



Efficient Semi-supervised Gross Target Volume of Nasopharyngeal Carcinoma Segmentation via Uncertainty Rectified Pyramid Consistency

Xiangde Luo¹, Wenjun Liao², Jieneng Chen³, Tao Song⁴, Yinan Chen^{4,5}, Shichuan Zhang⁶, Nianyong Chen², Guotai Wang^{1(✉)}, and Shaoting Zhang^{1,4}

¹ School of Mechanical and Electrical Engineering, University of Electronic Science and Technology of China, Chengdu, China

guotai.wang@uestc.edu.cn

² Department of Radiation Oncology, West China Hospital, Sichuan University, Chengdu, China

³ College of Electronics and Information Technology, Tongji University, Shanghai, China

⁴ SenseTime Research, Shanghai, China

⁵ West China Hospital-SenseTime Joint Lab, West China Biomedical Big Data Center, Sichuan University West China Hospital, Chengdu, China

⁶ Department of Radiation Oncology, Sichuan Cancer Hospital and Institute, University of Electronic Science and Technology of China, Chengdu, China

Abstract. Gross Target Volume (GTV) segmentation plays an irreplaceable role in radiotherapy planning for Nasopharyngeal Carcinoma (NPC). Despite that Convolutional Neural Networks (CNN) have achieved good performance for this task, they rely on a large set of labeled images for training, which is expensive and time-consuming to acquire. In this paper, we propose a novel framework with Uncertainty Rectified Pyramid Consistency (URPC) regularization for semi-supervised NPC GTV segmentation. Concretely, we extend a backbone segmentation network to produce pyramid predictions at different scales. The pyramid predictions network (PPNet) is supervised by the ground truth of labeled images and a multi-scale consistency loss for unlabeled images, motivated by the fact that prediction at different scales for the same input should be similar and consistent. However, due to the different resolution of these predictions, encouraging them to be consistent at each pixel directly has low robustness and may lose some fine details. To address this problem, we further design a novel uncertainty rectifying module to enable the framework to gradually learn from meaningful and reliable consensual regions at different scales. Experimental results on a dataset with 258 NPC MR images showed that with only 10% or 20% images labeled, our method largely improved the segmentation performance by

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-87196-3_30) contains supplementary material, which is available to authorized users.

© Springer Nature Switzerland AG 2021

M. de Bruijne et al. (Eds.): MICCAI 2021, LNCS 12902, pp. 318–329, 2021.

https://doi.org/10.1007/978-3-030-87196-3_30

leveraging the unlabeled images, and it also outperformed five state-of-the-art semi-supervised segmentation methods. Moreover, when only 50% labeled images, URPC achieved an average Dice score of 82.74% that was close to fully supervised learning. Code is available at: <https://github.com/HiLab-git/SSL4MIS>.

Keywords: Semi-supervised learning · Uncertainty rectifying · Pyramid consistency · Gross target volume · Nasopharyngeal carcinoma

1 Introduction

Nasopharyngeal Carcinoma (NPC) is one of the most common cancers in southern China, Southeast Asia, the Middle East, and North Africa [5]. The mainstream treatment strategy for NPC is radiotherapy, thus accurate target delineation plays an irreplaceable role for precise and effective radiotherapy. However, manual nasopharyngeal tumor contouring is tedious and laborious, since both the nasopharynx gross tumor volume (GTVnx) and lymph node gross tumor volume (GTVnd) need to be accurately delineated [13]. Recently, with a large amount of labeled data, deep learning has shown the potential for accurate GTV segmentation [13]. However, collecting a large labeled dataset for network training is difficult, as both time and expertise are needed to produce accurate annotation. In contrast, collecting a large set of unlabeled data is easier, which inspired us to develop a semi-supervised approach for NPC GTV segmentation by leveraging unlabeled data. What's more, semi-supervised learning (SSL) can largely reduce the workload of annotators for the development of deep learning models.

