Published in final edited form as:

Med Image Comput Comput Assist Interv. 2019 October ; 11765: 255–263.

doi:10.1007/978-3-030-32245-8_29.

Unsupervised Domain Adaptation via Disentangled Representations: Application to Cross-Modality Liver Segmentation

Junlin Yang¹, Nicha C. Dvornek³, Fan Zhang², Julius Chapiro³, MingDe Lin³, James S. Duncan^{1,2,3,4}

¹Department of Biomedical Engineering, Yale University, New Haven, CT, USA

²Department of Electrical Engineering, Yale University, New Haven, CT, USA

³Department of Radiology and Biomedical Imaging, Yale School of Medicine, New Haven, CT, USA

⁴Department of Statistics and Data Science, Yale University, New Haven, CT, USA

Abstract

A deep learning model trained on some labeled data from a certain source domain generally performs poorly on data from different target domains due to domain shifts. Unsupervised domain adaptation methods address this problem by alleviating the domain shift between the labeled source data and the unlabeled target data. In this work, we achieve cross-modality domain adaptation, i.e. between CT and MRI images, via disentangled representations. Compared to learning a one-to-one mapping as the state-of-art CycleGAN, our model recovers a many-to-many mapping between domains to capture the complex cross-domain relations. It preserves semantic feature-level information by finding a shared content space instead of a direct pixelwise style transfer. Domain adaptation is achieved in two steps. First, images from each domain are embedded into two spaces, a shared domain-invariant content space and a domain-specific style space. Next, the representation in the content space is extracted to perform a task. We validated our method on a cross-modality liver segmentation task, to train a liver segmentation model on CT images that also performs well on MRI. Our method achieved Dice Similarity Coefficient (DSC) of 0.81, outperforming a CycleGAN-based method of 0.72. Moreover, our model achieved good generalization to joint-domain learning, in which unpaired data from different modalities are jointly learned to improve the segmentation performance on each individual modality. Lastly, under a multi-modal target domain with significant diversity, our approach exhibited the potential for diverse image generation and remained effective with DSC of 0.74 on multi-phasic MRI while the CycleGAN-based method performed poorly with a DSC of only 0.52.

1 Introduction

Deep neural networks have been very successful in a variety of computer vision tasks, including medical image analysis. The majority of neural networks conduct training and

evaluation on images from the same distribution. However, real-world applications usually face varying visual domains. The distribution differences between training and test data, i.e. domain shifts, can lead to significant performance degradation. Data collection and manual annotation for every new task and domain are time-consuming and expensive, especially for medical imaging, where data are limited and are collected from different scanners, protocols, sites, and modalities. To solve this problem, domain adaptation algorithms look to build a model from a source data distribution that performs well on a different but related target data distribution [12]. In the context of medical image analysis, most prior studies on domain adaptation focus on aligning distributions of data from different scan protocols, scanners, and sites [4,9,11]. Related literature is relatively limited when it comes to different modalities.

In clinical practice, various imaging modalities may have valuable and complementary roles. For example, as a fast, less expensive, robust, and readily available modality, computed tomography (CT) plays a key role in the routine clinical examination of hepatocellular carcinoma (HCC), but has the disadvantages of radiation exposure and relatively low soft-tissue contrast. Magnetic resonance imaging (MRI) provides higher soft-tissue contrast for lesion detection and characterization but is more expensive, time-consuming, less robust, and more prone to artifacts. In practice, both multiphase contrast-enhanced MRI and CT may be used in the diagnosis and follow-up after treatment of HCC, and often require the same image analysis tasks, such as liver segmentation. While MRI acquisitions include more complex quantitative information than CT useful for liver segmentation, they are often less available clinically than CT images [8]. Thus, it would be helpful if we could learn a liver segmentation model from the more accessible CT data that also performs well on MRI images.

Given the significant domain shift, cross-modality domain adaptation is quite difficult (see Fig. 1). One promising approach utilizes CycleGAN, a pixel-wise style transfer model, for cross-modality domain adaptation in a segmentation task [3]. Compared to feature-based domain adaptation, it does not necessarily maintain the semantic feature-level information. More importantly, the cycle-consistency loss implies a one-to-one mapping between source domain and target domain and leads to lack of translated output diversity, generating very similar images. It thus fails to represent the complex real-world data distribution in the target domain and likely degrades the performance of segmentation or other follow-up analysis [5].

Our goal is to achieve domain adaptation between CT and MRI while maintaining the complex relationship between the two domains. Our model assumes the mapping to be many-to-many and learns it by disentangling the representation into content and style. A shared latent space is assumed to be found for both domains that preserves the semantic content information. Our main contributions are listed as follows. This is the first to achieve unsupervised domain adaptation for segmentation via disentangled representations in the field of medical imaging. Our model decomposes images across domains into a domain-invariant content space, which preserves the anatomical information, and a domain-specific style space, which represents modality information. We validated the superior performance of our model on a liver segmentation task with cross-modality domain adaptation and compared it to the state-of-art CycleGAN. We also demonstrated the generalizability of our

model to joint-domain learning and robust adaptation to a multi-modal target domain with large variety.

2 Methodology

2.1 Assumptions

Let $x_1 \in \chi_1$ and $x_2 \in \chi_2$ be images from two domains, which differ in visual appearance but share common semantic content. We assume there exists a mapping, potentially many-to-many instead of deterministic one-to-one, between χ_1 and χ_2 . Each image x_i from χ_i can be embedded into and generated from a shared semantic content space $c \in \mathcal{C}$ that is domain-invariant and a style code $s_i \in \mathcal{S}_i$ that is domain-specific ($i = 1, 2$) [2]. Specifically, MRI and CT of the abdomen from HCC patients can be considered as images from different domains χ_1 and χ_2 , since they exhibit quite different visual appearance with the same anatomical structure shared behind them. Therefore, a shared domain-invariant space that preserves the anatomical information and a domain-specific style code for each modality can be found to recover the underlying mapping between MRI and CT. Due to the many-to-many assumption, the relatively complex underlying distribution of the target domain can be recovered.

2.2 Model

Our Domain Adaptation via Disentangled Representations (DADR) pipeline consists of two modules: Disentangled Representation Learning Module (DRLModule) and Segmentation Module (SegModule) (see Fig. 2). Of note, the DRL Module box in the DADR pipeline at the top is expanded in the left large box.

DRLModule.—The module consists of two main components, a variational autoencoder (VAE) for reconstruction and a generative adversarial network (GAN) for adversarial training. We train the VAE component for in-domain reconstruction, where reconstruction loss is minimized to encourage the encoders and generators to be inverses to each other. The GAN component for cross-domain translation is trained to encourage the disentanglement of the latent space, decomposing it into content and style subspaces [6,7]. Similar to Huang's work [2], the DRLModule consists of several jointly trained encoders $E_{c1}, E_{c2}, E_{s1}, E_{s2}$, generators G_1, G_2 and discriminators D_1, D_2 , where $c_i = E_{c_i}(x_i) \sim p(c_i)$ and $s_i = E_{s_i}(x_i) \sim p(s_i)$ for $i = 1, 2$. Specifically, the generators are trying to fool the discriminators by successful cross-domain generation with swapped style code. Due to the disentangled style code $s_i \in \mathcal{S}_i$, the underlying mapping is assumed to be many-to-many. We have $p(c_1) = p(c_2)$ upon convergence, which is the shared content space that preserves anatomical information. The overall loss function is defined as the weighted sum of the three components:

$$L_{total} = \alpha L_{recon} + \beta L_{adv} + \gamma L_{latent} \quad (1)$$

a. In-domain reconstruction, $L_{recon} = L_{recon}^1 + L_{recon}^2$

$$L_{recon}^i = \mathbb{E}_{x_i \sim X_i} \|G_i(E_{ci}(x_i), E_{si}(x_i)) - x_i\|_1 \quad (2)$$

- b. Cross-domain translation, $L_{adv} = L_{adv}^{1 \rightarrow 2} + L_{adv}^{2 \rightarrow 1}$

$$L_{adv}^{1 \rightarrow 2} = \mathbb{E}_{c_1 \sim p(c_1), s_2 \sim p(s_2)} [\log(1 - D_2(x_1 \rightarrow 2))] + \mathbb{E}_{x_2 \sim X_2} [\log(D_2(x_2))] \quad (3)$$

- c. Latent space reconstruction, $L_{latent} = L_{recon}^{c_1} + L_{recon}^{s_1} + L_{recon}^{c_2} + L_{recon}^{s_2}$

$$L_{recon}^{c_1} = \|E_{c_2}(x_1 \rightarrow 2) - c_1\|_1, L_{recon}^{s_2} = \|E_{s_2}(x_1 \rightarrow 2) - s_2\|_1 \quad (4)$$

Domain Adaptation with Content-Only Images.—Once the disentangled representations are learned, content-only images can be reconstructed by using the content code c_i without style code s_i . For both CT and MR, their content codes are embedded in a shared latent space that incorporates the anatomical structure information and excludes the modality appearance information. We train a segmentation model on content-only images from CT domain and apply it directly on content-only images from MR domain.

Joint-Domain Learning.—Joint-domain learning aims to train a single model with data from both domains that works on both domains and outperforms models trained and tested separately on each domain. Our framework can easily generalize to joint-domain learning by including content-only images from both domains for the training segmentation module.

Implementation Details.—The SegModule is a standard UNet [10]. Content encoders consist of convolutional layers and residual layers followed by batch normalization, while style encoders consist of convolutional layers, a global average pooling layer, and a fully-connected layer. Generators take the style code (vector of length 8) and content code (feature map of $64 \times 64 \times 256$) as inputs. A multilayer perceptron takes the style code and generates affine transformation parameters. Residual blocks in the generator are equipped with an Adaptive Instance Normalization (AdaIN) layer to take affine transformation parameters from the style code. Discriminators are convolutional neural networks for binary classification. For the loss function, $\alpha = 25$, $\beta = 10$, $\gamma = 0.1$ in our experiments. Experiments were conducted on two Nvidia 1080ti GPUs. The training time each fold is ~5h for DRLModule and ~2h for SegModule. Testing is very quick.

3 Experiments and Results

3.1 Datasets and Experimental Setup

We tested our methods on unpaired CT slices of 130 patients from LiTS challenge 2017 [1] and multi-phasic MRI slices of 20 local patients with HCC (see Fig. 1). CT and MR were divided into 5 folds for subject-wise cross-validation. A supervised UNet [10] trained and tested on pre-contrast MR serves as upper bound of domain adaptation, while a supervised UNet trained on CT and tested on pre-contrast MRI, without domain adaptation, provides

the lower bound, which shows the relatively large domain shifts between CT and MRI (see Table 2).

For our DADR model, in experiment 1, 4 folds of CT and 4 folds of pre-contrast MR were used to train DRLModule, 4 folds of CT were used to train SegModule and 1 fold of pre-contrast MR were used to test SegModule. In experiment 2, it was the same as experiment 1, but 4 folds of pre-contrast MR were also used to train SegModule. In experiment 3, it was the same as experiment 1, except that pre-contrast MR were replaced with multi-phasic MR.

3.2 Results

Experiment 1: Segmentation with Domain Adaptation.—We evaluated Domain Adaptation with CycleGAN (DACGAN) and Domain Adaptation via Disentangled Representations (DADR) respectively and compared with UNet without Domain Adaptation (UNet w/o DA) (see Fig. 5 and Table 1).

DACGAN. For comparison purposes, we trained a CycleGAN model with unpaired CT and pre-contrast MR images and performed style transfer on test CT images to generate synthetic MR images. A standard UNet was trained on synthetic MR images and validated on real precontrast MR images. It achieved a DSC score of 0.72 (see Table 1) with subject-wise 5-fold cross-validation. Figure 3 shows two examples of synthetic MR generated by CycleGAN.

DADR. Our DRLModule embeds cross-domain images into a shared content space and generates content-only images. Figure 4 shows two examples of content-only images via Disentanglement Representations. We trained the UNet model on content-only CT and validated on content-only MR with 5-fold cross-validation and achieved a DSC score of 0.81 (see Table 1).

Experiment 2: Joint-Domain Learning.—Besides domain adaptation, our model can achieve joint-domain learning by feeding UNet with both content-only CT and content-only MR as training data. A single model that works on both CT and MR modality was obtained and outperformed two fully-supervised standard UNet models separately trained on each modality (see Table 2).

Experiment 3: Multi-modal Target Domain.—We considered multi-phasic MR with three phases as multi-modal target domain with complex underlying distribution and conducted domain adaptation and style transfer on it.

A Robust Domain Adaptation. A CycleGAN model was not able to handle such large variety in the multi-modal target domain e.g. multi-phasic MR. However, for our model, the shared content space provided a robust representation for anatomical information in multi-phasic MR; thus it remained effective even faced with the multi-modal target domain (see Fig. 6 and Table 3).

B Diverse Style Transfer. Furthermore, our model can realize diverse style transfer by changing the style code while preserving the anatomical structure with the same content

code. The style code can be randomly sampled in the style space or encoded from a reference image by style encoder E_{s_i} (see Fig. 7).

3.3 Analysis

We tested our model on unpaired CT and MRI data. It is noteworthy that it is a highly unbalanced cross-domain data, where CT is of better quality and the size of the CT dataset is about 6.5 times the size of MR. Experiment 1 shows that our model is superior to CycleGAN in terms of DSC score for cross-modality segmentation with domain adaptation. Experiment 2 shows a promising application of our model for joint-domain learning, which makes learning from unpaired medical images with different modalities a reality. Experiment 3 shows robustness of our model under multi-modal target domain with large diversity and the potential for diverse multi-modal style transfer. The disentangled representation is the key to fulfill the many-to-many mapping assumption and recover the complex relationship between two domains. It is also of vital importance to discover the shared latent space that preserves the semantic feature-level information.

4 Conclusions and Discussions

We proposed a cross-modality domain adaptation pipeline via disentangled representations, which may improve current clinical workflows and allow for robust intergration of multi-parametric MRI and CT data. Instead of one-to-one mapping, our model considers the complex mapping between CT and MR as many-to-many and preserves semantic feature-level information, thus ensuring robust cross-modality domain adaptation for a segmentation task. We validated and compared our model on CT and pre-phase MR from HCC patients with state-of-the-art methods. With multi-phasic MRI, our model demonstrated strong ability to handle multi-modal target domains with large variety. Furthermore, our model had good generalization towards joint-domain learning and showed the potential for multi-modal image generation, which can be further investigated in the future. In addition, we could focus on specific anatomical structures in the content space, such as liver and tumor, by including task-relevant loss to improve the results further.

Acknowledgments

This work was supported by NIH Grant 5R01 CA206180.

References

1. Christ P, Ettlinger F, Grün F, Lipkova J, Kaissis G: LiTS-liver tumor segmentation challenge. ISBI and MICCAI (2017)
2. Huang X, Liu M-Y, Belongie S, Kautz J: Multimodal unsupervised image-to-image translation In: Ferrari V, Hebert M, Sminchisescu C, Weiss Y (eds.) ECCV 2018. LNCS, vol. 11207, pp. 179–196. Springer, Cham (2018). 10.1007/978-3-030-01219-911
3. Jiang J, et al.: Tumor-aware, adversarial domain adaptation from CT to MRI for lung cancer segmentation In: Frangi AF, Schnabel JA, Davatzikos C, Alberola-López C, Fichtinger G (eds.) MICCAI 2018. LNCS, vol. 11071, pp. 777–785. Springer, Cham (2018). 10.1007/978-3-030-00934-286

4. Kamnitsa K, et al.: Unsupervised domain adaptation in brain lesion segmentation with adversarial networks In: Niethammer Ms., Niethammer M, et al. (eds.) IPMI 2017. LNCS, vol. 10265, pp. 597–609. Springer, Cham (2017). 10.1007/978-3-319-59050-9_47
5. Lee H-Y, Tseng H-Y, Huang J-B, Singh M, Yang M-H: Diverse image-to-image translation via disentangled representations In: Ferrari V, Hebert M, Sminchisescu C, Weiss Y (eds.) ECCV 2018. LNCS, vol. 11205, pp. 36–52. Springer, Cham (2018). 10.1007/978-3-030-01246-5_3
6. Mathieu MF, Zhao JJ, Zhao J, Ramesh A, Sprechmann P, LeCun Y: Disentangling factors of variation in deep representation using adversarial training. In: Advances in Neural Information Processing Systems, pp. 5040–5048 (2016)
7. Narayanaswamy S, et al.: Learning disentangled representations with semi-supervised deep generative models. In: Advances in Neural Information Processing Systems, pp. 5925–5935 (2017)
8. Oliva MR, Saini S: Liver cancer imaging: role of CT, MRI, US and PET. Cancer Imaging 4(Spec No A), S42 (2004) [PubMed: 18215974]
9. Perone CS, Cohen-Adad J: Promises and limitations of deep learning for medical image segmentation. J. Med. Artif. Intell 2 (2019)
10. Ronneberger O, Fischer P, Brox T: U-Net: convolutional networks for biomedical image segmentation In: Navab N, Hornegger J, Wells WM, Frangi AF (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). 10.1007/978-3-319-24574-2_28
11. Valindria VV, et al.: Domain adaptation for MRI organ segmentation using reverse classification accuracy. arXiv preprint arXiv:1806.00363 (2018)
12. Wang M, Deng W: Deep visual domain adaptation: a survey. Neurocomputing 312, 135–153 (2018)

Fig. 1. Images and histograms of liver (yellow) and whole image (blue). From left to right: CT, multiphasic MRI sequence at three time points (pre-contrast, 20 s post-contrast i.e. arterial phase, 70 s post-contrast i.e. portal venous phase) (Color figure online)

Fig. 2. Left: Framework for Disentangled Representation Learning Module. Solid line: in-domain reconstruction, Dotted line: cross-domain translation. Right: Pipeline of Domain Adaptation via Disentangled Representations (DADR)

Fig. 3.

Two examples of style transfer with CycleGAN, Left to right: CT, generated MR, MR

Fig. 4.
Two examples of content-only images via disentangled representations, Left to right: CT, content-only CT, content-only MR, MR

Fig. 5.

Two examples of segmentation results. Left to right: pre-contrast MR, liver mask, predictions from UNet w/o DA, DACGAN, DADR

Fig. 6.

Two examples of content-only images via disentangled representations for multi-phasic MR.
first row: original images, second row: content-only images.

Fig. 7.

Three examples of multi-modal style generation with reference, from left to right: CT, three pairs of reference MR and generated MR (pre-contrast, arterial, portal venous phase)

Table 1.

Comparison of Segmentation Results

Method	DICE	std
UNet w/o DA	0.26	0.07
DACGAN	0.72	0.05
DADR	0.81	0.03

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2.

Joint-domain Learning with DADR

Method	CT tested DSC	MR tested DSC
CT trained	0.901(0.020)	0.260(0.072)
MR trained	0.134(0.091)	0.869(0.044)
Joint CT & MR	0.912(0.012)	0.891(0.040)

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 3.

DA on Multimodal Target Domain

Method	DICE	std
DACGAN	0.52	0.06
DADR	0.74	0.04

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript