

Towards Cross-Modality Medical Image Segmentation with Online Mutual Knowledge Distillation

Kang Li,¹ Lequan Yu,^{1*} Shujun Wang,¹ Pheng-Ann Heng^{1,2}

¹Department of Computer Science and Engineering, The Chinese University of Hong Kong
²Guangdong Provincial Key Laboratory of Computer Vision and Virtual Reality Technology, Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, China
 {kli, lqyu, sjwang, pheng}@cse.cuhk.edu.hk

Abstract

The success of deep convolutional neural networks is partially attributed to the massive amount of annotated training data. However, in practice, medical data annotations are usually expensive and time-consuming to be obtained. Considering multi-modality data with the same anatomic structures are widely available in clinic routine, in this paper, we aim to exploit the prior knowledge (*e.g.*, shape priors) learned from one modality (*aka.*, assistant modality) to improve the segmentation performance on another modality (*aka.*, target modality) to make up annotation scarcity. To alleviate the learning difficulties caused by modality-specific appearance discrepancy, we first present an Image Alignment Module (IAM) to narrow the appearance gap between assistant and target modality data. We then propose a novel Mutual Knowledge Distillation (MKD) scheme to thoroughly exploit the modality-shared knowledge to facilitate the target-modality segmentation. To be specific, we formulate our framework as an integration of two individual segmentors. Each segmentor not only explicitly extracts one modality knowledge from corresponding annotations, but also implicitly explores another modality knowledge from its counterpart in mutual-guided manner. The ensemble of two segmentors would further integrate the knowledge from both modalities and generate reliable segmentation results on target modality. Experimental results on the public multi-class cardiac segmentation data, *i.e.*, MM-WHS 2017, show that our method achieves large improvements on CT segmentation by utilizing additional MRI data and outperforms other state-of-the-art multi-modality learning methods.

Introduction

Modern clinical practice usually involves multiple imaging techniques, *e.g.*, magnetic resonance imaging (MRI) and computed tomography (CT), to acquire a more comprehensive view of particular tissues or organs in disease diagnosis and surgical planning (Cao et al. 2017). For instance, CT and MRI are extensively used to provide clear anatomical information of cardiac structures (Zhuang et al. 2019). As different imaging modalities are based on different physical imaging principles, they normally emphasize organs or

tissues with distinct visual contrast. For example, in MRI brain scans, MRI-T2 has better visual effects in edema, while MRI-T1c highlights gross tumor core (Woo, Stone, and Prince 2014). Therefore, multi-modality learning is progressively developed in medical imaging domain. In recent years, deep learning based methods have achieved promising performance in many medical image analysis tasks, relying on the massive amount of annotated data (Litjens et al. 2017). However, large annotated datasets are usually expensive and time-consuming to acquire in medical domain, since the annotations have to be marked by professional medical experts under strict and meticulous inspection. Considering different imaging modalities reflect the same anatomic structures, multi-modality learning has become a promising direction to make up the annotation scarcity in automatic medical image analysis. In this paper, we aim to investigate the effectiveness of *utilizing the prior knowledge (e.g., shape priors) learned from one modality (aka. assistant modality) to improve the segmentation performance on another modality (aka. target modality)*, where no assistant-modality data is required in testing phase.

In general, one intuitive way is joint-training, where the network is trained with both assistant-modality and target-modality data. Another straightforward approach is to fine-tune a deep convolutional neural network (DCNN) learned from assistant-modality data with the target-modality data, so that it can transfer the prior knowledge learned from assistant modality to target-modality tasks. Nevertheless, the shared cross-modality information can not be well exploited in either joint-training or fine-tuning, since the representative modality-shared features are hard to directly learn from multi-modality data with large appearance discrepancy. Valindria et al. (2018) proposed to assign each modality data with its specific feature extractor to alleviate the negative influences from distinct multi-modality appearances. They presented four dual-stream architectures with well-designed parameter sharing strategies to explore cross-modality information, and demonstrated X-shape architecture is more suitable for multi-modality image segmentation. However, these dual-stream architectures incorporate human intervention in determining parameter sharing, and would require special adjustments to generalize to other medical

*Corresponding Author.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

analysis tasks. In our problem setting, two modality information are involved in the network learning, where the valuable modality-shared knowledge should be thoroughly explored to enhance the generalization of networks, while the redundant modality-specific appearances should be aligned to ease network learning.

In this paper, we propose a novel cross-modality image segmentation framework based on knowledge distillation to utilize assistant-modality (*e.g.*, MRI) priors to improve the segmentation performance on target-modality data (*e.g.*, CT). To alleviate the difficulties caused by modality-specific appearance discrepancy and facilitate the learning of modality-shared knowledge, we first present an Image Alignment Module (IAM) to reduce the appearance gap by translating the assistant-modality images to synthetic target-modality images via adversarial learning (Goodfellow et al. 2014). As the core workhorses in our framework, we develop a novel Mutual Knowledge Distillation (MKD) scheme to better exploit the modality-shared knowledge from both assistant and target modality data for superior segmentation results. We formulate this scheme as the cohort of two individual segmentors, which directly learn feature representations from synthetic target-modality data and real target-modality data, respectively. Besides explicitly learned from segmentation annotations, two segmentors are also mutually guided by its counterpart outputs implicitly via knowledge distillation (Hinton, Vinyals, and Dean 2015). Armed with the mutual guidance, both segmentors are able to explicitly learn the knowledge from one modality and implicitly learn the knowledge from another modality concurrently. Additionally, the ensemble of two segmentors would further integrate the knowledge learned from two modalities and perform more reliable segmentation in target modality. The whole framework is optimized in end-to-end online manner, so that the IAM could receive timely feedback from MKD scheme to generate more reasonable synthetic target-modality images with well-preserved shape priors. We extensively evaluate our method on the MM-WHS 2017 Challenge dataset (Zhuang et al. 2019). Our framework brings large dice improvements (3.06%) for cardiac segmentation on CT volumes by utilizing additional MRI data, substantially outperforming other multi-modality learning methods. The major contributions of this paper are as follows.