Recently, SSL has been widely used for medical image computing to reduce the annotation efforts [9, 12, 16, 19]. Bai et al. [1] developed an iterative framework where in each iteration, pseudo labels for unannotated images are predicted by the network and refined by a Conditional Random Field (CRF), then the new pseudo labels are used to update the network. After that, the perturbation-based methods have achieved increasing attention in semi-supervised learning [2, 4, 12, 19]. These methods add small perturbations to unlabeled samples and enforce the consistency between the model's predictions on the original data and the perturbed data. Meanwhile, the mean teacher-based [21] self-ensembling methods [7, 9, 26, 27] were introduced for semi-supervised medical image segmentation. Following [7], some recent works [9, 26, 27] used uncertainty map to guide the student model to learn more stably. In [11, 17, 29], an adversarial training strategy was used as regularization for SSL, which aims to minimize the adversarial loss to encourage the prediction of unlabeled data is anatomical plausible. Luo et al. [15] proposed a dual-task consistency framework for SSL by representing segmentation as a pixel-wise classification task and a level set regression task simultaneously, the difference between which was minimized during training. Despite their higher performance than learning from available labeled images only, existing methods are limited by high computational cost and complex training strategies in practice. For example, the co-training-based methods need to

train several models at the same time [20], and the uncertainty estimation-based frameworks need multiple forward passes [27]. Self-training-based approaches need to select and refine pseudo labels and update models’ parameters in several rounds [1], which is time-consuming.

In this work, we propose a novel efficient semi-supervised learning framework for the segmentation of GTVnx and GTVnd by further utilizing the unlabeled data. Inspired by deep supervision [10], Our method leverages a network that gives a pyramid (i.e., multi-scale) prediction, and encourages the predictions at multiple scales to be consistent for a given input, which is a simple yet efficient idea for SSL. A standard supervised loss at multiple scales is used for learning from labeled images. For unlabeled images, we encourage the multi-scale predictions to be consistent, which serves as a regularization. Since the ground truth of unlabeled images is unknown, the model may produce some unreliable prediction or noise which may cause the model to collapse and lose details. To overcome these problems, some existing works [3, 27] have leveraged model uncertainty to boost the stability of training and obtain better results. However, they typically estimate the uncertainty of each target prediction with Monte Carlo sampling [8], which needs massive computational costs as it requires multiple forward passes to obtain the uncertainty in each iteration. Differently from these methods, we estimate the uncertainty via the prediction discrepancy among multi-scale predictions, which just needs a single forward pass. With the guidance of the estimated uncertainty, we automatically emphasize the reliable predictions (low uncertainty) and weaken the unreliable ones (high uncertainty) when calculating the multi-scale consistency. Meanwhile, we introduce the uncertainty minimization [30] to reduce the prediction variance during training. Therefore, the proposed framework has high efficiency for semi-supervised segmentation by taking advantage of the unlabeled images. Our method was extensively evaluated on a clinical Nasopharyngeal Carcinoma dataset. Results show our method largely improved the segmentation performance by leveraging the unlabeled images, and it outperformed five state-of-the-art semi-supervised segmentation methods. Moreover, when only half of the training images are labeled, URPC achieves a very close result compared with fully supervised learning (the mean of Dice was 82.74% vs 83.51%).

2 Methods

The proposed URPC for semi-supervised segmentation is illustrated in Fig. 1. We add a pyramid prediction structure at the decoder of a backbone network and refer to it as PPNet. PPNet learns from the labeled data by minimizing a typical supervised segmentation loss directly. In addition, the PPNet is regularized by a multi-scale consistency between the pyramid predictions to deal with unlabeled data. The PPNet naturally leads to uncertainty estimation in a single forward pass by measuring the discrepancy between these predictions, and we propose to use this uncertainty to rectify the pyramid consistency considering the different spatial resolutions in the pyramid. To describe this work precisely, we first define

