



HHS Public Access

Author manuscript

Med Image Comput Assist Interv. Author manuscript; available in PMC 2019 September 01.

Published in final edited form as:

Med Image Comput Assist Interv. 2018 September ; 11071: 777–785. doi:
10.1007/978-3-030-00934-2_86.

Tumor-aware, Adversarial Domain Adaptation from CT to MRI for Lung Cancer Segmentation

Jue Jiang¹, Yu-Chi Hu¹, Neelam Tyagi¹, Pengpeng Zhang¹, Andreas Rimner², Gig S. Mageras¹, Joseph O. Deasy¹, and Harini Veeraraghavan^{#1,*}

¹Medical Physics, Memorial Sloan Kettering Cancer Center

²Radiation Oncology, Memorial Sloan Kettering Cancer Center

These authors contributed equally to this work.

Abstract

We present an adversarial domain adaptation based deep learning approach for automatic tumor segmentation from T2-weighted MRI. Our approach is composed of two steps: (i) a tumor-aware unsupervised cross-domain adaptation (CT to MRI), followed by (ii) semi-supervised tumor segmentation using Unet trained with synthesized and limited number of original MRIs. We introduced a novel target specific loss, called tumor-aware loss, for unsupervised cross-domain adaptation that helps to preserve tumors on synthesized MRIs produced from CT images. In comparison, state-of-the art adversarial networks trained without our tumor-aware loss produced MRIs with ill-preserved or missing tumors. All networks were trained using labeled CT images from 377 patients with non-small cell lung cancer obtained from the Cancer Imaging Archive and unlabeled T2w MRIs from a completely unrelated cohort of 6 patients with pre-treatment and 36 on-treatment scans. Next, we combined 6 labeled pre-treatment MRI scans with the synthesized MRIs to boost tumor segmentation accuracy through semi-supervised learning. Semi-supervised training of cycle-GAN produced a segmentation accuracy of 0.66 computed using Dice Score Coefficient (DSC). Our method trained with only synthesized MRIs produced an accuracy of 0.74 while the same method trained in semi-supervised setting produced the best accuracy of 0.80 on test. Our results show that tumor-aware adversarial domain adaptation helps to achieve reasonably accurate cancer segmentation from limited MRI data by leveraging large CT datasets.

1 Introduction

MRI-guided radiotherapy is an emerging technology for improving treatment accuracy over conventional CT-based radiotherapy due to better soft-tissue contrast in MR compared to CT images. Real-time and accurate tumor segmentation on MRI can help to deliver high dose to tumors while reducing normal tissue dose. However, as MRI-guided radiotherapy is not used in standard-of-care, only very few MRIs are available for training. Therefore, we developed an adversarial domain adaptation from large CT datasets for tumor segmentation on MRI.

* veerarah@mskcc.org.

Although deep neural networks excel in learning from large amounts of (labeled) data, their accuracy is reduced when applied to novel datasets or domains[1]. Differences between source and target domain distribution is called *domain shift*. Typically used fine-tuning methods require prohibitively large labeled data in the target domain. As an alternative, domain adaptation methods attempt to minimize domain shift either by feature sharing[2] or by learning to reconstruct the target from source domain[3, 4]. In essence, domain adaptation methods learn the marginal distributions [5] to transform source to target domain.

The problems of domain shift are exacerbated in medical images, where imaging modalities capture physical properties of the underlying anatomy differently (eg. CT vs. MRI). For example, whereas bones appear hyper-dense on CT and dark on MRI, tumors appear with similar contrast as normal soft-tissue on CT but have a distinct appearance on MRI (Fig. 1 (a) and (b)). Consequently, learning the marginal distributions of the domains alone may not be sufficient. Cross-domain adaptation of highly different modalities, has been applied in medical image analysis for image synthesis using paired images [6] and unpaired images [7], as well as for segmentation [8, 9]. However, all aforementioned approaches aim to only synthesize images that match the marginal but not the structure-specific conditional distribution such as tumors. Therefore, segmentation/classification using such synthetic images will lead to lower accuracy.