- We present an effective multi-modality learning framework to utilize prior knowledge of one modality to enhance the segmentation performance of another modality.
- We propose a novel Mutual Knowledge Distillation scheme to better exploit modality-shared knowledge with mutual guidance of segmentor outputs.
- We also present an Image Alignment Module to reduce cross-modality appearance discrepancy to promote the learning of modality-shared knowledge.
- We extensively evaluate our method on public multi-class cardiac segmentation challenge data (Zhuang et al. 2019). Our method outperforms other state-of-the-art multi-modality learning methods.

Related Work

Multi-modality Learning in Medical Imaging

Multi-modality learning has been widely studied in medical field. Many approaches (Fidon et al. 2017; Kamnitsas et al. 2017; Guo et al. 2018; Nie et al. 2016) proposed different feature fusion strategies to utilize complementary visual reflections of multi-modality data for comprehensive segmentation. For example, Kamnitsas et al. (2017) presented a dual pathway architecture to incorporate cross-modality local and global features for brain lesion segmentation from multi-channel MRI data, while Fidon et al. (2017) designed a nested architecture to jointly utilize different modality of MRI brain tumor data and adapt their framework to handle scalable inputs. Guo et al. (2018) compared the effectiveness of feature-level, classifier-level and decision-level feature fusion for multi-modality tumor segmentation. However, these methods also require paired multi-modality data in testing stage to perform thorough feature fusion, while we aim to exploit the priors of assisted modality to promote the performance on another modality by enhancing model generalization ability, where only target-modality data is required in the testing phase.

To utilize the prior knowledge between different modalities, Zheng (2015) proposed the marginal space learning method to transfer assisted modality knowledge to target modality. Moeskops et al. (2016) investigated how to utilize multi-modality information under multi-task learning frameworks. Recently, Valindria et al. (2018) developed dual-stream encoder-decoder architectures, which assign each modality with specific branch and extract cross-modality features with delicately designed parameter sharing strategies. With the recent advance in image translation (Goodfellow et al. 2014; Zhu et al. 2017), many researchers were inspired to tackle medical problems with modality translation. Zhao et al. (Zhao et al. 2019) integrated registration between different modality MR images for data augmentation. Jiang et al. (2018) and Zhang et al. (2018b) proposed to eliminate the appearance gap of different modalities with generative adversarial network (GAN) for multi-modal segmentation. Jiang et al. (2018) proposed to train a GAN to generate meaningful synthetic MRI images first, and take synthetic images with real ones together as segmentor inputs for lung cancer segmentation, while Zhang et al. (2018b) proposed to jointly optimize GAN and segmentor, where the updated synthetic information could be exploited. Different from these methods which directly fuse multi-modality knowledge by joint training, our framework takes two branches for synthetic data and real data respectively, and integrates cross-modality knowledge in a mutual-guided way.

Knowledge Distillation

Our work is also related to knowledge distillation, which was first proposed for network compression (Hinton, Vinyals, and Dean 2015). Hinton et al. (2015) proposed to use the output class probabilities of a static cumbersome model as soft targets to teach a lightweight student model. Zagoruyko and Komodakis (2017) further forced the stu-

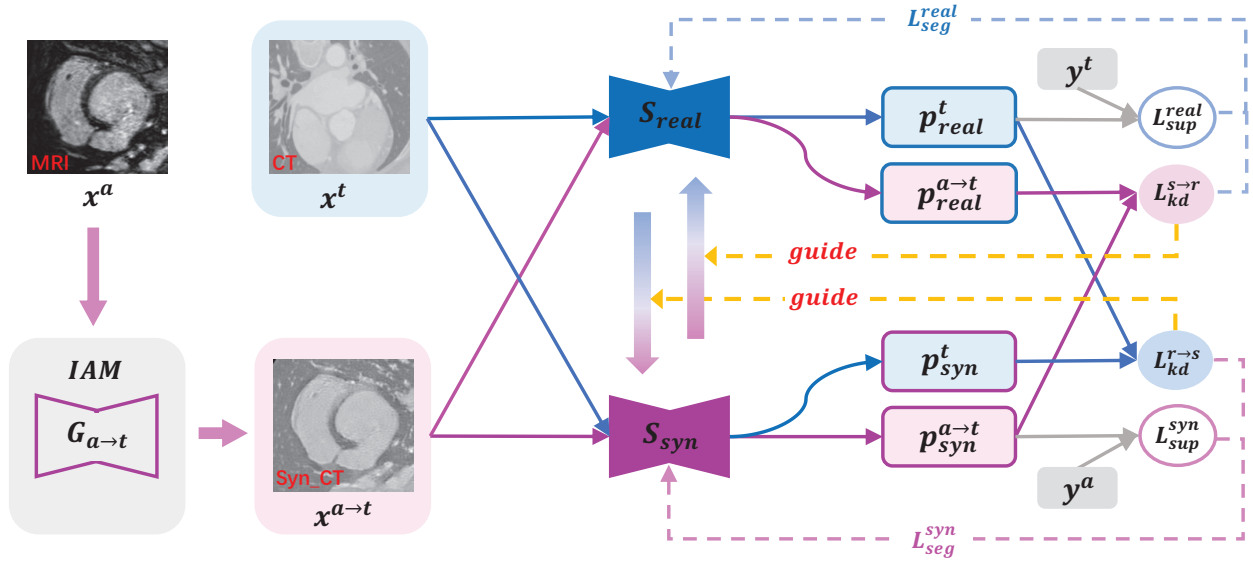


Figure 1: Overview of our framework, where magenta and blue represent the data flow of assistant modality (e.g., MRI) and target modality (e.g., CT) respectively. The generator $G_{a \rightarrow t}$ performs assistant-to-target translation and outputs synthesized CT data. The magenta-to-blue transition arrow represents the knowledge transfer from S_{syn} to S_{real} guided by $L_{kd}^{s \rightarrow r}$, while the blue-to-magenta transition arrow demonstrates the knowledge transfer from S_{real} to S_{syn} guided by $L_{kd}^{r \rightarrow s}$. In this way, both segmentors are mutually guided. In testing, we feed x_t into both segmentors and adopt the ensemble prediction as final results.

dent network to mimic middle-level attention maps of the teacher network besides its output probability maps. Since classical knowledge distillation approaches always require a powerful pre-trained teacher model, mutual learning was proposed to encourage an ensemble of networks to learn from each other cooperatively throughout the training process (Zhang et al. 2018a). Based on this, Wu et al. (2019b) utilized mutual learning to capture complementary features in semi-supervised classification. Wu et al. (2019a) applied mutual learning between contour extraction and edge extraction for saliency detection. We share a similar philosophy with mutual learning, but aim to design a cross-modality learning framework to utilize the shape priors from different modalities, which results in a different methodology.

