

Co-training with High-Confidence Pseudo Labels for Semi-supervised Medical Image Segmentation

Zhiqiang Shen^{1,2}, Peng Cao^{1,2*}, Hua Yang³, Xiaoli Liu⁴, Jinzhu Yang^{1,2}, Osmar R. Zaiane⁵

¹College of Computer Science and Engineering, Northeastern University, Shenyang, China

²Key Laboratory of Intelligent Computing in Medical Image, Ministry of Education, Shenyang, China

³College of Photonic and Electronic Engineering, Fujian Normal University, Fuzhou, China

⁴DAMO Academy, Alibaba Group, China

⁵Alberta Machine Intelligence Institute, University of Alberta, Edmonton, Alberta, Canada

xxszqyy@gmail.com, caopengneu@gmail.com, qsx20200801@student.fjnu.edu.cn,

liuxiaoli.lxl@alibaba-inc.com, yangjinzhu@cse.neu.edu.cn, zaiane@cs.ualberta.ca

Abstract

Consistency regularization and pseudo labeling-based semi-supervised methods perform co-training using the pseudo labels from multi-view inputs. However, such co-training models tend to converge early to a consensus, degenerating to the self-training ones, and produce low-confidence pseudo labels from the perturbed inputs during training. To address these issues, we propose an **Uncertainty-guided Collaborative Mean-Teacher (UCMT)** for semi-supervised semantic segmentation with the high-confidence pseudo labels. Concretely, UCMT consists of two main components: 1) collaborative mean-teacher (CMT) for encouraging model disagreement and performing co-training between the sub-networks, and 2) uncertainty-guided region mix (UMIX) for manipulating the input images according to the uncertainty maps of CMT and facilitating CMT to produce high-confidence pseudo labels. Combining the strengths of UMIK with CMT, UCMT can retain model disagreement and enhance the quality of pseudo labels for the co-training segmentation. Extensive experiments on four public medical image datasets including 2D and 3D modalities demonstrate the superiority of UCMT over the state-of-the-art. Code is available at: <https://github.com/Senyh/UCMT>.

1 Introduction

Semantic segmentation is critical for medical image analysis. Great progress has been made by deep learning-based segmentation models relying on a large amount of labeled data [Chen *et al.*, 2018; Ronneberger *et al.*, 2015]. However, labeling such pixel-level annotations is laborious and requires expert knowledge especially in medical images, resulting in that labeled data are expensive or simply unavailable. Unlabeled data, on the contrary, are cheap and relatively easy to

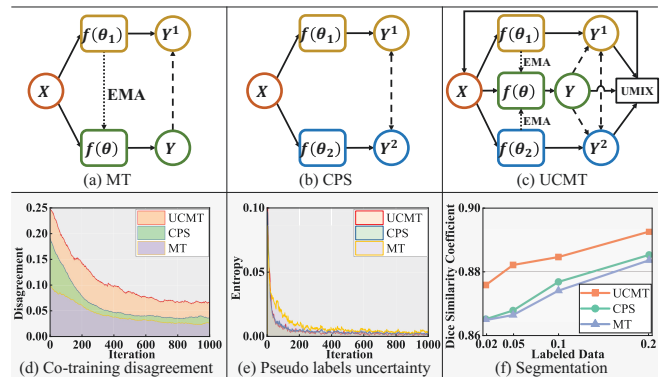


Figure 1: Illustration of the architectures and curves for co-training based semi-supervised semantic segmentation. (a) Mean-teacher [Tarvainen and Valpola, 2017], (b) Cross pseudo supervision [Chen *et al.*, 2021], (c) Uncertainty-guided collaborative mean-teacher, (d) the disagreement variation between the co-training sub-networks in terms of dice loss (*w.r.t.* number of iterations) of two branches (Y^1 and Y in MT; Y^1 and Y^2 in CPS; Y^1 and Y^2 in UCMT), (e) the uncertainty variation of the pseudo labels in terms of entropy *w.r.t.* number of iterations, and (f) the performance of MT, CPS, and UCMT on semi-supervised skin lesion segmentation under different ratio of labeled data. EMA: exponential moving average.

obtain. Under this condition, semi-supervised learning (SSL) has been the dominant data-efficient strategy through exploiting information from a limited amount labeled data and an arbitrary amount of unlabeled data, so as to alleviate the label scarcity problem [Van Engelen and Hoos, 2020].

Consistency regularization [Tarvainen and Valpola, 2017] and pseudo labeling [Lee and others, 2013] are the two main methods for semi-supervised semantic segmentation. Currently, combining consistency regularization and pseudo labeling via cross supervision between the sub-networks, has shown promising performance for semi-supervised segmentation [Lee and others, 2013; Ouali *et al.*, 2020; Ke *et al.*, 2020; Chen *et al.*, 2021; Liu *et al.*, 2022]. One critical limitation of these approaches is that the sub-networks tend to converge early to a consensus situation causing the co-training model degenerating to the self-training [Yu *et al.*, 2019b]. Disagreement between the sub-networks is crucial

*Corresponding author

for co-training, where the sub-networks initialized with different parameters or trained with different views have different biases (i.e., disagreement) ensuring that the information they provide is complementary to each other. Another key factor affecting the performance of these approaches is the quality of pseudo labels. Intuitively, high quality pseudo labels should have low uncertainty [Grandvalet and Bengio, 2004]. More importantly, these two factors influence each other. Increasing the degree of the disagreement between the co-training sub-networks by different perturbations or augmentations could result in their opposite training directions, thus increasing the uncertainty of pseudo labels. To investigate the effect of the disagreement and the quality of pseudo labels for co-training based semi-supervised segmentation, which has not been studied in the literature, we conduct a pilot experiment to illustrate these correlations. As shown in Figure 1, compared with mean-teacher (MT) [Tarvainen and Valpola, 2017] [Figure 1 (a)], cross pseudo supervision (CPS) [Chen *et al.*, 2021] [Figure 1 (b)] with the higher model disagreement [(d)] and the lower uncertainty [Figure 1 (e)] produces higher performance [Figure 1 (f)] on semi-supervised segmentation. Note that the dice loss of two branches are calculated to measure the disagreement. The question that comes to mind is: *how to effectively improve the disagreement between the co-training sub-networks and the quality of pseudo labels jointly in a unified network for SSL.*

