



# Reciprocal Learning for Semi-supervised Segmentation

Xiangyun Zeng<sup>1,2,3</sup>, Rian Huang<sup>1,2,3</sup>, Yuming Zhong<sup>1,2,3</sup>, Dong Sun<sup>1,2,3</sup>,  
Chu Han<sup>4</sup>, Di Lin<sup>5</sup>, Dong Ni<sup>1,2,3</sup>, and Yi Wang<sup>1,2,3</sup>(✉)

<sup>1</sup> National-Regional Key Technology Engineering Laboratory for Medical Ultrasound, Guangdong Provincial Key Laboratory of Biomedical Measurements and Ultrasound Imaging, School of Biomedical Engineering, Health Science Center, Shenzhen University, Shenzhen, China

<sup>2</sup> Medical UltraSound Image Computing (MUSIC) Lab, Shenzhen University, Shenzhen, China

<sup>3</sup> Marshall Laboratory of Biomedical Engineering, Shenzhen University, Shenzhen, China  
onewang@szu.edu.cn

<sup>4</sup> Department of Radiology, Guangdong Provincial People's Hospital, Guangdong Academy of Medical Sciences, Guangzhou, Guangdong, China

<sup>5</sup> College of Intelligence and Computing, Tianjin University, Tianjin, China

**Abstract.** Semi-supervised learning has been recently employed to solve problems from medical image segmentation due to challenges in acquiring sufficient manual annotations, which is an important prerequisite for building high-performance deep learning methods. Since unlabeled data is generally abundant, most existing semi-supervised approaches focus on how to make full use of both limited labeled data and abundant unlabeled data. In this paper, we propose a novel semi-supervised strategy called reciprocal learning for medical image segmentation, which can be easily integrated into any CNN architecture. Concretely, the reciprocal learning works by having a pair of networks, one as a student and one as a teacher. The student model learns from pseudo label generated by the teacher. Furthermore, the teacher updates its parameters autonomously according to the reciprocal feedback signal of how well student performs on the labeled set. Extensive experiments on two public datasets show that our method outperforms current state-of-the-art semi-supervised segmentation methods, demonstrating the potential of our strategy for the challenging semi-supervised problems. *The code is publicly available at <https://github.com/XYZach/RLSSS>.*

**Keywords:** Semi-supervised learning · Reciprocal learning · Segmentation · Meta learning · Deep learning

---

X. Zeng, R. Huang and Y. Zhong—Contribute equally to this work.

© Springer Nature Switzerland AG 2021

M. de Bruijne et al. (Eds.): MICCAI 2021, LNCS 12902, pp. 352–361, 2021.

[https://doi.org/10.1007/978-3-030-87196-3\\_33](https://doi.org/10.1007/978-3-030-87196-3_33)

## 1 Introduction

Accurate and robust segmentation of organs or lesions from medical images is of great importance for many clinical applications such as disease diagnosis and treatment planning. With a large amount of labeled data, deep learning has achieved great success in automatic image segmentation [7, 10]. In medical imaging domain, especially for volumetric images, reliable annotations are difficult to obtain as expert knowledge and time are both required. Unlabeled data, on the other hand, are easier to acquire. Therefore, semi-supervised approaches with unlabeled data occupying a large portion of the training set are worth exploring.

Bai et al. [1] introduced a self-training-based method for cardiac MR image segmentation, in which the segmentation prediction for unlabeled data and the network parameters were alternatively updated. Xia et al. [14] utilized co-training for pancreas and liver tumor segmentation tasks by exploiting multi-viewpoint consistency of 3D data. These methods enlisted more available training sources by creating pseudo labels, however, they did not consider the reliability of the pseudo labels which may leads to meaningless guidance. Some approaches to semi-supervised learning were inspired by the success of self-ensembling method. For example, Li et al. [5] embedded the transformation consistency into  $H$ -model [3] to enhance the regularization for pixel-wise predictions. Yu et al. [16] designed an uncertainty-aware mean teacher framework, which can generate more reliable predictions for student to learn. To exploit the structural information for prediction, Hang et al. [2] proposed a local and global structure-aware entropy regularized mean teacher for left atrium segmentation. In general, most teacher-student methods update teacher’s parameters using exponential moving average (EMA), which is an useful ensemble strategy. However, the EMA focuses on weighting the student’s parameters at each stage during training process, without evaluating the quality of parameters explicitly. It is more expected that the teacher model could purposefully update the parameters through a parameter evaluation strategy, so as to generate more reliable pseudo-labels.

