

## Article

# Semi-Supervised Medical Image Segmentation with Co-Distribution Alignment

Tao Wang <sup>1,2,3</sup>, Zhongzheng Huang <sup>2</sup>, Jiawei Wu <sup>4</sup> , Yuanzheng Cai <sup>1</sup>, Zuoyong Li <sup>1,\*</sup>

<sup>1</sup> Fujian Provincial Key Laboratory of Information Processing and Intelligent Control, College of Computer and Control Engineering, Minjiang University, Fuzhou 350108, China; twang@mju.edu.cn (T.W.)

<sup>2</sup> College of Computer and Data Science, Fuzhou University, Fuzhou 350108, China

<sup>3</sup> The Key Laboratory of Cognitive Computing and Intelligent Information Processing of Fujian Education Institutions, Wuyi University, Wuyishan 354300, China

<sup>4</sup> School of Electrical and Mechanical Engineering, Fujian Agriculture and Forestry University, Fuzhou 350002, China

\* Correspondence: fzulzytdq@126.com

**Abstract:** Medical image segmentation has made significant progress when a large amount of labeled data are available. However, annotating medical image segmentation datasets is expensive due to the requirement of professional skills. Additionally, classes are often unevenly distributed in medical images, which severely affects the classification performance on minority classes. To address these problems, this paper proposes Co-Distribution Alignment (Co-DA) for semi-supervised medical image segmentation. Specifically, Co-DA aligns marginal predictions on unlabeled data to marginal predictions on labeled data in a class-wise manner with two differently initialized models before using the pseudo-labels generated by one model to supervise the other. Besides, we design an over-expectation cross-entropy loss for filtering the unlabeled pixels to reduce noise in their pseudo-labels. Quantitative and qualitative experiments on three public datasets demonstrate that the proposed approach outperforms existing state-of-the-art semi-supervised medical image segmentation methods on both the 2D CaDIS dataset and the 3D LGE-MRI and ACDC datasets, achieving an mIoU of 0.8515 with only 24% labeled data on CaDIS, and a Dice score of 0.8824 and 0.8773 with only 20% data on LGE-MRI and ACDC, respectively.



**Citation:** Wang, T.; Huang, Z.; Wu, J.; Cai, Y.; Li, Z. Semi-Supervised Medical Image Segmentation with Co-Distribution Alignment. *Bioengineering* **2023**, *10*, 869. <https://doi.org/10.3390/bioengineering10070869>

Academic Editor: Giuseppe Baselli

Received: 13 June 2023

Revised: 13 July 2023

Accepted: 14 July 2023

Published: 21 July 2023



**Copyright:** © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

**Keywords:** medical image segmentation; semi-supervised learning; distribution alignment; co-training

## 1. Introduction

Currently, there are various new technologies and devices that assist in clinical diagnostic work [1–3], among which medical image segmentation plays an important role in clinical auxiliary diagnosis [4]. Recently, researchers have made great efforts in medical image segmentation [5–7] and achieved excellent performance with a large amount of labeled data. However, the annotation of medical data are typically dependent on medical professionals, and annotating large datasets is time-consuming.

To address this problem, semi-supervised (SS) medical image segmentation leverages a large amount of unlabeled data in conjunction with a small amount of labeled data to improve model performance. Particularly, unlabeled data are relatively affordable, as the laborious annotation process can be avoided. Recently, consistency regularization methods [8–12] have received great attention in SS medical image segmentation. The primary difference of various consistency regularization methods lies in their intended objectives. For example, the majority of perturbation consistency methods [10] tend to maintain a consistent unlabeled prediction with various augmentations. Uncertainty-aware methods [8,12] force consistent predictions in reliable regions of two corresponding models. Multi-task based methods [9,11] design a multi-task framework to guarantee an invariant relation of unlabeled data among different tasks. Despite the fact that these consistency

regularization methods have obtained encouraging results, most of the previous methods neglect two essential problems: class imbalance and the mismatch of class distributions between labeled and unlabeled data. For the first problem, there have been some effective and efficient solutions [13,14] to deal with class imbalance in the fully supervised scenario. However, these methods are unsuitable for the semi-supervised case because long-tail samples and noisy samples are usually difficult to identify in unlabeled data. For the second problem, ReMixMatch [15] proposed a coefficient transform to align labeled and unlabeled class distributions for SS classification. One of their key limitations, however, is that it only considers the empirical ground-truth class distribution, which could be highly imbalanced or even biased when the labeled data are scarce. In addition, estimating the labeled distributions may lead to an unaffordable computational cost for dense prediction tasks such as image segmentation.

In this paper, we focus on maintaining consistent distributions of labeled and unlabeled data under a co-training framework, unlike existing consistency regularization methods. However, it is non-trivial to achieve class distribution consistency. In particular, the empirical class distribution could be highly imbalanced or even biased when the data are sparsely labeled. Figure 1a shows the vanilla Distribution Alignment (DA) [15] in which an overall output class distribution on unlabeled data are maintained to align model outputs to the empirical ground-truth class distribution. More specifically, DA maintains a running average of predictions on unlabeled data. When the model outputs a prediction for an unlabeled sample, the distribution alignment scales the prediction by the empirical class distribution on labeled data over the average predictions on unlabeled data, so as to obtain an output that is aligned to the ground-truth class distribution. On the other hand, Figure 1b shows the cross-pseudo supervision [16] based on co-training [17] that uses two parallel networks with identical architecture but different initializations, and then uses the output of one network to supervise the other one. Inspired by ReMixMatch and cross-pseudo supervision, we propose a novel Co-Distribution Alignment (Co-DA) method to overcome class imbalance and the mismatch of class distributions by integrating the above methods into a unified learning framework. Specifically, Co-DA transforms the model output according to the ratio of the class-specific marginal distribution on labeled data over the average model predictions on unlabeled data for that same class and supervise the model from the other view as pseudo-labels. Different from ReMixMatch, Co-DA uses an exponential moving average (EMA) to simplify the estimation of the labeled distribution. More importantly, instead of relying solely on a single empirical ground-truth class distribution, we seek to fully exploit the model prediction for all classes, aiming to minimize the class-dependent distribution discrepancy between the model outputs on the labeled and the unlabeled data. Therefore, we preserve independent distributions for each class instead of an overall distribution as shown in Figure 1a,c. On the other hand, in contrast to typical co-training, Co-DA tends to keep the consistency between labeled and unlabeled distributions rather than the cross-consistency of predictions as shown in Figure 1b,c. Therefore, Co-DA is more computationally efficient for dense prediction tasks such as image segmentation and is less likely to be affected by errors at individual pixels within the co-training framework. To further reduce the impact from inaccuracies in the model prediction on unlabeled data, we design an over-expectation cross-entropy loss to filter out noises in pseudo-labels.

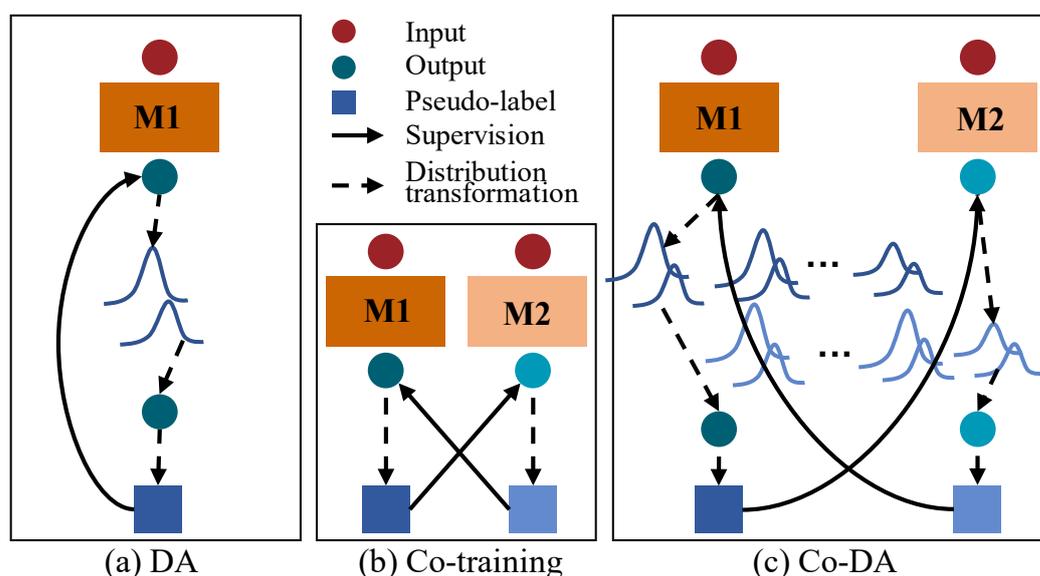
Our main contributions can be summarized as follows:

- To our knowledge, we are the first to solve the semi-supervised medical image segmentation problem with distribution alignment. In particular, we propose class-wise distribution alignment that utilizes the class-dependent output distribution instead of the overall empirical ground-truth class distribution, which could be highly imbalanced and biased when the labeled data are scarce.
- Our co-distribution alignment framework is more computationally efficient for dense prediction tasks involving a large number of pixels as compared to typical co-training

methods such as CPS [16]. More importantly, distribution alignment has better regularization, as the proposed method provides superior performance.

- To further reduce the impact from inaccurate predictions on unlabeled data, we propose a simple, yet effective over-expectation cross-entropy loss to filter out noises in pseudo-labels.
- Experimental evaluation results on three publicly available medical imaging datasets demonstrate the superior performance of our approach compared to the state-of-the-art methods. Moreover, ablation studies also verify the efficacy of the various components in Co-DA.
- Our method does not depend on a particular deep network architecture. Therefore, it can be used in conjunction with different models for medical image segmentation as a plug-and-play module to address the challenges of learning from imbalanced data and the distribution mismatch between labeled and unlabeled data.

The rest of the paper is organized as follows. Section 2 reviews recent literature in the areas of deep semi-supervised learning, semi-supervised medical image segmentation, co-training and distribution alignment methods. Section 3 describes the proposed method in detail, followed by experimental evaluation in Section 4 and closing remarks in Section 5.



**Figure 1.** Illustration of a comparison of Co-Distribution Alignment (Co-DA) with Distribution Alignment (DA) and co-training. (a) is the DA in ReMixMatch [15], (b) is co-training [16], (c) is the proposed Co-DA. M1 and M2 here stand for two differently initialized networks. The dark blue circles and squares denote the output and the pseudo-labels of M1. The light blue circles and squares denote the output and the pseudo-labels of M2.