In this paper, we focus on two major goals: maintaining model disagreement and the high-confidence pseudo labels at the same time. To this end, we propose the Uncertainty-guided Collaborative Mean Teacher (UCMT) framework that is capable of retaining higher disagreement between the co-training segmentation sub-networks [Figure 1 (d)] based on the higher confidence pseudo labels [Figure 1 (e)], thus achieving better semi-supervised segmentation performance under the same backbone network and task settings [Figure 1 (f)]. Specifically, UCMT involves two major components: 1) collaborative mean-teacher (CMT), and 2) uncertainty-guided region mix (UMIX), where UMIX operates the input images according to the uncertainty maps of CMT while CMT performs co-training under the supervision of the pseudo labels derived from the UMIX images. Inspired by the co-teaching [Han *et al.*, 2018; Yu *et al.*, 2019b; Chen *et al.*, 2021] for struggling with early converging to a consensus situation and degrading into self-training, we introduce a third component, the teacher model, into the co-training framework as a regularizer to construct CMT for more effective SSL. The teacher model acts as self-ensemble by averaging the student models, serving as a third part to guide the training of the two student models. Further, we develop UMIX to construct high-confident pseudo labels and perform regional dropout for learning robust semi-supervised semantic segmentation models. Instead of random region erasing or swapping [DeVries and Taylor, 2017; Yun *et al.*, 2019], UMIX manipulates the original image and its corresponding pseudo labels according to the epistemic uncertainty of the segmentation models, which not only reduces the uncertainty of the pseudo labels but also enlarges the training data distribution. Finally, by combining the strengths of UMIX with CMT, the proposed approach

UCMT significantly improves the state-of-the-art (sota) results in semi-supervised segmentation on multiple benchmark datasets. For example, UCMT and UCMT(U-Net) achieve 88.22% and 82.14% Dice Similarity Coefficient (DSC) on ISIC dataset under 5% labeled data, outperforming our baseline model CPS [Chen *et al.*, 2021] and the state-of-the-art UGCL [Wang *et al.*, 2022] by 1.41% and 9.47%, respectively.

In a nutshell, our contributions mainly include:

- We pinpoint the problem in existing co-training based semi-supervised segmentation methods: the insufficient disagreement among the sub-networks and the lower-confidence pseudo labels. To address the problem, we design an uncertainty-guided collaborative mean-teacher to maintain co-training with high-confidence pseudo labels, where we incorporate CMT and UMIX into a holistic framework for semi-supervised medical image segmentation.
- To avoid introducing noise into the new samples, we propose an uncertainty-guided regional mix algorithm, UMIX, encouraging the segmentation model to yield high-confident pseudo labels and enlarge the training data distribution.
- We conduct extensive experiments on four public medical image segmentation datasets including 2D and 3D scenarios. Comprehensive results demonstrate the effectiveness of each component of our method and the advantage of UCMT over the state-of-the-art.

2 Related Work

2.1 Semi-Supervised Learning

Semi-supervised learning aims to improve performance in supervised learning by utilizing information generally associated with unsupervised learning, and vice versa [Van Engelen and Hoos, 2020]. A common form of SSL is introducing a regularization term into the objective function of supervised learning to leverage unlabeled data. From this perspective, SSL-based methods can be divided into two main lines, i.e., pseudo labeling and consistency regularization. Pseudo labeling attempts to generate pseudo labels similar to the ground truth, for which models are trained as in supervised learning [Lee and others, 2013]. Consistency regularization enforces the model’s outputs to be consistent for the inputs under different perturbations [Tarvainen and Valpola, 2017]. Current state-of-the-art approaches have incorporated these two strategies and shown convincing performance for semi-supervised image classification [Sohn *et al.*, 2020; Zhang *et al.*, 2021]. Based on this line of research, we explore more effective consistency learning algorithms for semi-supervised segmentation.

2.2 Semi-Supervised Semantic Segmentation

Compared with image classification, semantic segmentation requires much more intensively and costly labeling for pixel-level annotations. Semi-supervised semantic segmentation inherits the main ideas of semi-supervised image classification. The combination of consistency regularization and

pseudo labeling, mainly conducting cross supervision between sub-networks using pseudo labels, has become the mainstream strategy for semi-supervised semantic segmentation in both natural images [Ouali *et al.*, 2020; Chen *et al.*, 2021] and medical images [Yu *et al.*, 2019a; Luo *et al.*, 2021; Wu *et al.*, 2021b; Wu *et al.*, 2022]. Specifically, these combined approaches enforce the consistency of the predictions under different perturbations, such as input perturbations [Li *et al.*, 2020b; Tu *et al.*, 2022], feature perturbations [Ouali *et al.*, 2020], and network perturbations [Tavainen and Valpola, 2017; Chen *et al.*, 2021; Wu *et al.*, 2021b; Wu *et al.*, 2022]. In addition, adversarial learning-based methods, rendering the distribution of model predictions from labeled data to be aligned with those from unlabeled data, can also be regarded as a special form of consistency regularization [Hung *et al.*, 2018; Li *et al.*, 2020a]. However, such cross supervision models may converge early to a consensus, thus degenerating to self-training ones. We hypothesize that enlarging the disagreement for the co-training models based on the high-confidence pseudo labels can improve the performance of SSL. Therefore, we propose a novel SSL framework, i.e., UCMT, to generate more accurate pseudo labels and maintain co-training for semi-supervised medical image segmentation.

