

Orthogonal Annotation Benefits Barely-supervised Medical Image Segmentation

Heng Cai¹, Shumeng Li¹, Lei Qi², Qian Yu³, Yinghuan Shi^{1,*}, Yang Gao¹
¹Nanjing University ²Southeast University ³Shandong Women's University

Abstract

Recent trends in semi-supervised learning have significantly boosted the performance of 3D semi-supervised medical image segmentation. Compared with 2D images, 3D medical volumes involve information from different directions, e.g., transverse, sagittal, and coronal planes, so as to naturally provide complementary views. These complementary views and the intrinsic similarity among adjacent 3D slices inspire us to develop a novel annotation way and its corresponding semi-supervised model for effective segmentation. Specifically, we firstly propose the orthogonal annotation by only labeling two orthogonal slices in a labeled volume, which significantly relieves the burden of annotation. Then, we perform registration to obtain the initial pseudo labels for sparsely labeled volumes. Subsequently, by introducing unlabeled volumes, we propose a dual-network paradigm named Dense-Sparse Co-training (DeSCO) that exploits dense pseudo labels in early stage and sparse labels in later stage and meanwhile forces consistent output of two networks. Experimental results on three benchmark datasets validated our effectiveness in performance and efficiency in annotation. For example, with only 10 annotated slices, our method reaches a Dice up to 86.93% on KiTS19 dataset. Our code and models are available at <https://github.com/HengCai-NJU/DeSCO>.

1. Introduction

Medical image segmentation is one of the most critical vision tasks in medical image analysis field. Thanks to the development of deep learning-based methods [8, 11, 28, 32], segmentation performance has now been substantially improved. However, the current promising performance is at

*Corresponding author: Yinghuan Shi. Heng Cai, Shumeng Li, Yinghuan Shi and Yang Gao are with the State Key Laboratory for Novel Software Technology and National Institute of Healthcare Data Science, Nanjing University, China. This work is supported by the NSFC Program (62222604, 62206052, 62192783), CAAI-Huawei Mind-Spore (CAAI-XSJLJJ-2021-042A), China Postdoctoral Science Foundation Project (2021M690609), Jiangsu Natural Science Foundation Project (BK20210224), and CCF-Lenovo Bule Ocean Research Fund.

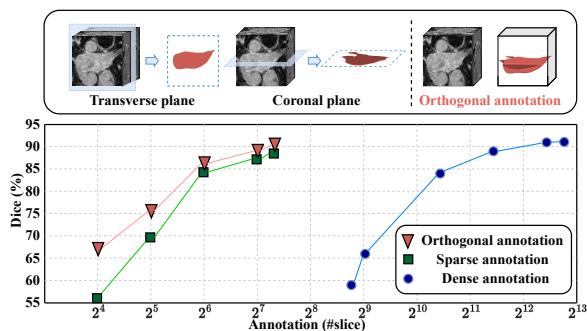


Figure 1. The upper figure illustrates our annotation method, each volume with annotations is labeled with only two orthogonal slices. The lower figure shows the comparison between the efficiency and effectiveness of our orthogonal annotation and other manners, including conventional dense annotation and previous sparse annotation which labels slices in one plane. All trained on LA [42] dataset with supervised setting. For sparse annotation and our orthogonal annotation, we train the models only on labeled voxels through partial cross-entropy and partial Dice loss.

the cost of large-scale manually precisely labeled dataset, which is prohibitively expensive and laborious to achieve. What's worse, different radiologists might provide different annotations even for a same image. Therefore, exploring ways to alleviate the requirement of quantity or quality of manual annotation is highly demanded. Mainstream methods typically follow two paradigms: 1) degrade annotation quality, *i.e.*, weakly-supervised segmentation, and 2) reduce annotation quantity, *i.e.*, semi-supervised segmentation.

Weakly-supervised segmentation methods usually utilize weak annotations, *e.g.*, image-level label [16, 17], scribble [20, 21], point [3] or partial slices [5, 18]. Unfortunately, most of them are either difficult to distinguish some fuzzy boundaries or with additional large computational burden [15]. What's more, weakly-supervised setting usually requires coarse annotation for every single image. This is still a heavy burden for radiologists. Besides, most current methods originally developed for 2D segmentation could not directly utilize 3D volumetric information [9].

Different from these weakly-supervised methods, semi-supervised methods train segmentation models with a small amount of manually labeled data and a large amount of unlabeled data, which have achieved remarkable perfor-

mance with an impressive deduction on demand for annotation [6, 19]. Despite their success, we notice most current semi-supervised segmentation methods still require full 3D annotation for each labeled volume. In fact, segmentation targets in adjacent slices of 3D volume are highly similar in both appearance and location, leading it redundant to label every slice. Although the sparse annotation is discussed in recent work [18], we notice these conventional methods still neglect the complementary views between different directions in 3D volume.

It is known that 3D medical volumes naturally contains different directions (e.g., transverse, coronal planes) which provide complementary information from different views. And recent trends in semi-supervised learning [7, 40] have revealed that learning from complementary view is indeed beneficial. Thus, we wonder *whether a novel annotation method coupled with its corresponding model could be investigated by introducing this complementary relation into 3D semi-supervised medical image segmentation.*

In this paper, for labeled volume, we innovatively investigate a novel sparse annotation way—*orthogonal annotation, i.e., only to label two slices in its orthogonal direction (e.g., transverse and coronal direction in Figure 1).* We believe our annotation way has two merits: 1) it could largely force the model to learn from complementary views with two diversely initialized labeled slices, 2) it helps greatly reduce the label costs with fully utilizing the inter-slice similarity. Following very recent work [18], we name the setting as Barely-supervised Segmentation.

To incorporate our orthogonal annotation, the most intuitive thought about training strategy of a segmentation model is that only the voxels on the labeled slices contribute to the training. However, directly learning from this sparse annotation is unstable and the training is apt to collapse (shown in Sec. 4). Thus, we apply registration to spread supervision signals from slice to volume, where the result of label propagation can serve as the dense pseudo label for training. By performing registration, we obtain two sets of pseudo labels for volumes from orthogonal directions. Yet, the obtained pseudo labels are not promising enough to directly train a segmentation model using current existing semi-supervised methods, which is mainly due to the accumulation of error in the registration process.

