

# Hierarchical Consistency Regularized Mean Teacher for Semi-supervised 3D Left Atrium Segmentation

Shumeng Li<sup>\*1,2</sup>, Ziyuan Zhao<sup>\*1,2</sup>, Kaixin Xu<sup>1</sup>, Zeng Zeng<sup>†1</sup>, Cuntai Guan<sup>2</sup>

**Abstract**—Deep learning has achieved promising segmentation performance on 3D left atrium MR images. However, annotations for segmentation tasks are expensive, costly and difficult to obtain. In this paper, we introduce a novel hierarchical consistency regularized mean teacher framework for 3D left atrium segmentation. In each iteration, the student model is optimized by multi-scale deep supervision and hierarchical consistency regularization, concurrently. Extensive experiments have shown that our method achieves competitive performance as compared with full annotation, outperforming other state-of-the-art semi-supervised segmentation methods.

**Clinical relevance**— The proposed method can efficiently utilize limited labeled data and abundant unlabeled data simultaneously for 3D left atrium segmentation, and achieve comparable performance compared with full annotation, which can be immediately implemented to any medical image segmentation tasks and extended into other clinical applications.

## I. INTRODUCTION

Deep learning has been widely applied in the field of medical image processing, in which, segmentation techniques play an important role in diagnostic analysis and clinical treatment [1]. Segmenting left atrium (LA) structures from magnetic resonance (MR) images accurately can provide a pre-operative assessment of its anatomy, which is essential for treating various cardiovascular diseases, such as atrial fibrillation [2]. Despite the promising performance of deep learning on LA segmentation [3], [4], most of the existing methods rely on a great number of expensive annotations delineated by experienced experts, and the labeling process is tedious and laborious.

To ease the manual labeling burden, significant effort has been devoted to utilizing available annotations and resources efficiently to improve the segmentation performance [5], [6]. Noting that plenty of unlabeled data can be easily collected, many works focus on leveraging limited data with annotations and abundant data without annotations together for improving the segmentation performance in a semi-supervised learning (SSL) manner [7], [8], [9], [10]. In recent years, the mean teacher (MT) model [11] has achieved great success in SSL, which enforces the consistency of the predictions with different perturbed inputs between student and teacher models thereby improving the robustness of the student model. However, the consistency is only applied

to the final predictions. Following [11], many perturbation-based SSL methods have been proposed, obtaining promising performance on LA segmentation [12], [13], [14]. These methods focus on uncertainty estimations to exploit rich unlabeled data from different perspectives with high computational complexity. For instance, UA-MT [12] estimates the uncertainty with multiple forward passes via Monte Carlo Dropout [15]. Another line of research [16], [17] tends to build task-level regularization by adding additional auxiliary tasks. In the absence of expert knowledge, these methods easily suffer from negative transfer.

In this work, a novel mean teacher-based SSL framework is presented for 3D LA segmentation, in which, we encourage the prediction consistency at different scales between the student and teacher networks. Concretely, we first extend the backbone framework to generate multi-scale predictions. Based on that, we inject multi-level deep supervision to different layers with labeled data to optimize the student network thereby producing more reliable targets. Meanwhile, we enforce the prediction consistency between student and teacher networks hierarchically to further regularize the student network with unlabeled data. Consequently, the proposed framework can efficiently take advantage of both labeled and unlabeled data for accurate SSL LA segmentation. Our framework is extensively evaluated on MICCAI 2018 Atrial Segmentation Challenge dataset [18]. Experimental results demonstrate that the proposed approach can effectively leverage unlabeled data for LA segmentation, and outperform many state-of-the-art SSL techniques.