2.3 Uncertainty-Guided Semi-Supervised Semantic Segmentation

Model uncertainty (epistemic uncertainty) can guide the SSL models to capture information from the pseudo labels. Two critical problems for leveraging model uncertainty are how to *obtain* and *exploit* model uncertainty. Recently, there are mainly two strategies to estimate model uncertainty: 1) using Monte Carlo dropout [Gal and Ghahramani, 2016], and 2) calculating the variance among different predictions [Zheng and Yang, 2021]. For semi-supervised semantic segmentation, previous works exploit model uncertainty to re-weight the training loss [Yu *et al.*, 2019a] or selecting the contrastive samples [Wang *et al.*, 2022]. However, these methods require manually setting a threshold to neglect the low-confidence pseudo labels, where the fixed threshold is hard to determine. In this paper, we *obtain* the epistemic uncertainty by the entropy of the predictions of CMT for the same input and *exploit* the uncertainty to guide the region mix for gradually exploring information from the unlabeled data.

3 Methodology

Before introducing our method, we first define the notations used in this work. The training set $\mathcal{D} = \{\mathcal{D}_L, \mathcal{D}_U\}$ contains a labeled set $\mathcal{D}_L = \{(X_i, Y_i)_{i=1}^N\}$ and an unlabeled set $\mathcal{D}_U = \{(X_j)_{j=N+1}^M\}$, where X_i/X_j denotes the i th/ j th labeled/unlabeled image, Y_i is the ground truth of the labeled image, and N and $M - N$ are the number of labeled and unlabeled samples, respectively. Given the training data \mathcal{D} , the goal of semi-supervised semantic segmentation is to learn a model $f(\cdot; \theta)$ performing well on unseen test sets.

3.1 Overview

To avoid the co-training degrading to the self-training, we propose to encourage model disagreement during training and

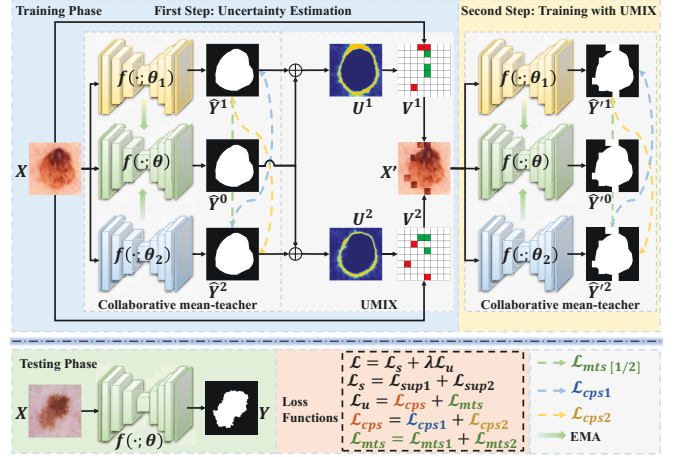


Figure 2: Overview of the proposed UCMT. CMT includes three sub-networks, i.e., the teacher sub-network ($f(\cdot; \theta)$) and the two student sub-networks ($f(\cdot; \theta_1)$ and $f(\cdot; \theta_2)$). UMIX constructs each new samples X' by replacing the top k most uncertain regions (red grids in V^1 and V^2) with the top k most certain regions (green grids in V^2 and V^1) in the original image X .

ensure pseudo labels with low uncertainty. With this motivation, we propose uncertainty-guided collaborative mean-teacher for semi-supervised image segmentation, which includes 1) collaborative mean-teacher, and 2) uncertainty-guided region mix. As shown in Figure 1 (d), CMT and UCMT gradually enlarge the disagreement between the co-training sub-networks. Meanwhile, CMT equipped with UMIX guarantees low-uncertainty for the pseudo labels. With the help of these conditions, we can safely maintain the co-training status to improve the effectiveness of SSL for exploring unlabeled data. Figure 2 illustrates the schematic diagram of the proposed UCMT. Generally, there are two steps in the training phase of UCMT. In the first step, we train CMT using the original labeled and unlabeled data to obtain the uncertainty maps; Then, we perform UMIX to generate the new samples based on the uncertainty maps. In the second step, we re-train CMT using the UMIX samples. Details of the training process of UCMT are shown in Algorithm 1. Although UCMT includes three models, i.e., one teacher model and two student models, only the teacher model is required in the testing stage.

3.2 Collaborative Mean-Teacher

Current consistency learning-based SSL algorithms, e.g., Mean-teacher [Tavainen and Valpola, 2017] and CPS [Chen *et al.*, 2021], suggest to perform consistency regularization among the pseudo labels in a multi-model architecture rather than in a single model. However, during the training process, the two-network SSL framework may converge early to a consensus and the co-training degenerate to the self-training [Yu *et al.*, 2019b]. To tackle this issue, we design the collaborative mean teacher (CMT) framework by introducing a "arbitrator", i.e., the teacher model, into the co-training architecture [Chen *et al.*, 2021] to guide the training of the two student models. As shown in Figure 2, CMT consists of one teacher model $f(\cdot; \theta)$ and two student models $f(\cdot; \theta_1)$ and $f(\cdot; \theta_2)$,

Algorithm 1 UCMT algorithm

Input: $\mathcal{D}_{\mathcal{L}} = \{(X_i, Y_i)\}_{i=1}^N$, $\mathcal{D}_{\mathcal{U}} = \{X_j\}_{j=N+1}^M$
Parameter: $\theta, \theta_1, \theta_2$
Output: $f(\cdot; \theta)$