Therefore, to leverage 1) the volumes with inaccurate pseudo labels and 2) the rest unlabeled volumes, we propose a simple yet effective end-to-end framework namely Dense-Sparse Co-training (DeSCO), which consists two segmentation models of a same structure. At the beginning of training, the models mainly learn from dense pseudo labels with a learning preference on voxels with more confident pseudo labels, *i.e., voxels near to registration source slice*, and exploit unlabeled volumes through cross-supervision. After the models have been improved through training, we

gradually get rid of pseudo label until the supervised loss solely comes from sparse annotation. Meanwhile, the role of cross-supervision is gradually emphasized correspondingly. Because in the process of reaching consensus through cross-supervision, the mistake introduced by previous training on inaccurate pseudo labels could be revised. Overall, our contributions are three folds:

- A new annotation way that only labels two orthogonal slices for a labeled 3D volume, which greatly reduces the annotation burden.
- A novel barely-supervised 3D medical image segmentation framework to steadily utilize our high-efficient sparse annotation with coupled segmentation method.
- A dense-sparse co-training paradigm to learn from dense pseudo label and sparse label while leveraging unlabeled volumes to reduce noise by reaching consensus through cross-supervision.

Extensive experiments on three public datasets validate that our barely-supervised method is close to or even better than its upper bound, *i.e., semi-supervised methods with fully annotated labeled volumes.* For example, on KiTS19, compared to Mean Teacher [36] that uses 320 labeled slices with a Dice of 84.98%, we only uses 10 labeled slices yet obtains a Dice of 86.93%.

2. Related Work

Semi-supervised Learning. Semi-supervised learning methods leverage a few labeled data and abundant unlabeled data to reduce the dependency on laborious annotation, and show their effectiveness on various tasks [22, 27, 45]. For example, Π -model [33] minimized the inconsistency between the outputs of differently disturbed inputs in order to improve the consistency of the model. Temporal Ensembling [33] and Mean Teacher [36] improved Π -model by leveraging the average of historic outputs to produce pseudo labels of higher quality. FixMatch [35] and FlexMatch [44] selected pseudo labels of higher quality with threshold. Some methods explored uncertainty by using novel consistency regularization, *e.g.,* UPS [31] and NP-Match [38]. The consistency regularization way used in these methods has greatly inspired current mainstream semi-supervised medical image segmentation methods.

Semi-supervised Medical Image Segmentation. To free radiologists from the burden of heavy annotation, many deep learning-based semi-supervised methods have been proposed for medical image segmentation in recent years.

Consistency regularization is a popular method to leverage unlabeled data. Li *et al.* [19] made a consistency constraint between the outputs of original images and transformed images. Similarly, Bortsova *et al.* [6] proposed a

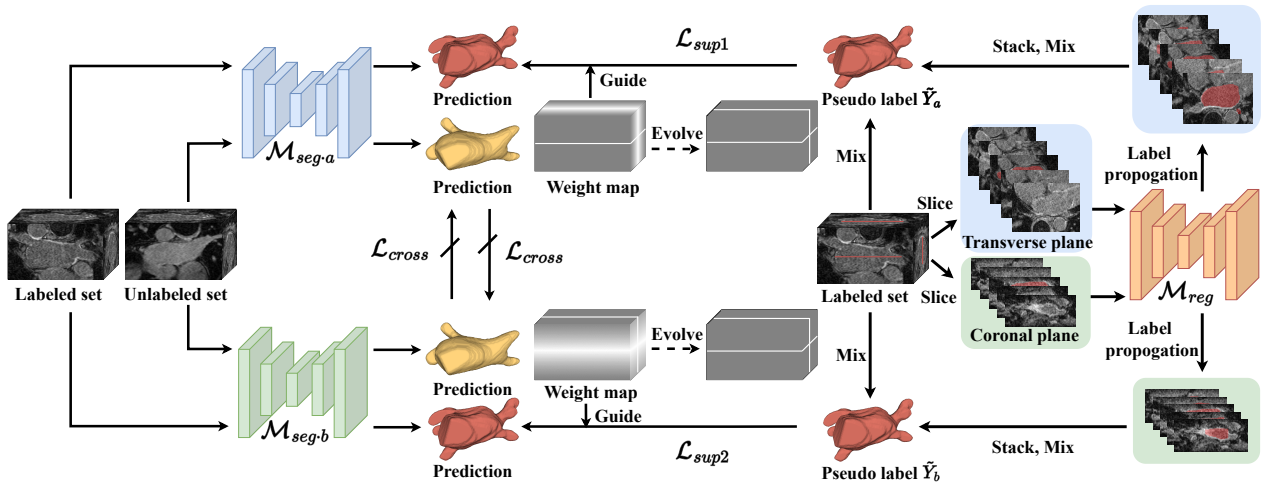


Figure 2. Overview of the proposed DeSCO paradigm. For a volume with orthogonal annotation, \mathcal{M}_{reg} propagates the orthogonal annotation into the whole volume in two directions, and the results serve as pseudo labels to supervise segmentation model \mathcal{M}_{seg-a} and \mathcal{M}_{seg-b} , respectively. For an unlabeled volume, \mathcal{M}_{seg-a} and \mathcal{M}_{seg-b} supervise each other with their outputs. Weight map guides \mathcal{L}_{sup} and whiter areas mean higher weight.

method that learns to predict consistently under different transformations. Xie *et al.* [41] enforced the consistency of images in feature space. Different from previous work, Luo *et al.* [24] imposed the consistency constraint at task level. Shi *et al.* [34] incorporated consistency regularization and entropy minimization by treating certain/uncertain region with different training strategies. Besides, Co-training is another widely-used framework for semi-supervised medical image segmentation. Zhou *et al.* [46] applied multi-planar fusion to generate more accurate 3D pseudo label to train a co-training model. Xia *et al.* [40] introduced multi-view co-training into semi-supervised medical image segmentation and domain adaptation. And some works [2, 43] improving the quality of pseudo labels also showed great effectiveness. However, when applied to 3D segmentation tasks, as a heavy burden, most of them still require a dozen pixel-wise fully annotation slice-by-slice.

Weakly-supervised Segmentation. Weakly-supervised segmentation methods give pixel-level output with coarse annotation including image-level annotation [16, 17, 39], bounding box [12, 29], scribble [20, 21, 37] and even some points [3]. Among them, image-level annotation is the most popular setting and has been extensively studied. Many image-level weakly-supervised segmentation methods are based on class activation map (CAM). However, the area suggested by CAM is incapable of directly training a segmentation model as it only covers the most discriminative part of large objects and is liable to over-covers small objects. Other annotation methods suffer from the same problem more or less. The common problem in these annotation methods is that they discard the information of the boundaries between objects, which is crucial for segmentation tasks. Our annotation method takes into consideration

the precise delineation of objects and utilizes volumetric information of 3D medical images.

Remark. Our setting could be regarded as a variant of semi-supervised segmentation by requiring much fewer annotated slices. Our proposed orthogonal annotation has its merits in preserving disparity to boost the performance in semi-supervised scenario. Also, compared with weakly-supervised segmentation, our method only requires very few annotated slices while yielding promising results.

