

Constructing and Exploring Intermediate Domains in Mixed Domain Semi-supervised Medical Image Segmentation

Qinghe Ma¹, Jian Zhang¹, Lei Qi², Qian Yu³, Yinghuan Shi^{1,*}, Yang Gao¹
¹Nanjing University ²Southeast University ³Shandong Women's University

Abstract

Both limited annotation and domain shift are prevalent challenges in medical image segmentation. Traditional semi-supervised segmentation and unsupervised domain adaptation methods address one of these issues separately. However, the coexistence of limited annotation and domain shift is quite common, which motivates us to introduce a novel and challenging scenario: **Mixed Domain Semi-supervised medical image Segmentation (MiDSS)**. In this scenario, we handle data from multiple medical centers, with limited annotations available for a single domain and a large amount of unlabeled data from multiple domains. We found that the key to solving the problem lies in how to generate reliable pseudo labels for the unlabeled data in the presence of domain shift with labeled data. To tackle this issue, we employ *Unified Copy-Paste (UCP)* between images to construct intermediate domains, facilitating the knowledge transfer from the domain of labeled data to the domains of unlabeled data. To fully utilize the information within the intermediate domain, we propose a symmetric *Guidance training strategy (SymGD)*, which additionally offers direct guidance to unlabeled data by merging pseudo labels from intermediate samples. Subsequently, we introduce a *Training Process aware Random Amplitude MixUp (TP-RAM)* to progressively incorporate style-transition components into intermediate samples. Compared with existing state-of-the-art approaches, our method achieves a notable 13.57% improvement in Dice score on Prostate dataset, as demonstrated on three public datasets. Our code is available at <https://github.com/MQinghe/MiDSS>

*Corresponding author. Qinghe Ma, Jian Zhang, Yinghuan Shi and Yang Gao are with the State Key Laboratory for Novel Software Technology and National Institute of Healthcare Data Science, Nanjing University, China. This work was supported by the NSFC Program (62222604, 62206052, 62192783), Jiangsu Natural Science Foundation (BK20210224), and Shandong Natural Science Foundation (ZR2023MF037).

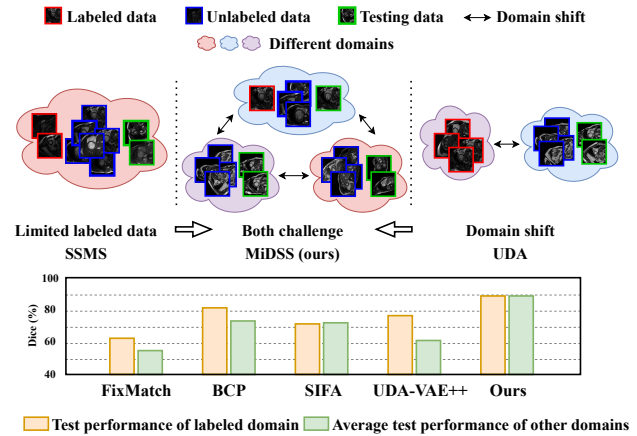


Figure 1. The upper figure illustrates SSMS, UDA, and MiDSS. The lower figure shows the comparison between different methods on the labeled domain (BIDMC) [23] and other domains.

1. Introduction

Semi-supervised medical image segmentation (SSMS) has aroused lots of attention in recent years, due to its advantage of improving performance while reducing the labeling burden [13, 20, 21, 43]. Surprisingly, the performance of semi-supervised methods [12, 22] is sometimes close to or even better than that of its upper bound—supervised methods [8, 32]. With 1) only a few labeled data as guidance, and 2) the development of unsupervised learning techniques, e.g., contrastive learning [18], SSMS is becoming a promising direction to utilize unlabeled data. In traditional SSMS tasks, it is typically assumed that labeled and unlabeled data come from the same distribution [1, 3, 42]. In this case, the model applies the information gained from labeled data directly to predict pseudo labels for unlabeled data.

Despite their success, we notice a significant yet easily underestimated issue: *During clinical data collection, do we really check and guarantee, whether each unlabeled sample belongs to the same distribution as labeled samples?* For example, given a certain amount of labeled data collected and well annotated from a typical center, since unlabeled data are easily accessed and collected, when the

physicians have a massive number of unlabeled data, they possibly lack sufficient time to verify whether the data belongs to the same distribution [10, 29].

As this dilemma mentioned above, which refers to “domain shift” [16, 38], could cause a large performance degeneration, we should pay more attention to this issue during our development of a general SSMS model. A more practical setting could be: given labeled data from a consistent distribution, the unlabeled data is allowed to come from the same distribution or even multiple different distributions. It is important to note that this issue could not be directly solved by current unsupervised domain adaptation (UDA) methods. Although UDA is promising to address domain shift, these methods [6, 11, 30, 41] demand abundant labeled source domain data, which is the primary reason for the existence of SSMS.

Despite the significance of the domain shift issue in SSMS, according to our best knowledge, our work is first to attempt to investigate this issue, which we termed as **Mixed Dmains Semi-supervised medical image Segmentation** (abbreviated as **MiDSS** in our paper). As illustrated in Fig. 1, confronted with the coexistence of limited annotation and domain shift, it is challenging to effectively utilize a large quantity of unlabeled medical images from multiple domains. Additionally, from the lower figure in Fig. 1, we found that achieving satisfactory results in our setting is difficult when relying solely on SSMS or UDA methods.

The key to addressing the problem lies in *how to generate reliable pseudo labels for unlabeled data, especially in the presence of domain shift*. However, the domain gap among training samples results in the model producing poor pseudo labels for unlabeled data. Thus, our key insight is first to construct several intermediate domains among original domains to narrow the domain gap, and then emphasize knowledge transfer approaches upon these constructed domains. CutMix [44], also referred to as Copy-Paste (CP), is a simple yet highly effective technique for generating intermediate samples by embedding patches from labeled samples into unlabeled samples along a single direction. To fully leverage the data generation capability of CP, we employ a Unified Copy-Paste (UCP) approach. UCP generates a great quantity of diverse intermediate samples by unified applying both embedding directions between labeled and unlabeled samples, narrowing the domain gap and effectively mitigating the issue of error accumulation.

Besides, the intermediate domains may not align perfectly with the domains of the unlabeled data. Solely training within the intermediate domains neglects the model performance on the unlabeled data, which are of great interest. The information from intermediate domains should be harnessed to promote the model to generate high quality pseudo labels for unlabeled data. Moreover, medical images of the same organ or lesion often exhibit similar

structures [37], with stylistic differences being the primary source of domain shift [24]. Local semantic mixing methods (*e.g.*, CP) fail to address these differences, leading to a lack of stylistic transition in intermediate domains. Additionally, it is crucial that an aggressive stylistic transition hinders the gradual construction of intermediate domains.