- 1: **for** $T \in [1, \text{numepochs}]$ **do**
- 2: **for** each minibatch B **do**
- 3: // i/j is the index for labeled/unlabeled data
- 4: **step 1: uncertainty estimation**
- 5: $\hat{Y}_i^0 \leftarrow f(X_i; \theta), \hat{Y}_i^1 \leftarrow f(X_i; \theta_1), \hat{Y}_i^2 \leftarrow f(X_i; \theta_2)$
- 6: $\hat{Y}_j^0 \leftarrow f(X_j; \theta), \hat{Y}_j^1 \leftarrow f(X_j; \theta_1), \hat{Y}_j^2 \leftarrow f(X_j; \theta_2)$
- 7: $L \leftarrow \mathcal{L}_s(\hat{Y}_i^1, \hat{Y}_i^2, Y_i) + \lambda(T)\mathcal{L}_u(\hat{Y}_j^0, \hat{Y}_j^1, \hat{Y}_j^2)$
- 8: Update $f(\cdot; \theta), f(\cdot; \theta_1), f(\cdot; \theta_2)$ using optimizer
- 9: $U_i^{[1/2]} \leftarrow \text{Uncertain}(f(X_i; \theta_{[1/2]}), f(X_i; \theta))$
- 10: $U_j^{[1/2]} \leftarrow \text{Uncertain}(f(X_j; \theta_{[1/2]}), f(X_j; \theta))$
- 11: **step 2: training with UMIX**
- 12: $X'_i/Y'_i \leftarrow \text{UMIX}(X_i/Y_i, U_i^1, U_i^2; k, 1/r)$
- 13: $X'_j/\hat{Y}_j^0 \leftarrow \text{UMIX}(X_j/\hat{Y}_j^0, U_j^1, U_j^2; k, 1/r)$
- 14: **Repeat** 5-8 using X'_i, X'_j, Y'_i , and \hat{Y}_j^0
- 15: **end for**
- 16: **end for**
- 17: **return** $f(\cdot; \theta)$

where the teacher model is the self-ensemble of the average of the student models. These models have the same architecture but initialized with different weights for network perturbations. For labeled data, these models are all optimized by supervised learning. For unlabeled data, there are two critical factors: 1) co-training between the two student models, and 2) direct supervision from the teacher to the student models.

To explore both the labeled and unlabeled data, the total loss \mathcal{L} for training UCMT involves two parts, i.e., the supervised loss \mathcal{L}_s and the unsupervised loss \mathcal{L}_u .

$$\mathcal{L} = \mathcal{L}_s + \lambda\mathcal{L}_u, \quad (1)$$

where λ is a regularization parameter to balance the supervised and unsupervised learning losses. We adopt a Gaussian ramp-up function to gradually increase the coefficient, i.e., $\lambda(t) = \lambda_m \times \exp[-5(1 - \frac{t}{t_m})^2]$, where λ_m scales the maximum value of the weighted function, t denotes the current iteration, and t_m is the maximum iteration in training.

Supervised Learning Path. For the labeled data, the supervised loss is formulated as,

$$\mathcal{L}_s = \frac{1}{N} \sum_{i=1}^N \{ \mathcal{L}_{seg}(f(X_i; \theta_1), Y_i) + \mathcal{L}_{seg}(f(X_i; \theta_2), Y_i) \}, \quad (2)$$

where \mathcal{L}_{seg} can be any supervised semantic segmentation loss, such as cross entropy loss and dice loss. Note that we choose dice loss in our experiments as its compelling performance in medical image segmentation.

Unsupervised Learning Path. The unsupervised loss \mathcal{L}_u acts as a regularization term to explore potential knowledge for the labeled and unlabeled data. \mathcal{L}_u includes the cross pseudo supervision \mathcal{L}_{cps} between the two student models and

the mean-teacher supervision \mathcal{L}_{mts} for guiding the student models from the teacher, as follow:

$$\mathcal{L}_u = \mathcal{L}_{cps} + \mathcal{L}_{mts}. \quad (3)$$

1) Cross Pseudo Supervision. The aim of \mathcal{L}_{cps} is to promote two students to learn from each other and to enforce the consistency between them. Let $\mathcal{L}_{cps} = \mathcal{L}_{cps1} + \mathcal{L}_{cps2}$ for encouraging bidirectional interaction for the two student sub-networks $f(\cdot; \theta_1)$ and $f(\cdot; \theta_2)$. The losses of \mathcal{L}_{cps1} and \mathcal{L}_{cps2} are defined as:

$$\mathcal{L}_{cps[1/2]} = \frac{1}{M-N} \sum_{j=1}^{M-N} \mathcal{L}_{seg}(f(X_j; \theta_{[1/2]}), \hat{Y}_j^{[2/1]}), \quad (4)$$

where \hat{Y}_j^1 and \hat{Y}_j^2 are the pseudo segmentation maps for X_j predicted by $f(\cdot; \theta_1)$ and $f(\cdot; \theta_2)$, respectively.

2) Mean-Teacher Supervision. To avoid the two students co-training in the wrong direction, we introduce a teacher model to guide the optimization of the student models. Specifically, the teacher model is updated by the exponential moving average (EMA) of the average of the student models: $\theta^t = \alpha\theta^{t-1} + (1-\alpha)[\beta\theta_1^t + (1-\beta)\theta_2^t]$, where t represents the current training iteration and α is the EMA decay that controls the parameters' updating rate.

The loss of mean-teacher supervision $\mathcal{L}_{mts} = \mathcal{L}_{mts1} + \mathcal{L}_{mts2}$ is calculated from two branches:

$$\mathcal{L}_{mts[1/2]} = \frac{1}{M-N} \sum_{j=1}^{M-N} \mathcal{L}_{seg}(f(X_j; \theta_{[1/2]}), \hat{Y}_j^0), \quad (5)$$

where \hat{Y}_j^0 is the segmentation map derived from $f(X_j; \theta)$.