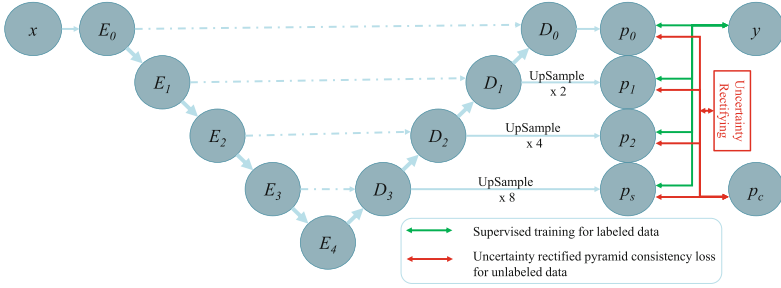


Fig. 1. Overview of the proposed Uncertainty Rectified Pyramid Consistency framework, which consists of a pyramid prediction network (PPNet) and an uncertainty rectifying module. It is based on a backbone of 3D UNet [6], where E s and D s are the blocks in the encoder and decoder of 3D UNet respectively. p_s is the prediction as scale s . The URPC is optimized by minimizing the supervised loss on the labeled data and the pyramid consistency loss on the unlabeled data. In addition, an uncertainty rectifying module is designed to reduce the impact of noise in the pyramid consistency and boost the stability of training.

some mathematical terms. Let D_l , D_u and $f_\phi(x)$ be the labeled set, unlabeled set and the PPNet’s parameters, respectively. Let $D = D_l \cup D_u$ be the whole provided dataset. We denote an unlabeled image as $x_i \in D_u$ and a labeled image pair as $(x_i, y_i) \in D_l$, where y_i is ground truth.

2.1 Multi-scale Prediction Network with Pyramid Consistency

To better exploit the prediction discrepancy of a single model at different scales, we firstly introduce the PPNet for the segmentation task, which can produce predictions with different scales. In this work, we employ 3D UNet [6] as a backbone and modify it to produce pyramid predictions by adding a prediction layer after each upsampling block in the decoder, where the prediction layer is implemented by $1 \times 1 \times 1$ convolution followed by a softmax layer. To introduce more perturbations in the network, a dropout layer and a feature-level noise addition layer are inserted before each of these prediction layers. For an input image x_i , PPNet $f_\phi(x)$ produces a set of multi-scale predictions $[p'_0, p'_1, \dots, p'_s, \dots, p'_{S-1}]$, where the p'_s is the prediction at scale s , and a smaller s corresponds to a higher resolution in the decoder, as shown in Fig. 1. S is the number of scales in the pyramid prediction. Then, we rescale these multi-scale predictions to the input size, and the corresponding results are denoted as $[p_0, p_1, \dots, p_s, \dots, p_{S-1}]$. For the labeled data, we use a supervised loss that is a combination of Dice and cross-entropy loss at multiple scales:

$$\mathcal{L}_{sup} = \frac{1}{S} \sum_{s=0}^{S-1} \frac{\mathcal{L}_{dice}(p_s, y_i) + \mathcal{L}_{ce}(p_s, y_i)}{2} \quad (1)$$

where y_i , \mathcal{L}_{dice} , \mathcal{L}_{ce} denote the ground truth of input x_i , the Dice loss and the cross entropy loss, respectively.

To efficiently leverage unlabeled data, we introduce a regularization by encouraging the multi-scale predictions of PPNet to be consistent. Concretely, we design a pyramid consistency loss to minimize the discrepancy (i.e., variance) among the predictions at different scales. First, we denote the average prediction across these scales as:

$$p_c = \frac{1}{S} \sum_{s=0}^{S-1} p_s \quad (2)$$

Then, the pyramid consistency loss is defined as:

$$\mathcal{L}_{pyc} = \frac{1}{S} \sum_{s=0}^{S-1} \|p_s - p_c\|_2 \quad (3)$$

where we encourage a minimized L_2 distance between the prediction at each scale and the average prediction.

2.2 Uncertainty Rectified Pyramid Consistency Loss

As the pyramid prediction at a range of scales has different spatial resolutions, even they can be resampled to the same resolution as the input, the resampled results still have different spatial frequencies, i.e., the prediction at the lowest resolution captures the low-frequency component of the segmentation, and the prediction at the highest resolution obtains more high-frequency components. Directly imposing a voxel-level consistency among these predictions can be problematic due to the different frequencies, such as loss of fine details or model collapse. Inspired by existing works [3, 24, 25, 27, 30], we introduce an uncertainty-aware method to address these problems. Unlike existing methods, our uncertainty estimation is a scale-level approach and only requires a single forward pass, which needs less computational cost and running time than exiting methods.

