

# Deep Symmetric Adaptation Network for Cross-modality Medical Image Segmentation

Xiaoting Han<sup>†</sup>, Lei Qi<sup>†</sup>, Qian Yu, Ziqi Zhou, Yefeng Zheng, Yinghuan Shi<sup>\*</sup>, Yang Gao

**Abstract**—Unsupervised domain adaptation (UDA) methods have shown their promising performance in the cross-modality medical image segmentation tasks. These typical methods usually utilize a translation network to transform images from the source domain to target domain or train the pixel-level classifier merely using translated source images and original target images. However, when there exists a large domain shift between source and target domains, we argue that this asymmetric structure could not fully eliminate the domain gap. In this paper, we present a novel deep symmetric architecture of UDA for medical image segmentation, which consists of a segmentation sub-network, and two symmetric source and target domain translation sub-networks. To be specific, based on two translation sub-networks, we introduce a bidirectional alignment scheme via a shared encoder and private decoders to simultaneously align features 1) from source to target domain and 2) from target to source domain, which helps effectively mitigate the discrepancy between domains. Furthermore, for the segmentation sub-network, we train a pixel-level classifier using not only original target images and translated source images, but also original source images and translated target images, which helps sufficiently leverage the semantic information from the images with different styles. Extensive experiments demonstrate that our method has remarkable advantages compared to the state-of-the-art methods in both cross-modality Cardiac and BraTS segmentation tasks.

**Index Terms**—Unsupervised Domain Adaptation, Medical Image Segmentation, Deep Symmetric Architecture.

## I. INTRODUCTION

Deep convolutional neural networks have made significant progress in medical image segmentation task. In the task, there is a common assumption that training and test images are drawn from the same data distribution. However, in numerous real-world applications, especially in healthcare field, due to different acquisition parameters or various imaging modalities, a large gap of data distributions between training and test sets possibly occurs. This distribution gap usually causes a drastic performance drop during the deployment of trained model.

Xiaoting Han and Lei Qi are the co-first authors. Corresponding Author: Yinghuan Shi.

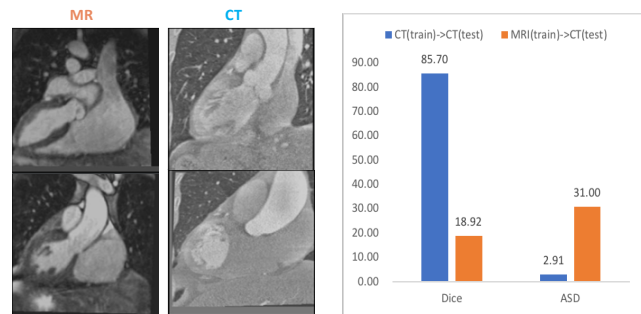
Xiaoting Han, Ziqi Zhou, Yinghuan Shi, and Yang Gao are with the State Key Laboratory for Novel Software Technology, the National Institute for Healthcare Data Science, and the Collaborative Innovation Center of Novel Software Technology and Industrialization, Nanjing University, China. (e-mail: hxt@smail.nju.edu.cn, zhouzq@smail.nju.edu.cn, syh@nju.edu.cn, gaoy@nju.edu.cn)

Lei Qi is with the School of Computer Science and Engineering, and the Key Lab of Computer Network and Information Integration (Ministry of Education), Southeast University, China. (e-mail: qilei@seu.edu.cn)

Qian Yu is with the School of Data and Computer Science, Shandong Women’s University, China. (e-mail: yuqian@sdwu.edu.cn)

Yefeng Zheng is with the Tencent Jarvis Lab, China. (e-mail: yefengzheng@tencent.com)

This work was supported by the National Key research and Development Plan (2019YFC0118300) and NSFC (61673203).



(a) Examples of CT and MR images. (b) Performance of different settings.

Fig. 1. Comparison of the models trained respectively on CT and MR images while evaluating on CT images. Note that, the smaller for “ASD” is better, and reverse on “Dice”.

For example, taking magnetic resonance imaging (MRI) and computed tomography (CT) in Fig. 1 as examples, we could observe there exists a large appearance variation between these two modalities. To quantitatively evaluate the performance drop caused by domain gap, we first train two respective segmentation models purely on CT or MR images, and then utilize these two models to segment a common CT image set, which denotes as “CT to CT” and “MRI to CT”, respectively. It is obvious that a drastic performance degeneration happens in “MRI to CT” compared to “CT to CT”, in terms of Dice and average surface distance (ASD). In this aspect, the model trained purely on MR images cannot directly generalize well on CT images due to this aforementioned domain gap.

To reduce the performance drop across different modalities, a trivial yet straightforward way is to fine-tune the model pre-trained on source images with labeled target images. However, the labeling process which requires pixel-wise manual delineation on target images is often labor-intensive to achieve, especially when we have a massive number of images for labeling. To tackle this issue, unsupervised domain adaptation (UDA)-based segmentation methods [1]–[4] have aroused considerable attention in recent years. In UDA segmentation, given a fully labeled source domain, our goal is to segment another target domain without labels in a cross-domain manner. Thanks to the recent development of UDA-based segmentation, the performance of cross-domain segmentation has now been greatly improved.

For these existing UDA-based segmentation methods, most of them focus on minimizing the differences between distributions of source and target domains from the *image-translation* or *feature-alignment* perspectives [5]. Specifically, for the image-translation perspective, a common way [6]–[9]

is to translate source domain images to the style of target domain images by image translation networks [10], [11]. The translated source images with their inherited ground truth labels from the source domain are utilized to train a target domain-oriented segmentation model in a supervised manner. However, due to the instability of generative adversarial network (GAN)-based methods [4], [12], [13], some inferior translations could destruct semantic information of original images, thus it may affect the segmentation accuracy.

For the feature-alignment perspective, several works [14]–[16] align distributions across domains in feature space to mitigate domain gap. Recent studies [17], [18] propose to project the feature space to other compact spaces, such as image space, since latent space of semantic segmentation might be over-complex and high-dimensional, which should simultaneously encode appearance, shape and context, *etc.* In this paper, we also apply adversarial learning in generated image space to mitigate distribution discrepancy.

Currently, the feature alignment is usually achieved by an independent network [8], [19]. Instead, we choose to bidirectionally align feature distributions across domains using two symmetric translation sub-networks, which achieves a significant performance improvement as compared to applying an independent network.