Therefore, to fully harness the information from intermediate domains and ensure comprehensive and stable knowledge transfer, we propose a Mixed Domains Semi-supervised medical image segmentation approach. In terms of training strategy, we design Symmetric Guidance training strategy (SymGD), which is composed of guidance from unlabeled data to intermediate samples and vice versa. This dual-perspective integration of pseudo labels enhances the precision of guidance. Concerning intermediate samples, we introduce a Training Process aware Random Amplitude MixUp module (TP-RAM) to promote smooth stylistic knowledge transfer. Our main contributions are three folds:

- We investigate a new yet underestimated semi-supervised medical image segmentation setting (namely MiDSS).
- A novel SymGD training strategy based on UCP, promotes the training on the unlabeled data with intermediate domains information.
- A TP-RAM module to make domain knowledge transfer comprehensive and stable.

Extensive experiments validate the effectiveness of our method on three public datasets. For example, on Prostate dataset, we achieved a remarkable Dice score improvement of 13.57% compared with all state-of-the-art methods.

2. Related Work

Semi-supervised Medical Image Segmentation. Manually annotating medical images is challenging and costly [49]. Semi-supervised medical image segmentation methods have shown promise in addressing limited-label segmentation tasks. Entropy minimization and consistency regularization are widely adopted techniques in this context. Yu *et al.* [43] encouraged consistent predictions under different perturbations. Luo *et al.* [26] proposed a dual-task network that predicts pixel-wise segmentation maps and geometry-aware level set representations. Li *et al.* [19] introduced a shape-aware strategy using a multi-task deep network that jointly predicts semantic segmentation and signed distance maps. Wu *et al.* [40] addressed challenges in SSMS by simultaneously enforcing pixel-level smoothness and inter-class separation. Miao *et al.* [27] pointed out the importance of algorithmic independence between two networks or branches in SSMS. Unfortunately, existing methods often face challenges when domain shift occurs, causing decreased performance due to the shared distribution assumption between labeled and unlabeled data.

Unsupervised Domain Adaptation. UDA methods train the model by leveraging abundant labeled source do-

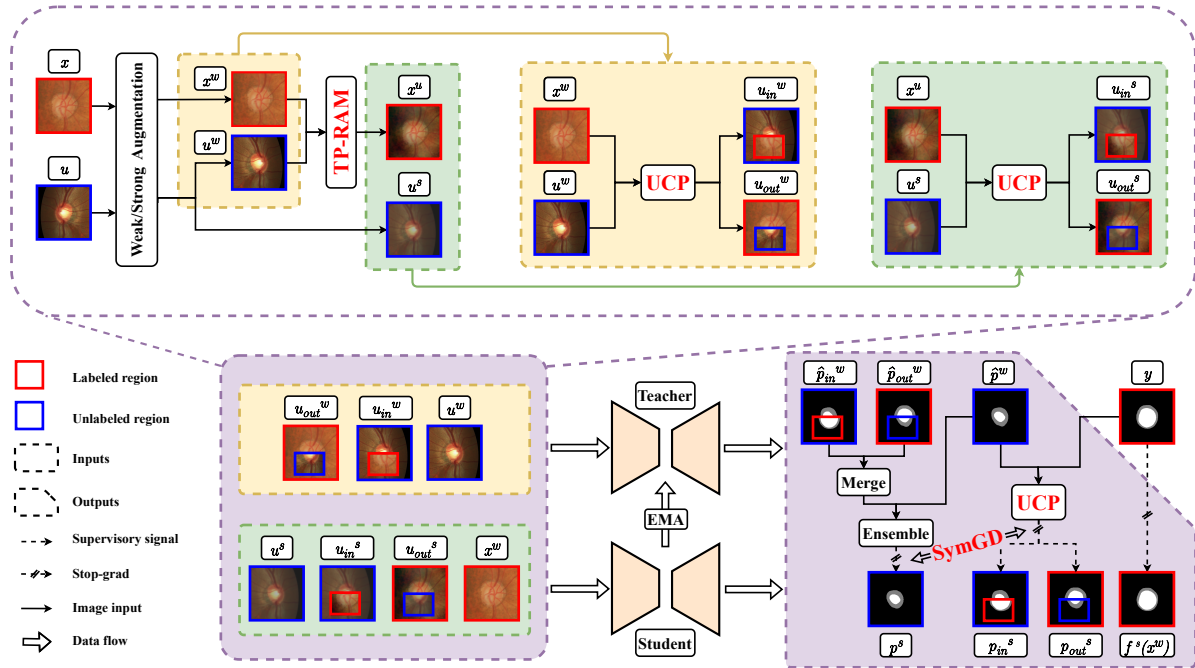


Figure 2. The overall framework of our method emphasizes domain knowledge transfer through data augmentation and training strategy. We generate intermediate samples through UCP between labeled data and unlabeled data. During training, we gradually introduce style transfer components to the intermediate samples, constructing intermediate domains at both semantic and stylistic levels. Further details about UCP and TP-RAM are provided in Fig. 3 and Fig. 5. Moreover, we design a symmetric guidance for model training. In addition to guiding from unlabeled data to intermediate samples, we merge unlabeled regions of intermediate samples to obtain the pseudo label of unlabeled data from another perspective. The integration of pseudo labels through ensemble from two perspectives guides the prediction of unlabeled data. Best viewed in color.

main data and unlabeled target domain data to achieve appealing performance in the target domain. In medical image segmentation tasks, domain shift issues arise from variations in device parameters, disease severity, imaging principles, *etc.* Thus, UDA plays a crucial role in addressing these issues. Adversarial learning based UDA methods achieve alignment at multiple levels to narrow domain gap, including input alignment [5, 30] and feature alignment [11, 15, 35, 47]. Self-training based UDA methods [46, 48] train the model by generating reliable pseudo labels for unlabeled data. Methods like [7, 41] promote domain knowledge transfer by generating intermediate domains. These methods typically rely on a sufficient amount of source domain data and focus on a single target domain, making it challenging to be effective in the MiDSS scenario.

Data Augmentation via Copy-Paste. Copy-paste (CP) [44] involves pasting the content of a particular region of one image onto the corresponding region of another image, creating a new image that retains the semantic information from both original images. Compared to other pixel-level fusion strategies, such as MixUp [45], CP excels in preserving and blending semantic information from source images. In contrast, MixUp combines global source images proportionally, potentially leading to ambiguity issues when

pixels from different classes are blended. Therefore, CP is a more suitable augmentation technique for medical image segmentation tasks. In line with this research direction, BCP [2] takes into account the dual embedding directions between unlabeled and labeled samples, randomly selecting one of them to generate intermediate samples. To mitigate domain shift, previous works [2, 14] employed CP to transfer knowledge from labeled data to unlabeled data by generating intermediate samples. Despite the progress, they fall short in fully leveraging information from intermediate domains, leading to sub-optimal transfer effects.