Therefore, we introduced a novel target specific loss, called tumor-aware loss, for unsupervised cross-domain adaptation that helps to preserve tumors on synthesized MRIs produced from CT images (Fig. 1(d)), which cannot be captured with just the cycle-loss (Fig. 1(c)).

2 Method

Our objective is to solve the problem of learning to segment tumors from MR images through domain adaptation from CT to MRI, where we have access to a reasonably sized labeled data in the source domain (X_{CT}, y_{CT}) but are provided with very limited number of target samples $X_{MRI} \ll X_{CT}$ and fewer labels y_{MR} . Our solution first employs tumor-aware unsupervised cross-domain adaptation to synthesize a reasonably large number of MRI from CT through adversarial training. Second, we combine the synthesized MRI with a fraction of real MRI with corresponding labels and train a Unet[10] for generating tumor segmentation as outlined in Fig. 2.

2.1 Step 1: MRI synthesis using tumor-aware unsupervised cross domain adaptation

The first step is to learn a mapping $G_{CT \rightarrow MRI}$ that synthesizes MRI from the CT images to fool a discriminator D_{MRI} using adversarial training [11]. Additionally, we compute an adversarial loss L_{adv}^{CT} for synthesizing CT from MRI by simultaneously training a network that learns a mapping $G_{MRI \rightarrow CT}$. The adversarial loss, L_{adv}^{MRI} , for synthesizing MRI from CT, and L_{adv}^{CT} , for synthesizing CT from MRI, are computed as:

$$\begin{aligned}
L_{adv}^{MRI}(G_{CT \rightarrow MRI}, D_{MRI}, X_{MRI}, X_{CT}) &= \mathbb{E}_{x_m \sim X_{MRI}} [\log(D_{MRI}(x_m))] \\
&+ \mathbb{E}_{x_c \sim X_{CT}} [\log(1 - (D_{MRI}(G_{CT \rightarrow MRI}(x_c))) \\
L_{adv}^{CT}(G_{MRI \rightarrow CT}, D_{CT}, X_{CT}, X_{MRI}) &= \mathbb{E}_{x_c \sim X_{CT}} [\log(D_{CT}(x_c))] \\
&+ \mathbb{E}_{x_m \sim X_{MRI}} [\log(1 - (D_{CT}(G_{MRI \rightarrow CT}(x_m)))
\end{aligned} \tag{1}$$

where x_c and x_m are real images sampled from the CT (X_{CT}) and MRI (X_{MRI}) domains, respectively. The total adversarial loss (Fig. 2 (purple ellipse)) is then computed as the summation of the two losses as $L_{adv} = L_{adv}^{MRI} + L_{adv}^{CT}$. We also compute a cycle consistency loss [5] to regularize the images synthesized through independent training of the two networks. By letting the synthesized images be $x'_m = G_{CT \rightarrow MRI}(x_c)$ and $x'_c = G_{MRI \rightarrow CT}(x_m)$, the cycle consistency loss L_{cyc} is calculated as:

$$\begin{aligned}
L_{cyc}(G_{CT \rightarrow MRI}, G_{MRI \rightarrow CT}, X_{CT}, X_{MRI}) &= \mathbb{E}_{x_c \sim X_{CT}} [\|G_{MRI \rightarrow CT}(x'_m) - x_c\|_1] \\
&+ \mathbb{E}_{x_m \sim X_{MRI}} [\|G_{CT \rightarrow MRI}(x'_c) - x_m\|_1].
\end{aligned} \tag{2}$$

The cycle consistency and adversarial loss only constrain the model to learn a global mapping that matches the marginal distribution but not the conditional distribution pertaining to individual structures such as the tumors. Therefore, a model trained using these losses does not need to preserve tumors, which can lead to either deterioration or total loss of tumors in the synthesized MRIs (Fig. 1(c)). Therefore, we introduced a tumor-aware loss that forces the network to preserve the tumors. To be specific, the tumor-aware loss is composed of a tumor loss (Fig. 2 (red ellipse)) and a feature loss (Fig. 2 (orange ellipse)). We compute the tumor loss by training two parallel tumor detection networks using simplified models of the Unet [10] for CT (U_{CT}) and the synthesized MRI (U_{MRI}). The tumor loss constrains the CT and synthetic MRI-based Unets to produce similar tumor segmentations, thereby, preserving the tumors and is computed as:

$$\begin{aligned}
L_{tumor} &= \mathbb{E}_{x_c \sim X_{CT}, y_c \sim y_{CT}} [\log P(y_c | G_{CT \rightarrow MRI}(x_c))] + \\
&\mathbb{E}_{x_c \sim X_{CT}, y_c \sim y_{CT}} [\log P(y_c | X_{CT})].
\end{aligned} \tag{3}$$