In this paper, we design a novel strategy named reciprocal learning for semi-supervised segmentation. Specifically, we make better use of the limited labeled data by using reciprocal learning strategy so that the teacher model can update its parameters with gradient descent algorithm and generate more reliable annotations for unlabeled set as the number of reciprocal learning step increases. We evaluate our approach on the pancreas CT dataset and the Atrial Segmentation Challenge dataset with extensive comparisons to existing methods. The results demonstrate that our segmentation network consistently outperforms the state-of-the-art method in respect to the evaluation metrics of Dice Similarity (Dice), Jaccard Index (Jaccard), 95% Hausdorff Distance (95 HD) and Average Symmetric Surface Distance (ASD). Our main contributions are three folds:

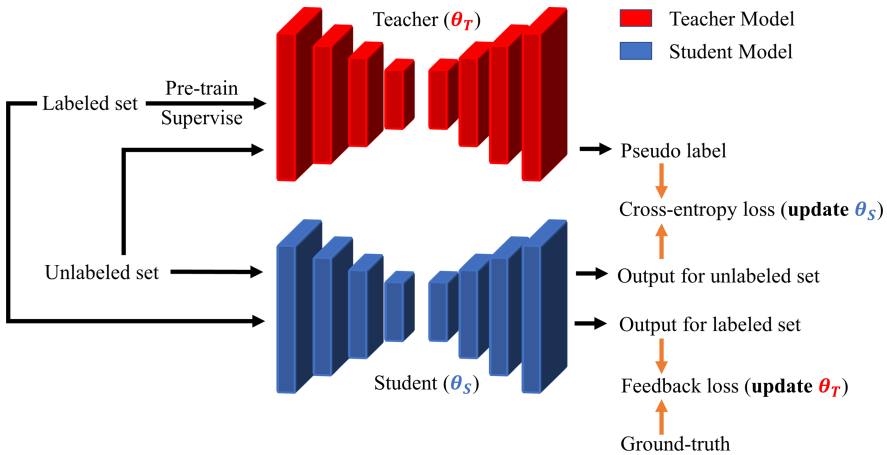
- We present a simple yet efficient reciprocal learning strategy for segmentation to reduce the labeling efforts. Inspired by the idea from *learning to learn*, we design a feedback mechanism for teacher network to generate more reliable

pseudo labels by observing how pseudo labels would affect the student. In our implementation, the feedback signal is the performance of the student on the labeled set. By reciprocal learning strategy, the teacher can update its parameters autonomously.

- The proposed reciprocal learning strategy can be utilized directly in any CNN architecture. Specifically, any segmentation network can be used as the backbone, which means there are still opportunities for further enhancements.
- Experiments on two public datasets show our proposed strategy can further raise semi-supervised segmentation quality compared with existing methods.

## 2 Methods

Figure 1 illustrates our reciprocal learning framework for semi-supervised segmentation. We deploy a meta-learning concept for teacher model to generate better pseudo labels by observing how pseudo labels would affect the student. Specifically, the teacher and student are trained in parallel: the student learns from pseudo labels generated by the teacher, and the teacher learns from the feedback signal of how well the student performs on the labeled set.



**Fig. 1.** The schematic illustration of our reciprocal learning framework for semi-supervised segmentation. In this paper, V-Net is used as the backbone. Again, we emphasize any segmentation network could be used as the backbone in our framework.

### 2.1 Notations

We denote the labeled set as  $(x_l, y_l)$  and the unlabeled set as  $x_u$ , where  $x$  is the input volume and  $y$  is the ground-truth segmentation. Let  $T$  and  $S$  respectively be the teacher model and the student model, and let their corresponding parameters be  $\theta_T$  and  $\theta_S$ . We denote the soft predictions of teacher network on the  $x_u$  as  $T(x_u; \theta_T)$  and likewise for the student.