In this paper, we put forward a novel deep symmetric adaptation network including a segmentation sub-network, and source/target domain translation sub-networks. In this network, we utilize two translation sub-networks via sharing the encoder and using private decoders to realize the bidirectional alignment (*i.e.*, from source to target and vice versa) of feature distributions between source and target domains, which can effectively alleviate the large discrepancy between domains. Moreover, all images from the translation sub-networks and raw source and target domains are used to train the segmentation sub-network, thus it helps sufficiently explore rich yet easily ignored semantic information. We evaluate our method on cross-modality Cardiac [20] and BraTS [21] segmentation tasks. Extensive experiments reveal the superiority of our method, compared to current state-of-the-art (SOTA) methods. Moreover, ablation study highlights the effectiveness of each module developed in our method. Main contributions of this paper are summarized as:

- We develop a novel symmetric adaptation network for cross-modality medical image segmentation, which consists of a segmentation sub-network, source and target domain translation sub-networks.
- We propose a bidirectional alignment scheme over source and target translation sub-networks. Besides, all images of different styles from source and target domains are utilized to train the segmentation module.
- We conduct extensive experiments in the cardiac segmentation task: “MRI to CT” and “CT to MRI”. We achieve a new state-of-the-art of 78.50% and 66.45% in mean Dice. We additionally evaluate our methods on the BraTS dataset, where our method outperforms other SOTA methods as well, achieving Dice of 67.18%.

## II. RELATED WORK

The existing UDA-based segmentation methods can be roughly divided into three categories according to different alignment perspectives, *i.e.*, 1) feature-alignment-based methods, 2) image-translation-based methods and 3) joint learning methods of feature alignment and image translation.

### A. Feature-alignment-based Methods

Many UDA methods pay attention to align distributions in the feature space, by reducing the distance metric between features extracted from source and target domains. Among them, maximum mean discrepancy (MMD) loss [22]–[24] is a popularly-used distance metric. As its extension, some statistics of feature distributions such as covariance [25], [26] are also utilized for alignment. Another prominent approach towards domain adaptation for image segmentation is adversarial learning [12], [27]. Several works use adversarial learning for latent features between the encoder and decoder, which resorts to learn domain-invariant representations [28]–[30].

Recent methods in literature [17], [18] perform adversarial learning in low-level space, *e.g.*, image space, to achieve a compact embedding, since latent space of image segmentation might be over-complex and high-dimensional, including various visual cues, such as appearance and context.

### B. Image-translation-based Methods

Inspired by GAN-based techniques [6], [12], some image-translation-based methods [7], [10], [31] are developed to convert source images into target style-like images. Afterwards, these generated images inheriting ground truth of source images can be used for supervised training of target segmentation network. For example, the PixelDA method [10] addresses domain adaptation by translating source images to target domain, thereby obtaining a simulated training set for target images. Huo *et al.* [32] first adopt a synthesis network to generate target-like images. Then they feed these generated images to a segmentation network to perform supervised training with source domain labels. However, due to instability of GAN-based methods, semantic information might be destructed during generating cross-domain images.

### C. Joint Learning Methods

Recently, some hybrid works integrate image translation and feature alignment to better mitigate the domain shift [33]–[35]. For example, Chang *et al.* [36] propose a domain invariant structure extraction (DISE) framework. This framework disentangles images into domain-invariant structure and domain-specific texture representation. It further translates image across domains and enables label transfer to improve the segmentation performance. Appearance adaptation networks (AAN) and representation adaptation networks (RAN) are combined in [34]. The former adapts source-domain images to the “style” of the target domain, and the latter attempts to learn domain-invariant representations. In medical image segmentation, Chen *et al.* [8], [9] present a novel unsupervised domain adaptation framework, namely synergistic image and

feature alignment (SIFA). This method considers the case of severe domain shift in cross-modality medical images, and synergistically merges the feature alignment and image translation into a unified network.

However, there are two disadvantages in the aforementioned methods. One is that aligning feature distributions via an independent network is difficult, especially for the case when distribution discrepancy is large. Another is that semantic information of original images might be destructed during generating cross-domain images. In response to these two shortcomings, we focus on the feature alignment using two symmetric translation sub-networks to make up for semantic knowledge for image-translation with source images. Furthermore, we explore more common semantic information using adversarial losses between not only original source and target images but also translated source and target images in the semantic space.

### III. METHOD

In this section, we present technical details of our method. Fig. 2 illustrates the overview of our framework, which consists of a shared encoder ( $E$ ), two domain-specific decoders ( $U_s, U_t$ ) and a classifier ( $C$ ). Among them, the shared encoder and a domain-specific decoder constitute a translation sub-network to reconstruct images and generate cross-domain images. The encoder and classifier form the segmentation sub-network. We focus on 1) bidirectionally aligning feature distributions, exploiting two symmetric translation sub-networks and 2) employing different styles of images to mine rich but easily underestimated semantic information. In the following, we introduce our method in detail from these two aspects. The symbols used in the following sections are summarized in Table I.

TABLE I  
SUMMARY OF SYMBOLS

Symbol	Notation
$\mathcal{X}^s, \mathcal{X}^t$	Image sets of source and target modalities (domains)
$x^s, x^t$	Samples of source and target modalities
$x^{s \rightarrow t}, x^{t \rightarrow s}$	Translated samples
$\mathcal{Y}^s, y^s$	Annotation set and sample label corresponding to the source domain
$E$	Shared encoder
$U_s, U_t$	Decoders for source and target domains
$Z_s, Z_t$	Embeddings of source and target domains
$D_t, D_s$	Discriminators for source and target domains
$D_{pi}$	Discriminator for segmentation map at different levels
$C_i$	Classifier for feature map at different levels

#### A. Translation Sub-network for Feature Alignment

We aim to address the problem of domain shift in medical image segmentation, where source images ( $\mathcal{X}^s$ ) with corresponding labels ( $\mathcal{Y}^s$ ), and target images ( $\mathcal{X}^t$ ) without labels, are given. Our goal is to achieve good performance in target domain by applying the model trained on source domain. However, due to the domain shift between source and target domains, the performance might significantly drop. Based on adversarial losses, the two translation sub-networks in our framework translate images 1) from source to target domain and 2) from target to source domain, respectively. We

also introduce reconstruction loss to better maintain semantic information. Furthermore, we share the encoder among the two translation sub-networks and the segmentation sub-network to achieve the bidirectional feature alignment, which is different from previous adaptation methods. The independent decoders for source and target domains mainly pay attention to domain-specific representations, which forces the shared encoder to merely learn domain-invariant feature representations. The detailed descriptions about adversarial loss and reconstruction loss of translation sub-networks are provided in the following.

1) *Adversarial Loss of Generated Space*: Specifically, the translation sub-network is composed of the shared encoder ( $E$ ) and domain-specific decoder ( $U_t$ ). The encoder takes source images ( $x^s$ ) as input and outputs source embedding ( $Z_s$ ). Then the target decoder ( $U_t$ ) takes source embedding ( $Z_s$ ) as input and generates target-like images to fool the target discriminator ( $D_t$ ), while the discriminator tries to distinguish translated images ( $x^{s \rightarrow t}$ ) from target images ( $x^t$ ). The target translation sub-network that acts as a generator, plays a minimax two-player game with the target discriminator ( $D_t$ ). They can be optimized via the adversarial loss as below:

$$\mathcal{L}_{adv}^t(E, U_t, D_t) = \mathbb{E}_{x^t \sim \mathcal{X}^t} [\log D_t(x^t)] + \mathbb{E}_{x^s \sim \mathcal{X}^s} [\log(1 - D_t(U_t(E(x^s))))], \quad (1)$$

where the discriminator maximizes the objective in the process of optimization. Meanwhile, the encoder and decoder encourage translated images to be indistinguishable from target images, hence they make the objective continuously reduced.