3.3 Uncertainty-Guided Mix

Although CMT can promote model disagreement for co-training, it also slightly increases the uncertainty of the pseudo labels as depicted in Figure 1. On the other hand, random regional dropout can expand the training distribution and improve the generalization capability of models [DeVries and Taylor, 2017; Yun *et al.*, 2019]. However, such random perturbations to the input images inevitably introduce noise into the new samples, thus deteriorating the quality of pseudo labels for SSL. One sub-network may provide some incorrect pseudo labels to the other sub-networks, degrading their performance. To overcome these limitations, we propose UMIX to manipulate image patches under the guidance of the uncertainty maps produced by CMT. The main idea of UMIX is constructing a new sample by replacing the top k most uncertain (low-confidence) regions with the top k most certain (high-confidence) regions in the input image. As illustrated in Figure 2, UMIX constructs a new sample $X' = \text{UMIX}(X, U^1, U^2; k, 1/r)$ by replacing the top k most uncertain regions (red grids in V^1 and V^2) with the top k most certain regions (green grids in V^2 and V^1) in X , where each region has size $1/r$ to the image size. To ensure the reliability of the uncertainty evaluation, we obtain the uncertain maps by integrating the outputs of the teacher and the student model instead of performing T stochastic forward passes designed by Monte

Carlo Dropout estimate model [Gal and Ghahramani, 2016; Yu *et al.*, 2019a], which is equivalent to sampling predictions from the previous and current iterations. This process can be formulated as:

$$U^m = \text{Uncertain}(f(X; \theta_m), f(X; \theta)) = - \sum_c P_c \log(P_c),$$

$$P_c = \frac{1}{2}(\text{Softmax}(f(X; \theta_m)) + \text{Softmax}(f(X; \theta))),$$
(6)

where $m = 1, 2$ denotes the index of the student models and c refers to the class index.

4 Experiments and Results

4.1 Experiments Settings

Datasets. We conduct extensive experiments on different medical image segmentation tasks to evaluate the proposed method, including skin lesion segmentation from dermoscopy images, polyp segmentation from colonoscopy images, and the 3D left atrium segmentation from cardiac MRI images.

Dermoscopy. We validate our method on the ISIC dataset [Codella *et al.*, 2018] including 2594 dermoscopy images and corresponding annotations. Following [Wang *et al.*, 2022], we adopt 1815 images for training and 779 images for validation.

Colonoscopy. We evaluate the proposed method on the two public colonoscopy datasets, including Kvasir-SEG [Jha *et al.*, 2020] and CVC-ClinicDB [Bernal *et al.*, 2015]. Kvasir-SEG and CVC-ClinicDB contain 1000 and 612 colonoscopy images with corresponding annotations, respectively.

Cardiac MRI. We evaluate our method on the 3D left atrial (LA) segmentation challenge dataset, which consists of 100 3D gadolinium-enhanced magnetic resonance images and LA segmentation masks for training and validation. Following [Yu *et al.*, 2019a], we split the 100 scans into 80 samples for training and 20 samples for evaluation.

Implementation Details

We use DeepLabv3+ [Chen *et al.*, 2018] equipped with ResNet50 as the baseline architecture for 2D image segmentation, whereas adopt VNet [Milletari *et al.*, 2016] as the baseline in the 3D scenario. All images are resized to 256×256 for inference, while the outputs are recovered to the original size for evaluation, in the 2D scenario. For 3D image segmentation, we randomly crop $80 \times 112 \times 112$ ($Depth \times Height \times Width$) patches for training and iteratively crop patches using a sliding window strategy to obtain the final segmentation mask for testing. We empirically set $\lambda_m = 1$, $k = 2$, $r = 16$, $\alpha = 0.99$ and $\beta = 0.99$ for our method in the experiments. We implement our method using PyTorch framework on a NVIDIA Quadro RTX 6000 GPU. We adopt AdamW as an optimizer with the fixed learning rate of $1e-4$. The batchsize is set to 16, including 8 labeled samples and 8 unlabeled samples. All 2D models are trained for 50 epochs, while the 3D models are trained for 1000 epochs¹.

¹Since UCMT performs the two-step training within one iteration, it is trained for half of the epochs.

Method	5% $\mathcal{D}_{\mathcal{L}}$	10% $\mathcal{D}_{\mathcal{L}}$
MT [Tarvainen and Valpola, 2017]	86.67	87.42
CCT [Ouali <i>et al.</i> , 2020]	83.97	86.43
CPS [Chen <i>et al.</i> , 2021]	86.81	87.70
UGCL(U-Net) [Wang <i>et al.</i> , 2022]	72.67	79.48
UCMT(U-Net) (ours)	82.14	83.33
CMT (ours)	87.86	88.10
UCMT (ours)	88.22	88.46

Table 1: Comparison with state-of-the-art methods on ISIC dataset. 5% $\mathcal{D}_{\mathcal{L}}$ and 10% $\mathcal{D}_{\mathcal{L}}$ of the labeled data are used for training, respectively. Results are measured by DSC.

Method	Kvasir-SEG		CVC-ClinicDB	
	15% $\mathcal{D}_{\mathcal{L}}$	30% $\mathcal{D}_{\mathcal{L}}$	15% $\mathcal{D}_{\mathcal{L}}$	30% $\mathcal{D}_{\mathcal{L}}$
AdvSemSeg	56.88	76.09	68.39	75.93
ColAdv	76.76	80.95	82.18	89.29
MT	87.44	88.72	84.19	84.40
CCT	81.14	84.67	74.20	78.46
CPS	86.44	88.71	85.34	86.69
CMT (ours)	88.08	88.61	85.88	86.83
UCMT (ours)	88.68	89.06	87.30	87.51

Table 2: Comparison with state-of-the-art methods on Kvasir-SEG and CVC-ClinicDB datasets in terms of DSC. 15% $\mathcal{D}_{\mathcal{L}}$ and 30% $\mathcal{D}_{\mathcal{L}}$ of the labeled data are individually used for training.

4.2 Comparison with State of the Arts

We compare the proposed method with state-of-the-art on the four public medical image segmentation datasets. We reimplement MT [Tarvainen and Valpola, 2017], CCT [Ouali *et al.*, 2020], and CPS [Chen *et al.*, 2021] by adopting implementations from [Chen *et al.*, 2021]. For other approaches, we directly use the results reported in their original papers.