3. Method

In the MiDSS scenario, the domain gap exists among images originating from K data centers $\{\mathcal{D}_i\}_{i=1}^K$. The training set comprises N labeled images $\{(x_i, y_i)\}_{i=1}^N$ from a single domain $\mathcal{D}_j \in \{\mathcal{D}_i\}_{i=1}^K$ and M unlabeled images $\{u_i\}_{i=1}^M$ from multiple domains $\mathcal{D}_1, \dots, \mathcal{D}_K$, where $M > N$. The image resolution of $H \times W \times D$ sequentially represents height and width and channel, and $y_i \in \{0, 1\}^{H \times W \times C}$ is the ground truth of x_i , where C is the number of class. We aim to train a model based on the aforementioned training set, capable of delivering outstanding segmentation performance across all domains with or without labeled data.

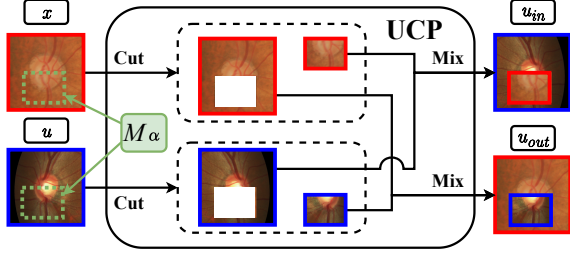


Figure 3. The illustration of UCP between images. **Cut** refers to splitting the image into two parts according to M_α , while **Mix** implies merging two parts of images back together.

Based on the mean teacher framework [34], for each labeled data x , teacher model f^t generates probability map p^w and pseudo label \hat{p} for weakly augmented unlabeled data u^w , guiding the student model f^s in making predictions p^s for strongly augmented unlabeled data u^s :

$$\begin{aligned} x^w &= \mathcal{A}^w(x); u^w = \mathcal{A}^w(u); u^s = \mathcal{A}^s(u), \\ p^w &= f^t(u^w); p^s = f^s(u^s); \hat{p} = \arg \max(p^w), \end{aligned} \quad (1)$$

where $\arg \max(\cdot)$ applied to probability distributions produces valid ‘‘one-hot’’ probability distributions, and f^t is updated by the EMA of f^s [34]. The weak augmentation \mathcal{A}^w includes cropping, rotation, flip, and elastic distortion, while the strong augmentation \mathcal{A}^s builds upon \mathcal{A}^w by incorporating non-geometric operations like color jitter and Gaussian blurring [33, 36]. The overall framework of our method is shown in Fig. 2.

3.1. Intermediate Samples Generation by UCP

Given a pair of labeled data (x^w, y^w) and unlabeled data (u^s, p^w, \hat{p}) , we simultaneously apply both embedding directions of CP. This enables unified Copy-Paste (UCP) between labeled and unlabeled data with distribution discrepancy. Through this approach, we obtain intermediate samples along with their pseudo-labels and probability maps:

$$\begin{aligned} u_{in}^s &= x^w \odot M_\alpha + u^s \odot (\mathbf{1} - M_\alpha), \\ p_{in}^w &= y^w \odot M_\alpha + p^w \odot (\mathbf{1} - M_\alpha), \\ \hat{p}_{in} &= y^w \odot M_\alpha + \hat{p} \odot (\mathbf{1} - M_\alpha), \\ u_{out}^s &= u^s \odot M_\alpha + x^w \odot (\mathbf{1} - M_\alpha), \\ p_{out}^w &= p^w \odot M_\alpha + y^w \odot (\mathbf{1} - M_\alpha), \\ \hat{p}_{out} &= \hat{p} \odot M_\alpha + y^w \odot (\mathbf{1} - M_\alpha), \end{aligned} \quad (2)$$

where $M_\alpha \in \{0, 1\}^{W \times H}$ is a randomly generated one-centered mask, indicating the region for CP operation. In this context, $\mathbf{1}$ indicates an all-one matrix, and \odot means element-wise multiplication. Referring to Fig. 3, the illustration explains UCP between images, with the same operation applied to both probability maps and pseudo labels.




Figure 4. The results depict the quality of pseudo labels in a SSMS method (FixMatch) with and without UCP, utilizing 40 labeled data from the BIDMC domain in Prostate dataset [23]. Each colored bar represents samples from different domains, with the bar height indicating the quality of pseudo labels generated by the model for unlabeled data from that domain.

3.2. Symmetric Guided Training with UCP

In the case of a labeled data (x^w, y^w) and unlabeled data (u^w, u^s) pair, we aim to facilitate the transfer of knowledge from the labeled domain while maximizing the extraction of latent information from the unlabeled data.

During the training phase, the teacher model firstly predicts u^w to generate probability map p^w and pseudo label \hat{p} for u^s . Then we employ UCP between (x^w, y^w) and (u^s, p^w, \hat{p}) to generate intermediate samples (u_{in}^s, u_{out}^s) , probability maps (p_{in}^w, p_{out}^w) and pseudo labels $(\hat{p}_{in}, \hat{p}_{out})$ by Eq. (2). We set weight map w to indicate whether the pseudo label is reliable:

$$w_i = \mathbf{1}(\max(p_i^w) \geq \tau), \quad (3)$$

where w_i is i^{th} pixel of w , and τ is a pre-defined confidence threshold used to filter noisy labels. The indicator function is denoted as $\mathbf{1}(\cdot)$. The pseudo labels \hat{p}_{in} and \hat{p}_{out} will be used as the supervision to guide the student model in predicting p_{in}^s and p_{out}^s for u_{in}^s and u_{out}^s :

$$\begin{aligned} \mathcal{L}_{in} &= \mathcal{L}_{ce}(\hat{p}_{in}, p_{in}^s, w_{in}) + \mathcal{L}_{dice}(\hat{p}_{in}, p_{in}^s, w_{in}), \\ \mathcal{L}_{out} &= \mathcal{L}_{ce}(\hat{p}_{out}, p_{out}^s, w_{out}) + \mathcal{L}_{dice}(\hat{p}_{out}, p_{out}^s, w_{out}), \end{aligned} \quad (4)$$

where \mathcal{L}_{ce} and \mathcal{L}_{dice} respectively represent the cross-entropy loss and dice loss, which are formulated as:

$$\begin{aligned} \mathcal{L}_{ce}(y, p, w) &= -\frac{1}{H \times W} \sum_{i=1}^{H \times W} w_i y_i \log p_i, \\ \mathcal{L}_{dice}(y, p, w) &= 1 - \frac{2 \times \sum_{i=1}^{H \times W} w_i p_i y_i}{\sum_{i=1}^{H \times W} w_i (p_i^2 + y_i^2)}, \end{aligned} \quad (5)$$

where p_i , y_i denote the probability of foreground and pseudo label of the i^{th} pixel, respectively.

The supervision provided by the fusion of \hat{p} with y^w to p_{in}^s and p_{out}^s effectively promotes the model training in the intermediate domains, facilitating adaptation to the domains

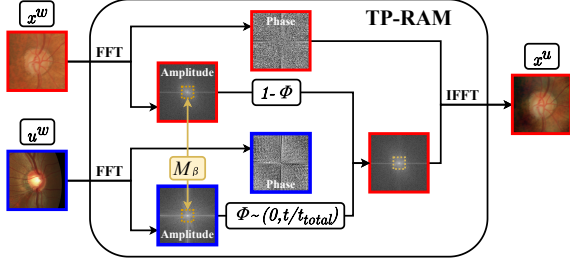


Figure 5. The illustration of TP-RAM. Fast Fourier transform (FFT) extracts the amplitude and phase maps of x^w and y^w . The low-frequency regions (determined by M_β) of the two phase maps are mixed. Through inverse Fast Fourier transform (IFFT), x^u is synthesized, preserving consistent semantics while introducing a different style compared to x^w .

of unlabeled data. With the assistance of UCP, the quality of pseudo labels for unlabeled data in each domain is significantly improved as shown in Fig. 4. However, this single-direction supervision lacks further utilization of information acquired in the intermediate domains, resulting in sub-optimal performance in the domains of unlabeled data. To address such limitation, we design a Symmetric Guidance training strategy (SymGD) to fully utilize the information from the intermediate domains. In addition to the consistency mentioned above, we also explore the direct guidance of unlabeled data through intermediate domain information.

Similarly, we employ M_α to mix x^w and u^w , generating intermediate samples u_{in}^w and u_{out}^w . The teacher model predicts them to generate pseudo labels \hat{p}_{in}^w and \hat{p}_{out}^w . We merge the unlabeled regions of the pseudo labels:

$$\hat{p}^{mg} = \hat{p}_{out}^w \odot M_\alpha + \hat{p}_{in}^w \odot (1 - M_\alpha). \quad (6)$$

We obtain the pseudo label for u^s from the pseudo labels of intermediate samples, then integrate it with p^w to provide more accurate guidance for the prediction of the student model on u^s . The weight map w^{ens} of \hat{p}^{mg} is determined as follows:

$$w^{ens} = (1 - (\hat{p} \oplus \hat{p}^{mg})) \odot w \odot w^{mg}, \quad (7)$$

where \oplus is the pixel-wise XOR operator to indicate the consistency between pseudo labels from two perspectives, and w and w^{mg} represent the weight map of \hat{p} and \hat{p}^{mg} respectively. The loss of such direction can be defined as:

$$\mathcal{L}_{sym} = \mathcal{L}_{ce}(\hat{p}^{mg}, p^s, w^{mg}) + \mathcal{L}_{dice}(\hat{p}^{mg}, p^s, w^{mg}), \quad (8)$$

To generate reliable pseudo labels for unlabeled data from the perspective of intermediate samples, it is necessary that the model performs well in the intermediate domain. Therefore, the guidance from \hat{p}^{mg} to p^s maintains a relatively low weight in the early stages of training and rapidly increases in the middle to later stages.

3.3. Style Transition in Intermediate Samples

The intermediate sample generation method based on UCP primarily focuses on local semantic blending while overlooking the transition of stylistic differences between different domains. Distribution information (*i.e.*, style) is typically represented in low-frequency components, while edge information is often found in high-frequency components [24]. By interpolating low-frequency information between data from different domains, amplitude MixUp generates samples with a new style. Then we mix these samples with u_w by UCP to introduce style-transition components in the intermediate domains.

However, arbitrary style transfer is not conducive to the stable model learning of the intermediate domain. We aim for the intermediate samples to exhibit a gradual transition in style between domains during the training process without compromising the diversity introduced by random MixUp. Therefore, we propose Training Progress aware Random Amplitude Mixup (TP-RAM).

Overall, as exhibited in Fig. 5, we obtain the frequency space signals of each channel of the image x through Fast Fourier transformation \mathcal{F} as follows:

$$\mathcal{F}(x)(u, v) = \sum_{h=0}^{H-1} \sum_{w=0}^{W-1} x(h, w) e^{-j2\pi(\frac{h}{H}u + \frac{w}{W}v)}, j^2 = -1. \quad (9)$$

Let $\mathcal{F}^A, \mathcal{F}^P$ be the amplitude and phase components of \mathcal{F} . For a pair of labeled and unlabeled data pair (x^w, u^w) , we obtain their amplitude ($\mathcal{F}^A(x^w), \mathcal{F}^A(u^w)$) and phase ($\mathcal{F}^P(x^w), \mathcal{F}^P(u^w)$). After that, we blend the low-frequency information of $\mathcal{F}^A(u^w)$ into $\mathcal{F}^A(x^w)$, and obtain a new image x^u through inverse Fast Fourier transformation \mathcal{F}^{-1} as follows:

$$x^u = \mathcal{F}^{-1}[M_\beta \odot \mathcal{F}^A(u^w) + (1 - M_\beta) \odot \mathcal{F}^A(x^w), \mathcal{F}^P(x^w)], \quad (10)$$

where M_β is a mask with values of 0 except in the central $2\beta W \times 2\beta H$ region, where the values range from 0 to Φ , indicating the mixing ratio for low-frequency amplitude. Φ is a parameter that increases during model training: $\Phi(t) = \frac{t}{t_{total}}$, where t denotes the current training iteration and t_{total} is the maximum training step, controlling the gradual enhancement of low-frequency components of unlabeled data in the intermediate samples.

3.4. Loss Function

For each pair of labeled and unlabeled data, the overall training objective function consists of two parts. The supervised loss is trivial, while the unsupervised loss is calculated through consistency regularization of symmetric guidance:

$$\mathcal{L}_{total} = \mathcal{L}_s + \lambda(\mathcal{L}_{in} + \mathcal{L}_{out} + \lambda\mathcal{L}_{sym}), \quad (11)$$

where λ is a weight coefficient decided by a time-dependent Gaussian warming-up function: $\lambda(t) = e^{-5(1-t/t_{total})}$.

Table 1. Comparison of different methods on Fundus dataset. #L represents the number of labeled samples. In the line of Upper bound, * denotes the Upper bound using all training samples in a domain as labeled data. \uparrow indicates that a higher value corresponds to better performance, while \downarrow suggests the opposite. The best performance is marked as **bold**, and the second-best is underlined.

Task		Optic Cup / Disc Segmentation								
Method	#L	DC \uparrow				DC \uparrow	JC \uparrow	HD \downarrow	ASD \downarrow	
		Domain 1	Domain 2	Domain 3	Domain 4	Avg.	Avg.	Avg.	Avg.	
U-Net		20	59.54 / 73.89	71.28 / 74.23	50.87 / 64.29	35.61 / 63.30	61.63	52.65	48.28	28.86
UA-MT	MICCAI'19	20	59.35 / 78.46	63.08 / 74.45	35.24 / 47.73	36.18 / 55.43	56.24	47.00	48.64	31.35
FDA	CVPR'20	20	76.99 / 89.94	<u>77.69</u> / 89.63	78.27 / 90.96	64.52 / 74.29	80.29	71.05	16.23	8.44
SIFA	TMI'20	20	50.67 / 75.30	64.44 / 80.69	61.67 / 83.77	55.07 / 70.67	67.78	54.77	20.16	10.93
FixMatch	NeurIPS'20	20	<u>81.18</u> / <u>91.29</u>	72.04 / 87.60	<u>80.41</u> / <u>92.95</u>	74.58 / 87.07	<u>83.39</u>	73.48	11.77	<u>5.60</u>
CPS	CVPR'21	20	64.53 / 86.25	70.26 / 86.97	42.92 / 54.94	36.98 / 46.70	61.19	52.69	34.44	26.79
CoraNet	TMI'21	20	61.64 / 87.32	65.56 / 87.05	66.12 / 83.54	49.01 / 77.73	72.25	60.50	20.52	10.44
UDA-VAE++	CVPR'22	20	55.01 / 80.76	68.87 / 85.94	63.23 / 84.92	68.42 / 80.89	73.51	61.40	17.60	9.86
SS-Net	MICCAI'22	20	59.42 / 78.15	67.32 / 85.05	45.69 / 69.91	38.76 / 61.13	63.18	53.49	44.90	25.73
BCP	CVPR'23	20	71.65 / 91.10	77.19 / 92.00	72.63 / 90.77	<u>77.67</u> / <u>91.42</u>	83.05	<u>73.66</u>	<u>11.05</u>	5.80
CauSSL	ICCV'23	20	63.38 / 80.60	67.52 / 80.72	49.53 / 63.88	39.43 / 49.43	61.81	51.80	41.25	23.94
Ours	this paper	20	83.71 / 92.96	80.47 / <u>89.93</u>	84.18 / 92.97	83.71 / 93.38	87.66	79.10	8.21	3.89
Upper bound	*		85.53 / 93.41	80.55 / 90.90	85.44 / 93.04	85.61 / 93.21	88.46	80.35	7.41	3.70




Figure 6. Visual results from Fundus dataset. The first row shows results for a test sample from the labeled domain (Domain 1), and the second row displays results for one from another domain (Domain 4). Red and green represent the Optical Cup and Disc, respectively.

Noting that the coefficient λ^2 for \mathcal{L}_{sym} increases relatively slowly in the early stages of training because the quality of pseudo labels generated through merge depends on the effective training of \mathcal{L}_{in} and \mathcal{L}_{out} .

3.5. Discussion of Inference and Parameter Size

Following previous work [34], our method consists of two models, both implemented based on U-Net [28], without introducing any additional network structures. We promote domain knowledge transfer through a training strategy and data augmentation to mitigate error accumulation in SSMS caused by domain shift. During the testing phase, given a test image x^{test} , we obtain the probability map $p^{test} = f^s(x^{test})$ through the student model and subsequently derive the segmentation result. The inference parameters only involve those from the student model, resulting in efficiency comparable to U-Net [28].

4. Experiments

4.1. Datasets

Fundus dataset [39] comprises retinal fundus images collected from four medical centers, primarily for optic cup and disc segmentation tasks. Each image has been cropped to form a region of interest with a 800×800 bounding-box.

We resize and randomly crop images to 256×256 .

Prostate dataset [23] contains prostate T2-weighted MRI data (with segmentation mask) collected from six different data sources out of three public datasets. We randomly split the dataset into training and testing sets based on a ratio of 4 : 1, with each 2D slice resized and randomly cropped to 384×384 . Labeled samples are selected from consecutive slices within cases, with at most one case of intersection and no slice intersection with unlabeled samples.

M&Ms dataset [4] contains 320 subjects scanned by four different magnetic resonance scanner vendors, primarily for the left ventricle (LV), left ventricle myocardium (MYO), and right ventricle (RV) segmentation tasks. Only the end-systole and end-diastole phases are annotated. We divide the data with annotation into training and testing sets following the same criteria as Prostate dataset. Each slice is resized and randomly cropped to 288×288 .

There are 20 labeled data in Fundus dataset and M&Ms dataset, and 40 labeled data in Prostate dataset. We normalized the data to $[-1,1]$. The detailed dataset partition information is presented in the supplementary material.

4.2. Implementation Details

Our method is implemented based on Pytorch and an NVIDIA GeForce RTX 4090 GPU. We set $\beta = 0.01$,

Table 2. Comparison of different methods on Prostate dataset.

Task		Prostate Segmentation									
Method	#L	DC \uparrow						DC \uparrow	JC \uparrow	HD \downarrow	ASD \downarrow
		RUNMC	BMC	HCRUDB	UCL	BIDMC	HK	Avg.	Avg.	Avg.	Avg.
U-Net		31.11	35.07	20.04	38.18	19.41	26.62	28.41	23.24	95.11	65.84
UA-MT	MICCAI'19	29.44	4.68	12.49	39.42	17.94	18.22	20.37	14.88	112.07	77.58
FDA	CVPR'20	47.44	35.37	24.54	61.01	28.19	40.51	39.51	32.17	76.67	47.87
SIFA	TMI'20	72.67	70.37	64.08	73.49	71.62	65.16	69.57	56.78	29.43	<u>13.03</u>
FixMatch	NeurIPS'20	<u>83.58</u>	69.17	<u>73.63</u>	<u>79.21</u>	56.07	84.78	<u>74.41</u>	<u>65.96</u>	<u>24.18</u>	14.09
CPS	CVPR'21	29.83	9.21	11.84	43.84	13.51	14.56	20.47	15.12	115.96	78.51
CoraNet	TMI'21	69.43	31.16	16.29	69.33	24.66	22.16	38.84	31.48	67.91	44.98
UDA-VAE++	CVPR'22	68.73	69.36	65.49	67.19	63.29	65.15	66.54	52.80	34.20	15.48
SS-Net	MICCAI'22	29.10	13.49	14.20	51.96	23.83	13.23	24.30	18.74	109.54	71.13
BCP	CVPR'23	70.15	<u>71.97</u>	46.15	58.93	<u>74.21</u>	67.47	64.81	55.17	52.60	27.22
CauSSL	ICCV'23	24.10	27.46	16.94	27.23	15.28	14.56	20.93	15.48	114.62	73.30
Ours	this paper	88.76	86.35	87.61	88.34	88.62	88.20	87.98	80.21	10.36	4.20
Upper bound	*	88.52	88.61	85.71	88.61	88.98	89.49	88.32	80.71	10.05	4.12

Table 3. Comparison of different methods on M&Ms dataset.