3. Method

In this work, we propose a novel annotation method called *orthogonal annotation* and a coupled training method comprised of 1) a registration module and 2) DesCO segmentation model. The training procedure is illustrated in Figure 2. We provide detailed introductions to the problem setting, our orthogonal annotation, registration module and the proposed DeSCO paradigm in the following parts.

3.1. Problem Setting and Notations

We aim to apply our orthogonal annotation in barely-supervised segmentation setting, *e.g.*, only 10 orthogonal slices from 5 volumes for an entire 3D medical image dataset. Training set contains N volumes $\{X_1, X_2, \dots, X_N\}$ of shape $H \times W \times D$, where H, W and D represent the height, width and depth of each volume, respectively. These volumes are further divided into $\{X_1, X_2, \dots, X_l\}$ with annotations and $\{X_{l+1}, X_{l+2}, \dots, X_N\}$ without annotations. Our methods annotates two slices from orthogonal planes denoted by a and b . The labeled slices of $X_i (1 \leq i \leq l)$ are denoted as X_{ia}^m and X_{ib}^n , which means m^{th} slice in plane a and n^{th} slice in plane b . Their annotations are denoted as Y_{ia}^m and Y_{ib}^n .

3.2. Justification for Orthogonal Annotation

Our proposed orthogonal annotation (OA in short), which only labels two slices in orthogonal directions for a 3D volume, is highly efficient in annotation. Additionally, we now validate its effectiveness empirically and experimentally, which includes the discussion about 1) disparity, 2) feature representation, and 3) initial performance. We here take LA dataset [42] as an example, which contains 100 3D volumes detailed in Sec. 4.1.

OA preserves the disparity. As revealed by previous theoretical studies in semi-supervised learning [30], disparity (or disagreement) between views (or models) providing complementary information is the key to success. Now we explore the similarity of parallel slices and orthogonal slices from both image- and feature-level. The slices are chosen at random and the features are extracted from V-Net [28] trained with fully annotated data. We utilize the popularly-used Hilbert-Schmidt independence criterion (HSIC) [26] as the evaluation metric. Smaller value of HSIC, more independence of two variables. The results are shown in Figure 3, which reveal that parallel slices are more relevant than orthogonal slices in terms of both image and feature. It implies that parallel annotations have more redundancy of information. Besides, slices from different planes are inherently varied and complementary. Thus, for the same amount of annotation, orthogonal annotation could be superior to parallel annotation.

OA enhances consistency from different directions. We now visualize the features under different settings. Here are three models trained with only one labeled slice in each volume. The labeled slices are from transverse plane for models t_1 and t_2 and from coronal plane for model c . In other words, labeled slices used to train models t_1 and t_2 are parallel and those used to train models t_1 and c are orthogonal.

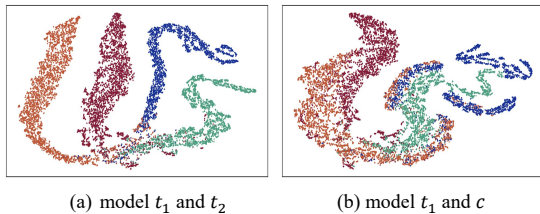


Figure 4. t-SNE result of feature for classification.

We extract the features before the classification layer of three models and use t-SNE to visualize them. The results are shown in Figure 4. Red and orange represent positive class while blue and aqua represent negative class. In par-

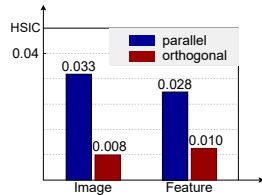


Figure 3. HSIC of parallel slices and orthogonal slices.

allel annotation, even the labeled slices come from a same plane, they learn different sub-patterns to distinguish foreground from background. Differently, the labeled slices to train t_1 and c share some common voxels (i.e., overlapping part) so that the features/predictions for the same class are closer. This result validates, by using orthogonal planes, different classes could have more consistent representation.

OA provides a promising initialization. To further explore the utility of the orthogonal annotation, we compare it to the parallel annotation in each plane. We trained models with three labeling manners: 1) one transverse plane slice and one coronal plane slice 2) two transverse plane slices 3) two coronal plane slices.

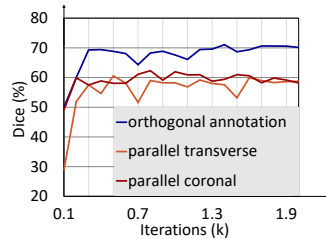


Figure 5. Dice coefficient on test set of models trained with different annotations.

The testing result is shown in Figure 5. The model trained with orthogonal annotations steadily and vastly outperform the other two models in its initial stage. It shows that simply leveraging orthogonal plane slices can promote the

performance greatly and this fact has always been neglected in medical image segmentation field.

These observations have revealed the superiority of our orthogonal annotation. More detailed settings of these experiments are provided in the supplementary material.

3.3. Registration Module

Registration seeks a spatial transformation to map an image to another image. For two slices: labeled slice (X_l, Y_l) and unlabeled slice X_u , registration obtains a spatial deformation field $\Phi(\cdot)$ by mapping from X_l to X_u . And pseudo label of X_u can be acquired by applying the same spatial deformation field to label Y_l , which can be formulated as $\hat{Y}_u = \Phi(Y_l)$.

As the segmentation targets in adjacent slices are usually similar in shape and size, it is feasible to propagate the label of a slice to its neighbouring slices through a registration module \mathcal{M}_{reg} . Note that registration module is an off-the-shelf tool that requires no training, so it introduces little extra computational cost except generating pseudo labels for only once.

For volume X with two orthogonal annotated slices (X_a^m, Y_a^m) and (X_b^n, Y_b^n) , we use the two slices to perform slice-by-slice label propagation and gradually generate two pseudo labels \hat{Y}_a and \hat{Y}_b for X through registration module \mathcal{M}_{reg} . First, we calculate spatial transformation fields that maps X_a^m to its closest slices X_a^{m-1} and X_a^{m+1} through \mathcal{M}_{reg} . Then the spatial transformation fields are applied to label Y_a^m to acquire pseudo labels \hat{Y}_a^{m-1} and

\hat{Y}_a^{m+1} , from which we further derive \hat{Y}_a^{m-2} and \hat{Y}_a^{m+2} . Through the same procedure, we obtain pseudo labels for all slices viewed in plane a , and concatenate them to compose pseudo label \tilde{Y}_a for volume X . Similarly, the pseudo label \tilde{Y}_b can be easily obtained. So here we have pseudo label sets $\{\tilde{Y}_{1a}, \tilde{Y}_{2a}, \dots, \tilde{Y}_{la}\}$ and $\{\tilde{Y}_{1b}, \tilde{Y}_{2b}, \dots, \tilde{Y}_{lb}\}$ for those volumes with orthogonal annotation.

