

Received May 30, 2019, accepted July 8, 2019, date of publication July 16, 2019, date of current version August 7, 2019.

Digital Object Identifier 10.1109/ACCESS.2019.2929258

PnP-AdaNet: Plug-and-Play Adversarial Domain Adaptation Network at Unpaired Cross-Modality Cardiac Segmentation

QI DOU^{1,2}, (Member, IEEE), CHENG OUYANG², CHENG CHEN¹, HAO CHEN¹, (Member, IEEE), BEN GLOCKER², XIAHAI ZHUANG⁴, AND PHENG-ANN HENG^{1,3}, (Senior Member, IEEE)

¹Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong

²Department of Computing, Imperial College London, London SW7 2AZ, U.K.

³T Stone Robotics Institute, The Chinese University of Hong Kong, Hong Kong

⁴School of Data Science, Fudan University, Shanghai 200433, China

Corresponding authors: Qi Dou (qi.dou@imperial.ac.uk) and Xiahai Zhuang (zxh@fudan.edu.cn)

This work was partially supported by HK RGC TRS project T42-409/18-R and in part by the CUHK T Stone Robotics Institute, CUHK, and in part by European Research Council (ERC) under the European Union's Horizon 2020 Research and Innovation Programme (grant agreement No 757173, project MIRA, ERC-2017-STG), and in part by Science and Technology Commission of Shanghai Municipality (17JC1401600).

ABSTRACT Deep convolutional networks have demonstrated state-of-the-art performance on various challenging medical image processing tasks. Leveraging images from different modalities for the same analysis task holds large clinical benefits. However, the generalization capability of deep networks on test data sampled from different distribution remains as a major challenge. In this paper, we propose a plug-and-play adversarial domain adaptation network (PnP-AdaNet) for adapting segmentation networks between different modalities of medical images, e.g., MRI and CT. We tackle the significant domain shift by aligning the feature spaces of source and target domains at multiple scales in an unsupervised manner. With the adversarial loss, we learn a domain adaptation module which flexibly replaces the early encoder layers of the source network, and the higher layers are shared between two domains. We validate our domain adaptation method on cardiac segmentation in unpaired MRI and CT, with four different anatomical structures. The average Dice achieved 63.9%, which is a significant recover from the complete failure (Dice score of 13.2%) if we directly test an MRI segmentation network on CT data. In addition, our proposed *PnP-AdaNet* outperforms many state-of-the-art unsupervised domain adaptation approaches on the same dataset. The experimental results with comprehensive ablation studies have demonstrated the excellent efficacy of our proposed method for unsupervised cross-modality domain adaptation. Our code is publically available at <https://github.com/carrenD/Medical-Cross-Modality-Domain-Adaptation>

INDEX TERMS Domain adaptation, adversarial learning, cardiac segmentation, medical imaging.

I. INTRODUCTION

Deep learning models, especially the convolutional neural networks (CNNs), have achieved remarkable successes during the past years, achieving state-of-the-art or even human-level performance on a variety of challenging medical imaging problems [1]–[3]. Typically, the deep networks are trained and tested on datasets where all the images are sampled from the same distribution. Despite the risk of over-fitting, the models are able to produce highly-accurate

predictions on new test data from the same domain. However, it has been frequently observed that established models under-perform when being tested on samples coming from a related but different new target domain [4]–[6]. For medical image computing in digital health field, the scenarios include the case that the test and training images come from different sites [7], [8] or different scanning protocols [9], [10] or even different imaging modalities [11], [12].

Different from natural images which are generally obtained by optical cameras, a typical situation in the medical field is the usage of various imaging modalities, capturing different physical properties. These different modalities play

The associate editor coordinating the review of this manuscript and approving it for publication was Yonghong Peng.




FIGURE 1. Illustration of severe domain shift existing in cross-modality images. The appearances of the cardiac structures (AA: Ascending aorta, LV-blood: Left ventricle blood cavity, LV-myo: Left ventricle myocardium) look significantly different on MRI and CT images, though segmentation masks look very similar.

complementary roles in clinical procedure of disease diagnosis and treatment. For example, Magnetic Resonance Imaging (MRI) and Computed Tomography (CT) have become indispensable tools for cardiac imaging. Specifically, MRI is ionizing radiation free and captures great contrast between soft tissues with high resolution in temporal space [13]. It features multi-parametric assessment of the myocardial contractivity and viability. By contrast, CT allows rapid imaging of the cardiac morphology, myocardial viability and coronary calcification, with great spatial resolution [14].

In practice, often the same image analysis task is required, such as segmentation or quantification of cardiac structures. Considering that annotation is prohibitively time-consuming and expensive (e.g., a whole heart segmentation from either MRI or CT takes up to 8 hours by a well-trained operator [15]), effectively adapting the model trained on one modality to the other holds clinical benefits. However, the appearances of cardiac MRI and CT are considerably different, with distinct contrasts and intensity histograms, as shown in Fig. 1. Unsupervised domain adaptation under such significant domain shift is very challenging, as well as yet to be explored.

Most methods for unsupervised domain adaptation focus on aligning the distributions in latent feature space, by minimizing measures of distances between the features extracted from the source and the target domain [16], [17]. For example, the Maximum Mean Discrepancy (MMD) is minimized together with a task-specific loss to learn domain-invariant and semantically-meaningful features in [18]. Long *et al.* [19] minimize MMD of domain features embedded in a reproducing kernel Hilbert space. Sun and Saenko [20] propose to align the feature covariances between domains. More recently, with the advancement of generative adversarial networks (GAN) [21] and its powerful extensions [22], [23], the latent feature spaces across domains can be implicitly aligned via adversarial learning. Notably, the DANN method is proposed to extract domain-invariant features by fully sharing weights of the CNN encoder between domains [17]. Tzeng *et al.* [16] introduce a more untied adversarial learning framework, named ADDA, where each domain has a dedicated domain-specific encoder before the last classification layer shared by both domains. Most recent studies further propose to train with auxiliary tasks on both domains, which serves as additional constraints

for feature alignment [24]. Alternatively, in terms of GAN based domain adaptation, another stream of solutions align input spaces of networks instead. They make use of unsupervised image-to-image translation, i.e., training the network with target-like synthetic source data, or testing with source-like target ones [12], [25], [26]. Most of them are based on a CycleGAN [23] foundation, where bi-directional image translations are learned by two GANs separately, and the consistency constraint between image transforms are enforced to preserve semantic information between transforms. Some other methods [27] combine feature alignment and image transformation together, however, the established framework rather sticks to improving target domain performance, and is not flexible between source and target domains.

For medical image computing, adversarial learning has presented inspiring efficacy on a wide variety of tasks [28]–[32]. In particular towards domain adaptation, early attempts have been made recently with the aim of generalizing the learned models to unseen target domains. For example, Zhang *et al.* [25] transformed unlabelled target X-Ray images to appear like those source radiographs and then directly tested the transformed images with the source model trained with labelled source data only. Similarly based on the CycleGAN, Jiang *et al.* [12] proposed a two-stage approach to first transform target CT images to resemble source MRI images, then conducted semi-supervised tumor segmentation with both synthetic image and a limited number of real MRI. Meanwhile, following the spirit of aligning latent feature spaces, there are a set of works aiming to extract domain-invariant representations from unpaired data. Degel *et al.* [33] minimized a segmentation loss in together with a domain adversarial training loss to encourage feature domain-invariance across ultrasound datasets for left atrium segmentation. Ren *et al.* [34] utilized adversarial learning to align the feature distribution of target images to the source domain for classifying histology images obtained in different staining procedures. These works have demonstrated that imposing alignment in feature space helps to generalize deep models to new data from a different domain. One of the most related work is Kamnitsas *et al.* [9], which conducted unsupervised domain adaptation by adversarial learning in multi-level feature space for brain lesion segmentation. The experimental setting is challenging as the unlabelled target domain contains a new MRI sequence which is unseen to the source model. By sharing encoders and aligning multi-level features, the method achieved promising results in target domain. However, designing flexible framework for unsupervised domain adaptation with minimal affect on the original source model, while mitigating drastic domain shift between modalities remains an open problem. Meanwhile, in terms of the highly challenging setting of cross-modality image segmentation from CT and MRI, to the best of our knowledge, related literature is limited. Valindria *et al.* [35] developed a joint learning method for multi-organ segmentation using unpaired MRI and CT. Zhang *et al.* [30] proposed