Efficient Uncertainty Estimation Based on Pyramid Predictions. As our PPNet obtains multiple predictions in a single forward pass, uncertainty estimation can be obtained efficiently by justing measuring their discrepancy without extra efforts. To be specific, we use the KL-divergence between the average prediction and the prediction at scale s as the uncertainty measurement:

$$\mathcal{D}_s \approx \sum_{j=0}^C p_s^j \cdot \log \frac{p_s^j}{p_c^j} \quad (4)$$

where p_s^j is the j th channel of p_s , and C is the class (i.e., channel) number. The approximated uncertainty shows the difference between the p_s and p_c . Note that for a given voxel in \mathcal{D}_s , a larger value indicates the prediction for that pixel at scale s is far from the other scales, i.e., with high uncertainty. As result, we obtain a set of uncertainty maps $\mathcal{D}_0, \mathcal{D}_1, \dots, \mathcal{D}_{S-1}$, where \mathcal{D}_s corresponds to uncertainty of p_s .

Uncertainty Rectifying. Based on the estimated uncertainty maps $\mathcal{D}_0, \mathcal{D}_1, \dots, \mathcal{D}_{S-1}$, we further extend the pyramid consistency \mathcal{L}_{pyc} to emphasize reliable parts and ignore unreliable parts of the predictions for stable unsupervised training. Specifically, for unlabeled data, we use the estimated uncertainty to automatically select reliable voxels for loss calculation. The rectified pyramid consistency loss is formulated as:

$$\mathcal{L}_{unsup} = \underbrace{\frac{1}{S} \sum_{s=0}^{S-1} \sum_v (p_s^v - p_c^v)^2 \cdot w_s^v}_{\text{uncertainty rectification}} + \underbrace{\frac{1}{S} \sum_{s=0}^{S-1} \|\mathcal{D}_s\|_2}_{\text{uncertainty minimization}} \quad (5)$$

where p_s^v and \mathcal{D}_s^v are the corresponding prediction and uncertainty values for voxel v . The consistency loss consists of two terms: the first is an uncertainty rectification (UR) term and the second is uncertainty minimization (UM) term. For a more stable training, we follow the policy in [30] and we use a voxel- and scale-wise weight w_s^v to automatically rectify the MSE loss rather than the threshold-based cut-off approaches [3, 27], as the threshold is hard to determine. The weight for a voxel v at scale s is defined as: $w_s^v = e^{-\mathcal{D}_s^v}$, it corresponds to voxel-wise exponential operation for $-\mathcal{D}_s$. According to this definition, for a given voxel at scale s , a higher uncertainty leads to a lower weight automatically. In addition, to encourage the PPNet to produce more consistent predictions at different scales, we use the uncertainty minimization term as a constraint directly. With this uncertainty rectified consistency loss, the PPNet can learn more reliable knowledge, which can then reduce the overall uncertainty of the model and produce more consistent predictions.

2.3 The Overall Loss Function

The proposed URPC framework learns from both labeled data and unlabeled data by minimizing the following combined objective function:

$$\mathcal{L}_{total} = \mathcal{L}_{sup} + \lambda \cdot \mathcal{L}_{unsup} \quad (6)$$

where \mathcal{L}_{sup} , \mathcal{L}_{unsup} are defined in Eq. 1 and Eq. 5, respectively. λ is a widely-used time-dependent Gaussian warming up function [21, 27] to control the balance between the supervised loss and unsupervised consistency loss, which is defined as $\lambda(t) = w_{max} \cdot e^{(-5(1 - \frac{t}{t_{max}})^2)}$, where w_{max} means the final regularization weight, t denotes the current training step and t_{max} is the maximal training step.