### 2.2 Reciprocal Learning Strategy

Figure 1 shows the workflow of our proposed reciprocal learning strategy. Firstly, the teacher model should be well pre-trained on labeled set  $(x_l, y_l)$  in a supervised manner. We use cross-entropy loss (CE) as loss function:

$$\mathcal{L}_{pre-train} = CE(y_l, T(x_l; \theta_T)). \tag{1}$$

Then we use the teacher’s prediction on unlabeled set as pseudo labels  $\hat{y}_u$  to train the student model. Specifically, Pseudo Labels (PL) trains the student model to minimize the cross-entropy loss on unlabeled set  $x_u$ :

$$\hat{y}_u \sim T(x_u; \theta_T), \tag{2}$$

$$\theta_S^{PL} = \arg \min_{\theta_S} CE(\hat{y}_u, S(x_u; \theta_S)). \tag{3}$$

After the student model updated, it’s expected to perform well on the labeled set and achieve a low cross-entropy loss, i.e.  $CE(y_l, S(x_l; \theta_S^{PL}))$ . Notice that the optimal student parameters  $\theta_S^{PL}$  always depend on the teacher parameters  $\theta_T$  via the pseudo labels (see Eq. (2) and (3)). Therefore, we express the dependency as  $\theta_S^{PL}(\theta_T)$  and further optimize  $\mathcal{L}_{feedback}$  with respect to  $\theta_T$ :

$$\min_{\theta_T} \mathcal{L}_{feedback}(\theta_S^{PL}(\theta_T)) = CE(y_l, S(x_l; \theta_S^{PL}(\theta_T))). \tag{4}$$

For each reciprocal learning step (including one update for the student using Eq. (3) and one update for the teacher using Eq. (4) respectively), however, solving Eq. (3) to optimize  $\theta_S$  until complete convergence is inefficient, as computing the gradient  $\nabla_{\theta_T} \mathcal{L}_{feedback}(\theta_S^{PL}(\theta_T))$  requires unrolling the entire student training process. Instead, a meta-learning approach [6] is utilized to approximate  $\theta_S^{PL}$  with one-step gradient update of  $\theta_S$ :

$$\theta_S^{PL} \approx \theta_S - \eta_S \nabla_{\theta_S} CE(\hat{y}_u, S(x_u; \theta_S)), \tag{5}$$

where  $\eta_S$  is the learning rate. In this way, the student model and the teacher model have an alternating optimization:

- (1) Draw a batch of unlabeled set  $x_u$ , then sample  $T(x_u; \theta_T)$  from the teacher model, and optimize with stochastic gradient descent (SGD):

$$\theta'_S = \theta_S - \eta_S \nabla_{\theta_S} CE(\hat{y}_u, S(x_u; \theta_S)). \tag{6}$$

- (2) Draw a batch of labeled set  $(x_l, y_l)$ , and reuse the student’s update to optimize with SGD:

$$\theta'_T = \theta_T - \eta_T \nabla_{\theta_T} \mathcal{L}_{feedback}(\theta'_S). \tag{7}$$

Optimize  $\theta_S$  with Eq. (6) can be simply computed via back-propagation. We now present the derivation for optimizing  $\theta_T$ . Firstly, by the chain rule, we have

$$\begin{aligned} \frac{\partial \mathcal{L}_{feedback}(\theta'_S)}{\partial \theta_T} &= \frac{\partial CE(y_l, S(x_l; \theta'_S))}{\partial \theta_T} \\ &= \frac{\partial CE(y_l, S(x_l; \theta'_S))}{\partial \theta_S} \bigg|_{\theta_S=\theta'_S} \cdot \frac{\partial \theta'_S}{\partial \theta_T} \end{aligned} \tag{8}$$

We focus on the second term in Eq. (8)

$$\begin{aligned} \frac{\partial \theta'_S}{\partial \theta_T} &= \frac{\partial}{\partial \theta_T} [\theta_S - \eta_S \nabla_{\theta_S} CE(\hat{y}_u, S(x_u; \theta_S))] \\ &= \frac{\partial}{\partial \theta_T} \left[ -\eta_S \cdot \left( \frac{\partial CE(\hat{y}_u, S(x_u; \theta_S))}{\partial \theta_S} \Big|_{\theta_S = \theta_S} \right)^\top \right] \end{aligned} \quad (9)$$

To simplify notations, we define *the gradient*

$$g_S(\hat{y}_u) = \left( \frac{\partial CE(\hat{y}_u, S(x_u; \theta_S))}{\partial \theta_S} \Big|_{\theta_S = \theta_S} \right)^\top \quad (10)$$