## 2. Related Work

In this section, we review related literature in semi-supervised medical image segmentation. We first review recent work in deep semi-supervised learning, and then more specifically in semi-supervised medical image segmentation, as well as progress in co-training and distribution alignment methods that are closely related to our approach.

### 2.1. Deep Semi-Supervised Learning

Recently, semi-supervised learning (SSL) has made remarkable progress in various machine learning tasks. SSL methods can be broadly categorized into pseudo-labeling methods, consistency regularization methods, entropy minimization methods and hybrid methods. Specifically, pseudo-labeling methods [18–21] aim at obtaining the pseudo-labels of unlabeled data by self-training. Consistency regularization methods [9,22–24] force similar predictions under different perturbations of unlabeled data to expand the decision

regions. Entropy minimization methods [25,26] tend to make the decision boundary follow low density regions with the help of unlabeled data. Hybrid methods [15,27,28] simultaneously combine some advantages of the above SSL methods. Our method belongs to both pseudo-labeling and consistency regularization methods, yet it addresses a critical problem that is largely ignored in existing methods. To be specific, most existing methods neglect class imbalance and the mismatch of class distributions between labeled data and unlabeled data, which are common in medical images and severely affect the performance of SSL methods.

## 2.2. Semi-Supervised Medical Image Segmentation

The distinctive appearance and class distribution characteristics of medical images pose unique challenges in applying SSL methods to them. In particular, medical image segmentation usually involves localizing objects with extreme shape and scale variations, and the class distribution could be highly skewed. Similar to generic SSL methods, common semi-supervised medical image segmentation methods include GAN-based, consistency regularization and pseudo-labeling methods. Specifically, GAN-based methods [29,30] attempt to use adversarial training to fool the discriminator with unlabeled data. For example, DCT-Seg [31] utilizes two models that are co-trained to generate pseudo-labels for each other. In addition, SS-Net [32] proposes a collaborative learning method to jointly improve the performance of disease grading and lesion segmentation with an attention mechanism. On the other hand, different from consistency regularization methods in generic SSL, for medical image segmentation, people usually design a strategy to keep consistency in local regions, e.g., regions of low uncertainty [12], regions of random category [33], etc. Finally, pseudo-labeling methods leverage an auxiliary model to generate pseudo-labels for unlabeled data. Unlike existing methods, our Co-DA focuses on exploring the discrepancy between labeled and unlabeled data distributions, which is of vital importance in SS medical image segmentation due to data scarcity and imbalance.

## 2.3. Co-Training

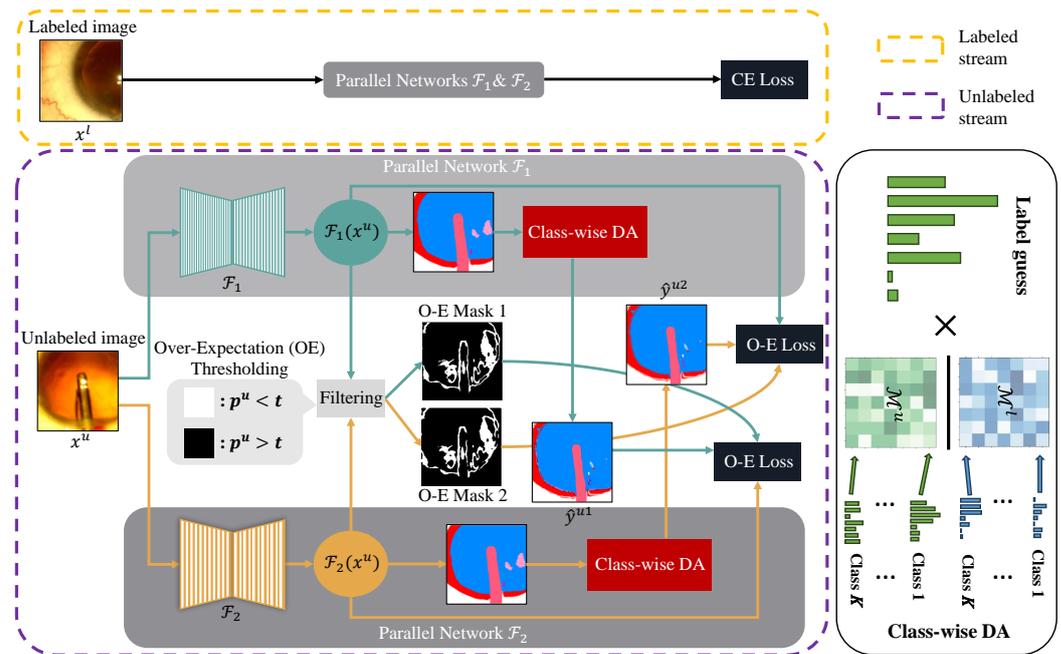
The original co-training algorithm [17] assumes that there are two naturally segmented views of the same instance, which are redundant and independent under certain conditions. Concretely, data from any of them are sufficient to train a strong learner, and the views are independent of each other. More specifically, the main steps of the co-training algorithms are divided into view acquisition, learner differentiation and label confidence estimation [34]. Recent progresses in co-training [31,35,36] primarily focus on maintaining the diversity across models with deep networks. In particular, cross-pseudo supervision [16] proposes to impose consistency on two segmentation networks perturbed with different initialization for the same input image. In contrast, Co-DA uses co-training to align the distributions of labeled and unlabeled predictions.

## 2.4. Distribution Alignment

Distribution alignment is widely used in domain adaptation. These methods work by aligning marginal distributions [37,38] or joint distributions [39–41] of different domains. In semi-supervised learning, domain alignment has also been considered to close the gap between predictions on labeled and unlabeled data. Compared to our proposed approach, the most relevant existing method is ReMixMatch [15], which uses a coefficient transformation to align the marginal distributions of labeled and unlabeled data. However, class distribution in medical images can be highly imbalanced, while ReMixMatch treats samples from all classes as a whole and neglects the variations in marginal distribution of each class. Different from ReMixMatch, in this work, we align marginal distributions in a class-wise manner that is more general and works better for imbalanced datasets with minority classes.

### 3. Our Approach

In this section, we describe the proposed approach in detail. We begin by revisiting the cross-pseudo supervision framework that our work is based on. Afterwards, we describe the three main components of Co-Distribution Alignment, i.e., marginal distribution estimation, class-wise distribution estimation and distribution transformation. In addition, we introduce an over-expectation cross-entropy loss that is used to further improve model performance by filtering out inaccurate pseudo-labels. The overall framework is presented in Figure 2.



**Figure 2.** Overview of Co-DA for semi-supervised medical image segmentation, which consists of a stream for labeled data and a stream for unlabeled data.  $\mathcal{F}_1$  and  $\mathcal{F}_2$  are two parallel networks with the same architecture but different initializations. For labeled data, the standard cross-entropy loss (CE Loss) is used. For unlabeled data,  $\hat{y}^{u1}$  and  $\hat{y}^{u2}$  are pseudo-labels given by Equation (8), and we use the aligned output distribution from one network to supervise the other one. For the over-expectation cross-entropy loss (O-E Loss) in Equation (10),  $p^u = \max(\mathcal{F}(x^u))$  denotes the class probability corresponding to the pseudo-labels,  $t$  stands for the threshold given by Equation (9) and the O-E Masks are the thresholded binary masks to filter out pixels with low confidence in their pseudo-labels. We use class-wise DA to align marginal distributions for each class to better deal with imbalanced datasets. Here,  $\mathcal{M}^l$  and  $\mathcal{M}^u$  denote the labeled and the unlabeled marginal distributions.

#### 3.1. Cross-Pseudo Supervision

Cross-pseudo supervision works by generating two views of the same input image with differently initialized segmentation networks. In this way, one network can discover the self-mistakes made by the other model [16,17]. Following [16], we also use two parallel networks  $\mathcal{F}_1$ ,  $\mathcal{F}_2$  with the same structure but different initialization for collaborative training. To define our problem more formally, and without loss of generality, we view the medical image segmentation problem as a pixel-wise classification one since our focus in this paper is the cross-alignment of marginal distributions. More specifically, let  $\mathcal{D}_l = \{(x_i^l, y_i^l)\}_{i=1}^m$  denote the labeled dataset and  $\mathcal{D}_u = \{x_j^u\}_{j=1}^n$  the unlabeled dataset. Here,  $x_i^l$  and  $x_j^u$  are the image pixels, and  $y_i^l$  is the class label for  $x_i^l$ . The loss function on the labeled dataset can be written as:

$$\mathcal{L}_s = -\frac{1}{m} \sum_{i=1}^m y_i^l \left[ \log \mathcal{F}_1(x_i^l) + \log \mathcal{F}_2(x_i^l) \right], \tag{1}$$

For the unlabeled data, we adopt the cross-pseudo supervision loss, which uses the pseudo-labels generated by one model to supervise the prediction of the other model, as follows:

$$\mathcal{L}_u = -\frac{1}{n} \sum_{j=1}^n \left[ \hat{y}_j^{u2} \log \mathcal{F}_1(x_j^u) + \hat{y}_j^{u1} \log \mathcal{F}_2(x_j^u) \right], \tag{2}$$

where  $\hat{y}_j^{u1}, \hat{y}_j^{u2}$  are the pseudo-labels generated by  $\mathcal{F}_1$  and  $\mathcal{F}_2$ , respectively. Here, pseudo-labels are the mode of prediction, i.e., the label derived from the most probable class.

Equation (2) shows that pseudo-labels determine our optimization objective. Therefore, most existing methods [16,36] focus on correcting or generating more accurate pseudo-labels to supervise unlabeled predictions. However, their methods neglect the fundamental challenges of learning from imbalanced data and the distribution mismatch between labeled and unlabeled data. In the next section, we introduce the proposed Co-DA that specifically addresses these problems.