3 Experiments and Results

3.1 Dataset and Implementations

The NPC dataset used in this work was collected from a local cancer center. A total number of 258 T1-weighted MRI images from 258 patients of NPC before

radiotherapy were acquired on several 3T Siemens scanners. The mean resolution of the dataset was $1.23 \text{ mm} \times 1.23 \text{ mm} \times 1.10 \text{ mm}$, and the mean dimension was $176 \times 286 \times 245$. The ground truth for GTVnx and GTVnd were obtained from manual segmentation by two experienced radiologists using ITK-SNAP [28]. The dataset was randomly split into 180 cases for training, 20 cases for validation, and 58 cases for testing. For the training images, only 18 (i.e., 10%) were used as labeled and the remaining 162 scans were used as unlabeled. For pre-processing, we just normalize each scan to zero mean and unit variance. In the evaluation stage, following existing work [13], we used the commonly-adopted Dice Similarity Coefficient (*DSC*) and the Average Surface Distance (*ASD*) as segmentation quality evaluation metrics.

The framework was implemented in PyTorch [18], using a node of a cluster with 8 TiTAN 1080TI GPUs. We used the SGD optimizer (weight decay = 0.0001, momentum = 0.9) with Eq. 6 for training our method. During the training processing, the poly learning rate strategy was used for learning rate decay, where the initial learning rate 0.1 was multiplied by $(1.0 - \frac{t}{t_{max}})^\gamma$ with $\gamma = 0.9$ and $t_{max} = 60k$. The batch size was set to 4, and each batch consists of two annotated images and two unannotated images. We randomly cropped $112 \times 112 \times 112$ sub-volumes as the network input and employed data augmentation to enlarge the dataset and avoid over-fitting, including random cropping, random flipping, and random rotation. The final segmentation results were obtained by using a sliding window strategy. Following [27], the w_{max} was set to 0.1 for all experiments. (Details of the NPC dataset is presented in supplementary materials.)

Table 1. Ablation study of the proposed URPC framework on the NPC MRI dataset, where 18 labeled and 162 unlabeled images were used for training. UR and UM denote the uncertainty rectification term and uncertainty minimization term, respectively.

Method	GTVnx		GTVnd		Mean	
	<i>DSC</i> (%)	<i>ASD</i> (voxel)	<i>DSC</i> (%)	<i>ASD</i> (voxel)	<i>DSC</i> (%)	<i>ASD</i> (voxel)
Baseline ($S = 1$)	71.94 ± 11.60	2.42 ± 1.65	66.27 ± 14.62	3.60 ± 3.12	69.10 ± 10.15	3.01 ± 1.76
$S = 2$	79.88 ± 6.91	1.79 ± 1.27	72.82 ± 15.55	2.85 ± 2.54	76.35 ± 9.48	2.32 ± 1.46
$S = 3$	79.09 ± 5.82	1.76 ± 0.97	75.08 ± 13.22	2.25 ± 2.27	77.09 ± 7.85	2.05 ± 1.24
$S = 4$	80.13 ± 6.37	1.82 ± 1.30	75.83 ± 12.93	2.65 ± 2.77	77.98 ± 8.00	2.24 ± 1.53
$S = 5$	79.10 ± 6.53	1.84 ± 1.14	75.73 ± 13.71	2.29 ± 2.61	77.42 ± 8.27	2.06 ± 1.43
$S = 4 + \text{UR}$	80.99 ± 5.50	1.70 ± 1.12	75.22 ± 13.86	3.05 ± 3.16	78.11 ± 8.06	2.38 ± 1.65
$S = 4 + \text{UR} + \text{UM}$	80.76 ± 5.72	1.69 ± 1.06	75.95 ± 12.74	2.20 ± 2.07	78.36 ± 7.66	1.95 ± 1.18

3.2 Evaluation of Our Proposed URPC on the NPC Dataset

Ablation Study. Firstly, to investigate the impact of different numbers of scales in the pyramid prediction of PPNet, as shown in Fig. 1, we set S of PPNet to 2, 3, 4, and 5, respectively, and UR and UM were not used at this stage. They were compared with the baseline of 3D UNet [6] without multi-scale predictions and therefore it only learns from labeled data. In contrast, the PPNet learns

