of FCN is recurrently updated to identify the selected permutation from the K options for each disarranged image, which can be seen a classification task with K categories; therefore, we employ the cross-entropy loss as $\mathcal{L}_{SS}$ to supervise the sub-task.

The Jigsaw puzzles transformation adopted in our approach has two differences, compared to the one in [14]. First, to increase the diversity of permutation, each of the tiles is randomly rotated by an angle $a \in \{0°, 90°, 180°, 270°\}$ besides the translation transformation. Second, to integrate the Jigsaw puzzles task into the end-to-end semi-supervised framework, the input of self-supervised sub-task is required to have the same size as that of the target segmentation task. Hence, instead of using the shared-weight neural network for each tile, the permuted tiles are first assembled to an image of the same size of the original image (i.e., $\{P'_1, ..., P'_Q\}$ shown in Fig. 1) and then fed as input to the neural network for the permutation classification.

2.2 Estimation of Self-loop Uncertainty for Unlabeled Data

The generation procedure of our self-loop uncertainty is presented in Algorithm 1. The self-supervised sub-task is able to recurrently optimize the neural network in an iteration, as the self-supervised signal can be self-driven without manual annotation. The different stages (i.e., $\{\theta_e^i, \theta_d\}, i \in \{1, ..., Q\}$) of self-supervised optimization are seen as different models, which enable the proposed self-loop uncertainty to approximate the ensemble uncertainty with a single neural network. The permutated images go through the FCN and yield a set of segmentation predictions $S_i, i \in \{1, ..., Q\}$. Since the calculated self-supervised loss (l) can explicitly represent the difficulty of puzzled image for neural network to restore, we formulate l as the confidence of corresponding segmentation result S (via $norm(.)$ and ω defined in Algorithm 1) to revise its contribution to the final pseudo-label. Our self-loop uncertainty thereby is the weighted average of the segmentation predictions produced by different stages of self-supervised optimization.

Uncertainty-Guided Loss. The set of segmentation predictions $\{S_1, ..., S_Q\}$ is presented in Fig. 1, where the red color represents the high score of foreground. The weight-averaged self-loop uncertainty y_{sl} can be used as the guidance to maintain the reliable predictions (i.e., high score) as target for the neural network to learn from unlabeled data. To achieve this goal, we adopt the mean squared error (MSE) loss as the uncertainty-guided loss $\mathcal{L}_{UG}$ for the network optimization with unlabeled data and pseudo-labels y_{sl}, which can be defined as:

$$\mathcal{L}_{UG}(S_x, y_{sl}) = \frac{\sum_{H \times W} \mathbb{I}(y_{sl} > th) \| S_x - y_{sl} \|^2}{\sum_{H \times W} \mathbb{I}(y_{sl} > th)} \tag{1}$$

where $\mathbb{I}$ is the indicator function; H and W are the image height and width, respectively; S_x is the segmentation prediction of input image x; and th is the threshold to select the high score target.

2.3 Objective Function

Assuming a batch contains N labeled data ($\{(x_j, y_j)\}_{j=1}^N$) and M unlabeled data $\{x_j\}_{j=N+1}^{N+M}$, where $x_j \in \mathbb{R}^{H \times W \times C}$ is the input image (H, W, and C are the height, width, and channel of the image, respectively) and $y_j \in \{0, 1\}^{H \times W}, j = 1, 2, \ldots, N$ is the ground-truth annotation, the objective function $\mathcal{L}$ for this batch can be formulated as:

$$\mathcal{L} = \sum_{j=1}^{N} \mathcal{L}_{SEG}(x_j, y_j) + \sum_{j=N+1}^{N+M} \mathcal{L}_{UG}(x_j, y_{sl}) + \sum_{j=1}^{N+M} \sum_{i=1}^{Q} \mathcal{L}_{SS}(T_{P_i'}(x_j), g_i). \tag{2}$$

During network optimization, for the unlabeled data, we first fixed the decoder of FCN and recurrently update the encoder with $\mathcal{L}_{SS}$ to generate y_{sl}.

Then, the weight of the whole FCN is optimized by $\mathcal{L}_{UG}$. In other words, an unsynchronized optimization of the encoder and decoder happens when using the unlabeled data. For the labeled data, on the other hand, the network is optimized with the $\mathcal{L}_{SEG}$ and $\mathcal{L}_{SS}$ simultaneously.

3   Experiments

MoNuSeg Dataset [13]**.** The dataset consists of diverse H&E stained tissue images captured from seven different organs (e.g., breast, liver, kidney, prostate, bladder, colon and stomach), which were collected from 18 institutes. The dataset has a public training set and a public test set for network training and evaluation, respectively. The training set contains 30 histopathological images with hand-annotated nuclei, while the test set consists of 14 images. The size of the histopathological images is 1000×1000 pixels.