Task		LV / MYO / RV Segmentation											
Method	#L	DC \uparrow								DC \uparrow	JC \uparrow	HD \downarrow	ASD \downarrow
		Vendor A		Vendor B		Vendor C		Vendor D		Avg.	Avg.	Avg.	Avg.
U-Net		57.29 / 37.85 / 34.65	73.44 / 64.20 / 53.58	55.83 / 48.47 / 44.84	63.85 / 52.25 / 49.85	53.01	44.30	38.07	22.88				
UA-MT	MICCAI'19	20 38.02 / 25.51 / 14.94	61.85 / 54.27 / 47.33	43.13 / 35.66 / 28.54	41.89 / 38.25 / 26.11	37.96	29.14	72.35	40.84				
FDA	CVPR'20	20 61.66 / 36.32 / 34.71	80.67 / 70.99 / 56.75	73.80 / 63.62 / 58.36	77.23 / 68.87 / 64.33	62.28	53.33	25.99	16.10				
SIFA	TMI'20	20 63.97 / 34.64 / 37.25	66.46 / 44.97 / 43.22	56.52 / 39.63 / 38.55	59.57 / 40.63 / 40.26	47.14	34.96	25.01	11.45				
FixMatch	NeurIPS'20	20 <u>87.26</u> / <u>77.78</u> / <u>77.14</u>	<u>91.06</u> / <u>82.78</u> / <u>79.07</u>	<u>87.84</u> / <u>80.07</u> / <u>78.03</u>	<u>90.86</u> / <u>81.75</u> / <u>81.84</u>	<u>82.96</u>	<u>73.99</u>	<u>6.21</u>	<u>3.51</u>				
CPS	CVPR'21	20 46.40 / 29.01 / 16.70	71.48 / 63.08 / 49.39	44.38 / 39.43 / 32.42	47.71 / 40.75 / 29.75	42.54	33.82	58.30	34.94				
CoraNet	TMI'21	20 65.70 / 27.79 / 22.16	63.32 / 48.63 / 46.56	64.89 / 48.59 / 45.30	68.38 / 55.88 / 46.79	50.33	40.54	32.98	19.22				
UDA-VAE++	CVPR'22	20 51.14 / 36.20 / 12.99	71.95 / 53.16 / 36.68	57.88 / 41.64 / 30.19	31.71 / 27.32 / 20.48	39.28	28.82	53.90	24.94				
SS-Net	MICCAI'22	20 48.98 / 27.02 / 20.65	59.67 / 52.69 / 40.86	54.92 / 41.40 / 39.62	55.04 / 52.07 / 35.88	44.07	35.89	49.50	32.55				
BCP	CVPR'23	20 85.91 / 73.82 / <u>78.08</u>	85.66 / 74.85 / 76.04	61.61 / 54.05 / 51.87	76.57 / 62.22 / 79.16	71.65	62.67	30.91	18.22				
CauSSL	ICCV'23	20 40.20 / 21.93 / 10.46	50.99 / 42.66 / 31.94	41.05 / 34.00 / 29.95	53.78 / 37.92 / 30.44	35.44	26.73	72.90	37.99				
Ours	this paper	20 87.77 / <u>76.36</u> / 80.65	91.48 / 83.68 / 81.46	89.25 / 82.65 / 82.27	90.91 / 82.34 / 82.86	84.31	75.18	5.15	2.42				
Upper bound	*	91.37 / 83.61 / 84.13	92.33 / 85.12 / 84.43	92.20 / 84.43 / 84.56	92.08 / 83.78 / 83.77	86.82	78.64	4.32	2.14				

$\tau = 0.95$ as the default value in experiments. We use Stochastic Gradient Descent (SGD) with a momentum of 0.9 and weight decay of 0.0001 as optimizer, with an initial learning rate of 0.03. The batch size is set as 8, including 4 labeled data and 4 unlabeled data. The iteration is set to 30,000 for Fundus dataset and 60,000 for Prostate and M&Ms dataset. In the testing stage, the final segmentation results are determined by the student model. We compare our method with other state-of-the-art (SOTA) methods, including supervised methods such as UA-MT [43], FixMatch [33], CPS [9], CoraNet [31], SS-Net [40], BCP [2], and CauSSL [27] as well as domain unsupervised adaptation methods like FDA [41], SIFA [6] and UDA-VAE++ [25]. In each experiment, a small amount of data from one domain is labeled (e.g., Domain 1 in Tab. 1), while the remainder serves as unlabeled data. For the upper bound, we employ UCP in FixMatch, and use all available training data from a certain domain as labeled data, providing the model with sufficient source domain information. We adopt the Dice coefficient (DC), Jaccard coefficient (JC), 95% Hausdorff Distance (HD), and Average Surface Distance (ASD) as evaluation metrics. Except for

SIFA, which uses ResNet blocks [17] for the generator and decoder, all other methods use U-Net [28] as the backbone.

4.3. Comparison with State-of-the-Art Methods

Results on Fundus dataset. We conduct experiments on Fundus dataset with only 20 labeled data, which is a dual object segmentation task with overlapping regions. As shown in Tab. 1, our method consistently outperforms other approaches in all metrics, and exhibits a minimal gap from the upper bound. Some SSMS methods exhibit performance degradation due to error accumulation, to the extent that their performance is even worse than training with labeled data alone. UDA methods face challenges in achieving effective knowledge transfer with limited labeling and multiple target domains. The segmentation result examples are shown in Fig. 6, demonstrating that our method achieves superior segmentation for samples from the same domain as labeled data and those from different domains.

Results on Prostate dataset. In the case of Prostate dataset, which is a single object segmentation task, our method demonstrates a more significant advantage and increased robustness when trained with 40 labeled data. The

Table 4. Ablation experiments on Fundus dataset.

Method	Task					Optic Cup / Disc Segmentation				
	UCP	VanillaGD	SymGD	TP-RAM	RAM	DC \uparrow				Avg.
						Domain 1	Domain 2	Domain 3	Domain 4	
#1	✓					82.67 / 93.11	72.08 / 88.86	82.97 / 92.78	80.84 / 92.94	85.78
#2	✓	✓				83.19 / 93.45	73.57 / 89.48	82.16 / 92.91	80.42 / 93.35	86.07
#3	✓		✓			83.21 / 93.48	76.13 / 89.06	83.04 / 92.87	83.63 / 93.69	86.89
#4	✓			✓		83.41 / 93.54	77.18 / 88.96	82.69 / 92.88	83.54 / 93.39	86.95
#5	✓		✓		✓	83.27 / 93.41	76.14 / 88.46	83.27 / 92.90	83.64 / 93.40	86.82
#6	✓		✓	✓		83.71 / 92.96	80.47 / 89.93	84.18 / 92.97	83.71 / 93.38	87.66

Table 5. Varying the low-frequency region size parameter β on Fundus dataset.