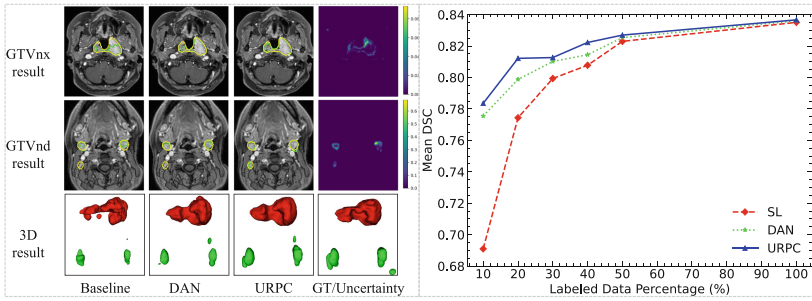
Table 2. Comparison between our method and existing methods on the NPC MRI dataset, when using 10% labeled data. * denotes p -value < 0.05 when comparing the proposed with the others.

Method	GTVnx		GTVnd		Mean		$T-T$ (h)
	DSC (%)	ASD (voxel)	DSC (%)	ASD (voxel)	DSC (%)	ASD (voxel)	
SL (10%)	71.94 \pm 11.60*	2.42 \pm 1.65*	66.27 \pm 14.62*	3.60 \pm 3.12*	69.10 \pm 10.15*	3.01 \pm 1.76*	73
SL (100%)	83.93 \pm 4.77*	1.35 \pm 0.73*	83.10 \pm 9.05*	1.48 \pm 1.73*	83.51 \pm 5.35*	1.41 \pm 0.94*	61
MT [21]	79.80 \pm 6.74*	1.70 \pm 1.17	69.78 \pm 16.34*	2.81 \pm 2.57*	74.79 \pm 9.15*	2.25 \pm 1.40*	76
ICT [22]	80.58 \pm 6.23	1.58 \pm 1.02*	72.62 \pm 13.47*	2.72 \pm 2.61*	76.59 \pm 7.98*	2.15 \pm 1.38*	78
EM [23]	79.85 \pm 6.32*	1.66 \pm 1.06	69.92 \pm 15.39*	3.14 \pm 2.82*	74.89 \pm 8.85*	2.40 \pm 1.54*	74
UAMT [27]	79.62 \pm 7.16*	1.67 \pm 1.05	71.98 \pm 15.66*	2.55 \pm 2.58*	75.78 \pm 9.67*	2.11 \pm 1.39*	95
DAN [29]	80.47 \pm 5.73	1.56 \pm 0.81*	74.62 \pm 12.83*	2.74 \pm 2.62*	77.55 \pm 7.39	2.15 \pm 1.26*	104
Ours	80.76 \pm 5.72	1.69 \pm 1.06	75.95 \pm 12.74	2.20 \pm 2.07	78.36 \pm 7.66	1.95 \pm 1.18	74

from both labeled data and unlabeled data. The quantitative results are shown in Table 1. It can be found that when S increases from 2 to 4, the performance of the proposed URPC improves gradually. However, we found that $S = 5$ achieved a lower performance than $S = 4$. That is because the resolution of p_4 is too small to preserve more details. It conforms to common sense and indicates that not more multi-scales predictions are better. Therefore, we used $S = 4$ for SSL in the following experiments. Secondly, to measure the contribution of the uncertainty rectifying module, we then turn on the UR term and UM term with $S = 4$ for training. From the last section of Table 1, we can see that both uncertainty rectifying (UR) term and uncertainty minimization (UM) term boost the model performance. What’s more, combining all sub-modules into a unified framework results in a better gain where the mean DSC and ASD were improved by 9.26% and 1.06 voxels than the baseline, demonstrating their effectiveness for semi-supervised segmentation.