Since  $g_S(\hat{y}_u)$  has dependency on  $\theta_T$  via  $\hat{y}_u$ , we apply the REINFORCE equation [13] to achieve

$$\begin{aligned} \frac{\partial \theta'_S}{\partial \theta_T} &= -\eta_S \cdot \frac{\partial g_S(\hat{y}_u)}{\partial \theta_T} \\ &= -\eta_S \cdot g_S(\hat{y}_u) \cdot \frac{\partial \log P(\hat{y}_u | x_u; \theta_T)}{\partial \theta_T} \\ &= \eta_S \cdot g_S(\hat{y}_u) \cdot \frac{\partial CE(\hat{y}_u, T(x_u; \theta_T))}{\partial \theta_T} \end{aligned} \quad (11)$$

Finally, we obtain the gradient

$$\begin{aligned} \nabla_{\theta_T} \mathcal{L}_{feedback}(\theta'_S) &= \eta_S \cdot (\nabla_{\theta'_S} CE(y_l, S(x_l; \theta'_S)))^\top \cdot \\ &\quad \nabla_{\theta_S} CE(\hat{y}_u, S(x_u; \theta_S)) \cdot \nabla_{\theta_T} CE(\hat{y}_u, T(x_u; \theta_T)). \end{aligned} \quad (12)$$

However, it might lead to overfitting if we rely solely on the student's performance to optimize the teacher model. To overcome this, we leverage labeled set to supervise teacher model throughout the course of training. Therefore, the ultimate optimal equation of the teacher model can be summarized as:  $\theta'_T = \theta_T - \eta_T \nabla_{\theta_T} [\mathcal{L}_{feedback}(\theta'_S) + \lambda CE(y_l, T(x_l; \theta_T))]$ , where  $\lambda$  is the weight to balance the importance of different losses.

## 3 Experiments

### 3.1 Materials and Pre-processing

To demonstrate the effectiveness of our proposed method, experiments were carried on two different public datasets.

The first dataset is the pancreas dataset [11] obtained using Philips and Siemens MDCT scanners. It includes 82 abdominal contrast enhanced CT scans, which have resolutions of  $512 \times 512$  pixels with varying pixel sizes and slice thickness between 1.5–2.5 mm. We used the soft tissue CT window range of  $[-125, 275]$  HU, and cropped the images centering at pancreas regions based on

the ground truth with enlarged margins (25 voxels)<sup>1</sup> after normalizing them as zero mean and unit variance. We used 62 scans for training and 20 scans for validation.

The second dataset is the left atrium dataset [15]. It includes 100 gadolinium-enhanced MR images, which have a resolution of  $0.625 \times 0.625 \times 0.625 \text{ mm}^3$ . We cropped centering at heart regions and normalized them as zero mean and unit variance. We used 80 scans for training and 20 scans for validation.

In this work, we report the performance of all methods trained with 20% labeled images and 80% unlabeled images as the typical semi-supervised learning experimental setting.

### 3.2 Implementation Details

Our proposed method was implemented with the popular library Pytorch, using a TITAN Xp GPU. In this work, we employed V-Net [9] as the backbone. More importantly, it's flexible that any segmentation network can be the backbone. We set  $\lambda = 1$ . Both the teacher model and the student model share the same architecture but have independent weights. Both networks were trained by the stochastic gradient descent (SGD) optimizer for 6000 iterations, with an initial learning rate  $\eta_T = \eta_S = 0.01$ , decayed by 0.1 every 2500 iterations. To tackle the issues of limited data samples and demanding 3D computations cost, we randomly cropped  $96 \times 96 \times 96$  (pancreas dataset) and  $112 \times 112 \times 80$  (left atrium dataset) sub-volumes as the network input and adopted data augmentation for training. In the inference phase, we only utilized the student model to predict the segmentation for the input volume and we used a sliding window strategy to obtain the final results, with a stride of  $10 \times 10 \times 10$  for the pancreas dataset and  $18 \times 18 \times 4$  for the left atrium dataset.