Performing label propagation suffers error accumulation as the pseudo label of each unlabeled slice is derived from its nearest pseudo label. Though error propagation can be greatly mitigated through morphology operations, it cannot be completely eliminated due to the limitation of registration itself. To alleviate the problem, we consider the importance of each slice in terms of their credibility.

3.4. Label Mixing

For each volume X with sparse orthogonal annotation, we have ground truth sparse annotation Y and dense registration pseudo label \hat{Y}_a and \hat{Y}_b . Now we generate pseudo label \tilde{Y}_a and \tilde{Y}_b by mixing up dense registration pseudo label and ground truth sparse annotation:

$$\tilde{Y}_a = \text{LabelMix}(\hat{Y}_a, Y), \quad (1)$$

where $\text{LabelMix}(\cdot, \cdot)$ denotes a function mixing them by replacing the unlabeled voxels in Y with pseudo label in \hat{Y}_a . Similarly, \tilde{Y}_b can be generated.

As mentioned in Sec. 3.3, different slices in \hat{Y} should be considered differently according to the quality of pseudo labels. We assume the error rate of label propagation is a certain value, and the credibility of slice i can be denoted as $w_i = \alpha^d$, where d is the distance from slice i to registration source slice and α ($0 \leq \alpha < 1$) is the decay rate. On these grounds, we set the voxel-wise weight for each voxel in \tilde{Y}_a :

$$W_a^i = \begin{cases} 1 & \text{if voxel } i \text{ is labeled in } Y \\ \alpha^d & \text{otherwise.} \end{cases} \quad (2)$$

Similarly for the voxel weight map of \tilde{Y}_b , defined as W_b .

3.5. Dense-Sparse Co-Training

Our DeSCO paradigm consists of two 3D segmentation networks \mathcal{M}_{seg-a} and \mathcal{M}_{seg-b} of a same structure. It leverages labeled and unlabeled volumes simultaneously in a mini-batch. As mentioned above, every volume X_i ($i \leq l$) has two pseudo labels, \hat{Y}_{ia} and \hat{Y}_{ib} . Segmentation network \mathcal{M}_{seg-a} is trained with \hat{Y}_{ia} and \mathcal{M}_{seg-b} is trained with \hat{Y}_{ib} , respectively. In this way, the two segmentation networks mainly learn from two different perspectives with the targets obtained from different plane registration, and the disparity in orthogonal annotation is well-preserved. The supervised loss contains weighted cross-entropy loss and weighted dice

loss, which are formulated as:

$$\mathcal{L}_{ce} = -\frac{1}{\sum_{i=1}^{H \times W \times D} w_i} \sum_{i=1}^{H \times W \times D} w_i y_i \log p_i, \quad (3)$$

and

$$\mathcal{L}_{dice} = 1 - \frac{2 \times \sum_{i=1}^{H \times W \times D} w_i p_i y_i}{\sum_{i=1}^{H \times W \times D} w_i (p_i^2 + y_i^2)}, \quad (4)$$

where w_i is i^{th} voxel of weight map W and p_i, y_i denote the probability of foreground and pseudo label of voxel i respectively.

At the early stage of training, segmentation models need to learn from dense pseudo label for steady improvement. Segmentation networks are gradually improved during training and can produce better segmentation results than the initial pseudo labels. Therefore, the initial pseudo labels are actually becoming the obstacle for continual improvement. Based on this fact, we should weaken the influence of pseudo label. To implement this, we gradually decrease the decay rate α , and down to zero at last, which means the networks get rid of pseudo labels and only learn from sparse annotation.

The supervised loss is the weighted sum of cross-entropy loss and dice loss:

$$\mathcal{L}_{sup} = \frac{1}{2} \mathcal{L}_{ce} + \frac{1}{2} \mathcal{L}_{dice}. \quad (5)$$

For those volumes without annotations, the two segmentation models teach each other with their predictions. In a training iteration of volume X_i ($i \geq l+1$), \mathcal{M}_{seg-a} takes X_i as input and generates its one-hot prediction, which will be partially selected as the pseudo label for \mathcal{M}_{seg-b} . Here we follow UAMT [43] by selecting those voxels whose uncertainty is lower than a threshold and produce mask M_{un} for better cross-supervision. Similarly, \mathcal{M}_{seg-b} regards masked one-hot prediction of \mathcal{M}_{seg-a} as the pseudo label. The loss is formulated as:

$$\mathcal{L}_{cross} = -\frac{1}{\sum_{i=1}^{H \times W \times D} m_i} \sum_{i=1}^{H \times W \times D} m_i \hat{y}_i \log p_i, \quad (6)$$

where m_i denotes the value in mask M_{un} , p_i and \hat{y}_i denote the probability of foreground and one-hot pseudo label predicted by two models.

For the late stage of training, the supervised loss mainly comes from the labeled slices and the information have been learnt by segmentation models, so keeping supervised loss weight high is to no avail. Oppositely, the cross-supervision weight should be increased for the two networks can correct the mistakes from noisy pseudo label and reach a consensus through cross-supervision. Thus, the overall objective is the weighted sum of \mathcal{L}_{sup} and \mathcal{L}_{cross} :

$$\mathcal{L} = (1 - \lambda) \mathcal{L}_{sup} + \lambda \mathcal{L}_{cross}, \quad (7)$$

where λ is a dynamic parameter gradually increasing to optimal cross-supervision weight λ_{oc} .

4. Experiments

4.1. Datasets

LA Dataset [42] contains 100 3D gadolinium-enhanced magnetic resonance images required using clinical whole-body MRI scanner and full annotations for left atrial cavity by radiologists. All scans have the same isotropic resolution of resolution of $0.625 \times 0.625 \times 0.625\text{mm}^3$ while their dimensions vary from each other.

KiTS19 Dataset [14] is a benchmark for kidney and kidney tumor segmentation. It contains 300 arterial phase abdominal CT where there are one or more kidney tumors and 210 of them have accessible kidney and kidney tumor delineations. The CTs have an average slice number of 216 and the slices thicknesses range from 1mm to 5mm. Most of slices in the transverse plane have the size of 512×512 .

LiTS Dataset [4] consists of 201 contrast-enhanced abdominal CT scans and corresponding segmentation of liver as well as liver tumor provided by various clinical sites. And 131 of segmentations are open to the public. There are high variations in pixel spacing and slice thickness.

4.2. Experiment Setting

For the selection of annotated slices, the principle is that the segmented target should be visible in the selected slice. For generalization, we simply select a fixed slice roughly located in the middle of the whole target.