Results on Dermoscopy. In Table 1, we report the results of our methods on ISIC and compare them with other state-of-the-art approaches. UCMT substantially outperforms all previous methods and sets new state-of-the-art of 88.22% DSC and 88.46% DSC under 5% and 10% labeled data. For fair comparison with UGCL [Wang *et al.*, 2022], replace the backbone of UCMT with U-Net. The results indicate that our UCMT(U-Net) exceeds UGCL by a large margin. Moreover, our CMT version also outperforms other approaches under the two labeled data rates. For example, CMT surpasses MT and CPS by 1.19% and 1.08% on 5% $\mathcal{D}_{\mathcal{L}}$ labeled data, showing the superiority of collaborative mean-teacher against the current consistency learning framework. By introducing UMIX, UCMT consistently increases the performance under different labeled data rates, which implies that promoting model disagreement and guaranteeing high-confident pseudo labels are beneficial for semi-supervised segmentation.

Results on Colonoscopy. We further conduct a comparative experiment on the polyp segmentation task from colonoscopy images. Table 2 reports the quantitative results on both Kvasir-SEG and CVC-ClinicDB datasets. Compared with the adversarial learning-based [Hung *et al.*, 2018; Wu *et al.*, 2021a] and consistency learning-based [Tarvainen and Valpola, 2017; Ouali *et al.*, 2020; Chen *et al.*, 2021] al-

Method	10% $\mathcal{D}_{\mathcal{L}}$				20% $\mathcal{D}_{\mathcal{L}}$			
	DSC	Jaccard	95HD	ASD	DSC	Jaccard	95HD	ASD
UA-MT [Yu <i>et al.</i> , 2019a]	84.25	73.48	13.84	3.36	88.88	80.21	7.32	2.26
SASSNet [Li <i>et al.</i> , 2020a]	87.32	77.72	9.62	2.55	89.54	81.24	8.24	2.20
LG-ER-MT [Hang <i>et al.</i> , 2020]	85.54	75.12	13.29	3.77	89.62	81.31	7.16	2.06
DUWM [Wang <i>et al.</i> , 2020]	85.91	75.75	12.67	3.31	89.65	81.35	7.04	2.03
DTC [Luo <i>et al.</i> , 2021]	86.57	76.55	14.47	3.74	89.42	80.98	7.32	2.10
MC-Net [Wu <i>et al.</i> , 2021b]	87.71	78.31	9.36	2.18	90.34	82.48	6.00	1.77
MT [Tarvainen and Valpola, 2017]	86.15	76.16	11.37	3.60	89.81	81.85	6.08	1.96
CPS [Chen <i>et al.</i> , 2021]	86.23	76.22	11.68	3.65	88.72	80.01	7.49	1.91
CMT (ours)	87.23	77.83	7.83	2.23	89.88	81.74	6.07	1.94
UCMT (ours)	88.13	79.18	9.14	3.06	90.41	82.54	6.31	1.70

Table 3: Comparison with state-of-the-art methods on LA dataset. 10% $\mathcal{D}_{\mathcal{L}}$ and 20% $\mathcal{D}_{\mathcal{L}}$ of the labeled data are used for training.

gorithms, the proposed methods achieve the state-of-the-art performance. For example, both CMT and UCMT outperform AdvSemSeg [Hung *et al.*, 2018] and ColAdv [Wu *et al.*, 2021a] by large margins on Kvasir-SEG and CVC-ClinicDB, except that ColAdv shows the better performance of 89.29% on CVC-ClinicDB under 30% labeled data. These results demonstrate that our uncertainty-guided collaborative mean-teacher scheme performs better than the adversarial learning and consistency learning schemes commonly used in the compared approaches. Notably, CMT and UCMT show better performance on the low-data regime, i.e., 15% $\mathcal{D}_{\mathcal{L}}$, and the performance between 15% $\mathcal{D}_{\mathcal{L}}$ and 30% $\mathcal{D}_{\mathcal{L}}$ labeled data is close. This phenomenon reflects the capacity of our method to produce high-quality pseudo labels from unlabeled data for semi-supervised learning, even with less labeled data.

Results on Cardiac MRI. We further evaluate the proposed method in the 3D medical image segmentation task. Table 3 shows the comparison results on the 3D left atrium segmentation from cardiac MRI. The compared approaches are all based on consistency learning and pseudo labeling, including uncertainty-aware [Yu *et al.*, 2019a; Wang *et al.*, 2020], shape-aware [Li *et al.*, 2020a], structure-aware [Hang *et al.*, 2020], dual-task [Luo *et al.*, 2021], and mutual training [Wu *et al.*, 2021b] consistency. It can be observed that UCMT achieves the best performance under both 10% and 20% $\mathcal{D}_{\mathcal{L}}$ in terms of DSC and Jaccard over the state-of-the-art methods. For example, compared with UA-MT [Yu *et al.*, 2019a] and MC-Net [Wu *et al.*, 2021b], UCMT shows 3.88% DSC and 0.43% DSC improvements on the 10% labeled data. The results demonstrate the superiority of our UCMT for 3D medical image segmentation.

4.3 Ablation Study

We conduct an ablation study in terms of network architectures, loss functions and region mix to investigate the effectiveness of each component and analyze the hyperparameters of the proposed method. There are three types of network architectures: 1) teacher-student (TS), 2) student-student (SS), and 3) student-teacher-student in the proposed CMT.

Effectiveness of Each Component. Table 4 reports the performance improvements over the baseline. It shows a trend that the segmentation performance improves when the components, including the STS (student-teacher-student),

Method	TS	SS	STS	\mathcal{L}_{cps}	\mathcal{L}_{mts}	U	DSC
Baseline							83.31
MT	✓				✓		86.67
CPS		✓		✓			86.81
CMTv1			✓		✓		86.84
CMTv2			✓	✓			87.48
CMTv3			✓	✓	✓		87.86
UCMT			✓	✓	✓	✓	88.22

Table 4: Ablation study of the different component combinations with only 5% labeled data on ISIC dataset. TS: teacher-student; SS: student-student; STS: student-teacher-student; \mathcal{L}_{cps} : cross pseudo supervision; \mathcal{L}_{mts} : mean-teacher supervision; U: UMIX.