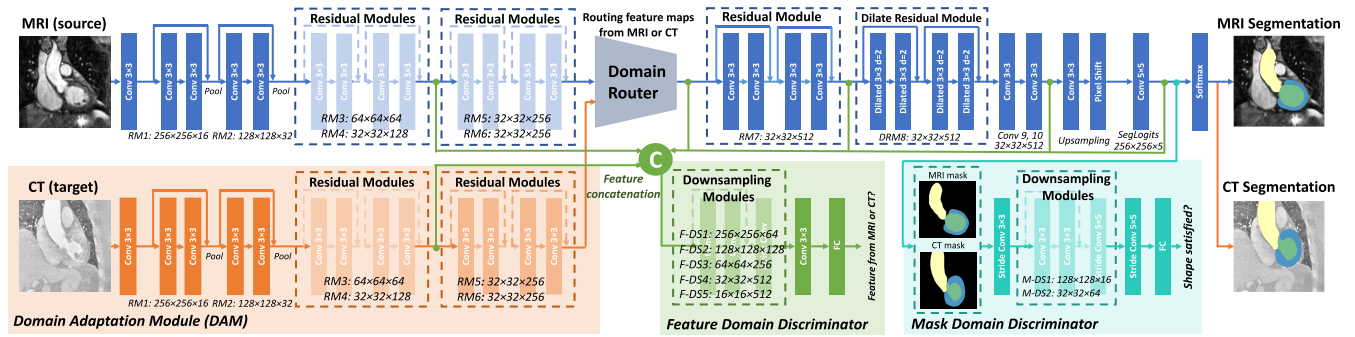


FIGURE 2. Overview of our proposed *PnP-AdaNet* (plug-and-play adversarial domain adaptation network), consisting of a source segmentation network, a domain adaptation module (DAM) and two discriminators. Multi-level activations and predicted segmentation masks are aggregated for alignment of the latent feature space. The domain router is for testing. It chooses which set of early layers to connect to higher layers for segmentation task. Specifically, when testing source data, it chooses to use original source early layers; when testing target data, it chooses to use DAM layers.

cross-modality image translation for improving cardiac segmentation with synthetic data. However, these works did not aim at our topic of unsupervised domain adaptation of CNNs, which is in principle much more difficult since annotation of target domain is completely unavailable.

In this paper, we study the challenging topic of unsupervised cross-modality domain adaptation on multi-class segmentation problem. We present a flexible plug-and-play adversarial domain adaptation network, called *PnP-AdaNet*, which effectively aligns the feature space of the target domain to that of the source domain. Specifically, the early encoders are replaced for target domain input, and higher layers are shared between domains. At adversarial learning, we build two domain discriminators, respectively connecting multi-level features and segmentation predictions for joint alignment of multi-level feature spaces and output spaces between domains. This paper is a substantial extension of our prior work [11], by further modifying the method with a significant performance boost, presenting a completely new series of ablation analysis of our proposed *PnP-AdaNet*, adding a comprehensive comparison with the state-of-the-art methods, and elaborating the literature review and discussion about unsupervised domain adaptation. Our main contributions are summarized as follows:

- We address the challenging yet crucial task of unsupervised cross-modality domain adaptation for medical image segmentation. A novel *PnP-AdaNet* is proposed to enable a flexible adaptation of the segmentation CNNs by plug-and-play feature encoders.
- Our model is learned with unpaired MRI and CT images via adversarial learning. To enhance supervision from discriminators, we aggregate multi-level features and segmentation mask predictions during training process.
- We extensively validate our method on multi-class cardiac segmentation with a public challenge data [36]. The mean Dice of four structures has been recovered from 13.2% to 63.9%, outperforming many state-of-the-art domain adaptation approaches. We also conduct comprehensive ablation studies on key method components. We release our code to facilitate the research community.

II. METHODS

Fig. 2 is an overview of our proposed *PnP-AdaNet* method. With a standard segmentation CNN learned on the source domain, we replace its early layers with a domain adaptation module while retain its higher layers, for testing on target domain data. Hence, we call our method as plug-and-play independent adaptation framework. The adaptation module maps the target images to the source domain in a latent feature space with aligned distribution. This process is trained with adversarial loss in an unsupervised way.

A. SEGMENTATION NETWORK WITHOUT SKIP CONNECTION

The essence of our proposed *PnP-AdaNet* is to establish an independent encoder for each domain and align their feature distributions in the latent space. Considering that only the early layers which compose the independent encoders are updated while those higher layers are fixed, the feature spaces at different layers need to be self-contained, i.e., not mixed-up with each other. This means that the network architectures using skip connections, e.g., the U-Net [37] and DenseNet [38], are not suitable choices. Otherwise, the plug-and-play setting would be problematic, because those domain-specific low-level features forwarded by skip-connections would affect the aligned high-level feature space (which is supposed to be shared across domains).

In this regard, we set up our segmentation model as a dilated network [39], which can extract representative features from large receptive fields, while also preserving the spatial acuity of feature maps. Residual connections inside local scopes are used for ease of gradients flow. Specifically, as illustrated in Fig. 2, the input image is firstly forwarded into a convolutional layer, then proceeded to a series of residual modules (termed as RM1-7, each consisting of the stacked 3×3 convolutions) and spatially down-sampled by a factor of eight. Next, to enlarge the receptive field for extracting the contextual information, 4 dilated convolutional layers are used in DRM8 with a dilation factor of 2. After another two convolutional layers Conv9 and Conv10, we conduct

upsampling to get dense predictions for the segmentation task. A 5×5 convolution operation is immediately followed to smooth out the activation maps. Finally, a softmax layer is utilized to obtain probabilistic pixel-wise predictions.

Formally, we denote the annotated dataset of the source domain by $X^s = \{(x_1^s, y_1^s), \dots, (x_i^s, y_i^s), \dots, (x_{N^s}^s, y_{N^s}^s)\}$, where x_i^s represents a sample pixel of the image and y_i^s is its category label of anatomical structure. We conduct supervised learning to establish a mapping M^s from the input image to the label space. The segmentation network of source domain is optimized by minimizing the hybrid loss \mathcal{L}_{seg} composed of multi-class cross-entropy loss and Dice coefficient loss [40]. Moreover, we denote $y_{i,c}^s$ as binary label regarding class $c \in C$ for sample x_i^s , its probability prediction is $\hat{p}_{i,c}^s$, and the label prediction is $\hat{y}_{i,c}^s$, the overall segmentation loss function is:

$$\mathcal{L}_{\text{seg}} = - \sum_{c \in C} \frac{\sum_{i=1}^{N^s} 2y_{i,c}^s \hat{y}_{i,c}^s}{\sum_{i=1}^{N^s} y_{i,c}^s y_{i,c}^s + \sum_{i=1}^{N^s} \hat{y}_{i,c}^s \hat{y}_{i,c}^s} - \lambda \sum_{i=1}^{N^s} \sum_{c \in C} w_c^s \cdot y_{i,c}^s \log(\hat{p}_{i,c}^s) + \beta \|W\|_2^2. \quad (1)$$

The first term is a Dice loss for multiple anatomical structures, and the second term is cross-entropy loss for individual pixels. The w_c^s is the weighting factor which copes with the issue of class imbalance. We combine the two complementary losses to tackle the challenging cardiac segmentation task. In practice, we also tried just using either one, but the performance was not as high as using both. The third term is the L2-regularization of the segmenter weights W . The λ and β are trade-off weights.

For the ease of notation, we will omit the subscript index i in following subsections, and directly use x^s and y^s to represent the sample and label from the source domain.

B. PLUG-AND-PLAY ADAPTATION MECHANISM

After obtaining the segmentation network trained on the source domain, we next aim to adapt it onto the target domain, in an unsupervised manner. In conventional transfer learning, it is common to update the last several layers of the pre-trained network towards a new given task with a new label space. The supporting assumption is that the early layers in the network extract low-level features which are universal for vision tasks. The higher layers are more task-specific and learn semantic-level features for conducting the defined predictions [41], [42]. In contrast, for domain adaptation, the defined task remains unchanged across domains. This means that the label space for source and target domains are identical, e.g., we segment the same anatomical structures from unpaired MRI/CT images. Not that images from different domains are not required to be co-registered. Basically, the distribution shift between the cross-modality domains are primarily low-level characteristics (e.g., gray-scale intensities) rather than high-level ones (e.g., geometric or semantic structures).