β	0.1	0.05	0.01	0.005
Avg. DC \uparrow	87.32	87.49	87.66	87.52

disparity from the upper bound is also minimal. Given the larger differences between domains in Prostate dataset, SSMS methods are notably influenced by more severe error accumulation. UDA methods also demonstrate poor performance in this scenario. The segmentation result examples are presented in the supplementary material.

Results on M&Ms dataset. We conducted experiments on M&Ms dataset using 20 labeled data. The results in Tab. 3 demonstrate that our approach is effective in multi-object segmentation tasks, achieving optimal performance. Due to the abundance of samples in M&Ms dataset, the substantial increase in labeled data significantly enhances the performance of the upper bound. The segmentation result examples are presented in the supplementary material.

It is worth noting that while BCP considers the distribution difference between labeled and unlabeled data in SSMS, our method still exhibits significant advantages in the scenario where there is distribution difference within unlabeled data and between unlabeled and labeled data. We also compare our method with more SSMS methods aimed at mitigating distribution differences, and detailed analysis results can be found in the supplementary material.

4.4. Ablation Study

In this paragraph, to validate the effectiveness of each component, we conduct a series of ablation experiments on Fundus dataset. The investigated settings are introduced as follows: 1) UCP: Generating intermediate samples based on UCP, 2) VanillaGD: Guiding p^s with \hat{p} , 3) SymGD: Guiding p^s with merge of \hat{p}_{out} and \hat{p}_{in} , 4) TP-RAM: Gradually increasing the range of value for β , 5) RAM: β is uniformly sampled from the distribution in the range from 0 to 1. The results are shown in Tab. 4.

The Efficiency of SymGD. Compared to #1, #3 introduces an additional guidance branch that guides the predictions of unlabeled data by integrating pseudo labels from dual perspectives. SymGD effectively leverages the inter-

mediate domain information, aiding the model in making more accurate predictions for unlabeled data. In #2, we use the pseudo label \hat{p} generated from the teacher model to guide the student model’s prediction p^s , and this straightforward additional consistency constraint has a limited impact on the model performance improvement.

The Efficiency of TP-RAM. Moreover, in #4, we gradually introduce stylistic transition components to intermediate samples, not limited to local semantic fusion. The comprehensive and stable construction of intermediate domains facilitates the transfer of domain knowledge from the domain of labeled data to the domains of unlabeled data. In #5, we replace the TP-RAM in our method with RAM, introducing stylistic transition arbitrarily for intermediate samples. Training with significant stylistic differences from labeled data in the early stages is detrimental to the stable construction of the intermediate domains, and compared to #2, it might even harm the model’s performance.

In #6, we incorporate two new modules to form a unified domain knowledge transfer framework, ultimately achieving the best performance in the MiDSS scenarios.

Size of Low-frequency Region in M_β . We also explore the relationship between the selection of low-frequency region size and model performance. Here, M_β represents the low-frequency region of an image with a size of $2\beta W \times 2\beta H$, $\beta \in \{0.1, 0.005, 0.01, 0.005\}$. Tab. 5 reveals that the model is quite robust to the choice of β . Overall, there is a slight trend of performance improvement followed by a decline, with the best performance achieved when $\beta = 0.01$.

5. Conclusion

In this paper, we explore a novel and practical setting, MiDSS, which involves the simultaneous challenges of limited labeled data and domain shift. To address the issue of declining pseudo label quality due to domain shift, we construct intermediate domains by UCP. Additionally, we introduce SymGD to enhance the utilization of intermediate domains information. Considering the stylistic differences between different domains, we design TP-RAM to introduce comprehensive and stable style transition components to intermediate domains. Extensive results on three datasets demonstrate the effectiveness of our method.

References

- [1] Wenjia Bai, Ozan Oktay, Matthew Sinclair, Hideaki Suzuki, Martin Rajchl, Giacomo Tarroni, Ben Glocker, Andrew King, Paul M Matthews, and Daniel Rueckert. Semi-supervised learning for network-based cardiac mr image segmentation. In *Medical Image Computing and Computer-Assisted Intervention*, pages 253–260. Springer, 2017. **1**
- [2] Yunhao Bai, Duowen Chen, Qingli Li, Wei Shen, and Yan Wang. Bidirectional copy-paste for semi-supervised medical image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11514–11524, 2023. **3, 7**
- [3] Heng Cai, Shumeng Li, Lei Qi, Qian Yu, Yinghuan Shi, and Yang Gao. Orthogonal annotation benefits barely-supervised medical image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3302–3311, 2023. **1**
- [4] Victor M Campello, Polyxeni Gkontra, Cristian Izquierdo, Carlos Martin-Isla, Alireza Sojoudi, Peter M Full, Klaus Maier-Hein, Yao Zhang, Zhiqiang He, Jun Ma, et al. Multi-centre, multi-vendor and multi-disease cardiac segmentation: the m&ms challenge. *IEEE Transactions on Medical Imaging*, 40(12):3543–3554, 2021. **6**
- [5] Agisilaos Chatsias, Thomas Joyce, Rohan Dharmakumar, and Sotirios A Tsaftaris. Adversarial image synthesis for unpaired multi-modal cardiac data. In *Simulation and Synthesis in Medical Imaging: Second International Workshop, SASHIMI*, pages 3–13. Springer, 2017. **3**
- [6] Cheng Chen, Qi Dou, Hao Chen, Jing Qin, and Pheng Ann Heng. Unsupervised bidirectional cross-modality adaptation via deeply synergistic image and feature alignment for medical image segmentation. *IEEE Transactions on Medical Imaging*, 39(7):2494–2505, 2020. **2, 7**
- [7] Lin Chen, Zhixiang Wei, Xin Jin, Huaian Chen, Miao Zheng, Kai Chen, and Yi Jin. Deliberated domain bridging for domain adaptive semantic segmentation. *Advances in Neural Information Processing Systems*, 35:15105–15118, 2022. **3**
- [8] Xuming Chen, Shanlin Sun, Narisu Bai, Kun Han, Qianqian Liu, Shengyu Yao, Hao Tang, Chupeng Zhang, Zhipeng Lu, Qian Huang, et al. A deep learning-based auto-segmentation system for organs-at-risk on whole-body computed tomography images for radiation therapy. *Radiotherapy and Oncology*, 160:175–184, 2021. **1**
- [9] 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. **7**
- [10] Sabyasachi Dash, Sushil Kumar Shakyawar, Mohit Sharma, and Sandeep Kaushik. Big data in healthcare: management, analysis and future prospects. *Journal of Big Data*, 6(1):1–25, 2019. **2**
- [11] Qi Dou, Cheng Ouyang, Cheng Chen, Hao Chen, and Pheng-Ann Heng. Unsupervised cross-modality domain adaptation of convnets for biomedical image segmentations with adversarial loss. *arXiv preprint arXiv:1804.10916*, 2018. **2, 3**
- [12] Yue Duan, Zhen Zhao, Lei Qi, Luping Zhou, Lei Wang, and Yinghuan Shi. Towards semi-supervised learning with non-random missing labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16121–16131, 2023. **1**
- [13] Deng-Ping Fan, Tao Zhou, Ge-Peng Ji, Yi Zhou, Geng Chen, Huazhu Fu, Jianbing Shen, and Ling Shao. Inf-net: Automatic covid-19 lung infection segmentation from ct images. *IEEE Transactions on Medical Imaging*, 2020. **1**
- [14] Jiashuo Fan, Bin Gao, Huan Jin, and Lihui Jiang. Ucc: Uncertainty guided cross-head co-training for semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9947–9956, 2022. **3**
- [15] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016. **3**
- [16] Hao Guan and Mingxia Liu. Domain adaptation for medical image analysis: a survey. *IEEE Transactions on Biomedical Engineering*, 69(3):1173–1185, 2021. **2**
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. **7**
- [18] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. **1**
- [19] Shuailin Li, Chuyu Zhang, and Xuming He. Shape-aware semi-supervised 3d semantic segmentation for medical images. In *Medical Image Computing and Computer-Assisted Intervention*, pages 552–561. Springer, 2020. **2**
- [20] Shumeng Li, Heng Cai, Lei Qi, Qian Yu, Yinghuan Shi, and Yang Gao. Pln: Parasitic-like network for barely supervised medical image segmentation. *IEEE Transactions on Medical Imaging*, 42(3):582–593, 2022. **1**
- [21] 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. **1**
- [22] Zekun Li, Lei Qi, Yinghuan Shi, and Yang Gao. Iomatch: Simplifying open-set semi-supervised learning with joint inliers and outliers utilization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15870–15879, 2023. **1**
- [23] Quande Liu, Qi Dou, and Pheng-Ann Heng. Shape-aware meta-learning for generalizing prostate mri segmentation to unseen domains. In *Medical Image Computing and Computer-Assisted Intervention*, pages 475–485. Springer, 2020. **1, 4, 6**
- [24] Quande Liu, Cheng Chen, Jing Qin, Qi Dou, and Pheng-Ann Heng. Feddg: Federated domain generalization on medical image segmentation via episodic learning in continuous frequency space. In *Proceedings of the IEEE/CVF Conference*