ISIC Dataset [3]**.** The ISIC dataset is widely-used to assess the segmentation accuracy of skin lesion areas of automatic segmentation algorithms. The dataset contains 2,594 dermoscopic images. The skin lesion area of each image has been manually annotated by the data provider. The image size varies from around 1000×1000 pixels to 4000×3000 pixels. We resize all the images to a uniform size of 512×512 pixels for network training and validation. The dataset is randomly separated to training and test sets according to the ratio of 75:25.

Evaluation Criterion. The F1 score, i.e., the unweighted average classification accuracy of the foreground and background tissues, which is widely-used in the area of nuclei [11,15,20] and skin lesion [10,12,18] segmentation, is adopted as the metric to evaluate the segmentation performance.

Baselines. Three popular uncertainty approaches—softmax probability map [1], Monte Carlo (MC) dropout [17,19], and uncertainty estimation via ensembling networks [8]—are involved as baselines in this study. Similar to [17], we set the dropout rate to 0.2 and forward the image through the neural network for ten times to generate MC dropout uncertainty. The ensemble uncertainty is generated by ensembling ten models trained with different network initializations. Consistent with the baselines, we generate ten permutations for an image to iteratively optimize the neural network and accordingly yield the self-loop uncertainty. The widely used ResUNet-18 [6,16] is used as the backbone for uncertainty estimation. For fair comparison, all the baselines are trained according to the same protocol.

3.1   Evaluation of Pseudo-Label Quality

Compared to skin lesion segmentation, which contains a single target in each image, the annotation of nucleus is more difficult and laborious. Hence, we mainly

use the MoNuSeg dataset to evaluate the quality of pseudo-label yielded by different approaches in this section.[1]

To quantitatively validate the accuracy of different pseudo-labels, we calculate the F1 score between the pseudo-labels and ground-truth and present the results in Table 1. The pseudo-labels are generated with different amounts (i.e., 20% and 50%) of labeled data D_L and the remaining training set is used as unlabeled data D_U. As shown in Table 1, our self-loop uncertainty outperforms all the baselines under different amounts of labeled data, which are +2.27% and +2.88% higher than the runner-up (i.e., MC Dropout) with 20% and 50% labeled data, respectively. The pesudo-labels yielded by uncertainty via ensembling models achieve lower accuracy among the baselines. The underlying reason may be that the MoNuSeg training set only contains 30 histopathological images, which make the amount of labeled data (i.e., 20% and 50%) insufficient to well train the neural network. Therefore, the ensembling of multiple unsatisfactory models cannot improve the accuracy of uncertainty estimation.

Table 1. F1 score (%) between ground-truth and the pseudo-labels generated by different uncertainty approaches with different amounts of labeled data. The superscript of SL is the number of permutations Q generated for self-supervised learning. (MC D.—MC Dropout, SL—Self-loop)

Amount of D_L	Softmax	MC D.	Ensemble	SL^3	SL^6	SL^{10}
20%	67.48	72.42	67.46	73.90	74.68	**75.24**
50%	69.53	73.58	70.01	76.51	76.77	**76.85**

Ablation Study. We conduct an ablation study to investigate the relationship between the number of permutation Q and the quality of pseudo-label. Q is set to 3, 6, 10, respectively, for the generation of self-loop uncertainty. As shown in Table 1, the self-loop uncertainty generated with a larger Q achieves the higher F1 score. However, the improvement of F1 score provided by increasing Q from 6 to 10 becomes marginal (e.g., +0.08% using 50% labeled data), which illustrates that Q may not be the larger the better for practical applications, when taking the computational cost into account.

3.2 Segmentation Performance Evaluation

To validate the effectiveness of pseudo-labels, we evaluate the performance of different semi-supervised frameworks on the test sets of MoNuSeg and ISIC. The semi-supervised approaches are trained with different portions (i.e., 20% and 50%) of labeled data. The evaluation results are listed in Table 2. The performance of fully-supervised approach with 100% labeled data is also assessed as the

---

[1] For visual comparison between pseudo-labels, please refer to *arxiv version*.

upper bound for the semi-supervised approaches. To validate the effectiveness of self-supervised sub-task, the self-loop uncertainty without $\mathcal{L}_{SS}$ is also involved for comparison. We pass the ten permutated images through the FCN without self-supervised optimization and yield the uncertainty by averaging the segmentation predictions. Due to lack of extra information exploited by self-supervised sub-task, the improvements yielded without $\mathcal{L}_{SS}$ significantly decrease.

Nuclei Segmentation. As shown in Table 2, the performance of fully-supervised approach significantly drops from 79.30% to 75.87% and 71.51%, respectively, with the reductions (i.e., −50% and −80%) of manual annotations. The application of pseudo-labels provides a consistent improvement to the segmentation accuracy. Among them, the proposed self-loop uncertainty yields the largest improvements, especially under the condition with 20% annotated data, i.e., +5.6% higher than the fully-supervised approach. Furthermore, we notice that our semi-supervised framework trained with 50% labeled data achieves comparable F1 score (79.10%) to that of 100% fully-supervised approach (79.30%), which demonstrates the potential of our approach for reducing the workload of manual annotations.