## II. RELATED WORK

Semi-supervised learning has achieved promising performance on medical image segmentation under the scarcity of labeled data. Bai *et al.* [7] involve a self-training process for cardiac MR image segmentation, where model updates and pseudo annotations are performed alternatively. Meanwhile, self-ensembling methods, particularly the mean teacher (MT) model [11], have received much attention on semi-supervised image segmentation. Specifically, the mean teacher (MT) structure [11] forces the consistency of the predictions with inputs under different perturbations between the student and teacher models and further boosts the model performance. Following that, Yu *et al.* [12] utilize uncertainty estimations with Monte Carlo sampling [15] to generate uncertainty maps for more reliable predictions. Beyond the prediction consistency, Li *et al.* [10] construct the transformation consistency into the self-ensembling model (TCSM) to improve the model generalization ability. Hang *et al.* [13] design the local

\* Contributed equally. † Corresponding author. This research is supported by Institute for Infocomm Research (I2R), Agency for Science, Technology and Research (A\*STAR), Singapore. <sup>1</sup> Institute for Infocomm Research (I2R), Agency for Science, Technology and Research (A\*STAR), Singapore. <sup>2</sup> Nanyang Technological University, Singapore. This work was done when Shumeng Li was an intern at I2R, A\*STAR.

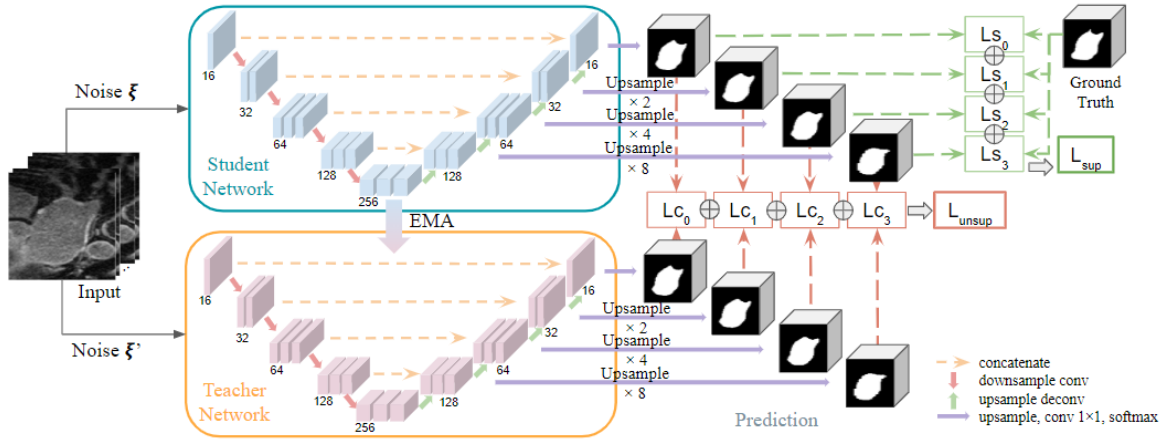


Fig. 1. Overview of our proposed framework. The student network is trained by minimizing  $L_{sup}$  weighted by the hierarchical supervised loss and  $L_{unsup}$  weighted by the hierarchical consistency loss.

and global structural consistency to extract structural information for semi-supervised LA segmentation. Li *et al.* [16] impose geometric constraints via an adversarial loss for shape-aware segmentation. Considering the consistency of different tasks, Luo *et al.* [17] apply a multi-task framework for segmentation and level set function regression together, enforcing the representation consistency at task level.

Despite the great success of semi-supervised segmentation, these methods suffer from high computational complexity, which increases the running time. It is observed that the consistency of existing methods is only applied to the final outputs of each voxel, leaving the abundant information from hidden layers unexplored. On the other hand, many works [19], [20] obtain multi-scale predictions from different layers for boosting the network performance. Inspired by this, we inject deep supervision into hidden layers for accelerating the convergence of model optimization, thereby improving the prediction accuracy. Moreover, we expect to leverage predictions at different scales to further regularize the mean teacher model hierarchically, while obtaining more reliable outputs via multi-level deep supervision with minimal extra computational cost.