Comparison with Other Semi-supervised Methods. We compared our method with only using 18 annotated images for supervised learning with 3D UNet, which is denoted as SL (10%). Similarly, SL (100%) denote supervised learning with all the training images annotated, which gives the performance upper bound. In addition, we further compared our methods with five state-of-the-art semi-supervised segmentation methods, including Mean Teacher (MT) [21], Interpolation Consistency Training (ICT) [22], Entropy Minimization (EM) [23], Uncertainty Aware Mean Teacher (UAMT) [27] and Deep Adversarial Network (DAN) [29]. Note that, for a fair comparison, all these methods were implemented by using 3D UNet [6] as the backbone and they are online available [14]. Table 2 shows the quantitative comparison of these methods. It can be found that compared with SL (10%), all semi-supervised methods improve the segmentation performance by a large margin, as they can learn from the unannotated data by a regularization loss during the training, and the DAN [29] achieve the best results among existing methods. Our framework (URPC) achieves better performance than these semi-supervised methods when using 10% labeled

data. These results show that our URPC has the capability to capture the rich information from the unlabeled data in addition to labeled data. What's more, our method is more efficient than existing methods and requires less training time ($T-T$) and computational cost, as it just needs to pass an input image once in an iteration. In Fig. 2(a), we visualize some 2D and 3D results of the supervised and semi-supervised methods when using 10% labeled data. Compared with supervised learning (SL) baseline and DAN [29], our method has a higher overlap ratio with the ground truth and reduces the false negative in both slice level and volume level, especially in GTVnd segmentation. We further visualized the estimated uncertainty (\mathcal{D}_0 in Eq. 4) in the last column of Fig. 2(a), showing that the uncertain region is mainly distributed near the boundary. We further performed a study on the data utilization efficiency of the URPC. Figure 2(b) shows the evolution curve of mean DSC of GTVnx and GTVnd segmentation obtained by SL, DAN [29] and URPC when using different numbers of labeled data. It can be found that URPC consistently outperforms SL and DAN [29], and when increasing the labeled ratio to 50%, URPC achieves the mean DSC of 82.74% which is very close to 83.51% obtained by SL (100%). These results demonstrate that the URPC has the capability to utilize the unlabeled data to bring performance gains. More results on 20% labeled data presented in supplementary materials showed that our method also outperforms these existing methods.



(a) Visualization comparison between different methods. (b) Performance under different ratios of labeled images.

Fig. 2. Comparison between different methods.

4 Conclusion

In this paper, we proposed a novel efficient semi-supervised learning framework URPC for medical image segmentation. A pyramid prediction network is employed to learn from the unlabeled data by encouraging to produce consistent predictions at multiple scales. An uncertainty rectifying module is designed to improve the stability of learning from unlabeled images and further boost model performance, where the uncertainty estimation can be obtained with a

single forward pass efficiently. We applied the proposed method to the segmentation of GTVnx and GTVnd, and the results demonstrated the effectiveness and generalization of URPC and also indicated the promising potential of our proposed approach for further clinical use. In the future, we will evaluate this framework on other segmentation tasks.

Acknowledgment. This work was supported by the National Natural Science Foundations of China [81771921, 61901084], and also by key research and development project of Sichuan province, China [20ZDYF2817]. We thank M.D. Mengwan Wu and Yuanyuan Shen from the Sichuan Provincial People’s Hospital for the data annotation and checking.

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. Cao, X., Chen, H., Li, Y., Peng, Y., Wang, S., Cheng, L.: Uncertainty aware temporal-ensembling model for semi-supervised abus mass segmentation. *TMI* **40**(1), 431–443 (2020)
4. Chaitanya, K., Karani, N., Baumgartner, C.F., Becker, A., Donati, O., Konukoglu, E.: Semi-supervised and task-driven data augmentation. In: Chung, A.C.S., Gee, J.C., Yushkevich, P.A., Bao, S. (eds.) IPMI 2019. LNCS, vol. 11492, pp. 29–41. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-20351-1_3
5. Chen, W., et al.: Cancer statistics in China, 2015. *CA: A Cancer J. Clin.* **66**(2), 115–132 (2016)
6. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9901, pp. 424–432. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46723-8_49
7. Cui, W., et al.: Semi-supervised brain lesion segmentation with an adapted mean teacher model. In: Chung, A.C.S., Gee, J.C., Yushkevich, P.A., Bao, S. (eds.) IPMI 2019. LNCS, vol. 11492, pp. 554–565. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-20351-1_43
8. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: representing model uncertainty in deep learning. In: ICML, pp. 1050–1059 (2016)
9. Hang, W., et al.: Local and global structure-aware entropy regularized mean teacher model for 3D left atrium segmentation. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12261, pp. 562–571. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59710-8_55
10. Lee, C.Y., Xie, S., Gallagher, P., Zhang, Z., Tu, Z.: Deeply-supervised nets. In: Artificial Intelligence and Statistics, pp. 562–570. PMLR (2015)