Methodology

In this section, we introduce our proposed cross-modality image segmentation framework in details. Given a set of labeled samples $\{x_i^a, y_i^a\}_{i=1}^N$ from assistant-modality data X^a , and a set of labeled samples $\{x_j^t, y_j^t\}_{j=1}^M$ from target-modality data X^t , we involve both X^a and X^t in network training, to improve the segmentation performance on target modality during testing. As illustrated in Figure 1, we design an Image Alignment Module (IAM) to eliminate the modality-specific appearance differences, and a Mutual Knowledge Distillation (MKD) scheme to thoroughly exploit modality-shared knowledge. The whole framework is optimized in online mutual learning manner.

Image Alignment Module

Since redundant modality-specific appearances would introduce bias and increase the difficulty in leveraging the valuable modality-shared prior knowledge from X^a , we propose to perform a transformation from assistant-modality image x^a to target-modality image x^t . In this way, the synthetic target-modality image $x^{a \rightarrow t}$ would acquire similar appearances to target-modality data with unaffected assistant-modality structures.

Inspired by generative adversarial networks (Goodfellow et al. 2014), we adopt a generator $G_{a \rightarrow t}$ and discriminator D_t to perform the assistant-to-target image translation $G_{a \rightarrow t} : X^a \rightarrow X^t$ by adversarial learning. To be specific, the generator $G_{a \rightarrow t}$ takes real assistant-modality image x^a as inputs and produces realistic synthetic target-modality image $x^{a \rightarrow t} = G_{a \rightarrow t}(x^a)$ to fool the discriminator D_t , while the discriminator D_t tries to distinguish the synthetic images from real ones. In this way, the $G_{a \rightarrow t}$ and the D_t are competing while improving each other. The parameters of $G_{a \rightarrow t}$ and D_t are optimized by adversarial loss as

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

where $G_{a \rightarrow t}$ tries to minimize the objective function by generating more realistic target-modality images, while D_t intends to maximize it by identifying synthetic target-modality images from real ones.

Since the above mapping alone is weakly-constrained, an inverse mapping $G_{t \rightarrow a} : X^t \rightarrow X^a$ is introduced to impose the cycle-consistency constrain (Zhu et al. 2017). Particularly, we attach another generator $G_{t \rightarrow a}$ and discriminator

D_a into this module and form a similar minimax game by adversarial loss in the following

$$\mathcal{L}_{adv}^a(G_{t \rightarrow a}, D_a) = \mathbb{E}_{x^a \sim X^a} [\log D_a(x^a)] + \mathbb{E}_{x^t \sim X^t} [\log(1 - D_a(G_{t \rightarrow a}(x^t)))] \quad (2)$$

The cycle-consistency constraint forces the synthetic images could be translated back to the input images, meaning the reconstructed assistant-modality image $\hat{x}^a = G_{t \rightarrow a}(G_{a \rightarrow t}(x^a))$ should be similar to x^a . And the reconstructed target-modality image $\hat{x}^t = G_{a \rightarrow t}(G_{t \rightarrow a}(x^t))$ should also be closer to x^t . Specifically, the cycle-consistency loss is defined as

$$\mathcal{L}_{cyc}(G_{a \rightarrow t}, G_{t \rightarrow a}) = \mathbb{E}_{x^a \sim X^a} \|\hat{x}^a - x^a\|_1 + \mathbb{E}_{x^t \sim X^t} \|\hat{x}^t - x^t\|_1 \quad (3)$$

Overall, the full objective functions to optimize $G_{a \rightarrow t}$ and $G_{t \rightarrow a}$ are summarized as

$$\mathcal{L}_{gan}^{a \rightarrow t} = \mathcal{L}_{adv}^t + \lambda_{cyc} \mathcal{L}_{cyc} \quad (4)$$

$$\mathcal{L}_{gan}^{t \rightarrow a} = \mathcal{L}_{adv}^a + \lambda_{cyc} \mathcal{L}_{cyc} \quad (5)$$

where λ_{cyc} is a hyperparameter to balance training. Noted that the generator $G_{a \rightarrow t}$ is co-optimized with the following Mutual Knowledge Distillation scheme to preserve shape invariance in end-to-end manner, which will be concretely described in Section . Through this module, the appearances of assistant-modality data are aligned with target modality, and the modality-shared prior knowledge could be better investigated in the following step.

Mutual Knowledge Distillation

To thoroughly exploit modality-shared knowledge, we develop a novel Mutual Knowledge Distillation (MKD) scheme to facilitate the network learning. We formulate it as a cohort of two segmentation networks, *i.e.*, synthetic segmentor S_{syn} and real segmentor S_{real} , where S_{syn} and S_{real} receive direct supervision from synthetic target-modality image $x^{a \rightarrow t}$ and real target-modality image x^t , respectively. During training, we not only encourage each segmentor to explicitly learn one modality knowledge from corresponding annotations, but also guide it to implicitly explore another modality knowledge from its peer segmentor outputs. With mutual guidance, both segmentors could explore feature representations of two modalities in explicit and implicit ways.