### III. METHODOLOGY

Fig. 1 illustrates the proposed mean teacher-based SSL framework. Our method follows the basic spirit of the mean teacher and takes 3D medical images under different perturbations as inputs and encourages the consistency of the segmentation predictions. The student network and the teacher network utilize the same backbone structure. We introduce additional prediction layers to estimate the quality of the hierarchical hidden representations. On one hand, deeply supervised learning with labeled data is performed by minimizing the multi-level segmentation loss from the different hidden layers and the final layer. On the other hand, the multi-level segmentation prediction maps of the teacher network are taken as the learning goal of the student network, which is regularized by hierarchical consistency to effectively take advantage of unlabeled data.

#### A. Multi-scale Deeply supervised Network

In our framework, 3D V-Net [21] is employed as the backbone. To extract hierarchical hidden representations, we add auxiliary layers after each block of the decoding stage to form a hierarchical supervised component (HS). Specifically, the auxiliary layers consist of an upsampling layer, a  $1 \times 1 \times 1$  convolution, and a softmax layer to obtain the multi-scale predictions. For the labeled dataset  $D_L = \{(x_i, y_i), i = 1, \dots, L\}$ , we consider the multi-scale predictions into model optimization for deeply supervised learning. We calculate the loss between predictions at different scales and ground truths respectively to obtain multi-scale deep supervision. Formally, the hierarchical supervised segmentation loss function is defined as:

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

where  $\mathcal{L}_{ce}$  and  $\mathcal{L}_{dice}$  are cross-entropy loss and dice loss, respectively.  $f^s(\cdot)$  denotes the segmentation prediction of HS at scale  $s$ .  $\alpha_s$  is the weights of HS component at scale  $s$  in the hierarchical network.

#### B. Hierarchical Consistency Regularized Mean Teacher

In SSL task, the training dataset consists of the labeled data  $D_L = \{(x_i, y_i), i = 1, \dots, L\}$  and the unlabeled data  $D_U = \{x_i, i = L + 1, \dots, L + U\}$ , where  $x_i \in \mathbb{R}^{H \times W \times D}$  is the input 3D volume and  $y_i \in \{0, 1\}^{H \times W \times D}$  is the 3D ground truth.  $L$  and  $U$  denote the number of 3D volumes of the labeled dataset and unlabeled dataset, respectively.

Following the design of the mean teacher [12], the student network and the teacher network share the same network backbone, and the teacher network weights  $\theta'$  are updated with the exponential moving average (EMA) of the student network weights  $\theta$ . The teacher network weights update strategy at step  $t$  is expressed as  $\theta'_t = \eta \theta'_{t-1} + (1 - \eta) \theta_t$ , where  $\eta$  is to control the EMA update rate, weighing the influence of the teacher network by the weights of current student network and the teacher network at step  $t - 1$ .

For the unsupervised dataset, we expect the segmentation results of the student network and teacher network to be consistent. To further enforce the consistency between the two networks, we design a hierarchical unsupervised consistency component (HU) to encourage the prediction consistent in the hidden feature space. In detail, based on the multi-scale feature maps generated from the student and teacher networks, the hierarchical consistency regularization is introduced. The regularization effectively utilizes unlabeled data, and the consistency constraint between the student network and the teacher network is determined as the mean squared error (MSE) loss of multi-scale predictions between them, which is denoted as:

$$\mathcal{L}_{unsup} = \sum_{s=1}^S \alpha_s \cdot \|f^s(x_i; \theta', \xi') - f^s(x_i; \theta, \xi)\|^2, \quad (2)$$

where  $(\theta, \xi)$  and  $(\theta', \xi')$  denote the different weights and input data with different perturbations of student and teacher networks, respectively.

The unsupervised regularization loss can be combined with the deeply supervised loss to optimize the network together. The training objective can be formulated as:

$$\min_{\theta} \sum_{i=1}^L \mathcal{L}_{sup}(x_i) + \lambda(t) \sum_{i=1}^{L+U} \mathcal{L}_{unsup}(x_i), \quad (3)$$

where  $\mathcal{L}_{sup}$  is the supervised term and  $\mathcal{L}_{unsup}$  is the unsupervised consistency regularization term, respectively.  $\lambda(t)$  is a time-dependent Gaussian weighting function that controls the supervised loss and the regularization term, which is defined as  $\lambda(t) = 0.1 \cdot e^{-5(1 - \frac{t}{t_{max}})^2}$ , where  $t$  and  $t_{max}$  denote the current and maximum training steps, respectively.

## IV. EXPERIMENTS

### A. Dataset and Experimental Settings

We evaluate our proposed semi-supervised method on the dataset of 2018 Atrial Segmentation Challenge for left atrium segmentation in 3D gadolinium-enhanced MR image scans (GE-MRIs) [18]. The dataset consists of 100 scans with segmentation masks and the resolution of the dataset is  $0.625 \times 0.625 \times 0.625 \text{ mm}^3$ . Following the previous work [12], [16], [13], [14], [17], we select 80 samples for training and 20 for testing and use the same data preprocessing methods for a fair comparison.

The proposed architecture is implemented in Pytorch and updated by the SGD optimizer for 6000 iterations in total by setting the initial learning rate of 0.01 and a stepped decay by 0.1 every 2500 iterations in our experiments. Hierarchical classifier mechanism is introduced to the V-Net backbone and the number of scales  $S$  is 4. The scale-wise weights  $\{\alpha_s, s = 0, \dots, 3\}$  for supervised loss and unsupervised consistency loss are specified as  $\{0.5, 0.4, 0.05, 0.05\}$ , empirically. Following Tarvainen and Valpola [11], the decay parameter  $\eta$  of exponential moving average (EMA) is 0.99. We set the batch size to 4 and each batch consists of two labeled volumes and two unlabeled volumes. For evaluation,

TABLE I  
SEGMENTATION RESULTS OF DIFFERENT APPROACHES.

Method	Scans Used	Metrics (%) $\uparrow$		Metrics (Voxel) $\downarrow$	
	L / U	Dice	Jaccard	ASD	95HD
V-Net	80 / 0	91.14	83.32	1.52	5.75
V-Net	16 / 0	86.03	76.06	3.51	14.26
ASDNet [8]	16 / 64	87.90	78.85	2.08	9.24
TCSM [10]	16 / 64	88.15	79.20	2.44	9.57
MT [11]	16 / 64	88.12	79.03	2.65	10.92
UA-MT [12]	16 / 64	88.88	80.21	2.26	7.32
SASSNet [16]	16 / 64	89.27	80.82	3.13	8.83
DTC [17]	16 / 64	89.42	80.98	2.10	7.32
LG-ER-MT [13]	16 / 64	89.62	81.31	2.06	7.16
DUWM [14]	16 / 64	89.65	81.35	<b>2.03</b>	7.04
<b>Ours</b>	16 / 64	<b>90.04</b>	<b>81.98</b>	2.18	<b>6.93</b>

four metrics employed to quantitatively evaluate the segmentation performance are Dice, Jaccard, the average surface distance (ASD), and the 95% Hausdorff Distance (95HD).

### B. Results and Discussions




Fig. 2. Example of 2D and 3D results of five semi-supervised segmentation methods, where GT denotes corresponding ground truth labels.

In our experiments, we use 16 out of 80 training samples as labeled dataset  $D_L$  and the rest 64 samples are served as unlabeled dataset  $D_U$ . The results in Table I compare our proposed approach and several other previous image segmentation methods. It is observed that there are significant improvements in segmentation performance for semi-supervised methods when compared to the supervised-only method (V-Net) for the same amount of labeled data. In addition, the results show that our method achieved up to 90.04% in Dice score. It is noticed that our approach is superior to other semi-supervised methods and obtains comparable segmentation results using only 1/5 labeled data as compared with full annotation.