Besides, we can also formulate the adversarial loss of translation sub-network from target to source as below:

$$\mathcal{L}_{adv}^s(E, U_s, D_s) = \mathbb{E}_{x^s \sim \mathcal{X}^s} [\log D_s(x^s)] + \mathbb{E}_{x^t \sim \mathcal{X}^t} [\log(1 - D_s(U_s(E(x^t))))]. \quad (2)$$

2) *Reconstruction Loss*: To ensure translated images ( $x^{s \rightarrow t}$ ) to preserve content information of original images and learn the style of target images, we add the target domain reconstruction loss. That is to say, we first use the encoder ( $E$ ) and target decoder ( $U_t$ ) to reconstruct target domain images, and then optimize them by the reconstruction loss

$$\mathcal{L}_{rec}^t(E, U_t) = \mathbb{E}_{x^t \sim \mathcal{X}^t} \|U_t(E(x^t)) - x^t\|_1. \quad (3)$$

Similarly, we also use the reconstruction loss in source translation sub-network, ensuring source decoder to learn distribution information of source domain as follows:

$$\mathcal{L}_{rec}^s(E, U_s) = \mathbb{E}_{x^s \sim \mathcal{X}^s} \|U_s(E(x^s)) - x^s\|_1. \quad (4)$$

So far, we obtain the optimization losses for target and source domain translation sub-networks, which can be calculated as follows:

$$\mathcal{L}_{gen}^t = \lambda_{rec} \mathcal{L}_{rec}^t + \lambda_{adv} \mathcal{L}_{adv}^t, \quad (5)$$

$$\mathcal{L}_{gen}^s = \lambda_{rec} \mathcal{L}_{rec}^s + \lambda_{adv} \mathcal{L}_{adv}^s, \quad (6)$$

where  $\lambda_{rec}$ ,  $\lambda_{adv}$  are hyper-parameters, which are used to adjust the weight of each component. The corresponding values in the experiment are 1.0, 0.01, respectively.

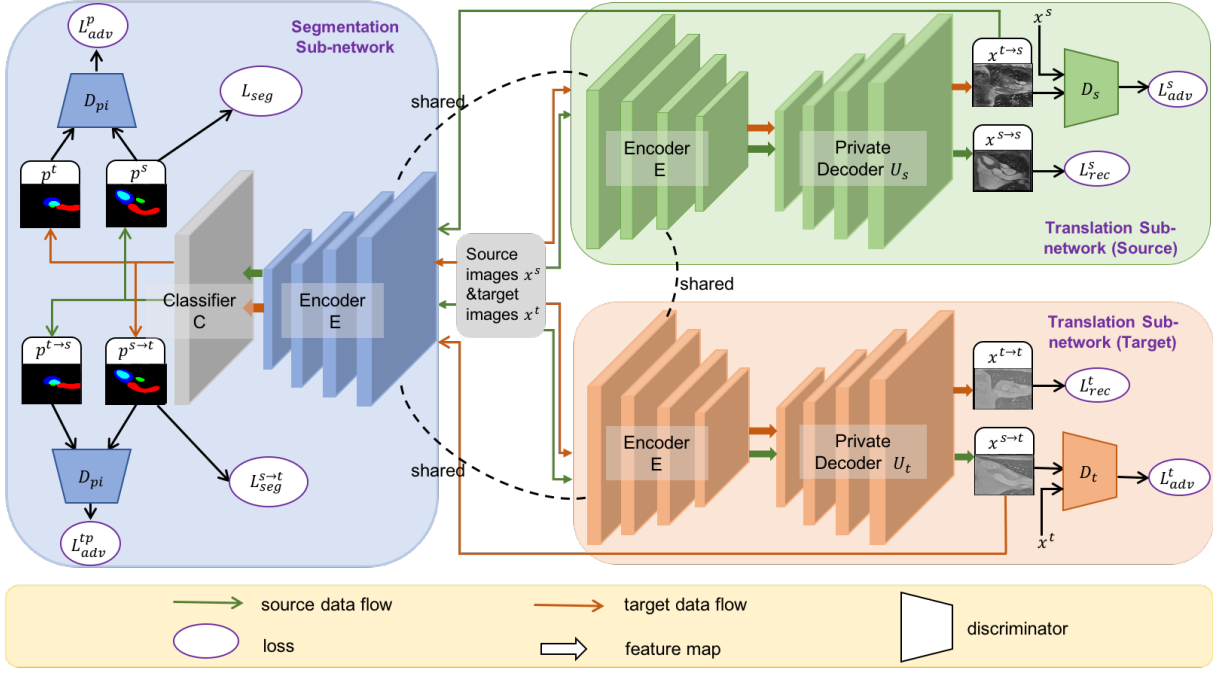


Fig. 2. An overview of our proposed method. The whole network is a completely symmetric framework, which is composed of a common encoder ( $E$ ) shared across domains, two domain-specific private decoders ( $U_s, U_t$ ) and a pixel-wise classifier ( $C$ ). Among them, the shared encoder and a domain-specific decoder constitute a translation sub-network for reconstructing images and generating cross-domain images. In addition, the encoder and classifier form the segmentation sub-network. The green and orange blocks in the figure respectively represent the source and target domain translation sub-network, and the blue block is the segmentation sub-network.

### B. Segmentation Sub-network for Semantic Mining

The previous methods [8], [9] utilize translated source images to learn the semantic information. Since translated source images might miss some semantic information in the image translation stage, caused by the instability of GAN-based methods [12], [13], we utilize both translated source images and original source images to train the segmentation sub-network. Furthermore, the raw target images and translated target images are employed to conduct the adversarial task, which can help the segmentation sub-network focus more on the common semantic knowledge of source and target domains. This sub-network includes two kinds of losses, *i.e.*, segmentation loss and adversarial loss.