### 3.3 Segmentation Performance

We compared results of our method with several state-of-the-art semi-supervised segmentation methods, including mean teacher self-ensembling model (MT) [12], uncertainty-aware mean teacher model (UA-MT) [16], shape-aware adversarial network (SASSNet) [4], uncertainty-aware multi-view co-training (UMCT) [14] and transformation-consistent self-ensembling model (TCSM) [5]. Note that we used the official code of MT, UA-MT, SASSNet, TCSM and reimplemented the UMCT which didn't release the official code. For a fair comparison, we obtained the results of our competitors by using the same backbone (V-Net) and re-training their networks to obtain the best segmentation results on the Pancreas dataset and the Left Atrium dataset.

---

<sup>1</sup> This study mainly focused on the challenging problem of semi-supervised learning for insufficient annotations. Several semi-supervised segmentation studies used cropped images for validations, e.g., UAMT [16] used cropped left atrium images, and [8] used cropped pancreas images. We followed their experimental settings.

The metrics employed to quantitatively evaluate segmentation include Dice, Jaccard, 95 HD and ASD. A better segmentation shall have larger values of Dice and Jaccard, and smaller values of other metrics.




**Fig. 2.** 2D visualization of our proposed semi-supervised segmentation method under 20% labeled images. The first two rows are the segmentation results of pancreas and the last two rows are the segmentation results of left atrium. Red and blue colors show the ground truths and the predictions, respectively. (Color figure online)

We first evaluated our proposed method on pancreas dataset. The first two rows of Fig. 2 visualize 12 slices of the pancreas segmentation results. Apparently, our method consistently obtained similar segmented boundaries to the ground truths. Table 1 presents the quantitative comparison of several state-of-the-art semi-supervised segmentation methods. Compared with using only 20% annotated images (the first row), all semi-supervised segmentation methods achieved greater performance proving that they could both utilize unlabeled images. Notably, our method improved the segmentation by 9.76% Dice and 12.00% Jaccard compared with the fully supervised baseline's results. Furthermore, our method achieved the best performance over the state-of-the-art semi-supervised methods on all metrics. Compared with other methods, our proposed method utilized the limited labeled data in a better way by using reciprocal learning strategy so that the teacher model could update its parameters autonomously and generate more reliable annotations for unlabeled data as the number of reciprocal learning step increases. The first two rows of Fig. 3 visualize the pancreas segmentation results of different semi-supervised segmentation methods in 3D. Compared with other methods, our method produced less false positive predictions especially in the case as shown in the first row in Fig. 3.

We also evaluated our method on the left atrium dataset, which is a widely-used dataset for semi-supervised segmentation. The last two rows of Fig. 2 visualize 12 segmented slices. Obviously, our results can successfully infer the ambigu-

**Table 1.** Quantitative comparison between our method and other semi-supervised methods on the pancreas CT dataset.

Method	# Scans used		Metrics			
	Labeled	Unlabeled	Dice [%]	Jaccard [%]	ASD [voxel]	95HD [voxel]
V-Net	12	0	72.16	57.67	2.80	12.34
V-Net	62	0	84.89	73.89	1.22	4.44
MT	12	50	79.14	66.04	2.21	8.71
SASS	12	50	79.55	66.32	2.29	8.53
UMCT	12	50	79.74	66.59	4.13	14.01
UAMT	12	50	80.04	67.52	2.99	10.96
TCSM	12	50	78.17	64.95	5.06	17.52
<b>Ours</b>	12	50	<b>81.92</b>	<b>69.67</b>	<b>1.82</b>	<b>6.36</b>


**Fig. 3.** Four cases of 3D visualization of different semi-supervised segmentation methods under 20% labeled images. The first two rows are the results of pancreas segmentation and the last two rows are the results of left atrium segmentation.

**Table 2.** Quantitative comparison between our method and other semi-supervised methods on the Left Atrium MRI dataset.

Method	# scans used		Metrics			
	Labeled	Unlabeled	Dice [%]	Jaccard [%]	ASD [voxel]	95HD [voxel]
V-Net	16	0	84.41	73.54	5.32	19.94
V-Net	80	0	91.42	84.27	1.50	5.15
MT	16	64	88.12	79.03	2.65	10.92
SASS	16	64	89.27	80.82	3.13	8.83
UMCT	16	64	89.36	81.01	2.60	7.25
UAMT	16	64	88.88	80.21	2.26	7.32
TCSM	16	64	86.26	76.56	2.35	9.67
<b>Ours</b>	16	64	<b>90.06</b>	<b>82.01</b>	<b>2.13</b>	<b>6.70</b>