Figure 2 exemplifies the visualization results of UA-MT [12], SASSNet [16], LG-ER-MT [13], DTC [17], our network and ground truth segmentation. It is observed that our method can obtain more reliable predictions which are closer to the ground truth segmentation labels. Remarkably, through our method, a more complete shape of the left atrium and more details are predicted, especially in the boundary area. This further demonstrates the effectiveness of our proposed strategy for 3D left atrium image segmentation.

TABLE II

EFFECTIVENESS OF CORE COMPONENT IN THE PROPOSED FRAMEWORK, WHERE HS AND HU RESPECTIVELY DENOTE THE HIERARCHICAL SUPERVISED TERM AND UNSUPERVISED CONSISTENCY TERM.

Method	Scans Used	Metrics (%) $\uparrow$		Metrics (voxel) $\downarrow$	
	L / U	Dice	Jaccard	ASD	95HD
V-Net	16 / 0	86.03	76.06	3.51	14.26
V-Net + HS	16 / 0	86.25	76.45	2.89	11.07
MT	16 / 64	88.12	79.03	2.65	10.92
MT + HU	16 / 64	88.16	79.09	2.95	10.62
MT + HS	16 / 64	89.32	80.87	2.70	8.73
MT + HU + HS	16 / 64	<b>90.04</b>	<b>81.98</b>	<b>2.18</b>	<b>6.93</b>

Furthermore, ablation studies are adopted to evaluate the effectiveness of the hierarchical supervised component (HS) and the hierarchical unsupervised consistency component (HU). We list the results in Table II. Compared with the V-Net trained on labeled data only, multi-scale deep supervision, *i.e.*, V-Net + HS can efficiently improve the segmentation performance. And it can also be employed for SSL to better utilize the labeled data. Then we introduce HU and HS to the mean teacher model, individually, *i.e.*, MT + HU and MT + HS, respectively. The results demonstrate the effectiveness of the two components. Lastly, combining two components together can achieve better performance.

## V. CONCLUSION

In a nutshell, a novel and simple approach is proposed for semi-supervised 3D medical image segmentation. We present a framework based on the mean teacher architecture and it encourages the consistency of predictions between the student and the teacher networks with the same input under different perturbations. We perform the hierarchical prediction consistency based on hidden feature space to regularize the mean teacher model while obtaining more reliable outputs via multi-scale deep supervision. Experimental results on the Atrial Segmentation Challenge dataset demonstrate the superiority of our framework over the other state-of-the-art semi-supervised learning methods. Besides, the proposed method is simple and versatile, which can be widely applied to other segmentation tasks.

## REFERENCES

- [1] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen Awm Van Der Laak, Bram Van Ginneken, and Clara I Sánchez, "A survey on deep learning in medical image analysis," *Medical image analysis*, vol. 42, pp. 60–88, 2017.
- [2] Roberto M Lang, Luigi P Badano, Victor Mor-Avi, Jonathan Afilalo, Anderson Armstrong, Laura Ernande, Frank A Flachskampf, Elyse Foster, Steven A Goldstein, Tatiana Kuznetsova, et al., "Recommendations for cardiac chamber quantification by echocardiography in adults: an update from the american society of echocardiography and the european association of cardiovascular imaging," *European Heart Journal-Cardiovascular Imaging*, vol. 16, no. 3, pp. 233–271, 2015.
- [3] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [4] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi, "V-net: Fully convolutional neural networks for volumetric medical image segmentation," in *2016 fourth international conference on 3D vision (3DV)*. IEEE, 2016, pp. 565–571.