1) *Segmentation Loss*: Since all images from source domain have corresponding labels, we feed source images ( $x^s$ ) into the encoder ( $E$ ) and pixel-level classifier ( $C$ ) to generate prediction maps. Then we optimize the whole segmentation sub-network by the segmentation loss. Inspired from the SIFA [8], we also apply the idea of deep supervision. We add an additional classifier for prediction of lower feature maps. To overcome the class imbalance issue between relative small-sized foreground and large-sized background, we employ a sum of soft Dice and weighted cross-entropy (CE) loss to train the segmentation sub-network as follows:

$$\mathcal{L}_{seg}^s(E, C_i) = \mathcal{L}_{CE}(C_i(E(x^s)), y^s) + \mathcal{L}_{Dice}(C_i(E(x^s)), y^s), i = 1, 2. \quad (7)$$

In addition, the translated images ( $x^{s \rightarrow t}$ ) maintain some content information of original images, and they can inherit the ground truth of source images, thus these images can be

used for supervised training of the segmentation sub-network. We feed the translated images ( $x^{s \rightarrow t}$ ) into segmentation sub-network and calculate the segmentation loss with source domain labels. The cross-domain segmentation loss, which is similar to source domain segmentation loss, is calculated as follows:

$$\mathcal{L}_{seg}^{s \rightarrow t}(E, C_i) = \mathcal{L}_{CE}(C_i(E(x^{s \rightarrow t})), y^s) + \mathcal{L}_{Dice}(C_i(E(x^{s \rightarrow t})), y^s), i = 1, 2. \quad (8)$$

2) *Adversarial Loss of Semantic Space*: It is known that segmentation outputs of source and target domains share a significant amount of similarity, *e.g.*, spatial layout and local context. Also, source domain can get corresponding labels, while there is no label information in target domain. Nevertheless, prediction maps can be utilized to aid the segmentation sub-network to mine the shared semantic knowledge between source and target domains. Specifically, we input target images into the segmentation sub-network, and the prediction maps, together with the source prediction maps, are feed into the discriminators ( $D_{pi}$ ), which aims to distinguish the prediction maps from source domain or target domain. Differently, the segmentation sub-network fools the discriminators ( $D_{pi}$ ). Therefore, the adversarial loss in the semantic space of real images is defined as:

$$\mathcal{L}_{adv}^p(E, C_i, D_{pi}) = \mathbb{E}_{x^s \sim \mathcal{X}^s} [\log D_{pi}(C_i(E(x^s)))] + \mathbb{E}_{x^t \sim \mathcal{X}^t} [\log(1 - D_{pi}(C_i(E(x^t))))]. \quad (9)$$

Similarly, translated source images and translated target images are available from the translation sub-networks, and translated source images inherit source labels. We add the cor-

responding adversarial loss in the semantic space of generated images, which can be described as:

$$\mathcal{L}_{adv}^{tp}(E, C_i, D_{pi}) = \mathbb{E}[\log D_{pi}(C_i(E(x^{s \rightarrow t}))) + \mathbb{E}[\log(1 - D_{pi}(C_i(E(x^{t \rightarrow s})))]. \quad (10)$$

To sum up, in order to mine more semantic information, the whole adversarial loss in the semantic space can be expressed as:

$$\mathcal{L}_{adv}^{sec}(E, C_i, D_{pi}) = \mathcal{L}_{adv}^p(E, C_i, D_{pi}) + \mathcal{L}_{adv}^{tp}(E, C_i, D_{pi}). \quad (11)$$

### C. The Overall Loss

In summary, the whole objective of our network for unsupervised domain adaptation is formulated as follows:

$$\mathcal{L} = \lambda_{gen}(\mathcal{L}_{gen}^t + \mathcal{L}_{gen}^s) + \lambda_{seg}(\mathcal{L}_{seg}^s + \mathcal{L}_{seg}^{s \rightarrow t}) + \lambda_{adv}^{sec} \mathcal{L}_{adv}^{sec}, \quad (12)$$

where  $\lambda_{gen}$ ,  $\lambda_{seg}$ , and  $\lambda_{adv}^{sec}$  are hyper-parameters adjusting the weight of each module. The corresponding values in our experiment are set to 1.0, 1.0 and 0.1, respectively.

### D. Network Architecture

The whole framework consists of the segmentation sub-network and two translation sub-networks. Among them, these three parts share an encoder. The segmentation sub-network, which is built on Deeplab-ResNet50, includes the encoder ( $E$ ) and classifiers ( $C_i$ ). Two domain-specific decoders have the same architecture, but they do not share any weights. They are composed of three residual blocks and three deconvolution layers. There are three kinds of discriminators, which are used to distinguish true and generated images of source domain, true and generated images of target domain, and different domains of output prediction maps. All of them have the same architecture, including four convolution layers, but do not share any parameters with each other.

## IV. EXPERIMENTS

In this section, we evaluate the effectiveness of our proposed method from both qualitative and quantitative perspectives. Specifically, we compare our method with several state-of-the-art methods and perform some ablation experiments to investigate the effect of various constraints on model performance.

### A. Datasets and Implementation Details

1) *Datasets*: In our experiment, we utilize two datasets (*i.e.*, Cardiac dataset [20] and BraTS dataset [21]) to validate the efficacy of our method. The first dataset is the Multi-Modality Whole Heart Segmentation (MMWHS) Challenge 2017 dataset [20], which consists of unpaired 20 MRI and 20 CT volumes from different clinical sites. We complete adaptation experiments both in the ‘‘MRI to CT’’ direction and in the ‘‘CT to MRI’’ direction. The label mask of cardiac images is manually delineated by the professional radiologists. The segmentation labels contain four cardiac structures, which are the left ventricle myocardium (LVM), left atrium blood

cavity (LAB), left ventricle blood cavity (LVB) and ascending aorta (AA). The other dataset is the Multi-Modality Brain Tumor Segmentation Challenge 2018 dataset [21], which is composed of four modalities of MRI imaging: FLAIR, T1, T1CE, and T2. In the experiment, T2 is treated as the source domain and the other modalities are considered as the target domains. We aim to segment the whole brain tumor of target domains in an unpaired way. For both datasets, we randomly select 80% of patient data as the training set and 20% as the test set. The image intensity is first normalized by subtracting the mean intensity and dividing by the standard deviation, and then transferred to range [-1, 1]. We perform the random crop, rotation and other augmentation operations in the data pre-processing stage.

2) *Implementation Details*: We implement our model with PyTorch on a Tesla V100 with 32 GB memory. The model on the Cardiac dataset is trained for  $2 \times 10^4$  iterations and the batch size is set to 8. The full training process takes 10 hours. On the BraTS dataset, we train the model for  $6 \times 10^4$  iterations with a batch size of 8, and the full training process takes 24 hours. For both tasks, the shared encoder ( $E$ ) and classifiers ( $C_i$ ) are trained with the Adam optimizer with learning rate of  $2 \times 10^{-4}$ . The two domain-specific decoders are trained with the Adam optimizer using the learning rate of  $1 \times 10^{-3}$ . Differently, the domain-specific discriminators ( $D_s$  and  $D_t$ ) are trained with the Adam optimizer with learning rate of  $1 \times 10^{-4}$ , and the discriminators of semantic space ( $D_{pi}$ ) are trained with the Adam optimizer with learning rate of  $5 \times 10^{-5}$ .

### B. Evaluation Metrics

There are three evaluation metrics in our experiment, including the Dice similarity coefficient (Dice), the average surface distance (ASD) and Hausdorff distance (HD). Dice is mainly used to calculate the similarity between prediction map and ground truth. Higher Dice value indicates better segmentation performance. ASD is used to calculate the average distances between the surface of the prediction mask and ground truth in 3D. And HD is introduced to measure the extreme distance between two sets of points. Lower ASD and HD values indicate better performance. Being consistent with previous works [8], [19], we choose Dice and ASD as evaluation metrics on the Cardiac dataset, and choose Dice and HD as evaluation metrics on the BraTS dataset.