Let us take a closer look at the pseudo-labels in Equation (2). We can conjecture that, ideally, if the pseudo-labels  $\hat{y}_j^{u1}, \hat{y}_j^{u2}$  follow the class distribution on labeled data, applying Equation (2) would already align the unlabeled class distribution to the labeled class distribution. However, due to class imbalance in medical image segmentation, the distribution of pseudo-labels could be biased toward majority classes. In this paper, we propose to address this problem by explicitly aligning the marginal distributions on unlabeled and labeled data after cross-pseudo supervision. In particular, our Co-DA considers  $K$  class-dependent distributions (where  $K$  is the number of classes) instead of a single overall empirical class distribution in naive DA. This approach allows us to capture and align finer class-wise statistics in the marginal distributions, as we describe in the next section.

### 3.2. Co-Distribution Alignment

In this section, we outline the proposed Co-Distribution Alignment in detail. Specifically, we first introduce an efficient method to estimate the class distribution on labeled data during training, and then discuss class-wise distribution estimation for capturing detailed statistics in labeled and unlabeled marginal distributions, followed by the distribution transformation to align the unlabeled distributions to labeled distributions.

#### 3.2.1. Marginal Distribution Estimation

The vanilla DA in ReMixMatch [15] needs to use the empirical ground-truth class distribution for aligning model predictions on unlabeled data to it. However, estimating ground-truth class distribution becomes computationally expensive for segmentation tasks where a single image contains a large number of pixels, and using only a small amount of labeled data is inappropriate, as the estimation may not well represent the overall class distribution. To address this issue, we use the exponential moving average (EMA) to estimate the aggregated class predictions on labeled data, which can be written as:

$$\mathcal{T}^l = \alpha \times \hat{\mathcal{T}}^l + (1 - \alpha) \times \mathbb{E}[\mathcal{F}(x^l; \Theta)], \tag{3}$$

where  $\mathcal{T}^l$  and  $\hat{\mathcal{T}}^l$  are the labeled distributions of the current iteration and the previous iteration,  $\alpha$  determines that  $\mathcal{T}^l$  is the average prediction of last  $1/(1 - \alpha)$  iterations and  $\Theta$  represents the parameters of the network. It should be noted that our work differs from ReMixMatch [15] in that we use EMA to estimate the class predictions on labeled data, which is more computationally efficient for dense prediction tasks. In addition, we note that the MLE of observations can be given by:

$$\begin{aligned} \Theta &= \arg \max_{\theta} \mathcal{F}(\mathcal{X}; \theta) \\ &= \arg \min_{\theta} - \sum_{i=1}^m \log \mathcal{F}(x_i^l; \theta) - \sum_{i=1}^n \log \mathcal{F}(x_i^u; \theta), \end{aligned} \tag{4}$$

where  $\mathcal{X}$  is the entire dataset that includes both labeled and unlabeled data. Obviously,  $\Theta$  is related to the likelihoods of both labeled and unlabeled data. Therefore, the labeled distribution estimated by Equation (3) considers both labeled and unlabeled data to estimate the labeled distribution and only costs manageable computations during training.

### 3.2.2. Class-Wise Distribution Estimation

An important limitation of the original DA is that class imbalance may cause the tail classes to diminish from the estimated distribution, as they only represent a small fraction of labeled data, ultimately leading to a biased distribution transformation. Unlike naive DA [15], Co-DA builds  $K$  independent distributions for each class, where  $K$  is the number of classes and each distribution is class-specific. Therefore, Co-DA is more robust to class imbalance due to the decoupled modeling process. Specifically, we maintain two matrices  $\mathcal{M}^l \in \mathbb{R}^{K \times K}$  and  $\mathcal{M}^u \in \mathbb{R}^{K \times K}$  for labeled and unlabeled distributions, respectively, in which row  $i$  is the distribution for the  $i$ -th class. For the  $i$ -th row of the labeled matrix  $\mathcal{M}_i^l$ , we use Equation (3) to approximate the marginal distribution according to the ground-truth class label as follows:

$$\mathcal{M}_i^l = \alpha \cdot \hat{\mathcal{M}}_i^l + (1 - \alpha) \cdot \mathbb{E}[\mathbb{I}_{[y^l=i]} \mathcal{F}(x^l)], \tag{5}$$

where  $\mathbb{I}_{[i]}(\cdot)$  is the indicator function. For the  $i$ -th row of the unlabeled matrix  $\mathcal{M}_i^u$ , we update the marginal distribution by EMA according to the model prediction following ReMixMatch [15]. The class membership of unlabeled data are obtained by the mode of prediction, i.e., the most probable class. However, we observe that the tail categories in the unlabeled distributions may be difficult to update due to their infrequent presence in a batch. In this case, we use the inverse transformation from labeled distribution to unlabeled distribution to update the unlabeled distribution. The overall update strategy of  $\mathcal{M}_i^u$  can be written as follows:

$$\mathcal{M}_i^u = \begin{cases} \hat{\mathcal{M}}_i^l \times \mathbb{E}[\hat{\mathcal{M}}_i^u / \hat{\mathcal{M}}_i^l], & \text{if } \{\hat{y}^u | \hat{y}^u = i\} = \emptyset, \\ \alpha \times \hat{\mathcal{M}}_i^u + (1 - \alpha) \times \mathbb{E}[\mathbb{I}_{[\hat{y}^u=i]} \mathcal{F}(x^u)], & \text{otherwise,} \end{cases} \tag{6}$$

where  $\hat{y}^u = \arg \max \mathcal{F}(x^u)$  and  $\hat{\mathcal{M}}_i^u$  is the unlabeled distribution in the previous iteration. This update strategy stems from the original distribution alignment proposed in ReMixMatch [15], but here, we use it to estimate the unlabeled distribution when the unlabeled data from a certain class (usually the minority classes) are absent in a batch. In addition, we use EMA to maintain a stable update of the distributions required for our Co-DA in the training process. It should be noted that, in practice, there are two models (i.e.,  $\mathcal{F}_1$  and  $\mathcal{F}_2$ ) in our method, and each model will be used to update their own  $\mathcal{M}^l$  and  $\mathcal{M}^u$ ; we omit their subscripts for notational simplicity.

### 3.2.3. Distribution Transformation

Following ReMixMatch, we align the distributions using coefficient transformation but with each class functioning independently. Besides, we use a temperature  $\tau$  to scale the labeled distribution following CReST [42]. Different from CReST, however, we choose an adaptive temperature according to the class-specific aggregated model prediction. More specifically, the temperature and the transformation of class  $i$  is given by:

$$\tau_i = 1 - \mathcal{M}_{ii}^l, \tag{7}$$

$$[\hat{y}^u = i] \implies \mathcal{F}'(x^u) = \text{Normalize}(\mathcal{F}(x^u) \otimes \mathcal{M}_i^{l^{\tau_i}} \circ \mathcal{M}_i^u) \tag{8}$$

where  $[\hat{y}^u = i]$  denotes that, according to the network, the  $i$ -th class is the most probable, and we therefore use the  $i$ -th row in  $\mathcal{M}^l$  and  $\mathcal{M}^u$ , i.e.,  $\mathcal{M}_i^l$  and  $\mathcal{M}_i^u$ , for the coefficient transformation. In Equation (8),  $\otimes$  and  $\oslash$  denote the element-wise product and division, respectively. In addition,  $\mathcal{F}(x^u)$  and  $\mathcal{F}'(x^u)$  denote the network prediction for  $x^u$  before and after distribution transformation.

We note that the temperature  $\tau_i$  further prevents the transformed distribution to be dominated by majority classes. Specifically,  $\tau_i \rightarrow 0$  makes the labeled distribution closer to uniform distribution, while a larger  $\tau_i$  results in a smaller shift.  $\tau_i$  not only makes the model more robust against noises in pseudo-labels in the early training stage, but also encourages the emergence of minority classes in the middle and late stages.

After distribution transformation, the network prediction after alignment,  $\mathcal{F}'(x^u)$ , is then used for cross-pseudo supervision according to Equation (2), instead of  $\mathcal{F}(x^u)$  as originally shown in Equation (2).

### 3.3. Over-Expectation Cross-Entropy Loss

In order to further reduce the negative impact of inaccurate pseudo-labels, we propose an over-expectation cross-entropy loss for learning from unlabeled data. Motivated by the definition of EMA,  $\mathcal{M}^l$  and  $\mathcal{M}^u$  can be regarded as the expectations of the labeled and unlabeled distributions, respectively. Intuitively, we can use the estimated aggregated model prediction for the  $i$ -th class as an adaptive threshold to filter out unlabeled samples below expectations to reduce noise. More specifically, the threshold  $t(i)$  for the  $i$ -th class is given by:

$$t(i) = \mathcal{M}_{ii}^u \tag{9}$$

As we will show in Section 4.5, this adaptive and dynamic threshold is consistently superior to different static threshold values and do not bring in additional hyperparameters. In this way, Equation (2) can be rewritten as:

$$\mathcal{L}_u = -\mathbb{E}[\mathbb{I}_{[p_i^{u2} > t(\hat{y}_i^{u2})]} \hat{y}_i^{u2} \log \mathcal{F}'_1(x_i^u) + \mathbb{I}_{[p_i^{u1} > t(\hat{y}_i^{u1})]} \hat{y}_i^{u1} \log \mathcal{F}'_2(x_i^u)], \tag{10}$$

where  $p_i^u$  stands for the probability corresponding to the pseudo-labels,  $\hat{y}_i^u$  stands for the class prediction of the pseudo-labels and we refer to our loss in Equation (10) as the over-expectation cross-entropy loss (O-E Loss). In a nutshell, this is an extension to the soft label cross-entropy loss that incorporates the threshold filtering in Equation (9). The logarithm comes from the cross-entropy between two probability distributions, and we refer readers to [43] for further background information. See Figure 2 for the O-E Masks 1 and 2 as an illustration for applying Equation (10) in practice. The complete training process of Co-DA is summarized in Algorithm 1.