In these regards, for our model, we design a plug-and-play adaptation mechanism, i.e., a set of early layers being replaced, while the higher layers are reused for a new target domain. The underlying intuition is that, the higher layers are closely correlated with the shared semantic labels, while the respective early-layer encoders perform distribution mappings in feature space for our unsupervised domain adaptation. Formally, the obtained source segmentation model M^s is regarded as a layer-wise feature extractor composing stacked transformations of $\{M_{l_1}^s, \dots, M_{l_n}^s\}$, with the subscript denoting the network layer index. The prediction of the semantic label is represented by:

$$\hat{y}^s = M^s(x^s) = M_{l_1:l_n}^s(x^s) = M_{l_n}^s \circ \dots \circ M_{l_1}^s(x^s). \quad (2)$$

For input $x^t \in X^t$ from a target domain, we propose a domain adaptation module (DAM) denoted by \mathcal{M} that maps x^t to the feature space aligned with the source domain. We denote the adaptation depth by d , i.e., the layers earlier than and including l_d are replaced by DAM when processing the target domain images. At the same time, the source model's higher layers are frozen during domain adaptation learning and reused for target inference. Hence, the prediction for the target input is:

$$\hat{y}^t = M_{l_{d+1}:l_n}^s \circ \mathcal{M}(x^t) = M_{l_n}^s \circ \dots \circ M_{l_{d+1}}^s \circ \mathcal{M}(x^t), \quad (3)$$

where $\mathcal{M}(x^t) = \mathcal{M}_{l_1:l_d}(x^t) = \mathcal{M}_{l_d} \circ \dots \circ \mathcal{M}_{l_1}(x^t)$ represents the DAM which is a stack of convolutional layers as the feature encoder. In practice, we set the DAM layer configurations the same as the replaced set of early layers of source model, i.e., $\{M_{l_1}^s, \dots, M_{l_d}^s\}$. This is a reasonable and safe implementation choice, as we can initialize the DAM with a pre-trained source encoder rather than random initializations. This contributes to a more stable optimization at adversarial training, especially in our unsupervised learning scenario.

Overall, we can find that the proposed plug-and-play domain adaptation mechanism is elegant and rather flexible at testing. During inference for the target domain, the DAM directly replaces the early d layers of the source domain network. The images of the target domain are processed and mapped to the feature space of source domain using this DAM. These adapted features are robust to the cross-modality domain shift, and can be correctly transformed into the label space via the fixed high-level layers. It is also necessary to mention that our plug-and-play domain adaptation procedure does not hurt the performance for the source domain. The early layer encoding path for the source domain is preserved and the higher layers are unaffected at learning. Therefore, our *PnP-AdaNet* can be flexibly tested for both target and source domain data, just by selecting the input path.

C. ADVERSARIAL LEARNING FOR FEATURE SPACE ALIGNMENT

We train our plug-and-play domain adaptation network via adversarial learning in an unsupervised manner. In the spirit

of GAN, a generator and a discriminator form a minimax two-player game. The generator aims to capture the distribution of the real data, while the discriminator should identify whether a presented sample comes from the real or learned distributions. In our *PnP-AdaNet*, the DAM serves as the generator which maps the input target image into the latent feature space of the source domain. The aim of the domain adaptation module is to encode representations which are aligned with those encoded from the source domain images. Hence, the fixed layers in the higher part of source network can be reused to make semantic-level predictions of segmentation masks. As we have zero annotation for the target domain, the adaptation process is implicitly supervised by the discriminators, i.e., forming an adversarial learning game.

In our framework, we propose to employ two discriminators. Specifically, the input to the first discriminator (i.e., the green part in Fig. 2) is an array of aggregated feature maps of the segmenter. This input has a high dimension with a relatively complicated distribution. A natural thinking here is that, we connect the output features of the DAM into the discriminator. However, the convolutional neural network has a hierarchical architecture, and the features at one certain layer rely on activations from its previous layers, and also, the features are proceeded to affect following layers. If we just monitor the encoded features immediately obtained from the DAM, the latent space alignment can be unstable. In other words, we have no idea whether those activations in the layers earlier to the adaptation depth are aligned. Also, the small shift that may still exist at the adaptation layer l_d would be magnified after composing them in higher layers. To overcome this problem, we aggregate the activations from multiple layers as input to the discriminator. The gradients from the discriminator can flow to the DAM via multiple paths, so that the feature space alignment can be more tightly supervised. This way of learning shares the spirit of deep supervision [43] to some extent.

In practice, we aggregate the activations from multiple levels of layers, and reshape them to the same resolution for channel concatenation. Formally, we refer to the feature maps in the selected frozen layers as the set of $F_H(\cdot)$ where $H = \{k, \dots, q\}$ being the set of selected layer indices. Similarly, we denote the selected feature maps of DAM by $\mathcal{M}_A(\cdot)$ with A being the selected layer set. In this way, the feature space of the target domain is $(\mathcal{M}_A(x^t), F_H(x^t))$ and the array $(M_A^s(x^s), F_H(x^s))$ is their counterpart for the source domain. Given the distribution of $(M_A^s(x^s), F_H(x^s)) \sim \mathbb{P}_{\text{feature}}^s$ and that of $(\mathcal{M}_A(x^t), F_H(x^t)) \sim \mathbb{P}_{\text{feature}}^t$, the distance between these two domain distributions which is to be minimized is represented as $W(\mathbb{P}_{\text{feature}}^s, \mathbb{P}_{\text{feature}}^t)$. For stabilized training, we use Wassertein distance [22] of the two distributions as:

$$W(\mathbb{P}_{\text{feature}}^s, \mathbb{P}_{\text{feature}}^t) = \inf_{\gamma \sim \prod(\mathbb{P}_{\text{feature}}^s, \mathbb{P}_{\text{feature}}^t)} \mathbb{E}_{(x,y) \sim \gamma} [\|x - y\|], \quad (4)$$

where $\prod(\mathbb{P}_{\text{feature}}^s, \mathbb{P}_{\text{feature}}^t)$ is the set of all joint distributions of $\gamma(x, y)$ whose marginals are respectively $\mathbb{P}_{\text{feature}}^s$ and $\mathbb{P}_{\text{feature}}^t$.

Aligning the latent feature space by directly inputting the high-dimensional activations to a discriminator is effective and essential. This might be fine for classification tasks, but may be sub-optimal for segmentation tasks which require pixel-wise predictions with fine structures. This implies that well-aligned output spaces, in our case segmentation mask, between source and target are also critical. Early studies using GANs for segmentation applications (not necessarily under the domain adaptation setting) commonly input the predicted segmentation masks to the discriminator. When the shape or structure of the predicted segmentation mask looks distorted (i.e., not looking like the real mask), the discriminator would impose a penalty. For the specific problem of domain adaptation at segmentation, we also consider that monitoring the shape of the predicted segmentation mask is important. This serves as a correction mechanism for imperfections for feature alignment in adversarial training.

To this end, we further include an auxiliary discriminator in our *PnP-AdaNet*, whose inputs are the predicted segmentation masks of the source and target domains. In this case, the input is more compact and the discriminator focus purely on outputs, compared with those of the first discriminator. We denote the segmentation predictions for the target and source domains by $\mathcal{S}(x^s) \sim \mathbb{P}_{\text{mask}}^s$ and $\mathcal{S}(x^t) \sim \mathbb{P}_{\text{mask}}^t$. Following the Eq. (4), we also employ the Wassertein distance between source and target distributions:

$$W(\mathbb{P}_{\text{mask}}^s, \mathbb{P}_{\text{mask}}^t) = \inf_{\gamma \sim \prod(\mathbb{P}_{\text{mask}}^s, \mathbb{P}_{\text{mask}}^t)} \mathbb{E}_{(x,y) \sim \gamma} [\|x - y\|]. \quad (5)$$

The detailed network architectures of the discriminators are illustrated in Fig. 2. For the model configuration, the feature discriminator is relatively deeper than the mask discriminator.

D. LOSS FUNCTIONS AND TRAINING STRATEGIES

In adversarial learning, the DAM is pitted against an adversary with the above two discriminators. We represent the first discriminator with \mathcal{D}_f which takes high-dimensional features as input, and denote the second discriminator by \mathcal{D}_m which takes compact predicted segmentation masks as input. The pair of \mathcal{D}_f and \mathcal{D}_m implicitly estimate the $W(\mathbb{P}_{\text{feature}}^s, \mathbb{P}_{\text{feature}}^t)$ and the $W(\mathbb{P}_{\text{mask}}^s, \mathbb{P}_{\text{mask}}^t)$, respectively. During the learning process, the discriminators would try to differentiate the inputs from the source and target domains. The domain adaptation with DAM aims to not only remove the domain-specific patterns in early layers, but also disallow their recovery at higher semantic-level layers. The generator \mathcal{M} (DAM) and the discriminators $\{\mathcal{D}_f, \mathcal{D}_s\}$ are jointly optimized via adversarial loss functions. Specifically, the loss for generator DAM is:

$$\mathcal{L}_{\mathcal{M}} = -\mathbb{E}_{(\mathcal{M}_A(x^t), F_H(x^t)) \sim \mathbb{P}_{\text{feature}}^t} [\mathcal{D}_f(\mathcal{M}_A(x^t), F_H(x^t))] - \mathbb{E}_{\mathcal{S}(x^t) \sim \mathbb{P}_{\text{mask}}^t} [\mathcal{D}_m(\mathcal{S}(x^t))]. \quad (6)$$

The losses for learning the discriminators \mathcal{D}_f and \mathcal{D}_m are:

$$\begin{aligned} \mathcal{L}_{\mathcal{D}_f} = & \mathbb{E}_{(\mathcal{M}_A(x^t), F_H(x^t)) \sim \mathbb{P}_{\text{feature}}^t} [\mathcal{D}_f(\mathcal{M}_A(x^t), F_H(x^t))] \\ & - \mathbb{E}_{(\mathcal{M}_A^s(x^s), F_H(x^s)) \sim \mathbb{P}_{\text{feature}}^s} [\mathcal{D}_f(\mathcal{M}_A^s(x^s), F_H(x^s))], \\ \text{s.t. } & \|\mathcal{D}_f\|_L \leq K, \end{aligned} \quad (7)$$

$$\begin{aligned} \mathcal{L}_{\mathcal{D}_m} = & \mathbb{E}_{\mathcal{S}(x^t) \sim \mathbb{P}_{\text{mask}}^t} [\mathcal{D}_m(\mathcal{S}(x^t))] - \mathbb{E}_{\mathcal{S}(x^s) \sim \mathbb{P}_{\text{mask}}^s} [\mathcal{D}_m(\mathcal{S}(x^s))], \\ \text{-s.t. } & \|\mathcal{D}_m\|_L \leq K, \end{aligned} \quad (8)$$

where K is the Lipschitz constraint of \mathcal{D}_f , \mathcal{D}_m . By alternately updating the generators and discriminators, the DAM becomes more effective to generate source-like features from target data for domain adaptation.

In practice, we first train the segmentation network on the source domain in a supervised manner with standard stochastic gradient descent. We employ the Adam optimizer with a batch size of 10 and a learning rate of 1×10^{-3} . After obtaining the segmenter, we train $\{\mathcal{M}, \mathcal{D}_f, \mathcal{D}_m\}$ with above adversarial loss for unsupervised domain adaptation. To begin with, we solely update the discriminators for 20k iterations with a batch size of 6 as a pre-training process. Next, we alternately update the generator and discriminators. Following the heuristic rules of training WGAN [22], we update the generator \mathcal{M} for one iteration, every 20 iterations updating the \mathcal{D}_f and \mathcal{D}_m . In adversarial learning, we use the RMSProp optimizer with a learning rate of 3×10^{-4} , and a stepped decay rate of 0.98 every 100 joint updates. Weight clipping for discriminator weights is 0.03. The discriminator and generator losses are scaled with a factor of 0.002. Dropout (rate = 0.25) and batch normalization are used in all the convolutional layers. We use Leaky-ReLU as the activation function and strided convolutions for down-sampling inside the discriminators.

III. DATASET AND EVALUATION METRICS

A. DATASET FOR CROSS-MODALITY ADAPTATION

Publicly available medical datasets which contain different modalities of images for the same anatomical structure are rare, hindering progress of investigating the challenging task of MRI and CT cross-modality domain adaptation. Fortunately, the challenge of *MICCAI 2017 Multi-Modality Whole Heart Segmentation (MM-WHS)* presented 20 MRI and 20 CT cardiac images with accurate manual segmentation annotations. The images are unpaired with the MRI data and CT data coming from different patients and different sites. We refer the readers to the original data description paper of Zhuang et al. [36] for more details about data acquisition such as the employed scanning protocols. For evaluating the domain adaptation on cardiac segmentation, we include following four structures: ascending aorta (AA), the left atrium blood cavity (LA-blood), the left ventricle blood cavity (LV-blood), and the myocardium of the left ventricle (LV-myo). We randomly split each modality of the data into training (16 subjects) and testing (4 subjects) subsets in experiments.

Our PnP-AdaNet is designed for unpaired cross-modality medical image segmentation, the MRI and CT images are not

co-registered in our experiments, as it is not necessary. In pre-processing, the MRI and CT images are reoriented (in view direction), resized and cropped centering at the heart region, such that the view of multi-modal images are roughly on the same page. We extract MRI and CT scans by 2D slices of size 256×256 at coronal plane during training, and obtain ~ 100 -300 2D slices for each MRI or CT. The model takes three adjacent slices as input, for using contextual information. The mask label of the middle slice is used as the ground truth. We conduct the intensity normalization in 3D for each modality, respectively. Augmentations of rotation, zooming and affine transformations are utilized to enlarge the training database.

B. EVALUATION METRICS ON SEGMENTATION PERFORMANCE

For the evaluation metrics, we follow the common practice to quantitatively evaluate segmentation methods [44]. The Dice coefficient ([%]) is used to assess the agreement between the predicted segmentation and ground truth for cardiac structures. We use the symmetric average surface distance (ASD)[voxel] to measure the segmentation performance from the perspective of boundary agreement. A higher Dice and a lower ASD indicate a better segmentation performance. Both metrics are presented in the format of $mean \pm std$, which shows the average performance with the cross-subject variations of the results. For some results, the N/A on the ASD means that at least one subject did not receive any correct prediction on the structure.

IV. EXPERIMENTS

A. EXPERIMENTAL SETTINGS

In our experiments, we first set the source domain as MRI and target domain as CT. We conducted extensive experiments to demonstrate the severe cross-modality domain shift and the effectiveness of domain adaptation strategies. Specifically, we designed the following experimental configurations:

- 1) training and testing the segmentation network on the source domain (i.e., *Seg-MRI*);
- 2) training and testing the segmentation network on annotated target domain images, as an upper bound (i.e., *Seg-CT*);
- 3) directly testing the source domain segmenter on target data, with no domain adaptation (i.e., *Seg-CT-noDA*);
- 4) our *PnP-AdaNet* for unsupervised domain adaptation.

In addition, we optimized the practical configurations of our *PnP-AdaNet*. Specifically, our ablation studies investigated: i) how the balancing ratio in loss functions between the two discriminators affects the domain adaptation performance, and ii) the importance of inputting multiple levels of feature maps to \mathcal{D}_f . Moreover, we implemented a series of existing popular unsupervised domain adaptation methods on our dataset, and for comparison on the task of unpaired cross-modality MRI to CT domain adaptation. Last but not least, we also report the segmentation results by reverting the adaptation direction, i.e., adapting from CT to MRI.




FIGURE 3. Results of different methods for CT image segmentations. Each row presents one typical example, from left to right: (a) raw CT images (b) ground truth labels (c) supervised training on CT (d) directly applying MRI segmenter on CT data (e) results of DANN (f) results of ADDA (g) results of CycleGAN (h) results of our proposed *PnP-AdaNet*. The structures of AA, LA-blood, LV-blood and LV-myocardium are indicated by yellow, red, green and blue colors, respectively.

TABLE 1. Quantitative evaluation with Dice of segmentation results on cardiac structures of different experimental configurations.

Experiment configurations	AA	LA-blood	LV-blood	LV-myocardium
Cascaded-FCN-MRI [45]	76.6±13.8	81.1±13.8	87.7±7.7	75.2±12.1
U-Net-MRI [37]	70.6±20.2	74.0±24.7	81.1±23.8	68.1±25.3
Cascaded-FCN-CT [45]	91.1±18.4	92.4±3.6	92.4±3.3	87.2±3.9
U-Net-CT [37]	94.0±6.2	91.0±5.2	91.0±4.3	86.1±4.2
Seg-MRI	80.3±1.8	78.1±8.0	88.3±2.1	70.8±2.1
Seg-CT	77.6±29.2	89.4±2.2	91.3±2.5	85.8±1.5
Seg-CT-noDA	31.5±23.9	2.7±0.8	3.4±5.8	15.3±17.2
Proposed <i>PnP-AdaNet</i>	74.0±7.3	68.9±5.2	61.9±10.7	50.8±7.0

B. RESULTS OF PnP-AdaNet FOR UNSUPERVISED DOMAIN ADAPTATION

Our employed segmentation network without skip connection serves as the basis for our subsequent domain adaptation procedures. Hence, we choose to first validate the performance of this dilated network architecture on the applied whole-heart segmentation task. Specifically, for *Seg-MRI* setting, our model achieves an average Dice of 79.4% across four structures. To see the comparison, we reference Payer et al. [45] which used two cascaded fully convolutional networks and achieved an average Dice of 80.2%, ranking the first in *MICCAI 2017 MM-WHS Challenge*. We also directly quote their reported U-Net results as a complementary comparison. As the detailed results listed in Table 1, our network's segmentation performance with standard training is comparable to the state-of-the-arts, which employed networks based on 3D convolutional kernels. Hence, we can safely regard our trained segmentation network as a standard baseline model for the cardiac segmentation task.

As for observing the severe domain shift inherent in cross-modality medical images, we first directly deploy the segmentation model trained on MRI domain to CT data. Unsurprisingly, the MRI segmenter completely fails on CT

images, with an average Dice of merely 13.2% across all the four structures. Specifically, the *Seg-CT-noDA* only receives a Dice of 2.7% for LA-blood and 3.4% for LV-blood. After domain adaptation, our *PnP-AdaNet* presents a great recovery of the segmentation performance on target CT data compared with *Seg-CT-noDA*. More specifically, our method has increased the average Dice across the four cardiac structures by 50.7%, achieving a score of 63.9%. As presented in the last column of Fig. 3, the predicted segmentation masks from *PnP-AdaNet* can successfully localize the cardiac structures and further capture their anatomical shapes. Notably, the segmentation performance on aorta has been significantly recovered after the adaptation DANN process, almost approaching the fully supervised upper bound.

C. ABLATION STUDY ON CONFIGURATIONS

In this subsection, we extensively investigate the configurations of our *PnP-AdaNet*. Specifically, we observe the domain adaptation performance by adjusting two key properties: i) the ratio balancing for losses of these two discriminators \mathcal{D}_m and \mathcal{D}_f , ii) the selected level of layers to input to the feature discriminator. We list the results of our experimented configurations in Table 2.

Specifically, we adjust the ratio of $\mathcal{D}_m/\mathcal{D}_f$ from 0.1 : 1 to 0.7 : 1, with a step of 0.1. As we gradually increase this ratio, we are actually emphasizing more on the shape regularization from the mask discriminator. From Table 2, the first row is solely using \mathcal{D}_f for only feature-level alignment without regularization on the output map, reaching an average segmentation Dice of 59.6%. Comparing the first row with the second row (rate = 0.1), we find that adding the mask discriminator \mathcal{D}_m as an auxiliary shape constrain is effective, which helps to directly suppress those noises or shape distortions in the outputs. Generally, the *PnP-AdaNet* is

TABLE 2. Domain adaptation results of ablation studies adjusting configurations of the proposed PnP-AdaNet.

Methods	AA		LA-blood		LV-blood		LV-myocardium		Mean	
	Dice	ASD	Dice	ASD	Dice	ASD	Dice	ASD	Dice	ASD
solely use \mathcal{D}_f	67.1±11.5	13.6±5.8	65.2±15.1	9.3±4.3	58.1±8.5	15.9±7.6	48.0±4.7	15.0±5.1	59.6±10.0	13.5±5.7
$\mathcal{D}_m/\mathcal{D}_f$ rate = 0.1	74.0±7.3	12.8±3.2	68.9±5.2	6.3±2.3	61.9±10.7	17.4±7.0	50.8±7.0	14.7±4.8	63.9±7.5	12.8±4.3
$\mathcal{D}_m/\mathcal{D}_f$ rate = 0.2	71.1±8.8	14.2±5.5	67.5±17.3	7.2±3.5	61.3±12.2	18.0±7.5	47.3±9.5	17.6±5.8	61.8±11.9	14.2±5.6
$\mathcal{D}_m/\mathcal{D}_f$ rate = 0.3	72.0±10.7	14.8±4.2	65.7±16.0	7.0±3.0	62.4±13.5	15.3±7.2	49.2±6.1	13.4±4.7	62.3±11.6	12.6±4.8
$\mathcal{D}_m/\mathcal{D}_f$ rate = 0.4	73.3±6.2	12.5±6.0	68.2±21.4	7.2±4.6	60.9±10.7	18.2±7.4	52.6±10.1	14.5±3.7	63.8±12.1	13.1±5.4
$\mathcal{D}_m/\mathcal{D}_f$ rate = 0.5	73.1±7.4	13.9±5.5	56.9±28.6	8.3±3.5	50.0±13.4	16.8±5.5	43.9±8.7	13.9±2.6	56.0±14.5	13.2±4.3
$\mathcal{D}_m/\mathcal{D}_f$ rate = 0.6	75.9±5.3	13.2±5.9	55.1±30.8	7.6±4.8	48.7±5.0	20.2±7.3	43.7±7.3	17.9±6.0	55.9±12.1	14.7±6.0
$\mathcal{D}_m/\mathcal{D}_f$ rate = 0.7	63.2±6.4	17.5±5.9	62.2±15.2	9.2±4.1	56.5±11.1	17.3±6.5	44.2±2.9	17.2±4.9	56.5±8.9	15.3±5.3
– low-level layer	72.1±5.5	12.6±4.2	47.8±32.5	9.9±6.3	50.5±12.6	16.0±4.9	48.7±8.6	15.0±3.3	54.8±14.8	13.4±4.7
– transaction layer	73.3±6.1	12.9±4.0	58.2±17.7	8.8±1.8	51.4±14.9	19.9±8.2	43.0±1.4	16.1±5.8	56.5±10.0	14.4±4.9
– high-level layer	70.4±7.5	13.5±3.2	69.4±10.6	6.6±1.8	53.6±14.2	21.8±8.0	43.0±5.0	18.2±6.9	59.1±9.3	15.0±5.0

not very sensitive to this ratio in a range of small-scale values. When adjusting the ratio from 0.1 to 0.4, the performance on segmenting CT images can stably achieve a Dice of over 60%. The change of the Dice score is not monotonic as we adjust the ratio. Meanwhile, it also shows that the ratio could not be set too high, otherwise, the segmentation accuracy on the target domain would decrease as observed in our experiments. The reason could be that straight-forward supervision of the mask is rather high-level. The mask quality also heavily relies on the achieved domain-invariance from earlier features. The two discriminators might also compete with each other when back-propagating their gradients to the generator (DAM), especially at the very beginning of training. For the sake of a high adaptation performance, monitoring the predicted masks in the compact semantic space should serve as an auxiliary supporting role.

As already observed that aligning the feature space is essential in our method, we furthermore investigate which layers to input their features to \mathcal{D}_f . Intuitively, aligning shallow layers encourages domain-invariance of low-level features which are the basis for all higher-level representations. The layer at the adaptation depth is obviously crucial as it is the dividing point in our plug-and-play setting. Connecting the discriminator with deeper layers, although they are not updated in the adaptation process, can help to retain the alignment and prevent recovery of domain-specific information within the high-level space. To observe the effectiveness of multi-level feature supervision, we conducted experiments by removing a low-level layer (RM4), the transaction layer (RM6) and a high-level layer (Conv10) from feature discriminator \mathcal{D}_f . From Table 2, we can observe that missing either level of the feature maps would reduce the performance by a noticeable margin. In practice, we also find that connecting the compact activations immediately before the softmax layer to \mathcal{D}_f is crucial. If we remove it from the feature discriminator input, a severe drop of performance will be seen. This observation again confirms the importance of monitoring the high-level compact spaces with strong semantic relevance.

D. COMPARISON WITH STATE-OF-THE-ART METHODS

As our investigated topic of unsupervised domain adaptation on cross-modality medical image segmentation is quite new in the field, there are few previous studies that could

be compared with. However, this problem deserves further careful explorations, given that CNNs are dominating current segmentation methods, and their generalization capability considerably matters. Different from the cross-site and cross-sequence domain shifts, we regard cross-modality shift between CT and MRI as one of the most challenging settings.

To promote and facilitate future studies on cross-modality domain adaptation, in addition to our own prior work of [11], we also implemented several state-of-the-art methods which are popular for unsupervised domain adaptation in natural computer vision. We used the same dataset and segmenter architecture settings for all the methods, with results listed in Table 3. Specifically, the approach of DANN [17] encourages domain-invariance in feature space. The source and target domains share the feature extractor. A domain classifier is connected to the output of the encoder to monitor domain-invariance. Another alternative approach is ADDA [16] which also aligns source and target domain distributions in the feature space. In this method, the source and target domains have their own feature encoders until the last softmax layer. A discriminator is used to differentiate which features come from which domain. A third comparison method we include in the comparison is CycleGAN [23], which produces impressive image-to-image transformations. We transform the MRI images to the appearance of CT, and train a segmenter with the transformed images and the MRI labels. We demonstrate typical examples of the generated CT images from MRI using CycleGAN in Fig. 4.


Observing performances of the domain adaptation methods, we can see that all of them are able to effectively recover the segmentation accuracy on target CT data. More specifically, in terms of different cardiac structures, all the methods perform better for ascending aorta and left atrium blood cavity than on the structures of left ventricle blood cavity and myocardium. The reason could be that these anatomical structures, especially the myocardium, have relatively complicated geometry, which increases the difficulty for unsupervised domain adaptation. In contrast, the aorta presents a more compact shape as well as clear boundaries, and reasonably, receiving the best recovery. Compared with DANN and ADDA which also used feature-level adaptation, our proposed *PnP-AdaNet* achieves better performance. The reason is that we neither share nor separate all the

TABLE 3. Comparison of different approaches on unpaired MRI to CT cross-modality domain adaptation of cardiac structure segmentation.

Methods	AA		LA-blood		LV-blood		LV-myocardium		Mean	
	Dice	ASD	Dice	ASD	Dice	ASD	Dice	ASD	Dice	ASD
Seg-CT-noDA	31.5±23.9	21.4±14.1	2.7±0.8	19.6±4.6	3.4±5.8	N/A	15.3±17.2	20.7±7.7	13.2±11.9	N/A
DANN [17]	39.0±35.1	16.2±5.8	45.1±23.6	9.2±2.9	28.3±11.8	12.1±1.6	25.7±13.2	10.1±3.2	34.5±20.9	11.9±3.4
ADDA [16]	47.6±15.2	13.8±3.0	60.9±13.2	10.2±6.5	11.2±13.1	N/A	29.2±16.4	13.4±5.0	37.2±14.5	N/A
CycleGAN [23]	73.8±7.4	11.5±2.9	75.7±4.3	13.6±3.6	52.3±21.0	9.2±3.9	28.7±13.3	8.8±4.3	57.6±11.5	10.8±3.7
Prior PnP-AdaNet [11]	74.8±6.2	27.5±7.6	51.1±11.2	20.1±4.5	57.2±12.4	29.5±11.7	47.8±5.8	31.2±10.1	57.7±8.9	27.1±8.5
Proposed PnP-AdaNet	74.0±7.3	12.8±3.2	68.9±5.2	6.3±2.3	61.9±10.7	17.4±7.0	50.8±7.0	14.7±4.8	63.9±7.5	12.8±4.3

TABLE 4. Results of reverting the domain adaptation direction, i.e., adapting the segmentation network from CT to MRI.

Experimental settings	AA		LA-blood		LV-blood		LV-myocardium		Mean	
	Dice	ASD	Dice	ASD	Dice	ASD	Dice	ASD	Dice	ASD
Seg-MRI-noDA	31.5±23.9	21.4±14.1	2.7±0.8	19.6±4.6	3.4±5.8	N/A	15.3±17.2	20.7±7.7	13.2±11.9	N/A
Proposed PnP-AdaNet	43.7±10.8	11.4±3.2	47.0±7.3	14.5±4.1	77.7±10.4	4.5±1.4	48.6±2.9	5.3±1.8	54.3±7.9	8.9±2.6

**FIGURE 4.** Examples of MRI to CT image-to-image translations with CycleGAN. Left to right: Original MRI image, generated CT image, MRI with segmentation ground truth, and generated CT with corresponding MRI ground truth.

feature encoding layers between source and target domains. Instead, our plug-and-play mechanism regards low-level features as domain-specific, while high-level feature encoders as sharable. Moreover, connecting multi-level features to the discriminator also plays a crucial role. The CycleGAN presents highly competitive results. We analyze that this can attribute to the pixel-wise supervision, particularly towards the segmentation task. The explicit supervision, although it maybe noisy, can encourage the predictions to consistently focus on the heart region. The good ASD results with CycleGAN also indicate clean masks. The deficiency of our feature-aligning method mainly appears at those unclear boundaries between neighboring structures, or wrong predictions on relatively homogeneous tissues but away from the ROI. With a very simple post-processing strategy, i.e., only remaining the largest 3D connected component for every class, we can reduce the ASD to 4.1, 5.4, 7.4 and 6.2 for the AA, LA-blood, LV-blood and LV-myocardium, respectively. Finally, compared with our own prior work of [11], the performance improvement comes from the usage of dual discriminators, with a more explicit constrain on the shape of segmentation.

E. REVERTING DOMAIN ADAPTATION DIRECTION

A natural question to ask is whether domain adaptation for the segmentation task is *symmetric to modality*, i.e. whether the reverse adaptation direction from CT to MRI can also

be achieved. To investigate this, we apply the same model setting, but only replacing the source domain as CT and target domain as MRI. The quantitative adaptation results are presented in Table 4. Directly using the CT segmenter on MRI data also unsurprisingly fails. Our proposed PnP-AdaNet is able to recover the average segmentation Dice to 54.3%. Notably, the best recovered structure in this reverse setting is the LV-blood, not the same structure (AA) as the direction of MRI to CT. As a more interesting observation, the Dice increases from complete failure (3.4%) to a considerably high value (77.7%). Results of the other three structures are not as promising. Compared with adaptation from MRI to CT, the reverse direction yields lower performance, which can be expected. Segmenting cardiac MRI itself is more difficult than segmenting cardiac CT. This is also notable from Table 1, where the CT segmentation Dice is higher than the result of MRI across all four structures. In these regards, transferring a CT segmenter to MRI seems more challenging. Our experimental results indicate that cross-modality domain adaptation can be achieved in both directions, however, the difficulty seems to be asymmetric.

V. DISCUSSIONS

In this paper, we try to tackle the domain adaptation problem under a very challenging as well as important setting of cross-modality medical datasets. As deep learning has become the de-facto standard for solving segmentation and detection tasks, investigating its generalization capability and robustness is essential. Existing successful practice using deep networks is to train and test the models with the same data source. However, it has been frequently revealed in very recent works, that the models would perform poorly on unseen datasets [7], [10], [26]. Resolving the domain adaptation issue holds great potentials for, applying trained deep learning models to wider clinical use, building more powerful networks using large-scale database combing images from multiple sites, and helping to understand how the networks capture the data distributions to make recognition predictions. In these regards, we explore effective unsupervised domain adaptation solution, under one of the most challenging

settings as the *multi-class segmentation of cross-modality medical images*.

The appearance differences between MRI and CT scans are apparent. Although the human eye could match the structures on the two modalities, from a perspective of image computing, the system “sees” distinct distributions of intensity values. For concrete example, see Fig. 1 again, the intensity range of myocardium and its contrast with the nearby tissues is very different on MRI and CT. In this case, the model that learns discriminative features from one modality is naturally inapplicable to the data manifold of another modality. Generally speaking, it is rather difficult for a model to generalize from one domain to another, where the measure of their support intersection is almost zero. On the other hand, multi-modality data have become indispensable tools in modern clinical routine. Cross-modality synthesis and segmentation has been gaining research popularity rapidly [46], [47]. Considering cardiology as currently the most typical scenario which uses both MRI and CT, we conduct our study on cardiac images.

Towards domain adaptation, we present a novel framework, i.e., *PnP-AdaNet*, which is very flexible and lightweight in practical use. We replace the early layers of an established network with a DAM at the testing process. The assumption is that, in cross-modality medical images for the same organ, those low-level features are domain-specific, while the semantic feature compositions at higher layers are shared and could be re-used across different modalities. To learn the DAM, we encourage the extracted features of the target domain to be aligned with the source feature in distributions. A notable setting in our *PnP-AdaNet* is incorporating multi-level adversarial learning, which makes the DAM tightly conditioned by the source distributions. Another hyper-parameter, which needs to be considered when using the plug-and-play strategy, is the adaptation depth d . In our prior work Dou *et al.* [11], we have experimentally investigated the influence of adjusting the choice of adaptation depth. Intuitively, a smaller d (i.e., shallower DAM with fewer parameters) might be less capable of learning effective feature mappings across domains. On the other hand, training a much deeper DAM solely with adversarial gradients would be difficult and unstable. A recent work studying a generalization upper-bound for unsupervised domain adaptation also indicated similar insights [48]. The paper derived that minimizing the generalization risk across domains requires a function with particular characteristics. It needs to be capable enough for reducing the Wasserstein distance to a desired degree for maintaining adaptation quality, meanwhile, also should control the complexity of reducing the generalization risk. Our experimental findings in [11] matched these insights, i.e., plugging the first six residual blocks yielded the highest performance. Hence, we suggest to select a middle-level layer as the adaptation depth, when embedding our plug-and-play strategy into other similar applications.

When thinking of solutions for domain adaptation, an alternative way is to take advantages of image-to-image

translation. The networks using cycle-consistency loss can generate plausible synthetic images, which then can either be employed for training or testing. The potential risk of pixel-space translation would be that the distortions and artifacts might be propagated or amplified in downstream processing. For example, Cohen *et al.* [49] found that CycleGAN based medical image translations models trained on imbalanced datasets would hide dangerous brain tumors in synthetic images. From our own experimental experiences, the cross-modality translation of non-tumor medical images is feasible using CycleGAN, and the reconstructed images also look quite realistic. It can bring positive effects when using synthetic images towards the scenario of domain adaptation, however, it cannot solve the problem perfectly. We analyze that the reasons are at least two-folds. One is that there still exists domain shift between the synthetic images and real images, although their appearances somehow look similar [30]. The other is that the generated images may on a pixel-level not match with the original image. Therefore, the training pairs of image and label are imperfect. In contrast, feature-level adaptations, such as ours, do not suffer from these problems, as the adaptation and task-specific training are not detached.

Although we achieved promising results, there are still limitations. The backbone of our network is a 2D CNN. In practice, we aggregate the adjacent three slices as the input channels to the models. This is common practice to overcome current single GPU memory constraints. It maybe beneficial to employ more carefully tailored networks for CT/MRI image analysis with 2.5D [50] or 3D CNNs [51]. In this work, we also tried to conduct the domain adaptation on the basis of 3D segmentation networks. However, we were faced with difficulties of memory consumption given the use of one segmenter, one generator and two discriminators. Moreover, optimizing 3D networks in unsupervised adversarial learning is also very challenging. In these regards, we chose to first down-grade the network to 2D, so that we can focus on the domain adaptation part which is the core of this paper. There are still not many works using 3D CNNs for GANs on medical applications. One recent work is Zhang *et al.* [30] which synthesized CT from MRI on cardiac data using an end-to-end 3D CNN. Another work is Pan *et al.* [28] which used a 3D cycle-consistent GAN to synthesize missing PET from MRI for neuroimage. To the best of our knowledge, the work of [9] was the first to employ 3D CNNs in domain adaptation for medical image segmentation task, which tackled different MRI sequences. Exploring the effectiveness of cross-modality domain adaptation approaches on the basis of 3D network is planned in our future work. In addition, we also plan to extend our method towards heterogeneous domain adaptation, for tackling a more challenging situation where the target domain is significantly different from the source domain in terms of not only low-level appearance but also high-level semantic structures. For example, we may adapt a neural network trained on a healthy cohort of MRI to

a severe-diseased cohort of CT data, which holds promising clinical significance.

VI. CONCLUSION

In conclusion, we investigate the challenging yet crucial task of unsupervised domain adaptation on cross-modality medical images. We propose a novel approach, called *plug-and-play adversarial domain adaptation network (PnP-AdaNet)*, which aligns the latent feature space of the target domain to that of the source domain. Extensive experiments with comparison with state-of-the-arts have validated the effectiveness of our method. Moreover, to facilitate future research on the cross-modality domain adaptation problem, we open-source our code using the well-organized cardiac segmentation challenge dataset. We believe that the cross-modality domain adaptation task will witness rapid development as an important topic in the scope of digital health.

ACKNOWLEDGMENT

(Qi Dou and Cheng Ouyang contributed equally to this work.)

REFERENCES

- [1] A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, and S. Thrun, "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, no. 7639, pp. 115–118, 2017.
- [2] B. E. Bejnordi, M. Veta, P. J. van Diest, B. Van Ginneken, N. Karsssemeijer, G. Litjens, J. A. Van Der Laak, M. Hermesen, Q. F. Manson, M. Balkenhol, O. Geessink, "Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer," *JAMA*, vol. 318, no. 22, pp. 2199–2210, Dec. 2017.
- [3] J. De Fauw, J. R. Ledsam, B. Romera-Paredes, S. Nikolov, N. Tomasev, S. Blackwell, H. Askham, X. Glorot, B. O'Donoghue, and D. Visentin, "Clinically applicable deep learning for diagnosis and referral in retinal disease," *Nature Med.*, vol. 24, no. 9, pp. 1342–1350, 2018.
- [4] A. Torralba and A. A. Efros, "Unbiased look at dataset bias," in *Proc. CVPR*, Jun. 2011, pp. 1521–1528.
- [5] V. Kalogeiton, V. Ferrari, and C. Schmid, "Analysing domain shift factors between videos and images for object detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 11, pp. 2327–2334, Nov. 2016.
- [6] T. Tommasi, M. Lanzi, P. Russo, and B. Caputo, "Learning the roots of visual domain shift," in *Proc. ECCV*, 2016, pp. 475–482.
- [7] E. Gibson, Y. Hu, N. Ghavami, H. U. Ahmed, C. Moore, M. Emberton, H. J. Huisman, and D. C. Barratt, "Inter-site variability in prostate segmentation accuracy using deep learning," in *Proc. MICCAI*, 2018, pp. 506–514.
- [8] A. Bentaieb and G. Hamarneh, "Adversarial stain transfer for histopathology image analysis," *IEEE Trans. Med. Imag.*, vol. 37, no. 3, pp. 792–802, Mar. 2018.
- [9] K. Kamnitsas, C. Baumgartner, C. Ledig, V. Newcombe, J. Simpson, A. Kane, D. Menon, A. Nori, A. Criminisi, B. Glocker, and D. Rueckert, "Unsupervised domain adaptation in brain lesion segmentation with adversarial networks," in *Proc. IPMI*, 2017, pp. 597–609.
- [10] N. Karani, K. Chaitanya, C. Baumgartner, and E. Konukoglu, "A lifelong learning approach to brain MR segmentation across scanners and protocols," in *Proc. MICCAI*, 2018, pp. 476–484.
- [11] Q. Dou, C. Ouyang, C. Chen, H. Chen, and P.-A. Heng, "Unsupervised cross-modality domain adaptation of convnets for biomedical image segmentations with adversarial loss," in *Proc. IJCAI*, 2018, pp. 691–697.
- [12] J. Jiang et al., "Tumor-aware, adversarial domain adaptation from CT to MRI for lung cancer segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assisted Intervent.* Cham, Switzerland: Springer, 2018, pp. 777–785.
- [13] C. Westbrook, C. K. Roth, and J. Talbot, *MRI in Practice*. Hoboken, NJ, USA: Blackwell Publishing, 2004.
- [14] C. Chartrand-Lefebvre, A. Cadrin-Chênevert, E. Bordeleau, P. Ugolini, R. Ouellet, J. Sablayrolles, and J. Prenovault, "Coronary computed tomography angiography: Overview of technical aspects, current concepts, and perspectives," *J.-Can. Assoc. Radiol.*, vol. 58, no. 2, p. 92, 2007.
- [15] X. Zhuang, "Challenges and methodologies of fully automatic whole heart segmentation: A review," *J. Healthcare Eng.*, vol. 4, no. 3, pp. 371–407, 2013.
- [16] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, "Adversarial discriminative domain adaptation," in *Proc. CVPR*, Jul. 2017, pp. 2962–2971.
- [17] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 2030–2096, May 2015.
- [18] E. Tzeng, J. Hoffman, N. Zhang, K. Saenko, and T. Darrell, "Deep domain confusion: Maximizing for domain invariance" Dec. 2014, *arXiv:1412.3474*. [Online]. Available: <https://arxiv.org/abs/1412.3474>
- [19] M. Long, Y. Cao, J. Wang, and M. I. Jordan, "Learning transferable features with deep adaptation networks," in *Proc. 32nd Int. Conf. Int. Conf. Mach. Learn.*, Jul. 2015, pp. 97–105.
- [20] B. Sun and K. Saenko, "Deep CORAL: Correlation alignment for deep domain adaptation," in *Proc. ECCV Workshop*, 2016, pp. 443–450.
- [21] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Proc. NIPS*, 2014, pp. 2672–2680.
- [22] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein gan," in *Proc. ICML*, 2017, pp. 214–223.
- [23] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. ICCV*, Oct. 2017, pp. 2242–2251.
- [24] F. M. Carlucci and A. D'Innocente, S. Bucci, B. Caputo, and T. Tommasi, "Domain generalization by solving jigsaw puzzles," Apr. 2019, *arXiv:1903.06864*. [Online]. Available: <https://arxiv.org/abs/1903.06864>
- [25] Y. Zhang, S. Miao, T. Mansi, and R. Liao, "Task driven generative modeling for unsupervised domain adaptation: Application to X-ray image segmentation," in *Proc. MICCAI*, 2018, pp. 599–607.
- [26] C. Chen, Q. Dou, H. Chen, and P.-A. Heng, "Semantic-aware generative adversarial nets for unsupervised domain adaptation in chest X-ray segmentation," in *Proc. MICCAI Workshop*, 2018, pp. 143–151.
- [27] J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. A. Efros, and T. Darrell, "CyCADA: Cycle-consistent adversarial domain adaptation," Dec. 2017, *arXiv:1711.03213*. [Online]. Available: <https://arxiv.org/abs/1711.03213>
- [28] Y. Pan, M. Liu, C. Lian, T. Zhou, Y. Xia, and D. Shen, "Synthesizing missing PET from MRI with cycle-consistent generative adversarial networks for Alzheimer's disease diagnosis," in *Proc. MICCAI*, 2018, pp. 455–463.
- [29] J. Fan, X. Cao, Z. Xue, P.-T. Yap, and D. Shen, "Adversarial similarity network for evaluating image alignment in deep learning based registration," in *Proc. MICCAI*, 2018, pp. 739–746.
- [30] Z. Zhang, L. Yang, and Y. Zheng, "Translating and segmenting multimodal medical volumes with cycle- and shape-consistency generative adversarial network," in *Proc. CVPR*, Jun. 2018, pp. 9242–9251.
- [31] J. M. Wolterink, T. Leiner, M. A. Viergever, and I. Išgum, "Generative adversarial networks for noise reduction in low-dose CT," *IEEE Trans. Med. Imag.*, vol. 36, no. 12, pp. 2536–2545, Dec. 2017.
- [32] L. Zhang, M. Pereañez, S. K. Piechnik, S. Neubauer, S. E. Petersen, and A. F. Frangi, "Multi-input and dataset-invariant adversarial learning (MDAL) for left and right-ventricular coverage estimation in cardiac MRI," in *Proc. MICCAI*, 2018, pp. 481–489.
- [33] M. A. Degel, N. Navab, and S. Albarqouni, "Domain and geometry agnostic cnns for left atrium segmentation in 3D ultrasound," in *Proc. MICCAI*, 2018, pp. 630–637.
- [34] J. Ren, I. Hacihaliloglu, E. A. Singer, D. J. Foran, and X. Qi, "Adversarial domain adaptation for classification of prostate histopathology whole-slide images," in *Proc. MICCAI*, 2018, pp. 201–209.
- [35] V. V. Valindria, N. Pawlowski, M. Rajchl, I. Lavdas, E. O. Aboagye, A. G. Rockall, D. Rueckert, and B. Glocker, "Multi-modal learning from unpaired images: Application to multi-organ segmentation in CT and MRI," in *Proc. IEEE WACV*, Mar. 2018, pp. 547–556.
- [36] X. Zhuang and J. Shen, "Multi-scale patch and multi-modality atlases for whole heart segmentation of MRI," *Med. Image Anal.*, vol. 31, pp. 77–87, Jul. 2016.
- [37] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. MICCAI*, 2015, pp. 234–241.
- [38] L. Yu, J.-Z. Cheng, Q. Dou, X. Yang, H. Chen, J. Qin, and P.-A. Heng, "Automatic 3D cardiovascular MR segmentation with densely-connected volumetric ConvNets," in *Proc. MICCAI*, 2017, pp. 287–295.

- [39] F. Yu, V. Koltun, and T. Funkhouser, "Dilated residual networks," in *Proc. CVPR*, Jul. 2017, pp. 636–644.
- [40] F. Milletari, N. Navab, and S.-A. Ahmadi, "V-Net: Fully convolutional neural networks for volumetric medical image segmentation," in *Proc. 4th Int. Conf. 3D Vis. (3DV)*, Oct. 2016, pp. 565–571.
- [41] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *Proc. ECCV*, 2014, pp. 818–833.
- [42] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *Proc. NIPS*, 2014, pp. 3320–3328.
- [43] C.-Y. Lee, S. Xie, P. Gallagher, Z. Zhang, and Z. Tu, "Deeply-supervised nets," in *Proc. Artif. Intell. Statist.*, Feb. 2015, pp. 562–570.
- [44] Q. Dou, L. Yu, H. Chen, Y. Jin, X. Yang, J. Qin, and P.-A. Heng, "3D deeply supervised network for automated segmentation of volumetric medical images," *Med. Image Anal.*, vol. 41, pp. 40–54, Oct. 2017.
- [45] C. Payer, D. Štern, H. Bischof, and M. Urschler, "Multi-label whole heart segmentation using CNNs and anatomical label configurations," in *Proc. Int. Workshop Stat. Atlases Comput. Models Heart, (MICCAI)*, 2017, pp. 190–198.
- [46] T. Joyce, A. Chatsias, and S. A. Tsafaris, "Robust multi-modal MR image synthesis," in *Proc. MICCAI*, 2017, pp. 347–355.
- [47] Y. Hiasa, Y. Otake, M. Takao, T. Matsuoka, K. Takashima, A. Carass, J. L. Prince, N. Sugano, and Y. Sato, "Cross-modality image synthesis from unpaired data using CycleGAN," in *Proc. Int. Workshop Simulation Synth. Med. Imag.*, 2018, pp. 31–41.
- [48] T. Galanti, S. Benaim, and L. Wolf, "Generalization bounds for unsupervised cross-domain mapping with WGANs," Jul. 2018, *arXiv:1807.08501*. [Online]. Available: <https://arxiv.org/abs/1807.08501>
- [49] J. P. Cohen, M. Luck, and S. Honari, "Distribution matching losses can hallucinate features in medical image translation," in *Proc. MICCAI*, 2018, pp. 529–536.
- [50] H. R. Roth, L. Lu, A. Seff, K. M. Cherry, J. Hoffman, S. Wang, J. Liu, E. Turkbey, and R. M. Summers, "A new 2.5D representation for lymph node detection using random sets of deep convolutional neural network observations," in *Proc. MICCAI*, 2014, pp. 520–527.
- [51] Q. Dou, H. Chen, L. Yu, L. Zhao, J. Qin, D. Wang, V. C. T. Mok, L. Shi, and P.-A. Heng, "Automatic detection of cerebral microbleeds from MR images via 3D convolutional neural networks," *IEEE Trans. Med. Imag.*, vol. 35, no. 5, pp. 1182–1195, May 2016.




CHENG CHEN received the B.E. degree in biomedical engineering from Zhejiang University, Hangzhou, China, in 2011, and the M.S. degree in biomedical engineering from The Johns Hopkins University, Baltimore, USA, in 2013. She is currently pursuing the Ph.D. degree in computer science and engineering with The Chinese University of Hong Kong. Her research interests include medical image analysis and deep learning.



HAO CHEN (S'14–M'17) received the bachelor's degree in information engineering from Beihang University, Beijing, China, in 2013, and the Ph.D. degree in computer science from The Chinese University of Hong Kong, Hong Kong, in 2017. His research interests include medical image analysis, deep learning, object detection, and segmentation.



BEN GLOCKER received the Ph.D. degree from TU Munich. He held a postdoctoral position with Microsoft and a Research Fellow with the University of Cambridge. He is currently a Senior Lecturer in machine learning for imaging with the Department of Computing, Imperial College London, and co-leading the Biomedical Image Analysis Group. His research interests include the intersection of medical image analysis, and artificial intelligence aiming to build computational tools for improving image-based detection and diagnosis of disease.



XIAHAI ZHUANG graduated from the Department of Computer Science, Tianjin University, and received the M.S. degree in computer science from Shanghai Jiao Tong University and the Ph.D. degree from University College London. He is currently an Associate Professor with the School of Data Science, Fudan University. His research interests include medical image analysis, image processing, and computer vision. Two of his works have been the finalists for MICCAI Young Scientist Awards.



PHENG-ANN HENG (M'92–SM'06) received the Ph.D. degree in computer science from Indiana University, Indianapolis, IN, USA. He is currently a Professor with the Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong, where he is also the Director of the Virtual Reality, Visualization, and Imaging Research Centre. He is also the Director of the Research Center for Human-Computer Interaction, Shenzhen Institute of Advanced Integration Technology, Chinese Academy of Sciences, Shenzhen, China. His research interests include virtual reality applications in medicine, visualization, medical imaging, human-computer interfaces, rendering and modeling, interactive graphics, and animation.



QI DOU (S'14–M'18) received the bachelor's degree in biomedical engineering from Beihang University, Beijing, China, in 2014, and the Ph.D. degree from the Department of Computer Science and Engineering, The Chinese University of Hong Kong, in 2018. She is currently a Postdoctoral Research Associate with the Biomedical Image Analysis Group, Department of Computing, Imperial College London. Her research interests include medical image analysis, deep learning, and machine learning.



CHENG OUYANG received the B.E. degree in biomedical engineering from Xi'an Jiaotong University, China, in 2016, and the M.S. degree in electrical and computer engineering from the University of Michigan, Ann Arbor, USA, in 2018. He is currently pursuing the Ph.D. degree computing with Imperial College London. His research interests include cardiac MRI analysis and reconstruction.