### C. Comparison with the State-of-the-art Methods

We now report the quantitative results of our method and the state-of-the-art (SOTA) methods in the Cardiac and BraTS segmentation tasks.

In Table II, we visualize the adaptation results in the task ‘‘MRI to CT’’ and ‘‘CT to MRI’’ of the Cardiac dataset. First, the result of ‘‘supervised training’’ can be regarded as the upper bound of unsupervised domain adaptation methods. As seen in Table II, it achieves Dice of 85.70% in the ‘‘MRI to CT’’ task and 84.59% in the reverse task. ‘‘No adaptation’’ denotes the evaluation results on CT images using the model trained on MR images or vice versa, which can be taken as the lower bound of unsupervised domain adaptation methods. When

TABLE II  
COMPARISON OF DIFFERENT METHODS ON THE CARDIAC DATASET. THE **BOLD** NUMBER HIGHLIGHTS THE BEST PERFORMANCE.

Cardiac MRI → Cardiac CT										
Method	Dice (%)					ASD (voxel)				
	AA	LAC	LVC	MYO	Average	AA	LAC	LVC	MYO	Average
Supervised training	76.14	89.82	91.37	85.47	85.70	3.39	3.64	2.44	2.19	2.91
No adaptation	37.02	33.71	0.27	4.65	18.92	28.64	8.56	58.50	28.29	31.00
PnP-AdaNet [37]	74.00	68.90	61.90	50.80	63.90	12.80	6.30	17.40	14.70	12.80
AdaOutput [2]	65.20	76.60	54.40	43.60	59.90	17.90	<b>5.50</b>	5.90	8.90	9.60
CycleGAN [6]	73.80	75.70	52.30	28.70	57.60	11.50	13.60	9.20	8.80	10.80
CyCADA [33]	72.90	77.00	62.40	45.30	64.40	9.60	8.00	9.60	10.50	9.40
SIFA [8]	81.30	79.50	73.80	61.60	74.10	7.90	6.20	5.50	8.50	7.00
DSFN [19]	<b>84.70</b>	76.90	79.10	62.40	75.80	N/A	N/A	N/A	N/A	N/A
Ours	79.92	<b>84.76</b>	<b>82.77</b>	<b>66.52</b>	<b>78.50</b>	<b>7.68</b>	6.65	<b>3.77</b>	<b>5.59</b>	<b>5.92</b>

Cardiac CT → Cardiac MRI										
Method	Dice (%)					ASD (voxel)				
	AA	LAC	LVC	MYO	Average	AA	LAC	LVC	MYO	Average
Supervised training	82.24	86.33	91.51	78.28	84.59	3.44	1.82	1.52	1.72	2.13
No adaptation	0.08	9.31	12.50	0.73	5.66	57.92	18.09	16.12	16.81	27.23
PnP-AdaNet [37]	43.70	47.00	77.70	48.60	54.30	11.40	14.50	4.50	5.30	8.90
AdaOutput [2]	60.80	39.80	71.50	35.50	51.90	5.70	8.00	4.60	4.60	5.70
CycleGAN [6]	64.30	30.70	65.00	43.00	50.70	5.80	9.80	6.00	5.00	6.60
CyCADA [33]	60.50	44.00	77.60	47.90	57.50	7.70	13.90	4.80	5.20	7.90
SIFA [8]	65.30	62.30	<b>78.90</b>	47.30	63.40	7.30	7.40	<b>3.80</b>	4.40	5.70
Ours	<b>71.29</b>	<b>66.23</b>	76.20	<b>52.07</b>	<b>66.45</b>	<b>4.44</b>	<b>7.30</b>	5.46	<b>4.25</b>	<b>5.36</b>




Fig. 3. Cardiac segmentation results of different unsupervised domain adaptation methods. The first two rows are CT images in the “MRI to CT” task, and the last two rows are MR images in the reverse task.

compared to “supervised training”, we can observe a drastic drop in performance due to the distribution discrepancy. For example, the Dice is reduced by 66.78% in the “MRI to CT” task and 78.93% in the “CT to MRI” task.

We choose several recently proposed SOTA methods for comparison. In particular, PnP-AdaNet [37] aligns the feature spaces of source and target domains in an unsupervised

manner; The AdaOutput [2] conducts the adversarial task between source and target prediction maps; The CycleGAN [6] performs data augmentation in the target domain via image-to-image translation methods; The CyCADA [33] adapts representations at both the image-level and feature-level; The SIFA [8] conducts synergistic alignment of domain from both image and feature perspectives; And DSFN [19] introduces

TABLE III  
COMPARISON OF DIFFERENT METHODS ON THE BRATS DATASET. THE **BOLD** NUMBER HIGHLIGHTS THE BEST PERFORMANCE.

Method	Dice (%)				Hausdorff Distance (mm)			
	T1	FLAIR	T1CE	Average	T1	FLAIR	T1CE	Average
Supervised training	73.17	85.58	72.56	77.10	9.47	4.57	9.18	7.74
No adaptation	4.17	65.16	6.28	27.70	55.67	28.00	49.77	39.56
AdaOutput [2]	42.60	67.80	33.10	47.80	23.40	16.60	28.80	22.90
CycleGAN [6]	38.10	63.30	42.10	47.80	25.40	17.20	23.20	21.90
CyCADA [33]	49.60	72.00	51.70	57.80	20.20	14.90	18.40	17.80
SIFA [8]	51.70	68.00	58.20	59.30	19.60	16.90	15.01	17.10
DSFN [19]	57.30	78.90	<b>62.20</b>	66.10	17.50	13.80	15.50	15.60
Ours	<b>57.70</b>	<b>81.79</b>	62.04	<b>67.18</b>	<b>14.24</b>	<b>8.62</b>	<b>13.70</b>	<b>12.19</b>




Fig. 4. Brain tumor segmentation results of different methods in the unsupervised domain adaptation task. The input images from top to bottom are from T1, FLAIR, and T1CE modality, respectively.

a dual-scheme fusion network to reduce the domain gap. Among these methods, PnP-AdaNet and CycleGAN are the feature-alignment-based method and image-translation-based method, respectively. CyCADA, SIFA and DSFN are the joint learning methods of feature alignment and image translation. In particular, our method also belongs to the joint learning method. Different from other joint learning methods, from the feature alignment perspective, we propose the bidirectional feature alignment using two symmetric translation sub-networks. From the image translation perspective, we exploit all styled images to train the segmentation sub-network via segmentation loss and adversarial loss.