---

#### Algorithm 1 Co-Distribution Alignment.

---

- 1: **Input:** Labeled training data  $\{(x_i^l, y_i^l)\}_{i=1}^m$ , unlabeled training data  $\{x_i^u\}_{i=1}^n$ , max iterations  $I_{max}$ , two parallel networks  $\mathcal{F}_1(\Theta_1)$  and  $\mathcal{F}_2(\Theta_2)$ , learning rate  $\eta$ , momentum  $\alpha$  for exponential moving average
  - 2: Initialize  $\mathcal{F}_1(\Theta_1)$  and  $\mathcal{F}_2(\Theta_2)$  with different parameters
  - 3: **for** iteration=1 to  $I_{max}$  **do**
  - 4:     Update the labeled distribution  $\mathcal{M}^l$  by Equation (5);
  - 5:     Update the unlabeled distribution  $\mathcal{M}^u$  by Equation (6);
  - 6:     Generate the pseudo-labels  $\hat{y}^{u1}$  and  $\hat{y}^{u2}$  w.r.t.  $\mathcal{F}_1(\Theta_1)$  and  $\mathcal{F}_2(\Theta_2)$  by Equation (8);
  - 7:     Calculate the supervised loss  $\mathcal{L}_s(x^l, y^l)$  according to Equation (1);
  - 8:     Calculate the unsupervised loss  $\mathcal{L}_u(x^u, \hat{y}^{u1}, \hat{y}^{u2})$  according to Equation (10);
  - 9:     Update  $\Theta_1 = \Theta_1 - \eta \cdot \nabla_{\Theta_1}(\mathcal{L}_s + \mathcal{L}_u)$ ;
  - 10:    Update  $\Theta_2 = \Theta_2 - \eta \cdot \nabla_{\Theta_2}(\mathcal{L}_s + \mathcal{L}_u)$ ;
  - 11: **end for**
  - 12: **Output:**  $\mathcal{F}_1(\Theta_1)$  and  $\mathcal{F}_2(\Theta_2)$ .
-

## 4. Experiments

In this section, we thoroughly verify the efficacy of the proposed Co-DA on three challenging public medical image segmentation datasets. Specifically, we first outline the experimental setup, including an introduction to the datasets, the evaluation metrics and our implementation details in Section 4.1, followed by quantitative and qualitative results on the three datasets in Section 4.2, Section 4.3 and Section 4.4, respectively. We also present ablation studies in Section 4.5 to demonstrate that our method provides a strong performance that is comparable to its fully supervised variant, and the individual components proposed in our method are all contributing to the performance of our model.

### 4.1. Experimental Setup

#### 4.1.1. CaDIS

We first evaluate the performance of our method on a 2D medical image segmentation task. The publicly available CaDIS dataset [44,45] consists of 4671 frames from 25 surgical videos, which are collected by experts and annotated at the pixel level. Following [44], we consider three progressively more difficult semantic segmentation tasks on this dataset. Specifically, task 1 contains 8 classes, with 4 for anatomical structures, 1 for all instruments and 3 for other objects that appear in frames; task 2 contains 17 classes, which divides the single instrument classes in task 1 into 10 more specific classes of instruments; task 3 contains 25 classes, where instruments are further subdivided according to the handles and parts of certain instruments.

To thoroughly demonstrate the efficacy of the proposed method, we compare our method with state-of-the-art competing methods on all three tasks with varying levels of labeled data. This task poses some unique challenges for segmenting the anatomical structures, surgical instruments and other objects (i.e., surgical tapes, hands and eye retractors) simultaneously. In particular, classes are unevenly distributed, and some objects are either thin or small, or both.

#### 4.1.2. Late Gadolinium Enhancement MRI

To further demonstrate the performance of our method on 3D medical image segmentation tasks, we also evaluate our method on the Late Gadolinium Enhancement MRI (LGE-MRI) dataset [46]. This dataset is a collection of 154 3D LGE-MRIs acquired from the Left Atrial Segmentation Challenge, which contains data from 60 patients with atrial fibrillation prior to and post-ablation treatment. The goal is to perform Left Atrium (LA) segmentation, and the images have an isotropic resolution of  $0.625 \times 0.625 \times 0.625$  mm<sup>3</sup>. In particular, it is challenging to segment the LA in the top and the bottom slices, corresponding to the pulmonary veins and the mitral valve, respectively.

#### 4.1.3. ACDC

In addition, we evaluate the performance of our method on another 3D medical image segmentation task using the ACDC (Automated Cardiac Diagnosis Challenge) 2017 dataset [47]. This dataset consists of MRI images from 100 patients with expert annotations. Among these, 2 are used as the validation set, 20 are used as the test set and the rest are used as the training set. The ACDC dataset contains images from patients with normal cardiac anatomy as well as those with previous myocardial infarction, dilated cardiomyopathy, hypertrophic cardiomyopathy and an abnormal right ventricle. One unique challenge in this task lies in identifying the left ventricle, the myocardium and the right ventricle, which could all be very small at certain anatomical structures such as the apex.

#### 4.1.4. Evaluation Metrics

This section presents the evaluation metrics employed to assess the effectiveness of the proposed approach. Firstly, we follow [48] and use the mean Intersection over Union (mIoU) to evaluate the CaDIS dataset, given as follows:

$$mIoU = \frac{1}{c} \sum_{i=0}^c \frac{TP}{TP + FP + FN}, \tag{11}$$

where  $c$  is the number of classes, and TP, FP and FN denote the number of pixels that are true positives, false positives and false negatives, respectively. In addition, Dice score, Jaccard, the average surface distance (ASD) and the 95% Hausdorff Distance (95HD) are adopted to evaluate LGE-MRI [8]. Also, methods are evaluated with Dice score and 95HD for ACDC 2017. Specifically, Dice, Jaccard and ASD are given as follows:

$$Dice = 2TP / (2TP + FP + FN), \tag{12}$$

$$Jaccard = TP / (TP + FP + FN), \tag{13}$$

$$ASD = \frac{\left( \sum_{a \in S(A)} \min_{b \in S(B)} \|a - b\| + \sum_{b \in S(B)} \min_{a \in S(A)} \|b - a\| \right)}{|S(A)| + |S(B)|}, \tag{14}$$

where  $S(\cdot)$  denotes the set of surface voxels, and the two sets  $A$  and  $B$  refer to the ground-truth and network prediction, respectively.  $\|\cdot\|$  denotes the L2 distance. For the Hausdorff distance, it could be written as:

$$HD = \max(h(S(A), S(B)), h(S(B), S(A))), \tag{15}$$

where  $h(S(A), S(B)) = \max_{a \in S(A)} \left\{ \min_{b \in S(B)} \|a - b\| \right\}$ , denoting the maximum distance of voxels in  $A$  to the nearest voxel in  $B$ . The 95% Hausdorff Distance is based on the 95th percentile of the surface distance above in order to eliminate the impact from a small number of outliers.

#### 4.1.5. Implementation Details

For CaDIS, we use the publicly available split strategy [48] of 3550 frames for training and the remaining 1120 for validation. For 12%, 24% and 49% labeled data in our experiments, we randomly choose 424, 834 and 1729 labeled frames, respectively. In addition, we randomly crop the original image to  $352 \times 352$  and augment the data in the same way as in [48], including random horizontal flipping, Gaussian noise and color jittering. To mitigate the effects of data imbalance, we also follow [48] and use Repeat Factor Sampling in the labeled set. For LGE-MRI, we are consistent with [8]'s pre-processing scheme, dividing the original 100 sheets of data into 80 for training and 20 for validation. All data are centrally cropped in the heart region and randomly cropped to  $112 \times 112 \times 80$ . The data augmentation strategy used included random flips and rotations of 90, 180 and 270 degrees. For the ACDC 2017 dataset, we use random flip, rotation and scaling as data augmentation.

The proposed method is implemented with PyTorch [49]. We use an nVIDIA GeForce RTX 3090 GPU for training. We choose to train CaDIS and ACDC datasets with EfficientNet-B3-based [50] Unet [5] and LGE-MRI with Vnet [51] as backbone architectures. Both networks are updated using the SGD optimizer with a momentum of 0.9 and an initial learning rate of 0.0001, with the learning rate divided by  $0.001^{1/epoch}$  for each epoch, where epoch depends on the ratio of iterations to the number of labeled data in training, and the batch size is set to 8. For CaDIS, we set the number of iterations in one epoch to 10,000 when the ratio of labeled data are 12%, 20,000 when the ratio of labeled data are 24% and 50,000

when the ratio of labeled data are 49%; for training LGE-MRI, the number of iterations is set to 6000 uniformly; for ACDC, the number is 30,000 uniformly.

#### 4.2. Results on CaDIS

In order to better verify the effectiveness of the proposed method, we compare it with the following state-of-the-art methods: URPC [52], UAMT [8], CLD [53] and CPS [16]. Among these methods, URPC [52], UAMT [8] and CLD [53] are specifically designed for semi-supervised medical image segmentation, while CPS [16] is a generic semantic segmentation algorithm. We also experiment with a naive setup, i.e., to train the model with only labeled data, denoted as Baseline. As shown in Table 1, Co-DA outperforms other methods on different split settings. Specifically, in Task 1, our method is able to improve 0.1223, 0.1017 and 0.0329 on mean IoU with 12%, 24% and 49% labeled data, respectively. For Task 2 and Task 3, the corresponding improvements are 0.1304, 0.1228, 0.1832 and 0.0845, 0.1181, 0.1189. It should be noted that under Task 1 with 12% labeled data, the performance of our method is slightly lower than UAMT for semi-supervised medical image segmentation and CPS for generic image segmentation, which is probably due to the data distribution being relatively balanced when the amount of data and number of classes are small. However, as the problem gets more difficult when we have more imbalanced datasets with more classes, the performance gain obtained from our method becomes more prominent.

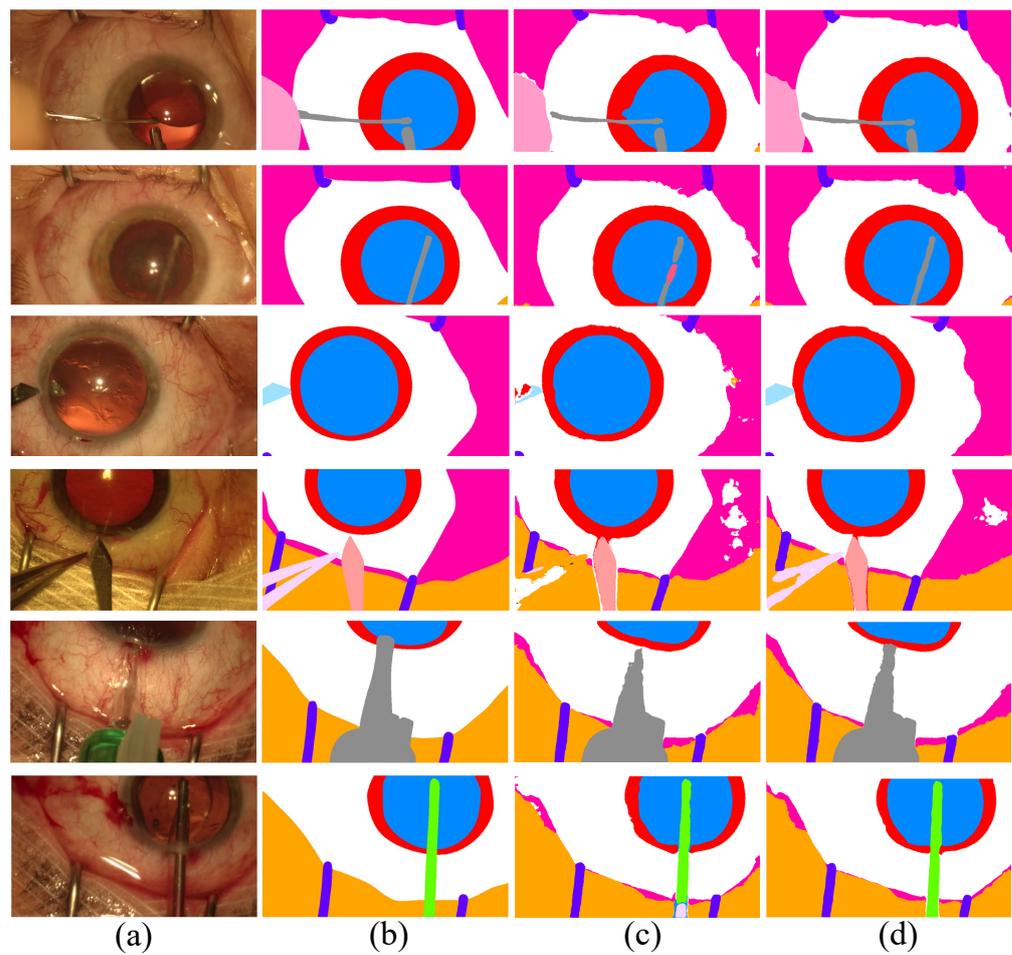
**Table 1.** Mean IoU results on CaDIS. Bold numbers represent the best performance.

Method	12% Labeled Data			24% Labeled Data			49% Labeled Data		
	Task 1	Task 2	Task 3	Task 1	Task 2	Task 3	Task 1	Task 2	Task 3
Baseline	0.5973	0.3719	0.2784	0.7498	0.5413	0.3859	0.8603	0.6066	0.5286
URPC [52]	0.6649	0.4449	0.3169	0.7486	0.5383	0.3886	0.8361	0.6328	0.5414
UAMT [8]	<b>0.7223</b>	0.2953	0.2180	0.7760	0.5288	0.4085	0.8811	0.6925	0.5815
CPS [16]	0.7222	0.4570	0.3525	0.8437	0.6562	0.4821	0.8874	0.7774	0.5837
CLD [53]	0.7120	0.4465	0.3056	0.8376	0.6474	0.4325	0.8857	0.7553	0.5999
<b>Co-DA (Ours)</b>	0.7196	<b>0.5023</b>	<b>0.3629</b>	<b>0.8515</b>	<b>0.6641</b>	<b>0.5040</b>	<b>0.8932</b>	<b>0.7898</b>	<b>0.6475</b>

In Figure 3, we show the qualitative results of Co-DA. The comparison between Co-DA and CLD [53] shows that Co-DA could obtain superior results, especially in the instrument regions. We present the confusion matrices of Co-DA and Baseline on CaDIS Task 1 in Appendix A.

#### 4.3. Results on Late Gadolinium Enhancement MRI

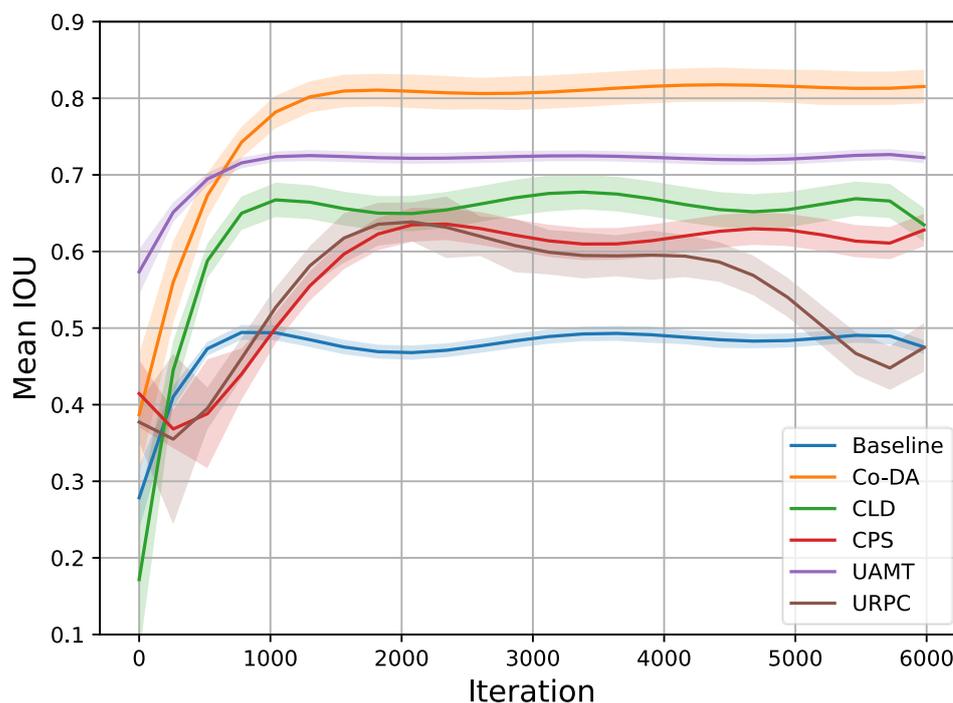
Similar to our setup in the CaDIS experiment, the quantitative results of the experiments on Late Gadolinium Enhancement MRI (LGE-MRI) with 10% and 20% labeled data are shown in Table 2, where we also present the experimental results for all the competing algorithms. Our proposed Co-DA consistently provides a strong performance among all algorithms. With 10% labeled data, Co-DA boosts 21.78% and 26.54% on Dice and Jaccard over the baseline, and reduces 2.86 voxels and 10.7 voxels on ASD and 95HD, respectively. With 20% labeled data, Co-DA consistently provides the best performance, improving 25.3% and 29.48% on Dice and Jaccard over the baseline, and reduces 5.54 voxels and 16.08 voxels on ASD and 95HD, respectively. We note that our method demonstrates the strongest performance on Dice score and Jaccard under all settings. Furthermore, we present the trends in mIoU with 20% data and different methods in Figure 4. It is evident that the proposed Co-DA trains more quickly than other methods, requiring fewer training iterations than CLD, CPS, URPC, etc. The performance of Co-DA also remains stable once peaked.



**Figure 3.** Comparison of the proposed Co-DA and another state-of-the-art method, CLD [53], on CaDIS with 49% labeled data. (a) is the input to the networks, (b) is the ground-truth of the corresponding input, (c) is the output of CLD and (d) is the output of Co-DA.

**Table 2.** Experimental results on LGE-MRI. Bold numbers represent the best performance.

Method	10% Labeled Data				20% Labeled Data			
	Dice (%)	Jaccard (%)	ASD (Voxels)	95HD (Voxels)	Dice (%)	Jaccard (%)	ASD (Voxels)	95HD (Voxels)
Baseline	62.41	46.67	10.27	36.01	62.94	49.70	9.23	30.85
URPC [52]	83.67	63.57	10.27	27.67	81.30	68.90	7.14	24.43
UAMT [8]	66.38	52.04	<b>6.64</b>	<b>21.27</b>	86.81	76.99	4.66	17.89
CPS [16]	69.38	47.77	17.09	27.09	73.29	59.42	8.73	26.39
CLD [53]	65.22	42.26	6.68	23.79	77.03	63.85	5.27	19.66
<b>Co-DA (Ours)</b>	<b>84.19</b>	<b>73.21</b>	7.41	25.31	<b>88.24</b>	<b>79.18</b>	<b>3.69</b>	<b>14.77</b>

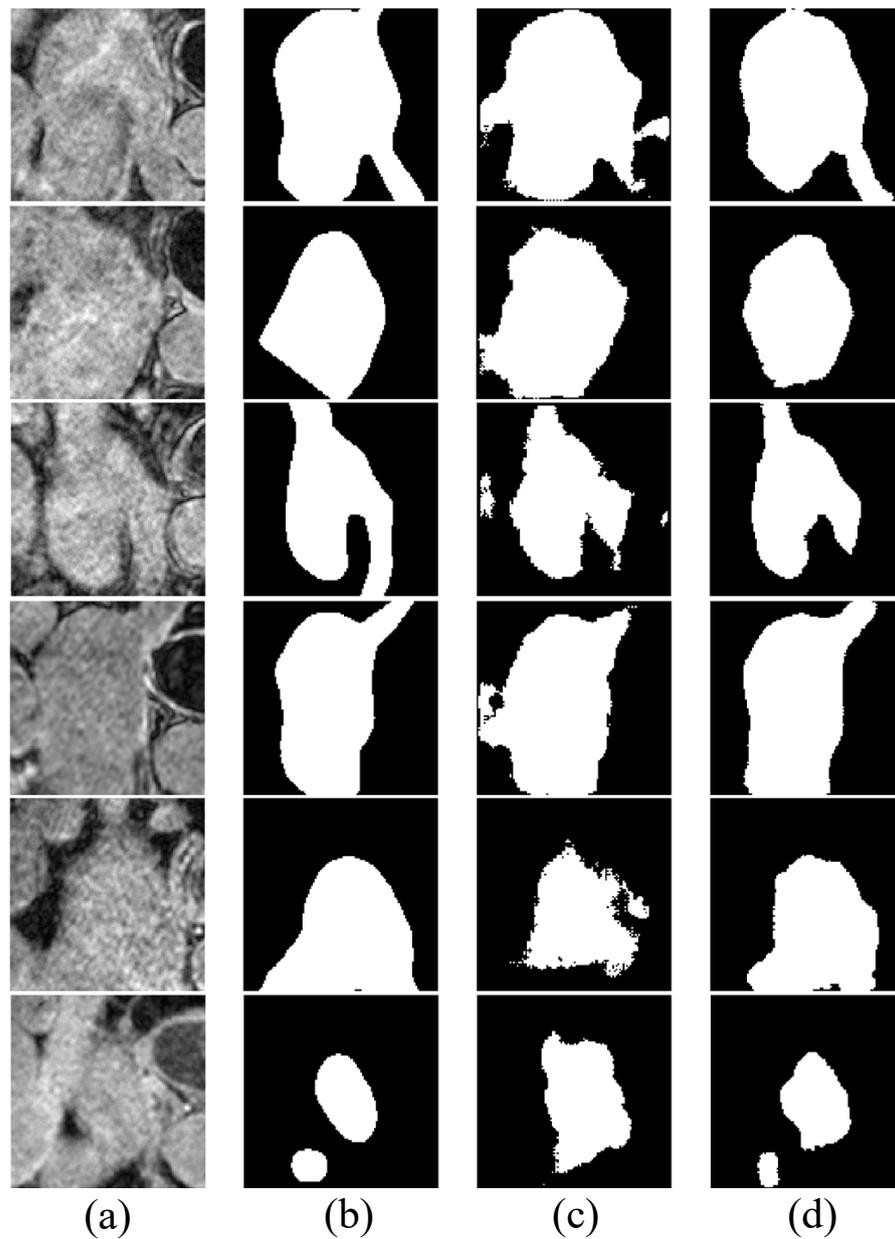


**Figure 4.** Trends in mIoU on the LGE-MRI dataset with 20% labeled data. It is clear that Co-DA not only provides the best performance, but also trains more quickly than most other methods, requiring a smaller number of training iterations. Additionally, the performance is stable once peaked.

Figure 5 shows some examples of the segmentation results on LGE-MRI, using the proposed Co-DA and another state-of-the-art method, CLD [53]. CLD tends to have difficulty in accurately segmenting the pathological regions and will partially misjudge the healthy regions, while our segmentation results of the pathological regions are much closer to the ground-truth. In addition, our method is less likely to generate spurious artefacts commonly found in the results obtained with CLD, showcasing the strong regularization ability of class-wise distribution alignment.

#### 4.4. Results on ACDC

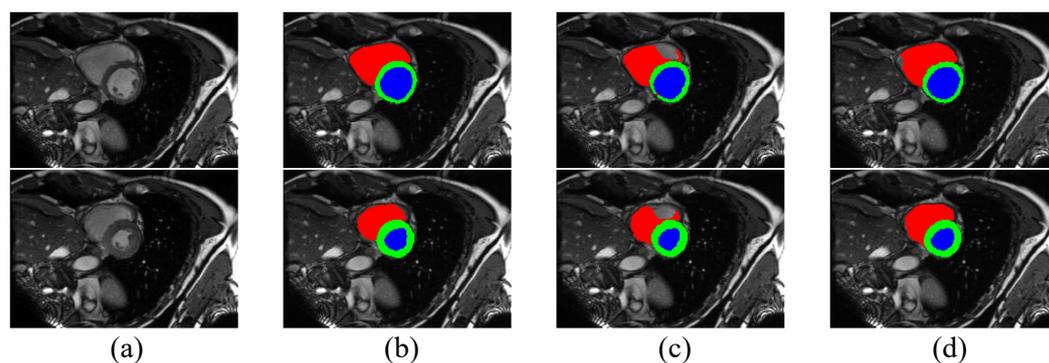
According to Table 3, the performance of our method on the ACDC dataset is better than or comparable to other competing algorithms, which shows that our method is also effective in this segmentation task. With 10% labeled data, our method is able to outperform the baseline with a 12.35% margin on Dice and a 6.18 voxels improvement on 95HD. With 20% labeled data, these performance boosts are 10.3% and 4.63 voxels, respectively. It is clear that our method provides the best performance in terms of 95HD under both settings. However, our method, along with other methods for semi-supervised medical image segmentation, displays a slightly poorer performance than CPS on Dice score, indicating the excellent performance of CPS and showing that distribution alignment for pseudo-labels does not work well enough on the ACDC dataset, which we will continue to explore in future work. Although our method does not provide the best Dice score, the performance in terms of 95HD is better than all other algorithms, demonstrating that Co-DA is able to provide high quality segmentation with accurate boundary delineation. We present example segmentation results of CLD and our proposed Co-DA in Figure 6.



**Figure 5.** Comparison of the proposed Co-DA and another state-of-the-art method, CLD [53], on 2D slices of MRI on LGE-MRI with 20% labeled data, where (a) is the input to the networks, (b) is the corresponding ground-truth, (c) is the output of CLD and (d) is the output of Co-DA.

**Table 3.** Experimental results on ACDC. Bold numbers represent the best performance.

Method	10% Labeled Data		20% Labeled Data	
	Dice (%)	95HD (Voxels)	Dice (%)	95HD (Voxels)
Baseline	70.59	12.11	77.43	9.57
URPC [52]	77.59	6.46	86.07	5.07
UAMT [8]	72.47	15.49	82.68	6.05
CPS [16]	<b>83.96</b>	8.75	<b>87.81</b>	5.95
CLD [53]	83.04	6.13	86.58	5.81
<b>Co-DA(Ours)</b>	82.94	<b>5.93</b>	87.73	<b>4.94</b>



**Figure 6.** Comparison of the proposed Co-DA and another state-of-the-art method, CLD [53], on the ACDC dataset with 10% labeled data, where (a) is the input to the networks, (b) is the corresponding ground-truth, (c) is the output of CLD and (d) is the output of Co-DA.

#### 4.5. Ablation Studies

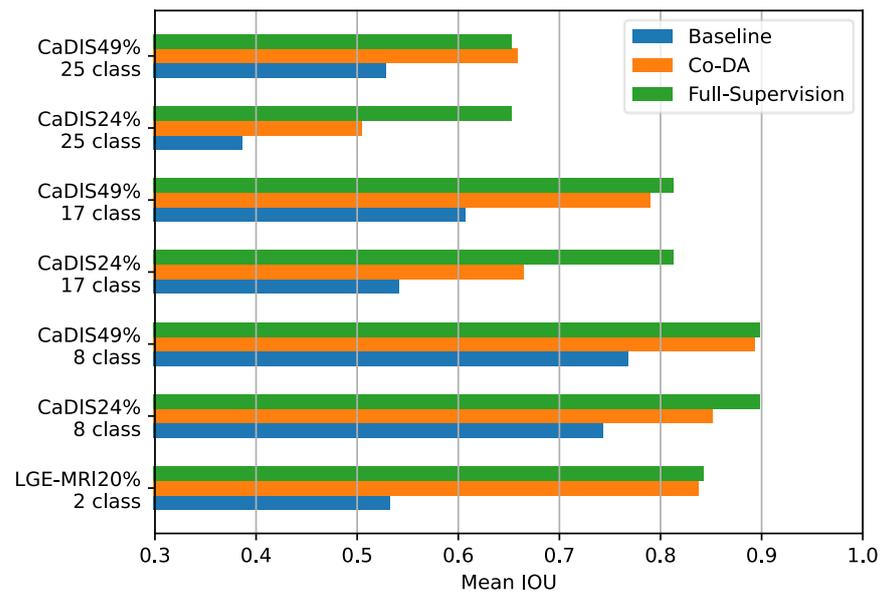
In this section, we perform additional experiments to verify the strong learning capacity Co-DA as compared to its fully supervised variant. In addition, we isolate the different components of our method to ensure that they all contribute to the final performance. Finally, we look into the efficacy of the dynamic threshold for the over-expectation cross-entropy loss and compare examples of pseudo-labels generated with and without Co-DA.

Firstly, we compare Co-DA with the fully supervised settings on both CaDIS and LGE-MRI and present the results in Figure 7. In many cases, the performance of Co-DA is very close to the fully supervised setting, demonstrating the outstanding efficacy of class-wise distribution alignment. Specifically, with 49% labeled data, Co-DA provides only slightly lower mIoU on CaDIS Tasks 1 and 2 as compared to the fully supervised setting. It even marginally outperforms the fully supervised case on Task 3, which means the refined pseudo-labels have a superior quality.

Secondly, to ensure the complementary effect of proposed components in our method, i.e., the co-distribution alignment and the over-expectation cross-entropy loss, we conduct a set of ablation experiments on CaDIS, and the results are presented in Table 4. The third row demonstrates that the over-expectation cross-entropy loss reduces the noise in pseudo-labels by dynamically varying the threshold to filter out low-confidence pixels, achieving a significant improvement of 0.0488 on Task 3 for mean IoU over the baseline. The fourth row shows that the proposed Co-DA can bring a gain of 0.1303 on Task 3, while there is a slight decrease in performance on Task 2. Empirically, we conjecture the reason is that Task 1 and Task 3 have more severe class imbalance than Task 2 [44]. Considering the overall stability of the proposed method, we apply the two newly designed modules to co-training simultaneously, obtaining improvements of 0.0058, 0.0124 and 0.0748 on the three tasks over the baseline, respectively.

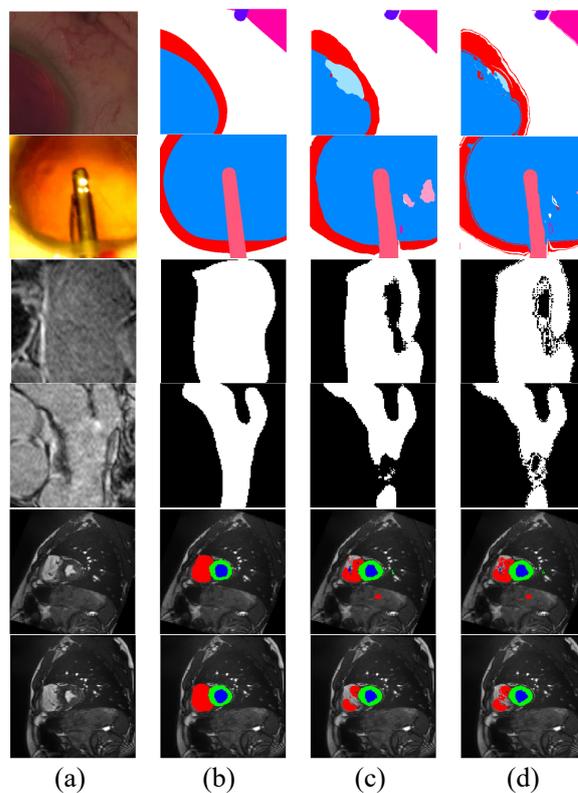
**Table 4.** Results from our ablation study on CaDIS. Performance shown in mean IoU. O-E denotes the proposed over-expectation cross-entropy loss.

Co-Training	DA	O-E	Co-DA	49% Labeled Data		
				Task 1	Task 2	Task 3
✓				0.8874	0.7774	0.5837
✓	✓			0.8893	0.7718	0.6438
✓		✓		0.8903	0.7824	0.6325
✓			✓	0.8907	0.7681	<b>0.7140</b>
✓		✓	✓	<b>0.8932</b>	<b>0.7898</b>	0.6585



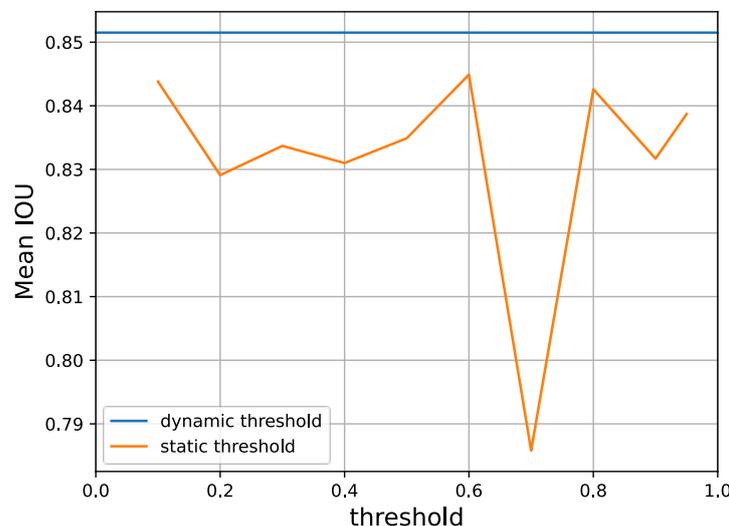
**Figure 7.** Comparing the baseline, Co-DA and the fully supervised method on the CaDIS dataset in terms of mIoU. In some cases, Co-DA performs comparably to the fully supervised method, demonstrating its strong learning capacity.

We also visualize examples of the pseudo-labels generated by the networks with and without Co-DA, and the results are shown in Figure 8. Clearly, Co-DA is able to improve the quality of pseudo-labels, resulting in an overall accuracy improvement and a reduction of the negative impact from incorrect pseudo-labels.



**Figure 8.** Visualization of pseudo-labels generated by Co-DA. (a) is the input image, (b) is the ground-truth, (c) are pseudo-labels generated by the network without Co-DA and (d) are the pseudo-labels refined by Co-DA.

Finally, according to Figure 9, different predefined thresholds for the over-expectation loss will also lead to different levels of performance. Yet, a dynamic threshold as provided by Equation (9) is consistently superior, as it is able to automatically select a suitable gating value for retaining reliable pseudo-labels.



**Figure 9.** A performance comparison of static thresholds versus the dynamic threshold that our method adopts, according to Equation (9) on Task 1 of the CaDIS dataset with 24% labeled data. It is clear that the dynamic threshold provides superior mean IoU.

## 5. Conclusions

In this paper, we propose a novel Co-Distribution Alignment (Co-DA) approach to partially labeled medical image segmentation tasks with imbalanced class distributions. Our key idea involves a class-dependent alignment between labeled distributions and unlabeled marginal distributions that is suitable for dense prediction tasks. Specifically, Co-DA simplifies the estimation of labeled distributions, extends the original DA to class-wise distribution alignment with cross-supervision and adopts adaptive temperature scaling for labeled distributions to avoid highly imbalanced estimations. In addition, we propose an over-expectation cross-entropy loss to reduce noises in pseudo-labels. Extensive experiments, including ablation studies, on three publicly available datasets demonstrate the consistently superior learning capacity of our approach.

**Author Contributions:** Conceptualization, T.W. and Z.H.; methodology, T.W. and J.W.; software, Z.H. and Y.C.; validation, T.W.; formal analysis, T.W. and Z.H.; investigation, Z.H.; resources, Y.C.; data curation, Z.H.; writing—original draft preparation, T.W. and Z.H.; writing—review and editing, Z.L.; visualization, Z.H. and J.W.; supervision, Z.L.; project administration, Z.L.; funding acquisition, Z.L. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research was funded by National Natural Science Foundation of China (61972187, 61703195), Fujian Provincial Natural Science Foundation (2022J011112, 2020J02024, 2020J01828), Research Project of Fashu Foundation (MFK23001) and The Open Program of The Key Laboratory of Cognitive Computing and Intelligent Information Processing of Fujian Education Institutions, Wuyi University (KLCCIP2020202).

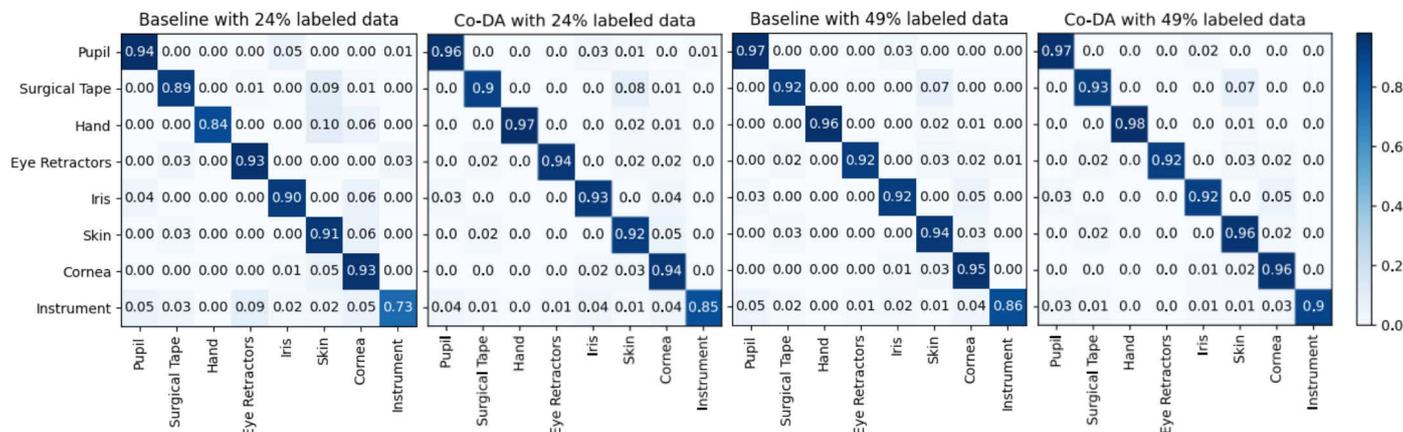
**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Conflicts of Interest:** The authors declare no conflict of interest.

### Appendix A

In Figure A1, we show the confusion matrices of Co-DA on CaDIS Task 1. We can clearly observe that under different split settings, Co-DA performs favorably in comparison to the Baseline.



**Figure A1.** The confusion matrices of our proposed Co-DA and Baseline (only with 24% and 49% labeled data) on CaDIS Task 1. The proposed Co-DA performs better than the Baseline on the majority of classes under both settings.

### References

- Hauser, P.V.; Chang, H.M.; Nishikawa, M.; Kimura, H.; Yanagawa, N.; Hamon, M. Bioprinting Scaffolds for Vascular Tissues and Tissue Vascularization. *Bioengineering* **2021**, *8*, 178. [CrossRef]
- Mustahsan, V.M.; Anugu, A.; Komatsu, D.E.; Kao, I.; Pentylala, S. Biocompatible Customized 3D Bone Scaffolds Treated with CRFP, an Osteogenic Peptide. *Bioengineering* **2021**, *8*, 199. [CrossRef]
- Cosoli, G.; Scalise, L.; De Leo, A.; Russo, P.; Tricarico, G.; Tomasini, E.P.; Cerri, G. Development of a Novel Medical Device for Mucositis and Peri-Implantitis Treatment. *Bioengineering* **2020**, *7*, 87. [CrossRef] [PubMed]
- Wang, W.; Yousaf, M.; Liu, D.; Sohail, A. A Comparative Study of the Genetic Deep Learning Image Segmentation Algorithms. *Symmetry* **2022**, *14*, 1977. [CrossRef]
- Ronneberger, O.; Fischer, P.; Brox, T. U-net: Convolutional networks for biomedical image segmentation. In Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention, Munich, Germany, 5–10 October 2015; Springer: Berlin/Heidelberg, Germany, 2015; pp. 234–241.
- Zhou, Z.; Rahman Siddiquee, M.M.; Tajbakhsh, N.; Liang, J. Unet++: A nested u-net architecture for medical image segmentation. In *Deep learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*; Springer: Berlin/Heidelberg, Germany, 2018; pp. 3–11.
- Hatamizadeh, A.; Tang, Y.; Nath, V.; Yang, D.; Myronenko, A.; Landman, B.; Roth, H.R.; Xu, D. Unetr: Transformers for 3d medical image segmentation. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, Waikoloa, HI, USA, 3–8 January 2022; pp. 574–584.
- Yu, L.; Wang, S.; Li, X.; Fu, C.W.; Heng, P.A. Uncertainty-aware self-ensembling model for semi-supervised 3D left atrium segmentation. In Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention, Shenzhen, China, 13–17 October 2019; Springer: Berlin/Heidelberg, Germany, 2019; pp. 605–613.
- Luo, X.; Chen, J.; Song, T.; Wang, G. Semi-supervised medical image segmentation through dual-task consistency. In Proceedings of the AAAI Conference on Artificial Intelligence, Virtual, 2–9 February 2021; Volume 35, pp. 8801–8809.
- Bortsova, G.; Dubost, F.; Hogeweg, L.; Katramados, I.; Bruijine, M.d. Semi-supervised medical image segmentation via learning consistency under transformations. In Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention, Shenzhen, China, 13–17 October 2019; Springer: Berlin/Heidelberg, Germany, 2019; pp. 810–818.
- Zhang, Y.; Zhang, J. Dual-task mutual learning for semi-supervised medical image segmentation. In Proceedings of the Chinese Conference on Pattern Recognition and Computer Vision, Beijing, China, 29 October–1 November 2021; Springer: Berlin/Heidelberg, Germany, 2021; pp. 548–559.
- Li, Y.; Chen, J.; Xie, X.; Ma, K.; Zheng, Y. Self-loop uncertainty: A novel pseudo-label for semi-supervised medical image segmentation. In Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention, Lima, Peru, 4–8 October 2020; Springer: Berlin/Heidelberg, Germany, 2020; pp. 614–623.
- Lin, T.Y.; Goyal, P.; Girshick, R.; He, K.; Dollár, P. Focal loss for dense object detection. In Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017; pp. 2980–2988.