On the other hand, the tumor feature loss L_{feat} forces the high-level features of X_{CT} and X_{CT}^{MRI} to be shared by using a constraint inspired by [12] as:

$$L_{feat}(x_c \sim X_{CT}) = \frac{1}{C \times H \times W} \left\| \phi_{CT}(x_c) - \phi_{MRI}(G_{CT \rightarrow MRI}(x_c)) \right\|^2. \tag{4}$$

where ϕ_{CT} and ϕ_{MRI} are the high-level features extracted from the U_{CT} and U_{MRI} respectively; C, H and W indicate the size of the feature. The total loss is then expressed as:

$$L_{total} = L_{adv} + \lambda_{cyc}L_{cyc} + \lambda_{tumor}L_{tumor} + \lambda_{feat}L_{feat} \quad (5)$$

Where λ_{cyc} , λ_{tumor} and λ_{feat} are the weighting coefficients for each lose. During training, we alternatively updated the domaine tranfer or generator network G, the discriminator D, and the tumor constraint network U with the following gradients,

$$-\Delta_{\theta_G}(L_{adv} + \lambda_{cyc}L_{cyc} + \lambda_{tumor}L_{tumor} + \lambda_{feat}L_{feat}), -\Delta_{\theta_D}(L_{adv}) \text{ and} \\ -\Delta_{\theta_U}(L_{tumor} + \lambda_{feat}L_{feat})$$

2.2 Step 2: Semi-supervised tumor segmentation from MRI

The synthesized MRI from the first step were combined with a small set of real MRI with labels (\tilde{X}_{MRI} and \tilde{y}_{MRI} in Fig. 2) to train a U-net [10] using Dice loss [13] (Fig. 2 (blue ellipse)) to generate tumor segmentation. Adversarial network optimization for MRI synthesis was frozen prior to semi-supervised tumor segmentation training to prevent leakage of MRI label information.

2.3 Network Structure and Implementation

The generators G and discriminators D for CT and MRI synthesis networks were implemented similar to that in [5]. We tied the penultimate layer in U_{MRI} and U_{CT} . The details of all networks are shown in the supplementary documents. Pytorch library [14] was used for implementing the proposed networks, which were trained on Nvidia GTX 1080Ti of 12 GB memory with a batch size of 1 during image transfer and batch size of 10 during semi-supervised segmentation. The ADAM algorithm [15] with an initial learning rate of 1e-4 was used during training. We set $\lambda_{cyc}=10$, $\lambda_{tumor}=5$ and $\lambda_{feat}=1$.

3 Experiments and Results

3.1 Ablation tests

We tested the impact of adding tumor-aware loss to the cycle loss (proposed vs. cycle-GAN [5] vs. masked-cycle-GAN [8]). Images synthesized using aforementioned networks were trained to segment using semi-supervised learning by combining with a limited number of real MRI. We call adversarial synthesis [8] that combined tumor labels as an additional channel with the original images as masked-cycle-GAN. We also evaluated the effect of adding a limited number of original MRI to the synthesized MRI on segmentation accuracy (tumor-aware with semi-supervised vs. tumor-aware with unsupervised training). We benchmarked the lowest achievable segmentation accuracy by training a network with only the pre-treatment (or week one) MRI.