For the implementation of our method, we adopt the SyNRA method in Ants [1] as our registration module and V-Net [28] as the backbone of our segmentation models.

For the training of segmentation models, we set the batch size as 2, one for labeled sample and one for unlabeled sample. We use Stochastic Gradient Descent (SGD) with momentum of 0.9 and weight decay of 0.0001 as optimizer. The learning rate is initialized as 0.01 and gradually decays to 0.0001. Training iteration is set to 6000 for all experiments. λ_{oc} is 0.8 for all three datasets. For the slice-wise weight decay rate, we initialize α to 0.95 and update every 1000 iterations according to cosine rampdown from [23].

For the evaluation and comparison of our method and other semi-supervised methods, we adopt four commonly used metrics, which are Dice coefficient, Jaccard coefficient, 95% Hausdorff Distance (95% HD) and Average Surface Distance (ASD). We compare our method with some classic and state-of-the-art (SOTA) methods, including Mean Teacher (MT) [36], Cross Pseudo Supervision (CPS) [10], Cross Teaching Between CNN and Transformer (CTBCT) [25] and Conservative-radical Network (CoraNet) [34]. All methods are implemented in 3D manners. MT, CPS, CoraNet and the CNN network of CTBCT

use V-Net [28] as backbone, the transformer network of CTBCT is implemented as UNETR, which is introduced in [13]. As those methods are not specifically designed for barely-supervised setting, for fairer comparisons, we designed two training settings for those methods. 1) Dense: we provided dense pseudo labels by registration as labeled training data. From the comparison between our method and other methods trained in this setting, we show that our method can better leverage the inaccurate pseudo labels while other methods suffers severe degradation. 2) Sparse: we provide original orthogonal labels without registration. And during training, the supervised loss only comes from labeled voxles. In the comparison between our method and other methods trained in this setting, we show that directly learning from few sparse annotations is almost unattainable, thus, generating initial pseudo labels is necessary and reasonable. All the experiments are implemented using PyTorch and an NVIDIA GeForce RTX 3090 GPU.

4.3. Compared with SOTA Methods

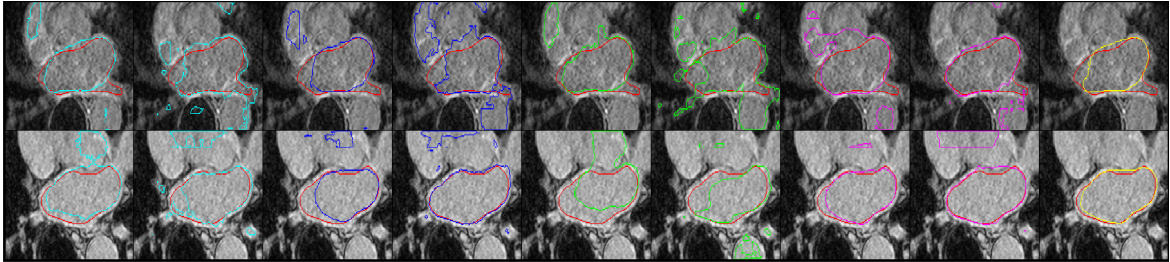
Results on LA Dataset. Following [43], we use 80 volumes for training and 20 for testing. 5 volumes of training data have orthogonal annotations on transverse plane and coronal plane, while the rest of the training data are used as unlabeled volumes. As there's no validation set in the experiments, so we report the results at the end of 6000 iterations of training for all methods. We randomly crop the original size $112 \times 112 \times 88$ into $112 \times 112 \times 80$ for training, and the patch size remains the same during testing. The prediction is generated by sliding the patch window with a stride of 18 on the coronal and sagittal plane and 4 on the transverse plane. The comparison results between our method and other methods are presented as mean \pm std of five cross-validation in Table 1. It is clear that our method outperforms other SOTA semi-supervised methods in all metrics except ASD. Compared to Mean Teacher trained with fully annotated volumes, our method uses only 2.3% annotations but there is only a 1.65% drop on Dice coefficient. We notice that these methods in Sparse setting work poorly, it is mainly because of the variety between slices and the lack of supervision signal on most parts of the whole volume. The segmentation result examples are shown in Figure 6.

Results on KiTS19 Dataset. We divided 210 samples into training set and test set, which have 190 and 20 samples respectively. Similar to LA, 5 volumes are randomly selected to label and the rest 185 volumes are unlabeled data. Patch size is set to $128 \times 128 \times 64$. The results are shown in Table 2. From the table we can see that our method outperforms all methods on all metrics. Note that our method even outperforms MT trained with 32 times the annotated slices. The segmentation result examples are shown in Figure 7.

Results on LiTS Dataset. For LiTS dataset, we also randomly select 5 volumes as labeled data, 95 volumes as

Table 1. Comparison on LA dataset segmentation.

Method	Venue	Setting	Scans Used		Metrics				
			L / U Volumes	Labeled Slices	Dice (%)	Jaccard (%)	HD (voxel)	ASD (voxel)	
Barely-supervised	MT [36]	NIPS'17	Dense	5 / 75	10	74.41±4.07	60.87±4.35	27.87±3.96	5.65±1.12
			Sparse	5 / 75	10	63.17±2.28	47.48±2.38	39.87±4.32	14.75±2.41
	CPS [10]	CVPR'21	Dense	5 / 75	10	74.02±0.41	60.62±1.37	30.32±1.73	5.86±0.92
			Sparse	5 / 75	10	62.24±2.55	46.12±2.87	45.03±2.77	17.83±2.29
	CTBCT [25]	MIDL'22	Dense	5 / 75	10	73.12±5.50	59.37±5.80	26.51±2.47	5.38±1.25
			Sparse	5 / 75	10	64.71±1.49	48.89±1.14	42.09±2.13	15.91±0.83
	CoraNet [34]	TMI'22	Dense	5 / 75	10	40.63±1.34	31.94±2.41	40.88±3.73	20.19±1.57
			Sparse	5 / 75	10	51.16±3.63	38.45±3.30	44.70±1.08	18.46±0.78
Ours	this paper	-	5 / 75	10	77.48±0.58	64.36±1.25	24.86±10.58	6.41±4.68	
Semi-supervised	MT [36]	NIPS'17	-	5 / 75	440	79.13±2.37	66.92±2.82	17.41±1.85	4.76±1.36