- on *Computer Vision and Pattern Recognition*, pages 1013–1023, 2021. 2, 5
- [25] Changjie Lu, Shen Zheng, and Gaurav Gupta. Unsupervised domain adaptation for cardiac segmentation: Towards structure mutual information maximization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2588–2597, 2022. 7
- [26] 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*, pages 8801–8809, 2021. 2
- [27] Juzheng Miao, Cheng Chen, Furui Liu, Hao Wei, and Pheng-Ann Heng. Causl: Causality-inspired semi-supervised learning for medical image segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21426–21437, 2023. 2, 7
- [28] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015. 6, 7
- [29] Vivek A Rudrapatna, Atul J Butte, et al. Opportunities and challenges in using real-world data for health care. *The Journal of Clinical Investigation*, 130(2):565–574, 2020. 2
- [30] Paolo Russo, Fabio M Carlucci, Tatiana Tommasi, and Barbara Caputo. From source to target and back: symmetric bi-directional adaptive gan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8099–8108, 2018. 2, 3
- [31] Yinghuan Shi, Jian Zhang, Tong Ling, Jiwen Lu, Yefeng Zheng, Qian Yu, Lei Qi, and Yang Gao. Inconsistency-aware uncertainty estimation for semi-supervised medical image segmentation. *IEEE Transactions on Medical Imaging*, 41(3):608–620, 2021. 7
- [32] Zhao Shi, Chongchang Miao, U Joseph Schoepf, Rock H Savage, Danielle M Dargis, Chengwei Pan, Xue Chai, Xiu Li Li, Shuang Xia, Xin Zhang, et al. A clinically applicable deep-learning model for detecting intracranial aneurysm in computed tomography angiography images. *Nature Communications*, 11(1):6090, 2020. 1
- [33] 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. 4, 7
- [34] 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. 4, 6
- [35] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schuster, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7472–7481, 2018. 3
- [36] Pratima Upreti and Bishesh Khanal. Fixmatchseg: Fixing fixmatch for semi-supervised semantic segmentation. *arXiv preprint arXiv:2208.00400*, 2022. 4
- [37] Jiacheng Wang, Xiaomeng Li, Yiming Han, Jing Qin, Liansheng Wang, and Zhou Qichao. Separated contrastive learning for organ-at-risk and gross-tumor-volume segmentation with limited annotation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2459–2467, 2022. 2
- [38] Qin Wang, Wen Li, and Luc Van Gool. Semi-supervised learning by augmented distribution alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1466–1475, 2019. 2
- [39] Shujun Wang, Lequan Yu, Kang Li, Xin Yang, Chi-Wing Fu, and Pheng-Ann Heng. Dofe: Domain-oriented feature embedding for generalizable fundus image segmentation on unseen datasets. *IEEE Transactions on Medical Imaging*, 39(12):4237–4248, 2020. 6
- [40] Yicheng Wu, Zhonghua Wu, Qianyi Wu, Zongyuan Ge, and Jianfei Cai. Exploring smoothness and class-separation for semi-supervised medical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention*, pages 34–43. Springer, 2022. 2, 7
- [41] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4085–4095, 2020. 2, 3, 7
- [42] Chenyu You, Yuan Zhou, Ruihan Zhao, Lawrence Staib, and James S Duncan. Simcvd: Simple contrastive voxel-wise representation distillation for semi-supervised medical image segmentation. *IEEE Transactions on Medical Imaging*, 41(9):2228–2237, 2022. 1
- [43] Lequan Yu, Shujun Wang, Xiaomeng Li, Chi-Wing Fu, and Pheng-Ann Heng. Uncertainty-aware self-ensembling model for semi-supervised 3d left atrium segmentation. In *Medical Image Computing and Computer-Assisted Intervention*, pages 605–613. Springer, 2019. 1, 2, 7
- [44] Sangdoon 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. 2, 3
- [45] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 3
- [46] Yifan Zhang, Ying Wei, Qingyao Wu, Peilin Zhao, Shuaicheng Niu, Junzhou Huang, and Mingkui Tan. Collaborative unsupervised domain adaptation for medical image diagnosis. *IEEE Transactions on Image Processing*, 29:7834–7844, 2020. 3
- [47] Ziyuan Zhao, Fangcheng Zhou, Kaixin Xu, Zeng Zeng, Cuntai Guan, and S Kevin Zhou. Le-uda: Label-efficient unsupervised domain adaptation for medical image segmentation. *IEEE Transactions on Medical Imaging*, 42(3):633–646, 2022. 3
- [48] 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. 3
- [49] Xiahai Zhuang et al. Challenges and methodologies of fully automatic whole heart segmentation: a review. *Journal of Healthcare Engineering*, 4:371–407, 2013. 2