\mathcal{L}_{cps} , \mathcal{L}_{mts} , and UMIX are introduced into the baseline, and again confirms the necessity of encouraging model disagreement and enhancing the quality of pseudo labels for semi-supervised segmentation. The semi-supervised segmentation model is boosted for two reasons: 1) \mathcal{L}_{cps} , \mathcal{L}_{mts} and the STS architecture that force the model disagreement in CMT for co-training, and 2) UMIX facilitating the model to produce high-confidence pseudo labels. All the components contribute to UCMT to achieve 88.22% DSC. These results demonstrate their effectiveness and complementarity for semi-supervised medical image segmentation. On the other hand, the two groups of comparisons between "TS (teacher-student) + \mathcal{L}_{mts} " (i.e., MT) vs. STS + \mathcal{L}_{mts} (i.e., CMTv1), and between "SS (student-student) + \mathcal{L}_{cps} " (i.e., CPS) vs. "STS + \mathcal{L}_{cps} " (i.e., CMTv2) show that the STS-based approaches yield the improvements of 0.17% and 0.64%, indicating the effectiveness of the STS component. However, the performance gaps are not significant because the STS architecture increases the co-training disagreement but decreases the confidences of pseudo labels. It can be easily found that the results are improved to 87.86% by "STS + \mathcal{L}_{cps} + \mathcal{L}_{mts} " (i.e., CMTv3) and the relative improvements of 1.55% and 1.41% DSC have been obtained by "STS + \mathcal{L}_{cps} + \mathcal{L}_{mts} + UMIX" (i.e., UCMT) compared with MT and CPS. The results demonstrate our hypothesis that maintaining co-training with high-confidence pseudo labels can improve the performance of semi-supervised learning.

Comparison of Different Data Augmentation. We fur-




Figure 3: Comparison of UMIX and CutMix on ISIC (a) and LA (b) dataset under 2%, 5%, 10%, and 20% $\mathcal{D}_{\mathcal{L}}$.

$1/r \backslash k$	1	2	3	4	5
1/16	87.95	88.22	88.12	87.96	88.08
1/4	87.65	88.15	87.80	87.54	87.90
1/8	87.86	87.92	88.03	88.01	87.87

Table 5: Investigation on how the top k and patch size affect the capacity of UMIX on ISIC dataset with 5% labeled data.

ther compare the proposed UMIX, component of our UCMT, with CutMix [Yun *et al.*, 2019] on ISIC and LA datasets with different labeled data to investigate their effects in semi-supervised segmentation. As illustrates in Figure 3, UMIX outperforms CutMix, especial in the low-data regime, i.e., 2% labeled data. The reason for this phenomenon is that CutMix performs random region mix that inevitably introduces noise into the new samples, which reduces the quality of the pseudo labels, while UMIX processes the image regions according to the uncertainty of the model, which facilitates the model to generate more confident pseudo labels.

Parameter Sensitivity Analysis. UMIX has two hyperparameters, i.e., the top k regions for mix and the size of the regions (patches) defined as the ratio of $1/r$ to the image size. We study the influence of these factors to UCMT on ISIC dataset with 5% $\mathcal{D}_{\mathcal{L}}$. It can be observed in Table 5 that reducing the patch size leads to a slight increase in performance. Moreover, varying the number of k does not bring us any improvement, suggesting that we can choose any value of K to eliminate outliers, thus bringing high-confidence pseudo labels for semi-supervised learning, which indicates the robustness of UMIX.

4.4 Qualitative Results

Figure 4 visualizes some example results of polyp segmentation, skin lesion segmentation, and left atrial segmentation. As shown in Figure 4 (a), the supervised baseline insufficiently segments some lesion regions, mainly due to the limited number of labeled data. Moreover, MT [Figure 4 (b)] and CPS [Figure 4 (c)] typically under-segment certain objects, which can be attributed to the limited generalization capability. On the contrary, our CMT [Figure 4 (e)] corrects these errors and produces smoother segment boundaries by gaining more effective supervision from unlabeled data. Besides, our complete method UCMT [Figure 4 (f)] further generates more accurate results by recovering finer segmentation details through more efficient training. These examples

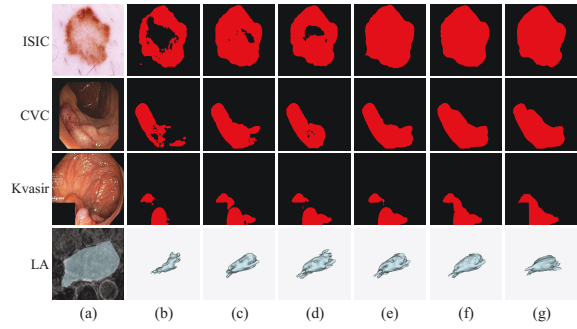


Figure 4: Qualitative examples on the four datasets. (a) images, (b) supervised baseline, (c) MT, (d) CPS, (e) CMT, (f) UCMT, and (g) ground truth.




Figure 5: Illustration of the uncertainty maps. (a) and (b) are the uncertainty maps of the original images and the UMIX images for the first example, while (c) and (d) are for the second example.

qualitatively verify the robustness of the proposed UCMT. In addition, to clearly give an insight into the procedure of the pseudo label generation and utilization in the co-training SSL method, we illustrate the uncertainty maps for two samples during the training in Figure 5. As shown, UCMT generates the uncertainty maps with high uncertainty [Figure 5 (a)/(c)] in the early training stage whereas our model produces relative higher confidence maps [Figure 5 (b)/(d)] from the UMIX images. During training, UCMT gradually improves the confidence for the input images. These results prove that UMIX can facilitate SSL models to generate high-confidence pseudo labels during training, guaranteeing that UCMT is able to maintain co-training in a more proper way.