3.2 Datasets

The image synthesis networks were trained using contrast-enhanced CT images with expert delineated tumors from 377 patients with non-small cell lung cancer (NSCLC) [16] available from The Cancer Imaging Archive (TCIA) [17], and an unrelated cohort of 6 patients scanned with T2w MRI at our clinic before and during treatment every week (n=7) with radiation therapy. Masked cycle-GANs used both tumor labels and the images as additional channels even for image synthesis training. Image regions enclosing the tumors were extracted and rescaled to 256×256 to produce 32000 CT image slices and 9696 T2w MR image slices. Only 1536 MR images from pre-treatment MRI were used for semi-supervised segmentation training of all networks. Segmentation validation was performed on the subsequent on-treatment MRIs (n=36) from the same 6 patients. Test was performed using 28 MRIs consisting of longitudinal scans (7,7,6) from 3 patients and pre-treatment scans from 8 patients not used in training. Tumor segmentation accuracy was evaluated by comparing to expert delineations using the Dice Score Coefficient (DSC), and the Hausdorff Distance 95%(HD95).

3.3 MR image synthesis results

Fig. 3 shows the representative qualitative results of synthesized MRI produced using only the cycle-GAN (Fig. 3(b)), masked cycle-GAN (Fig. 3(c)) and using our method (Fig. 3(d)). As seen, our method best preserves the anatomical details between CT and MRI. Quantitative evaluation using the Kullback - Leibler (KL) divergence computed from tumor regions between synthesized and original MRI, used for training, confirmed that our method resulted in the best match of tumor distribution with the lowest KL divergence of 0.069 compared with those obtained using the cycle-GAN (1.69) and masked cycle-GAN (0.32).

3.4 Segmentation results

Fig. 4 shows the segmentations generated using the various methods (yellow contours) for three representative cases from the test and validation sets, together with the expert delineations (red contours). As shown in Table 1, our approach outperformed cycle GAN irrespective of training without (unsupervised) or with (semi-supervised) labeled target data. Semi-supervised segmentation outperformed all methods in both test and validation datasets.

4 Discussion

In this work, we introduced a novel target-specific, tumor-aware loss for synthesizing MR images from unpaired CT datasets using unsupervised cross-domain adaptation. The tumor-aware loss forces the network to retain tumors that are typically lost when using only the cycle-loss and leads to accurate tumor segmentation. Although applied to lung tumors, our method is applicable to other structures and organs. Segmentation accuracy of our approach trained with only synthesized MRIs exceeded other methods trained in a semi-supervised manner. Adding small set of labeled target domain data further boosts accuracy. The validation set produced lower but not significantly different ($p=0.1$) DSC accuracy than the test set due to significantly smaller ($p=0.0004$) tumor volumes in validation (mean 37.66cc) when compared with the test set (mean 68.2cc). Our results showed that masked-cycle-GAN produced lower test performance compared to basic cycle-GAN, possibly due to poor

modeling from highly unbalanced CT and MR datasets. As a limitation, our approach only forces the synthesized MRIs to preserve tumors but not the MR intensity distribution within tumors. Such modeling would require learning the mapping for individual scan manufacturers, magnet strengths and coil placements which was outside the scope of this work. Additionally, synthesized images irrespective of the chosen method do not produce a one-to-one pixel mapping from CT to MRI similar to [8]. There is also room for improving the segmentation accuracy by exploring more advanced segmentation models, e.g. boundary-aware fully convolutional networks (FCN)[18].

5 Conclusions

In this work, we proposed a tumor-aware, adversarial domain adaptation method using unpaired CT and MR images for generating segmentations from MRI. Our approach preserved tumors on synthesized MRI and generated the best segmentation performance compared with state-of-the-art adversarial cross-domain adaptation. Our results suggest feasibility for lung tumor segmentation from MRI trained using MRI synthesized from CT.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