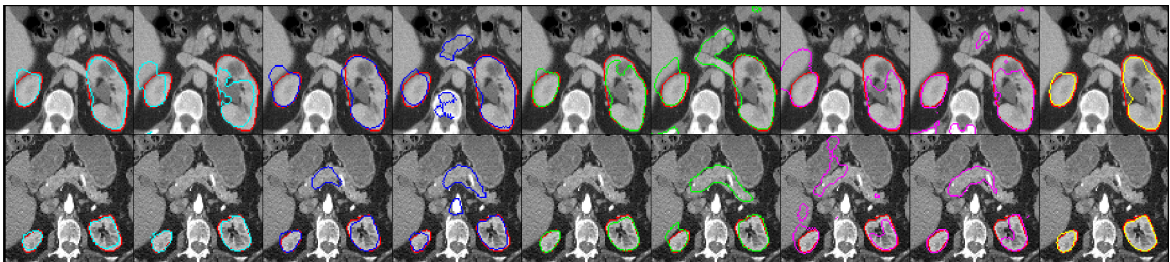


MT_{dense} MT_{sparse} CPS_{dense} CPS_{sparse} CTBCT_{dense} CTBCT_{sparse} CoraNet_{dense} CoraNet_{sparse} Ours

Figure 6. Visual segmentation examples from LA dataset. The red, cyan, blue, green, magenta, yellow curves denote the corresponding results of ground-truth, mean teacher [36], CPS [10], CTBCT [25], CoraNet [34] and ours, respectively.

Table 2. Comparison on KiTS19 dataset segmentation.

Method	Venue	Setting	Scans Used		Metrics				
			L / U Volumes	Labeled Slices	Dice (%)	Jaccard (%)	HD (voxel)	ASD (voxel)	
Barely-supervised	MT [36]	NIPS'17	Dense	5 / 185	10	78.16±1.99	66.96±1.41	17.15±5.93	5.40±2.25
			Sparse	5 / 185	10	32.16±13.78	20.41±11.04	67.42±12.73	32.35±9.83
	CPS [10]	CVPR'21	Dense	5 / 185	10	76.51±5.81	64.67±6.69	17.71±8.67	5.74±3.50
			Sparse	5 / 185	10	8.87±2.87	4.68±1.58	81.60±10.15	41.21±5.38
	CTBCT [25]	MIDL'22	Dense	5 / 185	10	80.65±0.83	69.99±2.27	14.88±6.07	4.28±1.88
			Sparse	5 / 185	10	17.87±0.87	9.88±0.55	76.06±8.57	37.83±5.17
	CoraNet [34]	TMI'22	Dense	5 / 185	10	43.08±28.61	33.27±23.87	24.17±3.88	8.61±2.11
			Sparse	5 / 185	10	28.64±25.05	18.90±19.44	72.62±22.57	35.22±13.43
Ours	this paper	-	5 / 185	10	86.93±2.78	78.33±3.21	11.61±2.15	3.28±0.51	
Semi-supervised	MT [36]	NIPS'17	-	5 / 185	320	84.98±4.68	76.51±5.63	16.51±0.18	4.61±0.83



MT_{dense} MT_{sparse} CPS_{dense} CPS_{sparse} CTBCT_{dense} CTBCT_{sparse} CoraNet_{dense} CoraNet_{sparse} Ours

Figure 7. Visual segmentation examples from KiTS19 dataset. The red, cyan, blue, green, magenta, yellow curves denote the corresponding results of ground-truth, mean teacher [36], CPS [10], CTBCT [25], CoraNet [34] and ours, respectively.

unlabeled data and the rest 31 as testing data. We adopt random cropping to get patches of $176 \times 176 \times 64$ for training and testing. From Table 3, our method uses only 10 labeled slices and obtains a Dice coefficient close to 90%. Compared with other methods, our method outperforms all methods by a large margin. And our method is only 1.53% lower than MT trained with full volumetric annotation on Dice metric, but performs much better on HD and ASD.

The segmentation result examples are shown in Figure 8.

4.4. Ablation Study

Effectiveness of Each Component: In order to better understand and evaluate the components of our method, we conduct an ablation experiment on the KiTS19 dataset. Now we introduce each setting of our experiment: 1) training from mixed dense pseudo label (Dense), 2) leverag-

Table 3. Comparison on LiTS dataset segmentation.

Method	Venue	Setting	Scans Used		Metrics				
			L / U Volumes	Labeled Slices	Dice (%)	Jaccard (%)	HD (voxel)	ASD (voxel)	
Barely-supervised	MT [36]	NIPS'17	Dense	5 / 95	10	81.76±4.82	69.73±7.01	28.87±13.85	8.57±4.03
			Sparse	5 / 95	10	56.82±25.76	43.83±28.78	74.03±34.02	31.41±17.58
	CPS [10]	CVPR'21	Dense	5 / 95	10	73.86±10.13	59.73±13.16	32.78±14.68	10.73±5.20
			Sparse	5 / 95	10	20.46±2.15	11.48±1.35	92.27±3.03	41.34±2.09
	CTBCT [25]	MIDL'22	Dense	5 / 95	10	79.68±5.45	66.95±7.41	30.46±13.36	9.18±4.12
			Sparse	5 / 95	10	40.07±7.95	25.52±6.24	71.83±11.41	29.49±5.44
	CoraNet [34]	TMI'22	Dense	5 / 95	10	80.17±1.77	68.97±3.16	19.42±7.47	4.34±0.50
			Sparse	5 / 95	10	36.84±8.20	23.13±6.03	95.89±1.94	43.25±2.03
Ours	this paper	-	5 / 95	10	89.24±1.37	81.10±2.28	10.05±2.42	2.27±0.45	
Semi-supervised	MT [36]	NIPS'17	-	5 / 95	320	90.77±1.91	83.62±3.11	18.32±7.17	4.95±1.81

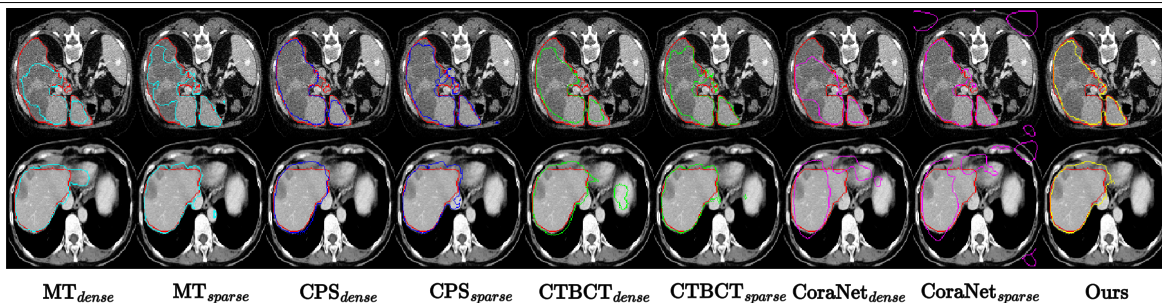


Figure 8. Visual segmentation examples from LiTS dataset. The red, cyan, blue, green, magenta, yellow curves denote the corresponding results of ground-truth, mean teacher [36], CPS [10], CTBCT [25], CoraNet [34] and ours, respectively.