Table 2. F1 score (%) yielded by different semi-supervised approaches on the two publicly available datasets.

	MoNuSeg			ISIC		
	20%	50%	100%	20%	50%	100%
Fully-supervised	71.51	75.87	**79.30**	81.49	84.86	**86.58**
Softmax [1]	73.65	76.18	-	82.81	85.11	-
MC Dropout [17]	75.31	77.98	-	83.68	85.74	-
Ensemble [8]	73.33	76.87	-	83.27	86.06	-
Self-loop w/o $\mathcal{L}_{SS}$	74.70	77.78	-	82.70	85.22	-
Self-loop (Ours)	77.11	79.10	-	84.92	86.17	-

Skin Lesion Segmentation. Similar trends of improvement are observed on the ISIC test set. Due to the extra information provided by the unlabeled data, the semi-supervised approaches outperform the fully-supervised one with limited annotated data (20% and 50%). The framework adopted our self-loop uncertainty as pseudo-labels achieves the highest F1 scores, i.e., 84.92% and 86.17% with 20% and 50% labeled data, respectively, and the latter is comparable to that of fully-supervied approach with 100% annotations (i.e., 86.17%). As ISIC has much more training data, compared to MoNuSeg, the ensemble-uncertainty-based framework achieves a comparable F1 score of 86.06% with 50% labeled data. However, it is worthwhile to mention that the generation of ensemble uncertainty requires 10 times of inferences during the test phase, as well as the MC

dropout. Conversely, the proposed self-loop uncertainty can be generated with a single inference, which significantly saves the computational cost.

4   Conclusion

In this paper, we proposed a semi-supervised approach to train neural networks with limited labeled data and a large quantity of unlabeled images for medical image segmentation. A novel pseudo-label (namely self-loop uncertainty), generated by recurrently optimizing the neural network with a self-supervised task, is adopted as the ground-truth for the unlabeled images to augment the training set and boost the segmentation accuracy.

References

1. Bai, W., et al.: Semi-supervised learning for network-based cardiac MR image segmentation. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) MICCAI 2017. LNCS, vol. 10434, pp. 253–260. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66185-8_29
2. Bortsova, G., Dubost, F., Hogeweg, L., Katramados, I., de Bruijne, M.: Semi-supervised medical image segmentation via learning consistency under transformations. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11769, pp. 810–818. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32226-7_90
3. Codella, N., et al.: Skin lesion analysis toward melanoma detection 2018: a challenge hosted by the International Skin Imaging Collaboration (ISIC). arXiv preprint arXiv:1902.03368 (2019)
4. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: representing model uncertainty in deep learning. In: International Conference on Machine Learning, pp. 1050–1059 (2016)
5. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. In: International Conference on Learning Representations (2018)
6. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)
7. Jungo, A., Reyes, M.: Assessing reliability and challenges of uncertainty estimations for medical image segmentation. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11765, pp. 48–56. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_6
8. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. In: Annual Conference on Neural Information Processing Systems, pp. 6402–6413 (2017)

9. Larsson, G., Maire, M., Shakhnarovich, G.: Colorization as a proxy task for visual understanding. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 840–849 (2017)
10. Li, Y., Shen, L.: Skin lesion analysis towards melanoma detection using deep learning network. Sensors **18**(2), 556 (2018)
11. Luna, M., Kwon, M., Park, S.H.: Precise separation of adjacent nuclei using a Siamese neural network. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11764, pp. 577–585. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32239-7_64
12. Nasr-Esfahani, E., et al.: Dense pooling layers in fully convolutional network for skin lesion segmentation. Comput. Med. Imaging Graph. **78**, 101658 (2019)
13. Naylor, P., Lae, M., Reyal, F., Walter, T.: Segmentation of nuclei in histopathology images by deep regression of the distance map. IEEE Trans. Med. Imaging **38**(2), 448–459 (2018)
14. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9910, pp. 69–84. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46466-4_5
15. Oda, H., et al.: BESNet: boundary-enhanced segmentation of cells in histopathological images. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11071, pp. 228–236. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00934-2_26
16. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
17. Sedai, S., et al.: Uncertainty guided semi-supervised segmentation of retinal layers in OCT images. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11764, pp. 282–290. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32239-7_32
18. Tang, Y., Yang, F., Yuan, S., Zhan, C.: A multi-stage framework with context information fusion structure for skin lesion segmentation. In: International Symposium on Biomedical Imaging, pp. 1407–1410 (2019)
19. 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: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11765, pp. 605–613. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_67
20. Zhou, Y., Onder, O.F., Dou, Q., Tsougenis, E., Chen, H., Heng, P.-A.: CIA-Net: robust nuclei instance segmentation with contour-aware information aggregation. In: Chung, A.C.S., Gee, J.C., Yushkevich, P.A., Bao, S. (eds.) IPMI 2019. LNCS, vol. 11492, pp. 682–693. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-20351-1_53