14. Li, Y.; Wang, T.; Kang, B.; Tang, S.; Wang, C.; Li, J.; Feng, J. Overcoming classifier imbalance for long-tail object detection with balanced group softmax. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 13–19 June 2020; pp. 10991–11000.
15. Berthelot, D.; Carlini, N.; Cubuk, E.D.; Kurakin, A.; Sohn, K.; Zhang, H.; Raffel, C. Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. *arXiv* **2019**, arXiv:1911.09785.
16. Chen, X.; Yuan, Y.; Zeng, G.; Wang, J. Semi-supervised semantic segmentation with cross pseudo supervision. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA, 20–25 June 2021; pp. 2613–2622.
17. Blum, A.; Mitchell, T. Combining labeled and unlabeled data with co-training. In Proceedings of the Eleventh Annual Conference on Computational Learning Theory, Madison, WI, USA, 24–26 July 1998; pp. 92–100.
18. Lee, D.H. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In Proceedings of the Workshop on Challenges in Representation Learning, ICML, Atlanta, GA, USA, 18 June 2013; Volume 3, p. 896.
19. Hu, Z.; Yang, Z.; Hu, X.; Nevatia, R. Simple: Similar pseudo label exploitation for semi-supervised classification. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA, 20–25 June 2021; pp. 15099–15108.
20. Li, Z.; Ko, B.; Choi, H.J. Naive semi-supervised deep learning using pseudo-label. *Peer-to-Peer Netw. Appl.* **2019**, *12*, 1358–1368. [[CrossRef](#)]
21. Zhang, B.; Wang, Y.; Hou, W.; Wu, H.; Wang, J.; Okumura, M.; Shinozaki, T. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. *Adv. Neural Inf. Process. Syst.* **2021**, *34*, 18408–18419.
22. Verma, V.; Kawaguchi, K.; Lamb, A.; Kannala, J.; Bengio, Y.; Lopez-Paz, D. Interpolation consistency training for semi-supervised learning. *arXiv* **2019**, arXiv:1903.03825.
23. Sajjadi, M.; Javanmardi, M.; Tasdizen, T. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. *Adv. Neural Inf. Process. Syst.* **2016**, *29*, 1–9.
24. Xie, Q.; Dai, Z.; Hovy, E.; Luong, T.; Le, Q. Unsupervised data augmentation for consistency training. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 6256–6268.
25. Grandvalet, Y.; Bengio, Y. Semi-supervised learning by entropy minimization. *Adv. Neural Inf. Process. Syst.* **2004**, *17*, 1–8.
26. Ma, N.; Bu, J.; Lu, L.; Wen, J.; Zhou, S.; Zhang, Z.; Gu, J.; Li, H.; Yan, X. Context-guided entropy minimization for semi-supervised domain adaptation. *Neural Netw.* **2022**, *154*, 270–282. [[CrossRef](#)]
27. Sohn, K.; Berthelot, D.; Carlini, N.; Zhang, Z.; Zhang, H.; Raffel, C.A.; Cubuk, E.D.; Kurakin, A.; Li, C.L. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 596–608.
28. Tarvainen, A.; Valpola, H. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Adv. Neural Inf. Process. Syst.* **2017**, *30*.
29. Sun, Y.; Zhou, C.; Fu, Y.; Xue, X. Parasitic GAN for semi-supervised brain tumor segmentation. In Proceedings of the IEEE International Conference on Image Processing, Taipei, Taiwan, 22–25 September 2019; IEEE: Piscataway, NJ, USA, 2019; pp. 1535–1539.
30. Peiris, H.; Chen, Z.; Egan, G.; Harandi, M. Duo-SegNet: Adversarial dual-views for semi-supervised medical image segmentation. In Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention, Strasbourg, France, 27 September–1 October 2021; Springer: Berlin/Heidelberg, Germany, 2021; pp. 428–438.
31. Peng, J.; Estrada, G.; Pedersoli, M.; Desrosiers, C. Deep co-training for semi-supervised image segmentation. *Pattern Recognit.* **2020**, *107*, 107269. [[CrossRef](#)]
32. Zhou, Y.; He, X.; Huang, L.; Liu, L.; Zhu, F.; Cui, S.; Shao, L. Collaborative learning of semi-supervised segmentation and classification for medical images. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 2079–2088.
33. Wu, Y.; Wu, Z.; Wu, Q.; Ge, Z.; Cai, J. Exploring Smoothness and Class-Separation for Semi-supervised Medical Image Segmentation. *arXiv* **2022**, arXiv:2203.01324.
34. Ning, X.; Wang, X.; Xu, S.; Cai, W.; Zhang, L.; Yu, L.; Li, W. A review of research on co-training. *Concurr. Comput. Pract. Exp.* **2021**, e6276. [[CrossRef](#)]
35. Qiao, S.; Shen, W.; Zhang, Z.; Wang, B.; Yuille, A. Deep co-training for semi-supervised image recognition. In Proceedings of the European Conference on Computer Vision, Munich, Germany, 8–14 September 2018; pp. 135–152.
36. Chen, M.; Du, Y.; Zhang, Y.; Qian, S.; Wang, C. Semi-supervised learning with multi-head co-training. In Proceedings of the AAAI Conference on Artificial Intelligence, Virtual, 22 February–1 March 2022; Volume 36, pp. 6278–6286.
37. Sun, B.; Feng, J.; Saenko, K. Return of frustratingly easy domain adaptation. In Proceedings of the AAAI Conference on Artificial Intelligence, Phoenix, AZ, USA, 12–17 February 2016; Volume 30.
38. Ghifary, M.; Balduzzi, D.; Kleijn, W.B.; Zhang, M. Scatter component analysis: A unified framework for domain adaptation and domain generalization. *IEEE Trans. Pattern Anal. Mach. Intell.* **2016**, *39*, 1414–1430. [[CrossRef](#)] [[PubMed](#)]
39. Wang, J.; Feng, W.; Chen, Y.; Yu, H.; Huang, M.; Yu, P.S. Visual domain adaptation with manifold embedded distribution alignment. In Proceedings of the 26th ACM International Conference on Multimedia, Seoul, Republic of Korea, 29 October 2018; pp. 402–410.
40. Shen, C.; Wang, X.; Wang, D.; Li, Y.; Zhu, J.; Gong, M. Dynamic joint distribution alignment network for bearing fault diagnosis under variable working conditions. *IEEE Trans. Instrum. Meas.* **2021**, *70*, 1–13. [[CrossRef](#)]

41. Saito, K.; Ushiku, Y.; Harada, T.; Saenko, K. Strong-weak distribution alignment for adaptive object detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 6956–6965.
42. Wei, C.; Sohn, K.; Mellina, C.; Yuille, A.; Yang, F. Crest: A class-rebalancing self-training framework for imbalanced semi-supervised learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA, 20–25 June 2021; pp. 10857–10866.
43. Shore, J.; Johnson, R. Properties of cross-entropy minimization. *IEEE Trans. Inf. Theory* **1981**, *27*, 472–482. [[CrossRef](#)]
44. Bouget, D.; Allan, M.; Stoyanov, D.; Jannin, P. Vision-based and marker-less surgical tool detection and tracking: A review of the literature. *Med. Image Anal.* **2017**, *35*, 633–654.
45. Trikha, S.; Turnbull, A.; Morris, R.; Anderson, D.; Hossain, P. The journey to femtosecond laser-assisted cataract surgery: New beginnings or a false dawn? *Eye* **2013**, *27*, 461–473. [[CrossRef](#)]
46. Xiong, Z.; Xia, Q.; Hu, Z.; Huang, N.; Bian, C.; Zheng, Y.; Vesal, S.; Ravikumar, N.; Maier, A.; Yang, X.; et al. A global benchmark of algorithms for segmenting the left atrium from late gadolinium-enhanced cardiac magnetic resonance imaging. *Med. Image Anal.* **2021**, *67*, 101832. [[CrossRef](#)] [[PubMed](#)]
47. Bernard, O.; Lalonde, A.; Zotti, C.; Cervenansky, F.; Yang, X.; Heng, P.A.; Cetin, I.; Lekadir, K.; Camara, O.; Ballester, M.A.G.; et al. Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: Is the problem solved? *IEEE Trans. Med. Imaging* **2018**, *37*, 2514–2525. [[CrossRef](#)] [[PubMed](#)]
48. Pissas, T.; Ravasio, C.S.; Cruz, L.D.; Bergeles, C. Effective semantic segmentation in Cataract Surgery: What matters most? In Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention, Strasbourg, France, 27 September–1 October 2021; Springer: Berlin/Heidelberg, Germany, 2021; pp. 509–518.
49. Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. Pytorch: An imperative style, high-performance deep learning library. *Adv. Neural Inf. Process. Syst.* **2019**, *32*, 1–12.
50. Tan, M.; Le, Q. Efficientnet: Rethinking model scaling for convolutional neural networks. In Proceedings of the International Conference on Machine Learning. PMLR, Long Beach, CA, USA, 10–15 May 2019; pp. 6105–6114.
51. Milletari, F.; Navab, N.; Ahmadi, S.A. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In Proceedings of the 2016 Fourth International Conference on 3D Vision (3DV), Stanford, CA, USA, 25–28 October 2016; IEEE: Piscataway, NJ, USA, 2016; pp. 565–571.
52. Luo, X.; Wang, G.; Liao, W.; Chen, J.; Song, T.; Chen, Y.; Zhang, S.; Metaxas, D.N.; Zhang, S. Semi-supervised medical image segmentation via uncertainty rectified pyramid consistency. *Med. Image Anal.* **2022**, *80*, 102517.
53. Lin, Y.; Yao, H.; Li, Z.; Zheng, G.; Li, X. Calibrating Label Distribution for Class-Imbalanced Barely-Supervised Knee Segmentation. *arXiv* **2022**, arXiv:2205.03644.

**Disclaimer/Publisher’s Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.