Table 4. Ablation study on KiTS19 dataset.

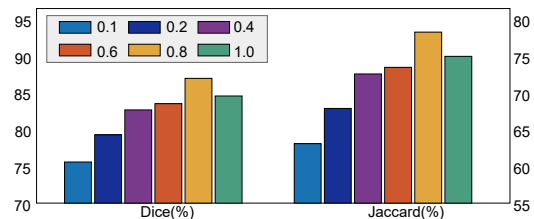
Methods	Dense	w_s	M_{un}	$D \rightarrow S$	$\lambda \uparrow$	Dice (%)	Jaccard (%)	HD (voxel)	ASD (voxel)
#1	✓					75.33±8.77	63.20±10.95	19.14±8.14	6.53±3.27
#2	✓	✓				77.69±5.07	65.75±6.57	17.90±6.34	5.65±2.47
#3	✓	✓	✓			78.31±3.94	66.48±4.67	18.35±5.59	5.43±1.62
#4	✓	✓	✓	✓		78.45±6.45	67.03±9.13	30.50±6.28	9.57±3.32
#5	✓	✓	✓		✓	84.86±2.01	74.99±2.16	12.21±2.75	3.34±0.95
#6	✓	✓	✓	✓	✓	86.93±2.78	78.33±3.21	11.61±2.15	3.28±0.51

ing slice-wise weight (w_s), 3) uncertainty-guided cross-supervision (M_{un}), 4) gradually decreasing the weight decay rate α to zero, *i.e.*, learning from dense to sparse ($D \rightarrow S$), and 5) gradually increasing cross-supervision weight to λ_{oc} ($\lambda \uparrow$). The results are shown in Table 4.

The process of learning from dense to sparse and a high cross-supervision weight in the late stage play the most critical role. Throughout the training procedure, the performance of the segmentation models is improving, and they can produce better results than the initial pseudo labels, so the high weighted supervised loss is gradually hampering the segmentation networks from improving. Simply leveraging slice-wise weight w_s is also effective.

Quantitative Analysis on Hyper-Parameter: We also conduct exhaustive experiments of different optimal cross-supervision weights λ_{oc} on the KiTS19 datasets. Specifically, we set the optimal cross-supervision weight as 0.1, 0.2, 0.4, 0.6, 0.8 and 1.0. We find that setting λ_{oc} as 0.8 serves best. Detailed results are shown in Figure 9.

It is shown that optimal cross-supervision weight should not be too large or too small, because small λ_{oc} makes the networks concentrate on learning from registration labels and neglect the fact that the networks are improving during the training procedure, while too large λ_{oc} makes the net-

Figure 9. Segmentation result on KiTS dataset with different optimal cross-supervision weight λ_{oc} .

work gradually lose the supervision from the ground truth labels, and the performance tends to decrease through further cross-supervision without constraint.

5. Conclusion

In this paper, we proposed orthogonal annotation for 3D medical image segmentation which is to label two orthogonal slices for a volume, and we verified its high-efficiency. Then we applied this annotation manner in barely-supervised segmentation setting. To better leverage the volumes with orthogonal annotation and the large number of unlabeled volumes, we designed a dense-sparse co-training paradigm, which learns from dense pseudo label first, then reduces noise and gains further promotion from sparse annotation in later stage. Our method could better synthesize the knowledge from two planes through cross-supervision. Large quantities of experiments have validated the effectiveness of our orthogonal annotation and DeSCO paradigm, *e.g.*, our method could achieve a close or better performance compared with semi-supervised method, but with only 2%~3% annotations.

References

- [1] Brian B Avants, Nicholas J Tustison, Gang Song, Philip A Cook, Arno Klein, and James C Gee. A reproducible evaluation of ANTs similarity metric performance in brain image registration. *Neuroimage*, 54(3):2033–2044, 2011. **6**
- [2] 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*, pages 253–260. Springer, 2017. **3**
- [3] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *European Conference on Computer Vision*, pages 549–565. Springer, 2016. **1, 3**
- [4] Patrick Bilic, Patrick Ferdinand Christ, Eugene Vorontsov, Grzegorz Chlebus, Hao Chen, Qi Dou, Chi-Wing Fu, Xiao Han, Pheng-Ann Heng, Jürgen Hesser, et al. The liver tumor segmentation benchmark (lits). *arXiv preprint arXiv:1901.04056*, 2019. **6**
- [5] Adeleh Bitarafan, Mahdi Nikdan, and Mahdieh Soleymani Baghshah. 3D image segmentation with sparse annotation by self-training and internal registration. *IEEE Journal of Biomedical and Health Informatics*, 25(7):2665–2672, 2020. **1**
- [6] Gerda Bortsova, Florian Dubost, Laurens Hogeweg, Ioannis Katramados, and Marleen de Bruijne. Semi-supervised medical image segmentation via learning consistency under transformations. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 810–818. Springer, 2019. **2**
- [7] Chen Chen, Carlo Biffi, Giacomo Tarroni, Steffen Petersen, Wenjia Bai, and Daniel Rueckert. Learning shape priors for robust cardiac MR segmentation from multi-view images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 523–531. Springer, 2019. **2**
- [8] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):834–848, 2017. **1**
- [9] Qihui Chen and Yi Hong. Scribble2d5: Weakly-supervised volumetric image segmentation via scribble annotations. *arXiv preprint arXiv:2205.06779*, 2022. **1**
- [10] Xiaokang Chen, Yuhui Yuan, Gang Zeng, and Jingdong Wang. Semi-supervised semantic segmentation with cross pseudo supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2613–2622, 2021. **6, 7, 8**
- [11] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3D U-Net: learning dense volumetric segmentation from sparse annotation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 424–432. Springer, 2016. **1**
- [12] Jifeng Dai, Kaiming He, and Jian Sun. Boxesup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1635–1643, 2015. **3**
- [13] Ali Hatamizadeh, Yucheng Tang, Vishwesh Nath, Dong Yang, Andriy Myronenko, Bennett Landman, Holger R Roth, and Daguang Xu. Unetr: Transformers for 3d medical image segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 574–584, 2022. **6**
- [14] Nicholas Heller, Niranjana Sathianathan, Arveen Kalapara, Edward Walczak, Keenan Moore, Heather Kaluzniak, Joel Rosenberg, Paul Blake, Zachary Rengel, Makinna Oestreich, et al. The kits19 challenge data: 300 kidney tumor cases with clinical context, ct semantic segmentations, and surgical outcomes. *arXiv preprint arXiv:1904.00445*, 2019. **6**
- [15] Sang Hyun Jo, In Jae Yu, and Kyung-Su Kim. Recurseed and certainmix for weakly supervised semantic segmentation. *arXiv preprint arXiv:2204.06754*, 2022. **1**
- [16] Jungbeom Lee, Eunji Kim, Sungmin Lee, Jangho Lee, and Sungroh Yoon. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5267–5276, 2019. **1, 3**
- [17] Jungbeom Lee, Eunji Kim, and Sungroh Yoon. Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4071–4080, 2021. **1, 3**
- [18] Shumeng Li, Heng Cai, Lei Qi, Qian Yu, Yinghuan Shi, and Yang Gao. PLN: Parasitic-like network for barely supervised medical image segmentation. *IEEE Transactions on Medical Imaging*, 2022. **1, 2**
- [19] Xiaomeng Li, Lequan Yu, Hao Chen, Chi-Wing Fu, and Pheng-Ann Heng. Semi-supervised skin lesion segmentation via transformation consistent self-ensembling model. *British Machine Vision Conference*, 2018. **2**
- [20] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3159–3167, 2016. **1, 3**
- [21] Xiaoming Liu, Quan Yuan, Yaozong Gao, Kelei He, Shuo Wang, Xiao Tang, Jinshan Tang, and Dinggang Shen. Weakly supervised segmentation of COVID19 infection with scribble annotation on CT images. *Pattern Recognition*, 122:108341, 2022. **1, 3**
- [22] Yuyuan Liu, Yu Tian, Yuanhong Chen, Fengbei Liu, Vasileios Belagiannis, and Gustavo Carneiro. Perturbed and strict mean teachers for semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4258–4267, 2022. **2**
- [23] Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations*, 2017. **6**