5 Conclusion

We present an uncertainty-guided collaborative mean-teacher for semi-supervised medical image segmentation. Our main ideas lies in maintaining co-training with high-confidence pseudo labels to improve the capability of the SSL models to explore information from unlabeled data. Extensive experiments on four public datasets demonstrate the effectiveness of this idea and show that the proposed UCMT can achieve state-of-the-art performance. In the future, we will investigate more deeply the underlying mechanisms of co-training for more effective semi-supervised image segmentation.

Ethical Statement

There are no ethical issues.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant 62076059 and the Natural Science Foundation of Liaoning Province under Grant 2021-MS-105.

References

- [Bernal *et al.*, 2015] Jorge Bernal, F Javier Sánchez, Gloria Fernández-Esparrach, Debora Gil, Cristina Rodríguez, and Fernando Vilariño. Wm-dova maps for accurate polyp highlighting in colonoscopy: Validation vs. saliency maps from physicians. *Computerized medical imaging and graphics*, 43:99–111, 2015.
- [Chen *et al.*, 2018] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018.
- [Chen *et al.*, 2021] 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.
- [Codella *et al.*, 2018] Noel CF Codella, David Gutman, M Emre Celebi, Brian Helba, Michael A Marchetti, Stephen W Dusza, Aadi Kalloo, Konstantinos Liopyris, Nabin Mishra, Harald Kittler, et al. Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic). In *2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018)*, pages 168–172. IEEE, 2018.
- [DeVries and Taylor, 2017] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- [Gal and Ghahramani, 2016] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016.
- [Grandvalet and Bengio, 2004] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. *Advances in neural information processing systems*, 17, 2004.
- [Han *et al.*, 2018] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. *Advances in neural information processing systems*, 31, 2018.
- [Hang *et al.*, 2020] 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*, pages 562–571. Springer, 2020.
- [Hung *et al.*, 2018] Wei Chih Hung, Yi Hsuan Tsai, Yan Ting Liou, Yen-Yu Lin, and Ming Hsuan Yang. Adversarial learning for semi-supervised semantic segmentation. In *29th British Machine Vision Conference, BMVC 2018*, 2018.
- [Jha *et al.*, 2020] Debesh Jha, Pia H Smedsrud, Michael A Riegler, Pål Halvorsen, Thomas de Lange, Dag Johansen, and Håvard D Johansen. Kvasir-seg: A segmented polyp dataset. In *International Conference on Multimedia Modeling*, pages 451–462. Springer, 2020.
- [Ke *et al.*, 2020] Zhanghan Ke, Di Qiu, Kaican Li, Qiong Yan, and Rynson WH Lau. Guided collaborative training for pixel-wise semi-supervised learning. In *European conference on computer vision*, pages 429–445. Springer, 2020.
- [Lee and others, 2013] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, page 896, 2013.
- [Li *et al.*, 2020a] 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*, pages 552–561. Springer, 2020.
- [Li *et al.*, 2020b] 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*, 32(2):523–534, 2020.
- [Liu *et al.*, 2022] 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.
- [Luo *et al.*, 2021] 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.
- [Milletari *et al.*, 2016] 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)*, pages 565–571. IEEE, 2016.
- [Ouali *et al.*, 2020] Yassine Ouali, Céline Hudelot, and Myriam Tami. Semi-supervised semantic segmentation

- with cross-consistency training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12674–12684, 2020.
- [Ronneberger *et al.*, 2015] 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.
- [Sohn *et al.*, 2020] 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.
- [Tarvainen and Valpola, 2017] 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.
- [Tu *et al.*, 2022] Peng Tu, Yawen Huang, Feng Zheng, Zhenyu He, Liujuan Cao, and Ling Shao. Guidedmix-net: Semi-supervised semantic segmentation by using labeled images as reference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2379–2387, 2022.
- [Van Engelen and Hoos, 2020] Jesper E Van Engelen and Holger H Hoos. A survey on semi-supervised learning. *Machine Learning*, 109(2):373–440, 2020.
- [Wang *et al.*, 2020] 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*, pages 542–551. Springer, 2020.
- [Wang *et al.*, 2022] Tao Wang, Jianglin Lu, Zhihui Lai, Jiajun Wen, and Heng Kong. Uncertainty-guided pixel contrastive learning for semi-supervised medical image segmentation. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI*, pages 1444–1450, 2022.
- [Wu *et al.*, 2021a] Huisi Wu, Guilian Chen, Zhenkun Wen, and Jing Qin. Collaborative and adversarial learning of focused and dispersive representations for semi-supervised polyp segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3489–3498, 2021.
- [Wu *et al.*, 2021b] Yicheng Wu, Minfeng Xu, Zongyuan Ge, Jianfei Cai, and Lei Zhang. Semi-supervised left atrium segmentation with mutual consistency training. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 297–306. Springer, 2021.
- [Wu *et al.*, 2022] Yicheng Wu, Zongyuan Ge, Donghao Zhang, Minfeng Xu, Lei Zhang, Yong Xia, and Jianfei Cai. Mutual consistency learning for semi-supervised medical image segmentation. *Medical Image Analysis*, 81:102530, 2022.
- [Yu *et al.*, 2019a] 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.
- [Yu *et al.*, 2019b] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor Tsang, and Masashi Sugiyama. How does disagreement help generalization against label corruption? In *International Conference on Machine Learning*, pages 7164–7173. PMLR, 2019.
- [Yun *et al.*, 2019] Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6023–6032, 2019.
- [Zhang *et al.*, 2021] 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.
- [Zheng and Yang, 2021] Zhedong Zheng and Yi Yang. Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. *International Journal of Computer Vision*, 129(4):1106–1120, 2021.