According to Table II, we can observe that 1) joint learning paradigm outperforms the independent image-translation or feature-alignment manner. The main reason is that image translation can help achieve better feature alignment in joint learning paradigm and vice versa. For example, PnP-AdaNet and CycleGAN perform worse than joint learning methods of feature alignment and image translation (*i.e.*, CyCADA, SIFA, and DSFN). As for AdaOutput, it adopts the alignment in the semantic space, which also belongs to the category of independent methods. Its test result is also in accordance with this observation. 2) Compared to the existing joint learning methods, our proposed method shows better performance,

which is attributed to the proposed bidirectional alignment and semantic mining schemes. As seen in Table II, compared to CyCADA, our method significantly outperforms CyCADA by 14.10% in “MRI to CT” and 8.95% in “CT to MRI”, which demonstrates the effectiveness of the shared encoder between segmentation sub-network and translation sub-networks. In addition, our method improves the Dice by 4.40% (“MRI to CT”) and 3.05% (“CT to MRI”) when compared to SIFA. The main reason might be that we fully exploit all styled images to incorporate more semantic knowledge. As for DSFN, our method outperforms DSFN by 2.70% in “MRI to CT” since we implement bidirectional feature alignment via two symmetric translation sub-networks. Some typical segmentation results of these adaptation methods on the Cardiac dataset are shown in Fig. 3.

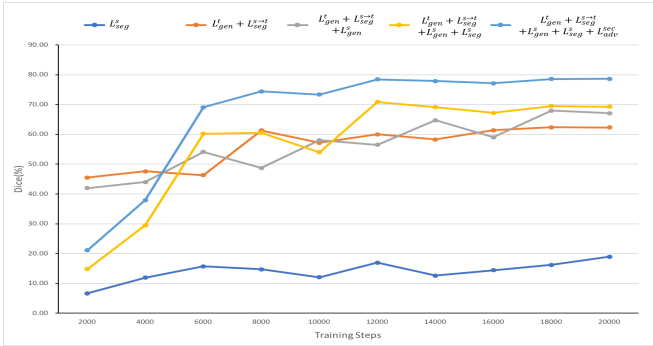
In Table III, we display the segmentation results on the BraTS dataset. Similar with that in the Cardiac dataset, we first report the experimental results of “supervised training” and “no adaptation” methods. Due to the domain gap, the Dice is reduced by 49.40% and HD is increased by 31.82. Since the appearance of FLAIR and T2 modality is similar, we achieve the Dice of 65.16% and the HD of 28.00, respectively, without adaptation. Differently, T1 and T1CE have a significant distribution discrepancy with T2 modality, the Dice

TABLE IV  
EXPERIMENTAL RESULTS OF ABLATION STUDY. THE **BOLD** NUMBER HIGHLIGHTS THE BEST PERFORMANCE.


Cardiac MRI → Cardiac CT										
Method	Dice (%)					ASD (voxel)				
	AA	LAC	LVC	Myo	Average	AA	LAC	LVC	Myo	Average
$\mathcal{L}_{seg}^s$	37.02	33.71	0.27	4.65	18.92	28.64	8.56	58.50	28.29	31.00
$\mathcal{L}_{gen}^t + \mathcal{L}_{seg}^{s \rightarrow t}$	71.11	71.21	56.37	46.67	61.34	<b>6.11</b>	6.84	9.77	8.93	7.91
$\mathcal{L}_{gen}^t + \mathcal{L}_{seg}^{s \rightarrow t} + \mathcal{L}_{gen}^s$	<b>82.11</b>	76.53	70.36	42.75	67.94	7.48	7.27	6.95	7.68	7.34
$\mathcal{L}_{gen}^t + \mathcal{L}_{seg}^{s \rightarrow t} + \mathcal{L}_{gen}^s + \mathcal{L}_{seg}^s$	81.04	77.03	69.42	51.59	69.77	7.82	<b>4.98</b>	10.09	6.24	7.28
$\mathcal{L}_{gen}^t + \mathcal{L}_{seg}^{s \rightarrow t} + \mathcal{L}_{gen}^s + \mathcal{L}_{seg}^s + \mathcal{L}_{adv}^{sec}(all)$	79.92	<b>84.76</b>	<b>82.77</b>	<b>66.52</b>	<b>78.50</b>	7.68	6.65	<b>3.77</b>	<b>5.59</b>	<b>5.92</b>

Cardiac CT → Cardiac MRI										
Method	Dice (%)					ASD (voxel)				
	AA	LAC	LVC	Myo	Average	AA	LAC	LVC	Myo	Average
$\mathcal{L}_{seg}^s$	0.08	9.31	12.50	0.73	5.66	57.92	18.09	16.12	16.81	27.23
$\mathcal{L}_{gen}^t + \mathcal{L}_{seg}^{s \rightarrow t}$	57.19	22.43	67.92	35.77	45.83	6.53	13.38	5.35	5.23	7.62
$\mathcal{L}_{gen}^t + \mathcal{L}_{seg}^{s \rightarrow t} + \mathcal{L}_{gen}^s$	56.50	47.61	75.12	35.30	53.63	6.17	10.73	4.89	5.25	6.76
$\mathcal{L}_{gen}^t + \mathcal{L}_{seg}^{s \rightarrow t} + \mathcal{L}_{gen}^s + \mathcal{L}_{seg}^s$	57.70	50.86	75.79	32.54	54.22	7.80	8.68	<b>4.64</b>	4.68	6.45
$\mathcal{L}_{gen}^t + \mathcal{L}_{seg}^{s \rightarrow t} + \mathcal{L}_{gen}^s + \mathcal{L}_{seg}^s + \mathcal{L}_{adv}^{sec}(all)$	<b>71.29</b>	<b>66.23</b>	<b>76.20</b>	<b>52.07</b>	<b>66.45</b>	<b>4.44</b>	<b>7.30</b>	5.46	<b>4.25</b>	<b>5.36</b>



(a) MRI to CT



(b) CT to MRI

Fig. 5. (a) The intermediate segmentation performance of the proposed network and its variant models after different training steps in “MRI to CT”. (b) The segmentation result boxplot with different ablation experimental settings in “CT to MRI”. The confidence intervals are generated based on the different volumes in the test sets.

are only 4.17% and 6.28%, respectively. From Table III, we can observe that the performance of all adaptation methods on this dataset shows the same trend as the Cardiac dataset. Some typical segmentation results of these adaptation methods for the BraTS dataset are illustrated in Fig. 4. As observed, our segmentation results are closer to supervised results and better than other methods.

#### D. Ablation Study

We perform ablation experiments on Cardiac dataset to verify the impact of various constraints on the entire network. The experimental results are displayed in Table IV. Overall, we focus on bidirectionally aligning feature distributions between source and target domains from the generated image space and learning common semantic knowledge in the semantic space.

1) *Effectiveness of unidirectional alignment*: We perform an experiment that only exploits unidirectional alignment via the “from source to target” translation sub-network and apply the translated segmentation loss ( $\mathcal{L}_{gen}^t + \mathcal{L}_{seg}^{s \rightarrow t}$ ). Compared to the result of “no adaptation”, the average Dice increases by 42.42% and the average ASD is reduced by 23.09 in the “MRI

to CT” task. Besides, in the “CT to MRI” task, the average Dice increases by 40.17% and the average ASD is reduced by 19.61.

2) *Effectiveness of bidirectional alignment*: We perform an experiment to conduct bidirectional feature alignment ( $\mathcal{L}_{gen}^t + \mathcal{L}_{seg}^{s \rightarrow t} + \mathcal{L}_{gen}^s$ ) by adding a symmetric “from target to source” translation sub-network. The two translation sub-networks consist of the shared encoder and private decoders, where two private decoders learn domain-specific representations by reconstruction losses and the shared encoder focuses on learning domain-invariant representations. In this way, our average Dice increases by 6.60% in “MRI to CT” and 7.80% in “CT to MRI”, compared to “ $\mathcal{L}_{gen}^t + \mathcal{L}_{seg}^{s \rightarrow t}$ ” (unidirectional alignment).

3) *Effectiveness of source segmentation*: Due to the instability of GAN-based methods [12], [13], translated source images might miss some semantic information of source images. As a result, we feed source images into the segmentation sub-network to make up for lost semantic information ( $\mathcal{L}_{gen}^t + \mathcal{L}_{seg}^{s \rightarrow t} + \mathcal{L}_{gen}^s + \mathcal{L}_{seg}^s$ ). The Dice increases by 1.83% in “MRI to CT” and 0.59% in “CT to MRI” when compared






Fig. 6. Comparison of the random initialization (blue) and the pre-trained initialization (orange).

TABLE V  
COMPARISON OF DIFFERENT PARAMETER SETTING AND OUR METHOD.  
THE **BOLD** NUMBER HIGHLIGHTS THE BEST PERFORMANCE.

Method	Cardiac MRI $\rightarrow$ Cardiac CT	
	Dice (%)	ASD (voxel)
Shared Decoders	44.93	10.50
Unshared Encoders	73.90	6.90
Unshared Discriminators	76.14	6.04
Ours	<b>78.50</b>	<b>5.92</b>

“Ours” : private decoders, shared encoders, and shared discriminators

to “ $\mathcal{L}_{gen}^t + \mathcal{L}_{seg}^{s \rightarrow t} + \mathcal{L}_{gen}^s$ ”.

4) *Effectiveness of adversarial loss in semantic space*: It is known that not only original source images are accessible to corresponding labels, but translated source images can inherit source labels. Being aware this fact, we use original images and generated images to respectively perform adversarial learning ( $\mathcal{L}_{gen}^t + \mathcal{L}_{seg}^{s \rightarrow t} + \mathcal{L}_{gen}^s + \mathcal{L}_{seg}^s + \mathcal{L}_{adv}^{sec}$ ) in the semantic space, which can exploit all styled images to explore more common semantic information. By this way, the average Dice increases to 78.50% (“MRI to CT”) and 66.45% (“CT to MRI”), and ASD is reduced to 5.92 (“MRI to CT”) and 5.36 (“CT to MRI”). Moreover, Fig. 5 shows the intermediate segmentation performance of the proposed network and its variant models after different training steps in the “MRI to CT” task and shows the “CT to MRI” segmentation result boxplot from different volumes in the MR test set. As observed, each module proposed in our method is useful for performance improvement.

## V. DISCUSSION

In this part, we mainly discuss different parameter settings in each module and our method. We report the experimental results in Fig. 6 and Table V.

1) *Initialization of network parameters*: In the experiment, there are two ways to initialize the network parameters, *i.e.*, 1) random initialization and 2) initialization with a pre-trained network. Specifically, for the latter, we use the pre-trained model ResNet50 on the ImageNet dataset [38]. In the experiment, we only reuse the parameters of layer1, layer2, layer3, layer4 in the pre-trained ResNet50 model, since our segmentation network replaces the last fully connected layer with a convolutional layer and changes the input channel to one. As verified by experiment, due to the initialization using pre-trained model, the encoder can extract more representative

features and the average Dice is improved to 80.73% compared with random initialization, as shown in Fig. 6. However, in order to make a fair comparison with other methods, we adopt the results of random initialization in this paper.

2) *Evaluation of domain-specific decoders*: In our proposed network, there are two translation sub-networks, which are composed of a common shared encoder and two private decoders. The private decoders are used to learn domain-specific representations via reconstruction losses and the shared encoder is used to learn domain-invariant representations via adversarial losses. If we share the decoders of translation sub-networks in our framework, it means that source and target domains apply a common translation network. The shared decoders could not learn domain-specific representations and could not encourage the shared encoder to learn the domain-invariant representations. As verified by experiment (Shared Decoders) in Table V, the result of shared decoders is poorer than private decoders in “ours”. It only achieves the Dice of 44.93% and the ASD of 10.50.

3) *Evaluation of shared encoders*: We share the encoder between the segmentation sub-network and two translation sub-networks. In this way, we achieve the bidirectional feature alignment. If the segmentation sub-network does not share the encoder with two translation sub-networks, it means that they are independent with each other. The network could not achieve the feature alignment from the generated image space. As observed in Table V, the Dice is reduced by 4.60% in “unshared encoders” when compared to shared encoders in “ours”.

4) *Effectiveness of the shared discriminators*: In our experiments, we use two sets of discriminators for adversarial learning in the semantic space. We share the parameters of these two sets of discriminators. Between them, one set of discriminators is used to distinguish prediction maps of real images which domain they come from. In this set, since the source domain is annotated, its label is a positive category and the target domain is a negative category. The other set of discriminators are used to distinguish which domain prediction maps of generated images come from. Since translated source images have corresponding labels, they are marked as a positive class, and translated target images are marked as a negative class. If the parameters are shared between the two sets of discriminators, it means the real and generated images use the same classifier. In addition, source images and translated source images are grouped into the same category, and target images and translated target images are grouped into the other category, which exactly meets the requirements of semantic consistency. As seen in Table V, sharing parameters of discriminators in “ours” improves the Dice by 2.36% when compared to not sharing parameters (*i.e.*, Unshared Discriminators).

## VI. CONCLUSION

In this paper, we proposed a novel symmetric network of unsupervised domain adaptation for medical image segmentation. Our model not only achieves the bidirectional feature alignment via two symmetric translation sub-networks,

but also sufficiently explores the semantic information from different style images. Through the experiments, our network shows significant advantages compared to the state-of-the-art methods on both the Cardiac and BraTS datasets.

Besides, our method can be further improved to better generalize in the target domain. For example, if we can get pseudo-labels of target domain or combine the self-training paradigm, more semantic information from target domain could be explored to achieve better segmentation performance. We will explore this direction in our future work.

## REFERENCES

- [1] K. Saito, K. Watanabe, Y. Ushiku, and T. Harada, "Maximum classifier discrepancy for unsupervised domain adaptation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3723–3732.
- [2] Y.-H. Tsai, W.-C. Hung, S. Schuler, K. Sohn, M.-H. Yang, and M. Chandraker, "Learning to adapt structured output space for semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7472–7481.
- [3] L. A. Gatys, A. S. Ecker, and M. Bethge, "Image style transfer using convolutional neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Jun. 2016, pp. 2414–2423.
- [4] T.-H. Vu, H. Jain, M. Bucher, M. Cord, and P. Pérez, "Dada: Depth-aware domain adaptation in semantic segmentation," in *Proceedings of the IEEE International Conference on Computer Vision*, Oct. 2019, pp. 7364–7373.
- [5] Y. Chen, W. Li, and L. Van Gool, "Road: Reality oriented adaptation for semantic segmentation of urban scenes," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7892–7901.
- [6] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2223–2232.
- [7] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo, "Stargan: Unified generative adversarial networks for multi-domain image-to-image translation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 8789–8797.
- [8] C. Chen, Q. Dou, H. Chen, J. Qin, and P. A. Heng, "Unsupervised bidirectional cross-modality adaptation via deeply synergistic image and feature alignment for medical image segmentation," *IEEE Transactions on Medical Imaging*, 2020.
- [9] C. Chen, Q. Dou, H. Chen, J. Qin, and P.-A. Heng, "Synergistic image and feature adaptation: Towards cross-modality domain adaptation for medical image segmentation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 865–872.
- [10] K. Bousmalis, N. Silberman, D. Dohan, D. Erhan, and D. Krishnan, "Unsupervised pixel-level domain adaptation with generative adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 3722–3731.
- [11] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb, "Learning from simulated and unsupervised images through adversarial training," *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2242–2251, 2017.
- [12] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [13] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1125–1134.
- [14] J. Solomon, F. De Goes, G. Peyré, M. Cuturi, A. Butscher, A. Nguyen, T. Du, and L. Guibas, "Convolutional wasserstein distances: Efficient optimal transportation on geometric domains," *ACM Transactions on Graphics (TOG)*, vol. 34, no. 4, pp. 1–11, 2015.
- [15] F. S. Saleh, M. S. Aliakbarian, M. Salzmann, L. Petersson, and J. M. Alvarez, "Effective use of synthetic data for urban scene semantic segmentation," in *Eur. Conf. Comput. Vis.* Munich, Germany: Springer, 2018, pp. 86–103.
- [16] E. Tzeng, J. Hoffman, T. Darrell, and K. Saenko, "Simultaneous deep transfer across domains and tasks," in *Proceedings of the IEEE International Conference on Computer Vision*, Dec. 2015, pp. 4068–4076.
- [17] X. Zhu, H. Zhou, C. Yang, J. Shi, and D. Lin, "Penalizing top performers: Conservative loss for semantic segmentation adaptation," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 568–583.
- [18] S. Sankaranarayanan, Y. Balaji, A. Jain, S. Nam Lim, and R. Chellappa, "Learning from synthetic data: Addressing domain shift for semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3752–3761.
- [19] D. Zou, Q. Zhu, and P. Yan, "Unsupervised domain adaptation with dual-scheme fusion network for medical image segmentation," in *Proceedings of the International Joint Conference on Artificial Intelligence*, 2020, pp. 3291–3298.
- [20] X. Zhuang and J. Shen, "Multi-scale patch and multi-modality atlases for whole heart segmentation of mri," *Medical image analysis*, vol. 31, pp. 77–87, 2016.
- [21] B. H. Menze, A. Jakab, S. Bauer, J. Kalpathy-Cramer, K. Farahani, J. Kirby, Y. Burren, N. Porz, J. Slotboom, R. Wiest *et al.*, "The multimodal brain tumor image segmentation benchmark (brats)," *IEEE transactions on medical imaging*, vol. 34, no. 10, pp. 1993–2024, 2014.
- [22] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A kernel two-sample test," *The Journal of Machine Learning Research*, vol. 13, no. 1, pp. 723–773, 2012.
- [23] M. Long, Y. Cao, J. Wang, and M. Jordan, "Learning transferable features with deep adaptation networks," in *International conference on machine learning*. PMLR, 2015, pp. 97–105.
- [24] K. Bousmalis, G. Trigeorgis, N. Silberman, D. Krishnan, and D. Erhan, "Domain separation networks," in *Advances in neural information processing systems*, 2016, pp. 343–351.
- [25] B. Sun and K. Saenko, "Deep coral: Correlation alignment for deep domain adaptation," in *European conference on computer vision*. Springer, 2016, pp. 443–450.
- [26] M. Mancini, L. Porzi, S. Rota Bulò, B. Caputo, and E. Ricci, "Boosting domain adaptation by discovering latent domains," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3771–3780.
- [27] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 2096–2030, 2016.
- [28] Y.-H. Chen, W.-Y. Chen, Y.-T. Chen, B.-C. Tsai, Y.-C. Frank Wang, and M. Sun, "No more discrimination: Cross city adaptation of road scene segmenters," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 1992–2001.
- [29] Y. Luo, L. Zheng, T. Guan, J. Yu, and Y. Yang, "Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2507–2516.
- [30] W. Hong, Z. Wang, M. Yang, and J. Yuan, "Conditional generative adversarial network for structured domain adaptation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 1335–1344.
- [31] A. Royer, K. Bousmalis, S. Gouws, F. Bertsch, I. Mosseri, F. Cole, and K. Murphy, "Xgan: Unsupervised image-to-image translation for many-to-many mappings," in *Domain Adaptation for Visual Understanding*. Springer, 2020, pp. 33–49.
- [32] Y. Huo, Z. Xu, H. Moon, S. Bao, A. Assad, T. K. Moyo, M. R. Savona, R. G. Abramson, and B. A. Landman, "Synseg-net: Synthetic segmentation without target modality ground truth," *IEEE transactions on medical imaging*, vol. 38, no. 4, pp. 1016–1025, 2018.
- [33] J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. Efros, and T. Darrell, "Cycada: Cycle-consistent adversarial domain adaptation," in *International conference on machine learning*, 2018, pp. 1989–1998.
- [34] Y. Zhang, Z. Qiu, T. Yao, D. Liu, and T. Mei, "Fully convolutional adaptation networks for semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 6810–6818.
- [35] Y. Li, L. Yuan, and N. Vasconcelos, "Bidirectional learning for domain adaptation of semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 6936–6945.
- [36] W.-L. Chang, H.-P. Wang, W.-H. Peng, and W.-C. Chiu, "All about structure: Adapting structural information across domains for boosting

- semantic segmentation,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 1900–1909.
- [37] Q. Dou, C. Ouyang, C. Chen, H. Chen, B. Glocker, X. Zhuang, and P.-A. Heng, “Pnp-adanet: Plug-and-play adversarial domain adaptation network at unpaired cross-modality cardiac segmentation,” *IEEE Access*, vol. 7, pp. 99 065–99 076, 2019.
- [38] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE Conference on Computer Vision and Pattern Recognition*. Ieee, 2009, pp. 248–255.