- [24] Xiangde Luo, Jieneng Chen, Tao Song, and Guotai Wang. Semi-supervised medical image segmentation through dual-task consistency. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 8801–8809, 2021. **3**
- [25] Xiangde Luo, Minhao Hu, Tao Song, Guotai Wang, and Shaoting Zhang. Semi-supervised medical image segmentation via cross teaching between cnn and transformer. In *Medical Imaging with Deep Learning*, 2022. **6, 7, 8**
- [26] Wan-Duo Kurt Ma, JP Lewis, and W Bastiaan Kleijn. The HSIC bottleneck: Deep learning without back-propagation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5085–5092, 2020. **4**
- [27] Peng Mi, Jiangang Lin, Yiyi Zhou, Yunhang Shen, Gen Luo, Xiaoshuai Sun, Liujuan Cao, Rongrong Fu, Qiang Xu, and Rongrong Ji. Active teacher for semi-supervised object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14482–14491, 2022. **2**
- [28] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-Net: Fully convolutional neural networks for volumetric medical image segmentation. In *4th International Conference on 3D Vision*, pages 565–571. IEEE, 2016. **1, 4, 6**
- [29] Youngmin Oh, Beomjun Kim, and Bumsub Ham. Background-aware pooling and noise-aware loss for weakly-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6913–6922, 2021. **3**
- [30] Siyuan Qiao, Wei Shen, Zhishuai Zhang, Bo Wang, and Alan Yuille. Deep co-training for semi-supervised image recognition. In *Proceedings of the European Conference on Computer Vision*, pages 135–152, 2018. **4**
- [31] Mamshad Nayeem Rizve, Kevin Duarte, Yogesh S Rawat, and Mubarak Shah. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. In *International Conference on Learning Representations*, 2021. **2**
- [32] 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*, pages 234–241. Springer, 2015. **1**
- [33] Laine Samuli and Aila Timo. Temporal ensembling for semi-supervised learning. In *International Conference on Learning Representations*, volume 4, page 6, 2017. **2**
- [34] Yinghuan Shi, Jian Zhang, Tong Ling, Jiwen Lu, Yefeng Zheng, Qian Yu, Lei Qi, and Yang Gao. Inconsistency-aware uncertainty estimation for semi-supervised medical image segmentation. *IEEE Transactions on Medical Imaging*, 41(3):608–620, 2022. **3, 6, 7, 8**
- [35] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in Neural Information Processing Systems*, 33:596–608, 2020. **2**
- [36] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in Neural Information Processing Systems*, 30, 2017. **2, 6, 7, 8**
- [37] Bin Wang, Guojun Qi, Sheng Tang, Tianzhu Zhang, Yunchao Wei, Linghui Li, and Yongdong Zhang. Boundary perception guidance: A scribble-supervised semantic segmentation approach. In *IJCAI International Joint Conference on Artificial Intelligence*, 2019. **3**
- [38] Jianfeng Wang, Thomas Lukasiewicz, Daniela Massiceti, Xiaolin Hu, Vladimir Pavlovic, and Alexandros Neophytou. NP-Match: When neural processes meet semi-supervised learning. In *International Conference on Machine Learning*, pages 22919–22934. PMLR, 2022. **2**
- [39] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12275–12284, 2020. **3**
- [40] Yingda Xia, Dong Yang, Zhiding Yu, Fengze Liu, Jinzheng Cai, Lequan Yu, Zhuotun Zhu, Daguang Xu, Alan Yuille, and Holger Roth. Uncertainty-aware multi-view co-training for semi-supervised medical image segmentation and domain adaptation. *Medical Image Analysis*, 65:101766, 2020. **2, 3**
- [41] Yutong Xie, Jianpeng Zhang, Zhibin Liao, Johan Verjans, Chunhua Shen, and Yong Xia. Pairwise relation learning for semi-supervised gland segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 417–427. Springer, 2020. **3**
- [42] 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*, 67:101832, 2021. **1, 4, 6**
- [43] 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*, pages 605–613. Springer, 2019. **3, 5, 6**
- [44] Bowen Zhang, Yidong Wang, Wenxin Hou, Hao Wu, Jindong Wang, Manabu Okumura, and Takahiro Shinozaki. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. *Advances in Neural Information Processing Systems*, 34:18408–18419, 2021. **2**
- [45] Dawei Zhou, Nannan Wang, Chunlei Peng, Yi Yu, Xi Yang, and Xinbo Gao. Towards multi-domain face synthesis via domain-invariant representations and multi-level feature parts. *IEEE Transactions on Multimedia*, 24:3469–3479, 2021. **2**
- [46] Yuyin Zhou, Yan Wang, Peng Tang, Song Bai, Wei Shen, Elliot Fishman, and Alan Yuille. Semi-supervised 3d abdominal multi-organ segmentation via deep multi-planar co-training. In *2019 IEEE Winter Conference on Applications of Computer Vision*, pages 121–140. IEEE, 2019. **3**