ous boundaries and have a high overlap ratio with the ground truths. A quantitative comparison is shown in Table 2. Compared with using only 20% labeled images (the first row), our method improved the segmentation by 5.65% Dice and 8.47% Jaccard, which were very close to using 100% labeled images (the second row). In addition, it can be observed that our method achieved the best performance than the state-of-the-art semi-supervised methods on all evaluation metrics, corroborating that our reciprocal learning strategy has the fully capability to utilize the limited labeled data. The last two rows of Fig. 3 visualize the left atrium segmentation results of different semi-supervised segmentation methods in 3D. Compared with other methods, our results were close to the ground truths and preserved more details and produced less false positives, which demonstrates the efficacy of our proposed reciprocal learning strategy.

We further conducted an ablation study to demonstrate the efficacy of the proposed reciprocal learning strategy. Specifically, we discarded our reciprocal learning strategy by fixing teacher model after it was well pretrained. The results degraded to 73.82%/86.82% Dice, 59.38%/77.27% Jaccard, 4.62/3.69 ASD and 17.78/12.29 95HD on pancreas/left atrium datasets, which shows our reciprocal learning contributes to the performance improvement.

## 4 Conclusion

This paper develops a novel reciprocal learning strategy for semi-supervised segmentation. Our key idea is to fully utilize the limited labeled data by updating parameters of the teacher and the student model in a reciprocal learning way. Meanwhile, our strategy is simple and can be used directly in existing state-of-the-art network architectures, where the performance can be effectively enhanced. Experiments on two public datasets demonstrate the effectiveness, robustness and generalization of our proposed method. In addition, our proposed reciprocal learning strategy is a general solution and has the potential to be used for other image segmentation tasks.

**Acknowledgements.** This work was supported in part by the National Key R&D Program of China (No. 2019YFC0118300), in part by the National Natural Science Foundation of China under Grants 62071305, 61701312 and 81971631, in part by the Guangdong Basic and Applied Basic Research Foundation (2019A1515010847), in part by the Medical Science and Technology Foundation of Guangdong Province (B2019046), in part by the Natural Science Foundation of Shenzhen University (No. 860-000002110129), and in part by the Shenzhen Peacock Plan (No. KQTD2016053112051497).

## 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](https://doi.org/10.1007/978-3-319-66185-8_29)

2. 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](https://doi.org/10.1007/978-3-030-59710-8_55)
3. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. arXiv preprint [arXiv:1610.02242](https://arxiv.org/abs/1610.02242) (2016)
4. 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](https://doi.org/10.1007/978-3-030-59710-8_54)
5. 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)
6. Liu, H., Simonyan, K., Yang, Y.: DARTS: differentiable architecture search. In: ICLR (2018)
7. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR, pp. 3431–3440 (2015)
8. Luo, X., Chen, J., Song, T., Wang, G.: Semi-supervised medical image segmentation through dual-task consistency. In: AAAI Conference on Artificial Intelligence (2021)
9. Milletari, F., Navab, N., Ahmadi, S.A.: V-Net: fully convolutional neural networks for volumetric medical image segmentation. In: 3DV, pp. 565–571. IEEE (2016)
10. 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](https://doi.org/10.1007/978-3-319-24574-4_28)
11. Roth, H.R., et al.: DeepOrgan: multi-level deep convolutional networks for automated pancreas segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9349, pp. 556–564. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-24553-9\\_68](https://doi.org/10.1007/978-3-319-24553-9_68)
12. Tarvainen, A., Valpola, H.: Mean teachers are better role models: weight-averaged consistency targets improve semi-supervised deep learning results. In: NIPS, pp. 1195–1204 (2017)
13. Williams, R.J.: Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach. Learn.* **8**(3–4), 229–256 (1992). <https://doi.org/10.1007/BF00992696>
14. Xia, Y., et al.: 3D semi-supervised learning with uncertainty-aware multi-view co-training. In: WACV, pp. 3646–3655 (2020)
15. Xiong, Z., et al.: A global benchmark of algorithms for segmenting the left atrium from late gadolinium-enhanced cardiac magnetic resonance imaging. *Med. Image Anal.* **67**, 101832 (2021)
16. 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](https://doi.org/10.1007/978-3-030-32245-8_67)