11. Li, S., Zhang, C., He, X.: Shape-aware semi-supervised 3D semantic segmentation for medical images. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12261, pp. 552–561. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59710-8_54
12. Li, X., Yu, L., Chen, H., Fu, C.W., Xing, L., Heng, P.A.: Transformation-consistent self-ensembling model for semisupervised medical image segmentation. *TNNLS* **32**(2), 523–534 (2020)
13. Lin, L., et al.: Deep learning for automated contouring of primary tumor volumes by MRI for nasopharyngeal carcinoma. *Radiology* **291**(3), 677–686 (2019)
14. Luo, X.: SSL4MIS (2020). <https://github.com/HiLab-git/SSL4MIS>
15. Luo, X., Chen, J., Song, T., Wang, G.: Semi-supervised medical image segmentation through dual-task consistency. In: AAAI, vol. 35, no. 10, pp. 8801–8809 (2021)
16. Ma, J., et al.: Active contour regularized semi-supervised learning for COVID-19 CT infection segmentation with limited annotations. *Phys. Med. Biol.* **65**(22), 225034 (2020)
17. Nie, D., Gao, Y., Wang, L., Shen, D.: ASDNet: attention based semi-supervised deep networks for medical image segmentation. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11073, pp. 370–378. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00937-3_43
18. Paszke, A., et al.: PyTorch: an imperative style, high-performance deep learning library. In: NeurIPS, pp. 8026–8037 (2019)
19. Peng, J., Pedersoli, M., Desrosiers, C.: Mutual information deep regularization for semi-supervised segmentation. In: MIDL, pp. 601–613. PMLR (2020)
20. Qiao, S., Shen, W., Zhang, Z., Wang, B., Yuille, A.: Deep co-training for semi-supervised image recognition. In: ECCV, pp. 135–152 (2018)
21. Tarvainen, A., Valpola, H.: Mean teachers are better role models: weight-averaged consistency targets improve semi-supervised deep learning results. In: NeurIPS, pp. 1195–1204 (2017)
22. Verma, V., Lamb, A., Kannala, J., Bengio, Y., Lopez-Paz, D.: Interpolation consistency training for semi-supervised learning. In: IJCAI, pp. 3635–3641 (2019)
23. Vu, T.H., Jain, H., Bucher, M., Cord, M., Pérez, P.: ADVENT: adversarial entropy minimization for domain adaptation in semantic segmentation. In: CVPR, pp. 2517–2526 (2019)
24. Wang, G., Aertsen, M., Deprest, J., Ourselin, S., Vercauteren, T., Zhang, S.: Uncertainty-guided efficient interactive refinement of fetal brain segmentation from stacks of MRI slices. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12264, pp. 279–288. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59719-1_28
25. Wang, G., Li, W., Aertsen, M., Deprest, J., Ourselin, S., Vercauteren, T.: Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks. *Neurocomputing* **338**, 34–45 (2019)
26. Wang, Y., et al.: Double-uncertainty weighted method for semi-supervised learning. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12261, pp. 542–551. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59710-8_53
27. 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

28. Yushkevich, P.A., et al.: User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability. *Neuroimage* **31**(3), 1116–1128 (2006)
29. Zhang, Y., Yang, L., Chen, J., Fredericksen, M., Hughes, D.P., Chen, D.Z.: Deep adversarial networks for biomedical image segmentation utilizing unannotated images. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) *MICCAI 2017*. LNCS, vol. 10435, pp. 408–416. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66179-7_47
30. Zheng, Z., Yang, Y.: Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. *IJCV* **129**(4), 1106–1120 (2021). <https://doi.org/10.1007/s11263-020-01395-y>