References

1. Tzeng E, Ho man J, Saenko K, Darrell T: Adversarial discriminative domain adaptation. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, 7 21–26, 2017 (2017) 2962–2971
2. Ganin Y, Lempitsky V: Unsupervised domain adaptation by backpropagation. In: Intl Conf. Machine Learning (ICML). (2015) 1180–1189
3. Shrivastava A, Pfister T, Tuzel O, Susskind J, Wang W, Webb R: Learning from simulated and unsupervised images through adversarial training. In: IEEE Conf. Computer Vision and Pattern Recognition (CVPR). Volume 3 (2017) 2107–2116
4. Yoo D, Kim N, Park S, Paek AS, Kweon IS: Pixel-level domain transfer In: European Conference on Computer Vision (ECCV), Springer (2016) 517–532
5. Zhu JY, Park T, Isola P, Efros A: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Intl. Conf. Computer Vision (ICCV). (2017) 2223–2232
6. Nie D, Trullo R, Lian J, Petitjean C, Ruan S, Wang Q, Shen D: Medical image synthesis with context-aware generative adversarial networks In: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI), Springer (2017) 417–425
7. Wolterink JM, Dinkla AM, Savenije MH, Seevinck PR, AT C, I I: Deep MR to CT synthesis using unpaired data In: Intl Workshop on Simulation and Synthesis in Medical Imaging, Springer (2017) 14–23
8. Chatsias A, Joyce T, Dharmakumar R, Tsaftaris SA: Adversarial image synthesis for unpaired multi-modal cardiac data In: Intl Workshop on Simulation and Synthesis in Medical Imaging, Springer (2017) 3–13
9. Huo Y, Xu Z, Bao S, Assad A, Abramson RG, B-A L: Adversarial synthesis learning enables segmentation without target modality ground truth In: IEEE Intl. Symposium on Biomedical Imaging. (2018)
10. Ronneberger O, Fischer P, Brox T: U-net: Convolutional networks for biomedical image segmentation In: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI), Springer (2015) 234–241

11. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y: Generative adversarial nets. In: Advances in Neural Information Processing Systems (NIPS). (2014) 2672–2680
12. Johnson J, Alahi A, Li FF: Perceptual losses for real-time style transfer and super-resolution In: European Conference on Computer Vision (ECCV), Springer (2016) 694–711
13. Milletari F, Navab N, Ahmadi SA: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 3D Vision (3DV), 2016 Fourth International Conference on, IEEE (2016) 565–571
14. Paszke A, Gross S, Chintala S, Chanan G, Yang E, DeVito Z, Lin ZM, Desmaison A, Antiga L, Lerer A: Automatic differentiation in PyTorch. (2017)
15. Kingma DP, Ba J: Adam: A method for stochastic optimization. Proceedings of the 3rd International Conference on Learning Representations (ICLR) (2014)
16. Aerts H, E., R.V., Leijenaar RT, Parmar C, Grossmann P, Carvalho S, Lambin P: Data from NSCLC-radiomics. The Cancer Imaging Archive (2015)
17. Clark K, Vendt B, Smith K, Freymann J, Kirby J, Koppel P, Moore S, Phillips S, Maffitt D, Pringle M: The cancer imaging archive (TCIA): maintaining and operating a public information repository. *Journal of digital imaging* 26(6) (2013) 1045–1057 [PubMed: 23884657]
18. Shen H, Wang R, Zhang J, McKenna SJ: Boundary-aware fully convolutional network for brain tumor segmentation In: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI), Springer (2017) 433–441




Fig. 1: MRI synthesized from a representative (a) CT image using (c) cycle-GAN[5] and (d) proposed method. The corresponding MRI scan for (a) is shown in (b). As shown, the proposed method (d) using tumor-aware loss helps to fully preserve tumor in the synthesized MRI compared with (c).