For explicit learning, the segmentors are optimized with the supervised segmentation loss with respect to the ground truth. Specifically, the synthetic segmentor S_{syn} mainly learns assistant-modality features from synthetic target-modality image $x^{a \rightarrow t}$ with assistant-modality annotation y^a , while the real segmentor S_{real} mainly exploits target-modality information from real target-modality image x^t with target-modality ground truth y^t . The supervised segmentation loss is formulated as the combination of cross-

entropy loss \mathcal{L}_{ce} and dice loss \mathcal{L}_{dice} :

$$\begin{aligned} p_{syn}^{a \rightarrow t} &= S_{syn}(x^{a \rightarrow t}), \\ p_{real}^t &= S_{real}(x^t), \\ \mathcal{L}_{sup}^{syn} &= \mathcal{L}_{ce}(y^a, p_{syn}^{a \rightarrow t}) + \mathcal{L}_{dice}(y^a, p_{syn}^{a \rightarrow t}), \\ \mathcal{L}_{sup}^{real} &= \mathcal{L}_{ce}(y^t, p_{real}^t) + \mathcal{L}_{dice}(y^t, p_{real}^t), \end{aligned} \quad (6)$$

where y^a and y^t are the segmentation ground truth of assistant modality and target modality, respectively.

For implicit learning, we propose to utilize the outputs of each segmentor as guidance for its counterpart, so that synthetic segmentor S_{syn} and real segmentor S_{real} are mutually guided by each other and interactively explored the dark knowledge (*i.e.*, output category similarity) from each other (Furlanello et al. 2018) (Hinton, Vinyals, and Dean 2015). To make real segmentor exploit the guidance from synthetic segmentor, we feed synthetic target-modality image $x^{a \rightarrow t}$ into real segmentor S_{real} to produce the probability map $p_{real}^{a \rightarrow t} = S_{real}(x^{a \rightarrow t})$, and encourage it to be similar to synthetic segmentor output $p_{syn}^{a \rightarrow t}$ with the synthetic-to-real knowledge distillation loss $\mathcal{L}_{kd}^{s \rightarrow r}$ (*i.e.*, the magenta filled circle in Figure 1). Since S_{syn} is directly supervised by synthetic target-modality annotations, the guidance (*i.e.*, $p_{syn}^{a \rightarrow t}$) is trustworthy. Furthermore, as S_{syn} is equipped with assistant-modality priors, this synthetic-to-real knowledge distillation scheme would transfer knowledge from S_{syn} to S_{real} and facilitate S_{real} to integrate the cross-modality information effectively. Similarly, we also feed real target-modality images x^t into synthetic segmentor S_{syn} to obtain the probability map $p_{syn}^t = S_{syn}(x^t)$. By narrowing down the differences between p_{syn}^t and real segmentor output p_{real}^t , S_{syn} takes the informative guidance from S_{real} by real-to-synthetic knowledge distillation loss $\mathcal{L}_{kd}^{r \rightarrow s}$ (*i.e.*, the blue filled circle in Figure 1). With the training going on, both segmentors are mutually guided to gradually explore the knowledge of another modality from its counterpart. We formulate the knowledge distillation loss as the cross-entropy between two probability maps, following previous work (Hinton, Vinyals, and Dean 2015):

$$\begin{aligned} \mathcal{L}_{kd}^{s \rightarrow r} &= \mathcal{L}_{ce}(p_{syn}^{a \rightarrow t}, p_{real}^{a \rightarrow t}), \\ \mathcal{L}_{kd}^{r \rightarrow s} &= \mathcal{L}_{ce}(p_{real}^t, p_{syn}^t). \end{aligned} \quad (7)$$

The objective functions to train S_{syn} and S_{real} are then defined as

$$\begin{aligned} \mathcal{L}_{seg}^{real} &= \mathcal{L}_{sup}^{real} + \lambda_{kd}^1 \mathcal{L}_{kd}^{s \rightarrow r}, \\ \mathcal{L}_{seg}^{syn} &= \mathcal{L}_{sup}^{syn} + \lambda_{kd}^2 \mathcal{L}_{kd}^{r \rightarrow s}, \end{aligned} \quad (8)$$

where λ_{kd}^1 and λ_{kd}^2 are hyperparameters. In summary, each individual segmentor is optimized by the supervised segmentation loss to explicitly learn one modality knowledge, as well as knowledge distillation loss to implicitly absorb another modality knowledge, and thus comprehensively investigates cross-modality information under peer guidance.

Online Mutual Training

Overall, our total objective to train the whole framework can be formulated as

$$\mathcal{L}_{obj} = \mathcal{L}_{gan}^{a \rightarrow t} + \mathcal{L}_{gan}^{t \rightarrow a} + \mathcal{L}_{seg}^{syn} + \mathcal{L}_{seg}^{real} \quad (9)$$

Algorithm 1 Training procedure of the proposed approach

Input: A batch of (x^t, y^t) from target-modality dataset X^t and (x^a, y^a) from assistant-modality dataset X^a

Output: The prediction p^t of input x^t

- 1: $\theta_{G_{a \rightarrow t}}, \theta_{G_{t \rightarrow a}}, \theta_{D_a}, \theta_{D_t}, \theta_{S_{syn}}, \theta_{S_{real}} \leftarrow$ initialize
 - 2: **while** not converge **do**
 - 3: $(x^t, y^t), (x^a, y^a) \leftarrow$ sampled from X^t and X^a
 - 4: Calculate $\mathcal{L}_{gan}^{a \rightarrow t}$ and \mathcal{L}_{sup}^{syn} as Eq. (4) and (6)
 - 5: Update $\theta_{G_{a \rightarrow t}} \stackrel{\pm}{\leftarrow} -\Delta_{\theta_{G_{a \rightarrow t}}} (\mathcal{L}_{gan}^{a \rightarrow t} + \mathcal{L}_{sup}^{syn})$
 - 6: Calculate \mathcal{L}_{adv}^t as Eq. (1)
 - 7: Update $\theta_{D_t} \stackrel{\pm}{\leftarrow} -\Delta_{\theta_{D_t}} \mathcal{L}_{adv}^t$
 - 8: Calculate $\mathcal{L}_{gan}^{t \rightarrow a}$ as Eq. (5)
 - 9: Update $\theta_{G_{t \rightarrow a}} \stackrel{\pm}{\leftarrow} -\Delta_{\theta_{G_{t \rightarrow a}}} \mathcal{L}_{gan}^{t \rightarrow a}$
 - 10: Calculate \mathcal{L}_{adv}^a as Eq. (2)
 - 11: Update $\theta_{D_a} \stackrel{\pm}{\leftarrow} -\Delta_{\theta_{D_a}} \mathcal{L}_{adv}^a$
 - 12: Calculate \mathcal{L}_{seg}^{syn} and \mathcal{L}_{seg}^{real} as Eq. (8)
 - 13: Update $\theta_{S_{syn}} \stackrel{\pm}{\leftarrow} -\Delta_{\theta_{S_{syn}}} \mathcal{L}_{seg}^{syn}$
 - 14: Update $\theta_{S_{real}} \stackrel{\pm}{\leftarrow} -\Delta_{\theta_{S_{real}}} \mathcal{L}_{seg}^{real}$
 - 15: Calculate p_{real}^t and p_{syn}^t
 - 16: **return** $p^t \leftarrow$ the ensemble of p_{real}^t and p_{syn}^t
-

The entire framework is trained in online manner, and we update all components alternatively in each iteration. Algorithm 1 presents detailed training procedures. We closely follow the settings in Zhu et al. (2017) to optimize $G_{a \rightarrow t}$, D_t , $G_{t \rightarrow a}$, and D_a , except that the generator $G_{a \rightarrow t}$ is optimized together with S_{syn} by $\mathcal{L}_{gan}^{a \rightarrow t}$ and \mathcal{L}_{sup}^{syn} with S_{syn} fixed. This strategy ensures geometric transformation consistency in $G_{a \rightarrow t}$ by constraining pixel-level semantic relationship of synthetic target-modality images. Due to that, $G_{a \rightarrow t}$ could generate meaningful synthetic target-modality data with unaffected shape priors.

For segmentation networks, we alternatively update the weights of S_{syn} and S_{real} by the combination of supervised segmentation loss and knowledge distillation loss. From S_{syn} view, the real-to-synthetic knowledge distillation loss provides chances to have a glimpse of real data, which guides S_{syn} to generalize towards a more reliable direction. From S_{real} view, the synthetic-to-real knowledge distillation loss brings additional knowledge of highly realistic synthetic data as data augmentation and directly enhance its generalization ability. The real-to-synthetic and synthetic-to-real knowledge distillation make S_{syn} and S_{real} mutually benefit from each other’s guidance in online manner. The ensemble model is able to explore more informative and comprehensive cross-modality knowledge. In testing, we feed x_t into both segmentors and adopt the ensemble prediction as final results.

Experiments

Dataset and Implementation Details

We evaluate the proposed method on the Multi-modality Whole Heart Segmentation Challenge 2017 (MM-WHS

2017) dataset, which contains unpaired 20 MRI and 20 CT volumes as the training data and the annotations of 7 cardiac substructures including the left ventricle blood cavity (LV), the right ventricle blood cavity (RV), the left atrium blood cavity (LA), the right atrium blood cavity (RA), the myocardium of the left ventricle (MYO), the ascending aorta (AA), and the pulmonary artery (PA) (Zhuang et al. 2019). Here, we take MRI as assistant modality and CT as target modality, as MRI images have better soft tissue contrast and would provide more detailed heart information for CT segmentation. We randomly split 20 CT volumes into two folds and perform two-fold cross validation. In each fold, we use all 20 MRI volumes and 10 CT volumes to train our network.

Following data preprocessing procedures of previous works in this dataset (Chen et al. 2019), the image slices are sampled in coronal view, cropped centered at heart region, and resized into 256×256 . Data augmentation methods applied here include random flipping and rotating. The architectures of generators and discriminators are closely following the setting in Zhu et al. (2017), and the segmentors adopt the baseline network architecture in (Valindria et al. 2018), consisting of Unet (2015) with residual blocks (2016). In training, Adam optimizer is used for the optimization of generators, discriminators and segmentors with the learning rate of 2×10^{-4} , except for segmentors which adopt a decay rate of 0.9 for every two epochs. The hyperparameters λ_{cyc} , λ_{kd}^1 and λ_{kd}^2 are empirically set as 10, 0.5 and 1, respectively.

Experimental Results

Comparison with other methods. To demonstrate the effectiveness of our method, we train a segmentor with only CT data as our baseline and compare our method with other multi-modality learning methods. The quantitative results in whole heart segmentation are shown in Table 1. In joint-training, the segmentor is trained with both CT and MRI data simultaneously, while in fine-tune, we pretrain a segmentor with MRI data and then update the network parameters by CT data. We compare with X-shape architecture, which achieves the best segmentation performance among all dual-stream architectures in (Valindria et al. 2018). We also compare with Jiang et al. (2018) and Zhang et al. (2018b), which employ GAN to alleviate multi-modality appearance discrepancy in network training. As S_{syn} and S_{real} acquire two modality prior knowledge, we input CT data in both segmentors. Considering complementary information, we adopt the ensemble prediction as final results. Note that all methods mentioned above are experimented with the same segmentor backbone for fair comparison. We use the Dice coefficient as evaluation metric.

It is observed that with assisted modality data, all multi-modality learning approaches achieve average dice improvements of all substructures in CT segmentation compared with baseline, which demonstrates the feasibility of utilizing assisted modality priors to promote the performance on target-modality segmentation. However, the mean dice improvements of fine-tune and joint-training are limited, *i.e.*, 0.63% and 0.37%, respectively. X-shape method only achieves comparable results with fine-tune, since separate branch could only avoid, not directly reduce appearance

Table 1: Quantitative comparison between our method and other multi-modality segmentation methods. Here we take CT as target modality and MR as assistant modality. The dice of all heart substructures and the average of them are reported here.

Method	Mean Dice	Dice of substructures of heart						
		MYO	LA	LV	RA	RV	AA	PA
Baseline	0.8706	0.8702	0.8922	0.9086	0.8386	0.846	0.9252	0.8134
Fine-tune	0.8769	0.8716	0.9040	0.9079	0.8443	0.8526	0.9274	0.8305
Joint-training	0.8743	0.8665	0.9076	0.9123	0.8278	0.8492	0.9302	0.8266
X-shape (2018)	0.8767	0.8719	0.8979	0.9094	0.8551	0.8444	0.9343	0.8240
Jiang et al. (2018)	0.8765	0.8723	0.9054	0.9073	0.8338	0.8525	0.9484	0.8156
Zhang et al. (2018b)	0.8850	0.8781	0.9112	0.9134	0.8514	0.8631	0.9430	0.8342
Ours (S_{syn})	0.8849	0.8738	0.9086	0.9169	0.8571	0.8633	0.9412	0.8333
Ours (S_{real})	0.8947	0.8893	0.9131	0.9197	0.8668	0.8758	0.9552	0.8432
Ours	0.9012	0.8934	0.9190	0.9267	0.8747	0.8814	0.9595	0.8538

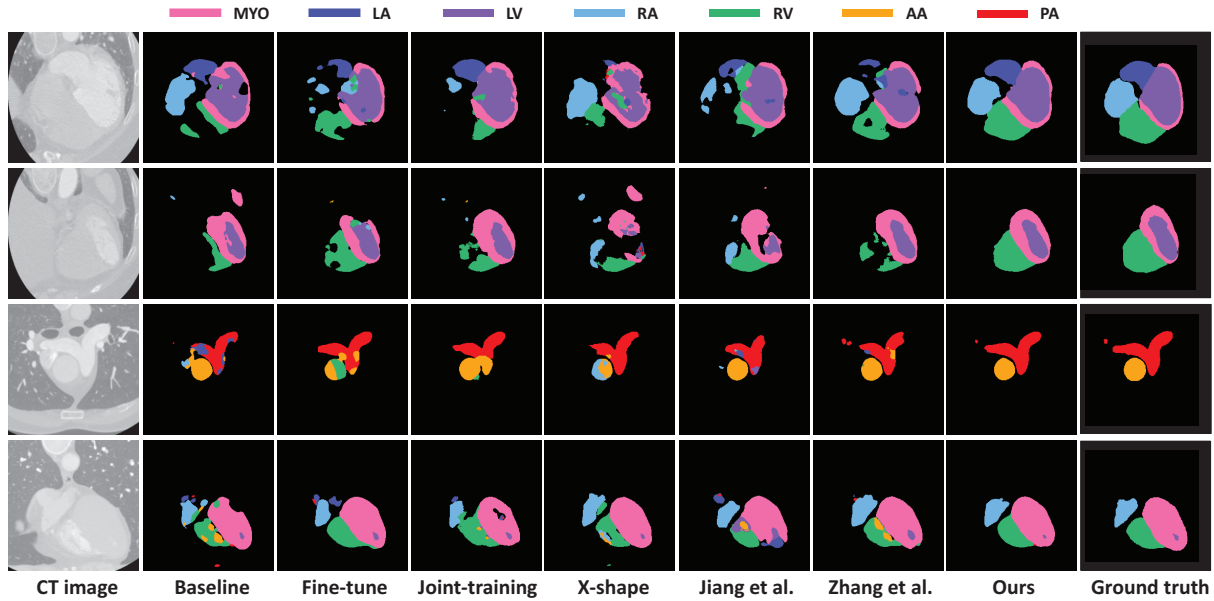


Figure 2: Visual comparison of segmentation results produced by different methods, where the legend for cardiac substructures is presented on the top. As we can see, our results (second last column) are closer to the ground truth (last column) than others.

gap. Although Jiang et al. (2018) proposed to use GAN for multi-modality appearance difference reduction, the improvement is still limited. The reason may be that they train the GAN and segmentor in offline manner, where the segmentor would be incompatible to the generated images. With the help of online training strategy, Zhang et al. (2018b) achieve 0.85% mean dice improvements than Jiang et al. (2018). Compared with Zhang et al. (2018b), our approach takes advantages of mutual knowledge distillation scheme and can interactively improves the cross-modality knowledge exploration of each segmentor. Therefore, our method (last row in Table 1) achieves the highest dice in all heart substructure segmentation and further outperforms Zhang et al. (2018b) with 1.62% in mean dice.

For our approach, both real segmentor S_{real} and synthetic segmentor S_{syn} enhance the mean dice with 2.41% and 1.43% compared to baseline. Since S_{real} is trained by

direct supervision from real CT annotations, while S_{syn} directly learn from synthetic CT data, the mean dice of S_{real} is slightly better than S_{syn} . During training, S_{real} learns explicit knowledge from CT data and implicit priors from MRI data. Meanwhile, S_{syn} absorbs explicit MRI knowledge and implicit CT priors. Considering the complimentary information, the ensemble segmentor acquires 3.06% improvements than baseline in mean dice, which demonstrates the effectiveness of our method. One intuitive explanation of why our mutual knowledge distillation performs better than joint-training is that joint training could be seen as “hard-parameter” sharing (same parameters) between real and synthetic segmentor, while our Mutual Knowledge Distillation scheme enables “soft-parameter” sharing, which provides more flexibility in feature extraction.

It is worth to mention that the effectiveness of our method is from the cross-modality mutual knowledge distillation,

Table 2: Ablation study of key components in our method, where the mean dice of all heart substructures is reported.

Methods	S_{syn}	S_{real}	S_{ens}
W/o IAM	0.8730	0.8745	0.8829
W/o $\mathcal{L}_{kd}^{s \rightarrow r}$	0.8825	0.8742	0.8879
W/o $\mathcal{L}_{kd}^{r \rightarrow s}$	0.5785	0.8904	0.7937
Ours	0.8849	0.8947	0.9012

Table 3: Performance comparison of different amounts of assisted modality data in training.

Setting	Mean Dice
5-MRI	0.8916
10-MRI	0.8945
20-MRI	0.9012

not solely gains from model ensemble. For single segmentor, our S_{real} outperforms Zhang et al. (2018b) with 1.27% in mean dice with the same segmentor architecture, as shown in Table 1. Moreover, for ensemble segmentor, we train two segmentors in Zhang et al. (2018b) with the combination of synthetic and real CT and use the ensemble prediction as the final results. This ensemble method achieve the mean dice of 88.92%, inferior to our final results by 1.2%, validating the effectiveness of proposed MKD regularization scheme.

Qualitative comparison. The qualitative comparison between different methods is presented in Figure 2. As observed, without assistant-modality prior knowledge, the network trained only with CT data would misidentify many substructures. After exploiting assisted modality information, the segmentation results are better but still scattered and fragmentary, especially in fine-tune and X-shape. Owing to the proposed Mutual Knowledge Distillation scheme, the predictions of our method have more smooth contours and are closer to the ground truth compared to others.

Ablation analysis. To investigate the effectiveness of the key components in our framework, we conduct ablation studies and the quantitative results are presented in Table 2. We first remove the Image Alignment Module (IAM) and directly train S_{syn} and S_{real} with raw CT and MRI data under the Mutual Knowledge Distillation scheme. As observed from Table 2, the mean dice of S_{syn} and S_{real} are similar to each other. The segmentation performance of S_{syn} , S_{real} and S_{ens} are degraded by 1.19%, 2.02% and 1.83%, respectively, showing the importance of IAM in our framework. We further present some synthesized CT images in Figure 3. As we can see, the heart structures are clear and well-preserved during translation. In addition, we also conduct experiments with only one-way knowledge distillation existing in our framework. As shown in Table 2, without synthetic-to-real knowledge distillation, the results of S_{syn} is comparable with mutual distillation, while the performance of S_{real} is weaken and compromised. Specifically, S_{real} and S_{ens} decrease 2.05% and 1.33% compared to mu-

tual distillation respectively. In another aspect, with the absence of real-to-synthetic knowledge distillation, the performance of S_{syn} is hugely deteriorated and further corrupts S_{ens} performance. As no real CT data is involved in any ways for training S_{syn} , the S_{syn} is completely isolated and lost. With our proposed mutual knowledge distillation, both segmentors could gain mutual benefits from each other’s training.

We also present the qualitative comparison among predictions of S_{syn} , S_{real} and our ensemble results in Figure 4. Since the cohort of S_{syn} and S_{real} integrates information from annotations and peer guidance from each other’s outputs, the ensemble segmentor could highlight the uniform predictions while correcting the misleading predictions. For those inconsistent predictions of S_{syn} and S_{real} (e.g., last row in Figure 4), S_{ens} comprehensively considers different opinions and gives back reliable predictions. Overall, the key components of our network are tightly incorporated and collaboratively devote to remarkable segmentation results.




Figure 3: Generated results of Image Alignment Module, where MRI structures in synthetic CT images are well-preserved with clear edges.

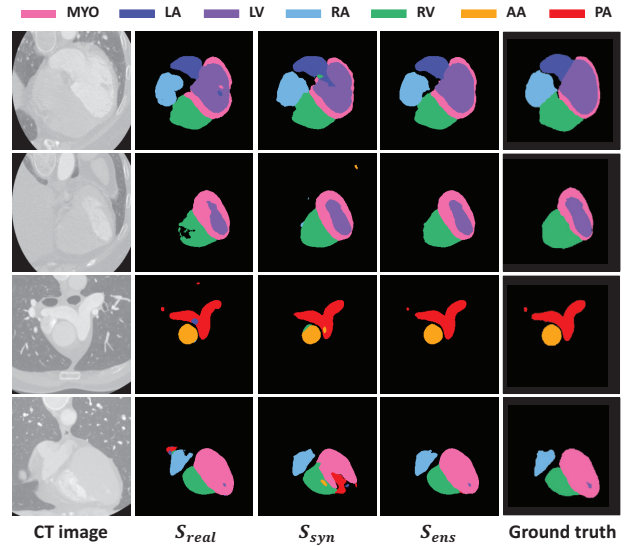


Figure 4: Visual comparison of segmentation results by S_{real} , S_{syn} and the ensemble segmentor, which could balance inconsistent predictions and produce reliable results.

Effect of different amounts of assistant-modality data. We further investigate the effects of different amounts of

Table 4: Quantitative comparison between our approach and other multi-modality segmentation methods in whole heart segmentation. Here we take MR as target modality and CT as assistant modality. The dice of all cardiac substructures and mean dice are reported.

Method	Mean Dice	Dice of substructures of heart						
		MYO	LA	LV	RA	RV	AA	PA
Baseline	0.8243	0.8259	0.8153	0.9347	0.8146	0.9140	0.7365	0.7288
Fine-tune	0.8217	0.8336	0.8447	0.9371	0.7893	0.8940	0.7015	0.7516
Joint-training	0.8306	0.8375	0.7893	0.9303	0.8507	0.9119	0.7487	0.7455
X-shape (2018)	0.8337	0.8321	0.8247	0.9347	0.8450	0.9181	0.7278	0.7536
Jiang et al. (2018)	0.8370	0.8309	0.8569	0.9356	0.8582	0.9016	0.7334	0.7422
Zhang et al. (2018b)	0.8375	0.8360	0.8169	0.9317	0.8501	0.9183	0.7435	0.7663
Ours (S_{syn})	0.8382	0.8405	0.8308	0.9337	0.8655	0.9136	0.7351	0.7483
Ours (S_{real})	0.8466	0.8507	0.8582	0.9427	0.8623	0.9091	0.7398	0.7633
Ours	0.8517	0.8525	0.8695	0.9437	0.8655	0.9121	0.7544	0.7640

assistant-modality data in the network training. We conduct experiments with 5, 10 and 20 MRI volumes as assistant-modality data in network training and quantitatively evaluate these settings on CT segmentation in Table 3. Like the previous, we use the same CT volumes and perform two-fold cross-validation, and randomly pick the corresponding number of MRI volumes as assisted modality data. Generally, it is observed that the performance of CT segmentation is enhanced with the increase of MRI data, as more MRI data implies more assisted prior knowledge.

Extensive experiments with CT as assistant modality.

To demonstrate the benefit of our proposed mutual knowledge distillation is bidirectional, we conduct another experiments to take CT as assistant modality, while MR as target modality. All substructure dice and their mean dice are reported in Table 4. Our method outperforms baseline and Zhang et al. (2018b) with 2.74% and 1.42% in mean dice, respectively. Our single segmentor S_{real} also achieves better performance than Zhang et al. (2018b) by 0.91% in average dice. The results show that our method could be bidirectionally applied to both modalities and have no constrain in target modality chosen.

Conclusion

In this paper, we propose a novel multi-modality image segmentation framework to demonstrate the effectiveness of exploiting assistant-modality prior knowledge to enhance the segmentation performance on target modality. In specific, we first design an Image Alignment Module to reduce appearance discrepancy by generating synthetic target-modality images via adversarial learning. More importantly, we present the Mutual Knowledge Distillation scheme to thoroughly exploiting the modality-shared knowledge by explicitly learning one modality knowledge from segmentation annotations and implicitly exploring another modality knowledge with peer guidance. The experimental results on public multi-modality cardiac segmentation dataset verify the effectiveness of our method and demonstrate superior performance against other state-of-the-art methods.

Acknowledgements

This work described in this paper was supported by the following grants from the Hong Kong Innovation and Technology Commission (Project No. ITS/426/17FP) and the National Natural Science Foundation of China (Project No. U1813204).

References

- Cao, X.; Yang, J.; Gao, Y.; Guo, Y.; Wu, G.; and Shen, D. 2017. Dual-core steered non-rigid registration for multi-modal images via bi-directional image synthesis. *Medical image analysis* 41:18–31.
- Chen, C.; Dou, Q.; Chen, H.; Qin, J.; and Heng, P.-A. 2019. Synergistic image and feature adaptation: Towards cross-modality domain adaptation for medical image segmentation. *arXiv preprint arXiv:1901.08211*.
- Fidon, L.; Li, W.; Garcia-Peraza-Herrera, L. C.; Ekanayake, J.; Kitchen, N.; Ourselin, S.; and Vercauteren, T. 2017. Scalable multimodal convolutional networks for brain tumour segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 285–293. Springer.
- Furlanello, T.; Lipton, Z. C.; Tschannen, M.; Itti, L.; and Anandkumar, A. 2018. Born again neural networks. *arXiv preprint arXiv:1805.04770*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.
- Guo, Z.; Li, X.; Huang, H.; Guo, N.; and Li, Q. 2018. Medical image segmentation based on multi-modal convolutional neural network: Study on image fusion schemes. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, 903–907. IEEE.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Identity mappings in deep residual networks. In *European conference on computer vision*, 630–645. Springer.

- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Jiang, J.; Hu, Y.-C.; Tyagi, N.; Zhang, P.; Rimner, A.; Mageras, G. S.; Deasy, J. O.; and Veeraraghavan, H. 2018. Tumor-aware, adversarial domain adaptation from ct to mri for lung cancer segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 777–785. Springer.
- Kamnitsas, K.; Ledig, C.; Newcombe, V. F.; Simpson, J. P.; Kane, A. D.; Menon, D. K.; Rueckert, D.; and Glocker, B. 2017. Efficient multi-scale 3d cnn with fully connected crf for accurate brain lesion segmentation. *Medical image analysis* 36:61–78.
- Litjens, G.; Kooi, T.; Bejnordi, B. E.; Setio, A. A. A.; Ciompi, F.; Ghafoorian, M.; Van Der Laak, J. A.; Van Ginneken, B.; and Sánchez, C. I. 2017. A survey on deep learning in medical image analysis. *Medical image analysis* 42:60–88.
- Moeskops, P.; Wolterink, J. M.; van der Velden, B. H.; Gilhuijs, K. G.; Leiner, T.; Viergever, M. A.; and Išgum, I. 2016. Deep learning for multi-task medical image segmentation in multiple modalities. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 478–486. Springer.
- Nie, D.; Wang, L.; Gao, Y.; and Shen, D. 2016. Fully convolutional networks for multi-modality isointense infant brain image segmentation. In *2016 IEEE 13Th international symposium on biomedical imaging (ISBI)*, 1342–1345. IEEE.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241. Springer.
- Valindria, V. V.; Pawlowski, N.; Rajchl, M.; Lavdas, I.; Aboagye, E. O.; Rockall, A. G.; Rueckert, D.; and Glocker, B. 2018. Multi-modal learning from unpaired images: Application to multi-organ segmentation in ct and mri. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 547–556. IEEE.
- Woo, J.; Stone, M.; and Prince, J. L. 2014. Multimodal registration via mutual information incorporating geometric and spatial context. *IEEE Transactions on Image Processing* 24(2):757–769.
- Wu, R.; Feng, M.; Guan, W.; Wang, D.; Lu, H.; and Ding, E. 2019a. A mutual learning method for salient object detection with intertwined multi-supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8150–8159.
- Wu, S.; Li, J.; Liu, C.; Yu, Z.; and Wong, H.-S. 2019b. Mutual learning of complementary networks via residual correction for improving semi-supervised classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6500–6509.
- Zagoruyko, S., and Komodakis, N. 2017. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *ICLR*.
- Zhang, Y.; Xiang, T.; Hospedales, T. M.; and Lu, H. 2018a. Deep mutual learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4320–4328.
- Zhang, Z.; Yang, L.; and Zheng, Y. 2018b. Translating and segmenting multimodal medical volumes with cycle- and shape-consistency generative adversarial network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9242–9251.
- Zhao, A.; Balakrishnan, G.; Durand, F.; Guttag, J. V.; and Dalca, A. V. 2019. Data augmentation using learned transformations for one-shot medical image segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8543–8553.
- Zheng, Y. 2015. Cross-modality medical image detection and segmentation by transfer learning of shapel priors. In *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*, 424–427. IEEE.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, 2223–2232.
- Zhuang, X.; Li, L.; Payer, C.; Stern, D.; Urschler, M.; Heinrich, M. P.; Oster, J.; Wang, C.; Smedby, O.; Bian, C.; et al. 2019. Evaluation of algorithms for multi-modality whole heart segmentation: An open-access grand challenge. *arXiv preprint arXiv:1902.07880*.