- [5] Veronika Cheplygina, Marleen de Bruijne, and Josien PW Pluim, "Not-so-supervised: a survey of semi-supervised, multi-instance, and transfer learning in medical image analysis," *Medical image analysis*, vol. 54, pp. 280–296, 2019.
- [6] Nima Tajbakhsh, Laura Jeyaseelan, Qian Li, Jeffrey N Chiang, Zhihao Wu, and Xiaowei Ding, "Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation," *Medical Image Analysis*, vol. 63, pp. 101693, 2020.
- [7] Wenjia Bai, Ozan Oktay, Matthew Sinclair, Hideaki Suzuki, Martin Rajchl, Giacomo Tarroni, Ben Glocker, Andrew King, Paul M Matthews, and Daniel Rueckert, "Semi-supervised learning for network-based cardiac mr image segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2017, pp. 253–260.
- [8] Dong Nie, Yaozong Gao, Li Wang, and Dinggang Shen, "Asdnet: attention based semi-supervised deep networks for medical image segmentation," in *International conference on medical image computing and computer-assisted intervention*. Springer, 2018, pp. 370–378.
- [9] Ziyuan Zhao, Xiaoman Zhang, Cen Chen, Wei Li, Songyou Peng, Jie Wang, Xulei Yang, Le Zhang, and Zeng Zeng, "Semi-supervised self-taught deep learning for finger bones segmentation," in *2019 IEEE EMBS International Conference on Biomedical Health Informatics (BHI)*, 2019, pp. 1–4.
- [10] Xiaomeng Li, Lequan Yu, Hao Chen, Chi-Wing Fu, Lei Xing, and Pheng-Ann Heng, "Transformation-consistent self-ensembling model for semisupervised medical image segmentation," *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [11] Antti Tarvainen and Harri Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2017, NIPS'17, p. 1195–1204.
- [12] Lequan Yu, Shujun Wang, Xiaomeng Li, Chi-Wing Fu, and Pheng-Ann Heng, "Uncertainty-aware self-ensembling model for semi-supervised 3d left atrium segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2019, pp. 605–613.
- [13] Wenlong Hang, Wei Feng, Shuang Liang, Lequan Yu, Qiong Wang, Kup-Sze Choi, and Jing Qin, "Local and global structure-aware entropy regularized mean teacher model for 3d left atrium segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2020, pp. 562–571.
- [14] Yixin Wang, Yao Zhang, Jiang Tian, Cheng Zhong, Zhongchao Shi, Yang Zhang, and Zhiqiang He, "Double-uncertainty weighted method for semi-supervised learning," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2020, pp. 542–551.
- [15] Alex Kendall and Yarin Gal, "What uncertainties do we need in bayesian deep learning? for computer vision?," in *Advances in Neural Information Processing Systems*, 2017, vol. 30.
- [16] Shuailin Li, Chuyu Zhang, and Xuming He, "Shape-aware semi-supervised 3d semantic segmentation for medical images," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2020, pp. 552–561.
- [17] Xiangde Luo, Jieneng Chen, Tao Song, and Guotai Wang, "Semi-supervised medical image segmentation through dual-task consistency," *AAAI Conference on Artificial Intelligence*, 2021.
- [18] Zhaohan Xiong, Qing Xia, Zhiqiang Hu, Ning Huang, Cheng Bian, Yefeng Zheng, Sulaiman Vesal, Nishant Ravikumar, Andreas Maier, Xin Yang, et al., "A global benchmark of algorithms for segmenting the left atrium from late gadolinium-enhanced cardiac magnetic resonance imaging," *Medical Image Analysis*, vol. 67, pp. 101832, 2021.
- [19] Qi Dou, Lequan Yu, Hao Chen, Yueming Jin, Xin Yang, Jing Qin, and Pheng-Ann Heng, "3d deeply supervised network for automated segmentation of volumetric medical images," *Medical image analysis*, vol. 41, pp. 40–54, 2017.
- [20] Ziyuan Zhao, Zeng Zeng, Kaixin Xu, Cen Chen, and Cuntai Guan, "Dsal: Deeply supervised active learning from strong and weak labelers for biomedical image segmentation," *IEEE Journal of Biomedical and Health Informatics*, pp. 1–1, 2021.
- [21] F Milletari, N Navab, SAV Ahmadi, and V-net, "Fully convolutional neural networks for volumetric medical image segmentation," in *Proceedings of the 2016 Fourth International Conference on 3D Vision (3DV)*, 2016, pp. 565–571.