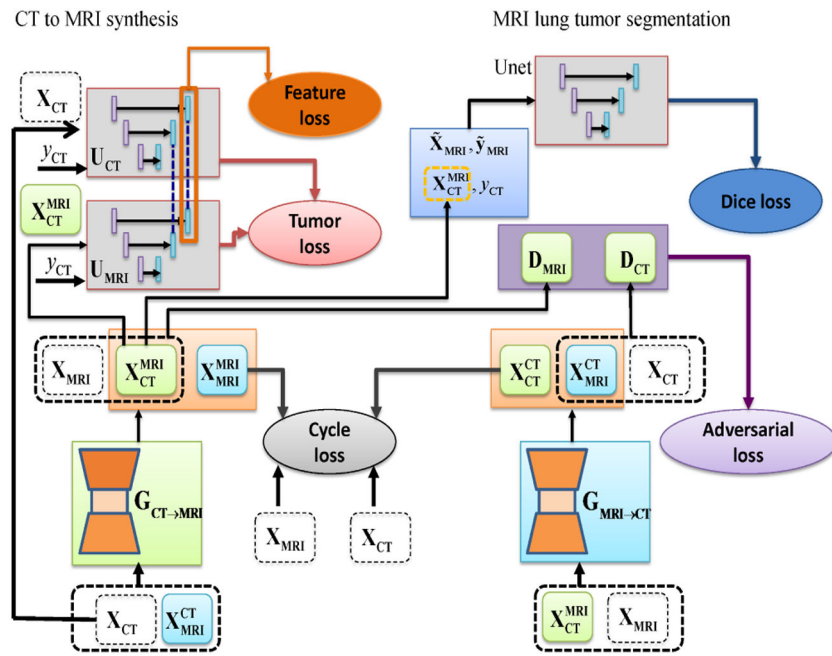


Fig. 2: Approach overview. X_{CT} and X_{MRI} are the real CT and MRI; X_{CT}^{MRI} and X_{MRI}^{CT} are the synthesized MR and CT images; y_{CT} is the CT image label; $G_{CT \rightarrow MRI}$ and $G_{MRI \rightarrow CT}$ are the CT and MRI transfer networks; \tilde{X}_{MRI} and \tilde{y}_{MRI} are a small sample set from the real MRI, used to train semi-supervised segmentation.




Fig. 3: MRI synthesized from CT using different deep learning methods. The red contour indicates the manually delineated tumor region in the NSCLC datasets [16]. (a) CT image; (b) cycle-GAN [5]; (c) Masked cycle-GAN [8]; (d) Proposed.

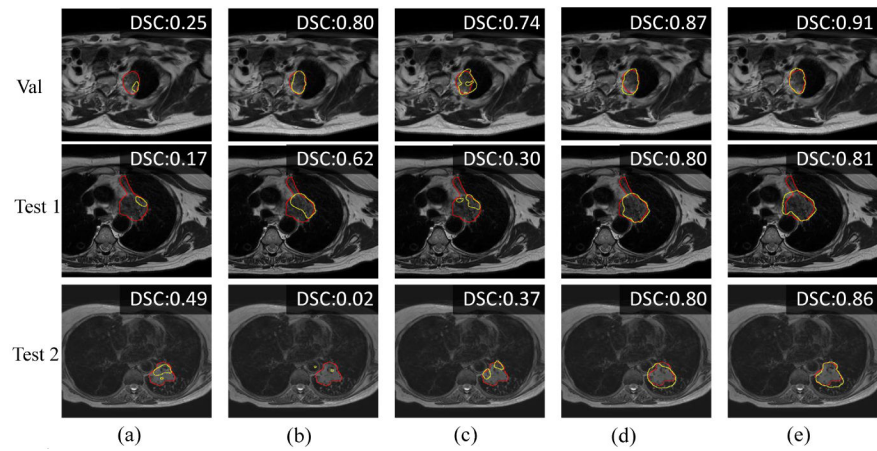


Fig. 4: Segmentation results on the representative examples from the validation and test set of different methods. The red contour stands for the expert delineations and the yellow contour stands for the segmentation results. (a) segmentation with only week 1 MRI; (b) segmentation using MRI synthesized by cycle-GAN [5]; (c) segmentation using MRI synthesized by masked cycle-GAN [8]; (d) tumor-aware unsupervised learning; (e) tumor-aware semi-supervised learning

Table 1:

Segmentation accuracy

Method	Validation		Test	
	DSC	HD95 mm	DSC	HD95 mm
Week one only	0.63±0.27	7.22±7.19	0.55±0.25	13.23±0.75
cycle-GAN [5]	0.57±0.24	11.41±5.57	0.66±0.16	11.91±4.44
masked cycle-GAN [8]	0.67±0.21	7.78±4.40	0.63±0.24	11.65±6.53
Tumor aware unsupervised	0.62±0.26	7.47±4.66	0.74±0.15	8.88±4.83
Tumor aware semi-supervised	0.70±0.19	5.88±2.88	0.80±0.08	7.16±4.52

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript