Published in final edited form as:

Med Image Comput Comput Assist Interv. 2019 October ; 11769: 221–229.

doi:10.1007/978-3-030-32226-7_25.

Integrating cross-modality hallucinated MRI with CT to aid mediastinal lung tumor segmentation

Jiang Jue¹, Hu Jason¹, Tyagi Neelam¹, Rimner Andreas², Berry L. Sean¹, Deasy O. Joseph¹, Veeraraghavan Harini^{1,*}

¹Medical Physics, Memorial Sloan Kettering Cancer Center

²Radiation Oncology, Memorial Sloan Kettering Cancer Center

Abstract

Lung tumors, especially those located close to or surrounded by soft tissues like the mediastinum, are difficult to segment due to the low soft tissue contrast on computed tomography images. Magnetic resonance images contain superior soft-tissue contrast information that can be leveraged if both modalities were available for training. Therefore, we developed a cross-modality educed learning approach where MR information that is educed from CT is used to hallucinate MRI and improve CT segmentation. Our approach, called cross-modality educed deep learning segmentation (CMEDL) combines CT and pseudo MR produced from CT by aligning their features to obtain segmentation on CT. Features computed in the last two layers of parallelly trained CT and MR segmentation networks are aligned. We implemented this approach on U-net and dense fully convolutional networks (dense-FCN). Our networks were trained on unrelated cohorts from open-source the Cancer Imaging Archive CT images (N=377), an internal archive T2-weighted MR (N=81), and evaluated using separate validation (N=304) and testing (N=333) CT-delineated tumors. Our approach using both networks were significantly more accurate (U-net $P < 0.001$; denseFCN $P < 0.001$) than CT-only networks and achieved an accuracy (Dice similarity coefficient) of 0.71 ± 0.15 (U-net), 0.74 ± 0.12 (denseFCN) on validation and 0.72 ± 0.14 (U-net), 0.73 ± 0.12 (denseFCN) on the testing sets. Our novel approach demonstrated that educing cross-modality information through learned priors enhances CT segmentation performance.

Keywords

Hallucinating MRI from CT for segmentation; lung tumors; adversarial cross-domain deep learning

1 Introduction

Precision medical treatments including image-guided radiotherapy require accurate target tumor segmentation [1]. Computed tomography (CT), the standard-of-care imaging modality lacks sufficient soft-tissue contrast, which makes visualizing tumor boundaries difficult, especially for those that are adjacent to soft-tissue structures. With the advent of new MRI

* veeraragh@mskcc.org.

simulator technologies, radiation oncologists can delineate target structures on MRI acquired in simulation position, which then have to be transferred using image registration to the planning CTs acquired at a different time in treatment position for radiation therapy planning [2]. Image registration itself is prone to errors and thus accurate segmentation on CT itself is more desirable for improving accuracy of clinical radiation treatment margins. More importantly, driven by the lack of simultaneously acquired CT and MR scans, current methods are restricted to CT alone. Therefore, we developed a novel approach, called cross-modality educed deep learning (CMEDL), that uses unpaired cross-domain adaptation between unrelated CT and MR datasets to hallucinate MR-like images or pseudo MR (pMR) from CT scans. The pMR image is combined with CT image to regularize training of a CT segmentation network. This is accomplished by aligning the features of the CT with the pMR features during training (Figure. 1).

Ours is not a method for data augmentation using cross-domain adaptation [3–5]. Our work is also unlike methods that seek to reduce the datashift differences between same imaging modalities [6–8]. Instead, our goal in this work is to maximize the segmentation performance in a single less informative imaging modality, namely, CT using learned information modeling the latent tissue relationships with a more informative modality, namely MRI. The key insight here is that the features dismissed as uninterpretable on CT can provide inference information when learning proceeds from a more informative modality such as MRI.

Our approach is most similar in its goal to compute shared representations for improving segmentations as in the work by [9], where several shared representations between CT and MRI were constructed using fully convolutional networks. Our approach, that is based on GANs for cross-modality learning, shares some similarities to [10] that also used a GAN as a backbone framework, and implemented dual networks for performing segmentations on both CT and MRI. However, our approach substantially differs from prior works in its use of the cross-modality tissue relations as priors to improve inference on the less informative source (or CT) domain. Though applied to segmenting lung tumors, this method is generally applicable to other structures and imaging modalities.

Our contributions in this work are as follows: (i) first, we developed a novel approach to generate segmentation on CT by leveraging more informative MRI through cross-modality priors. (ii) second, we implemented this approach on two different segmentation networks to study feasibility of segmenting lung tumors located in the mediastinum, an area where there is diminished contrast between tumor and the surrounding soft-tissue. (iii) third, we evaluated our approach on a large dataset of 637 tumors.

2 Methods

We use a supervised cross-modality and CT segmentation approach with a reasonably large number of expert segmented CT scans (X_{CT}, y_{CT}) and a few MR scans with expert segmentation ($\{X_{MR}, y_{MR}\}$, where, $N_{X_{MR}} \ll N_{X_{CT}}$). The cross-modality educed deep learning (CMEDL) segmentation consists of two sub-networks that are optimized alternatively. The first sub-network (Figure. 1 A) generates a pMR image given a CT image.

The second sub-network (Figure. 1 B), trains its CT segmentation network constrained using features from another network trained using pMRI. The alternative optimization enables the approach to regularize both the segmentation and pMR generation, such that the pMR is specifically tuned to increase segmentation accuracy. In other words, pMR acts as an informative regularizer for CT segmentation, while the gradients of segmentation errors serve to constrain the generated pMR images.

2.1 Cross-domain adaptation for hallucinating pseudo MR images:

A pair of conditional GANs[11] are trained with unpaired CT and T2-weighted (T2w) MR images arising from different sets of patients. The first GAN transforms CT into a pseudo MR (pMR) image ($G_{C \rightarrow M}$) while the second, transforms a MR image into its corresponding pseudo CT (pCT) ($G_{M \rightarrow C}$) image. The GANs are optimized using the standard adversarial loss ($L_{adv} = L_{adv}^{CT} + L_{adv}^{MR}$) and cycle consistency losses ($L_{cycl} = L_{cycl}^{CT} + L_{cycl}^{MR}$). In addition, we employed a contextual loss that was introduced for real-world images [12] in order to handle learning from image sets lacking spatial correspondence. The contextual loss facilitates such transformations by treating images as collection of features and computing a global similarity between all pairs of features between the two images ($\{g_{j \in N}, m_{i \in M}\}$) used in computing domain adaptation. The contextual similarity is expressed as:

$$CX(g, m) = \frac{1}{N} \sum_j \max_i CX(g_j, m_i), \quad (1)$$

where, N corresponds to the number of features. The contextual similarity is computed by normalizing the inverse of cosine distances between the features in the two images as described in[12]. The contextual loss is computed as:

$$L_{cx} = -\log(CX(f(G(X_{CT})), f(X_{MR}))). \quad (2)$$

The total loss for the cross-modality adaptation is then expressed as the summation of all the aforementioned losses. The pMR generated from this step is passed as an additional input for training the CT segmentation network.

2.2 Segmentation combining CT with pMR images

Our approach for combining the CT with pMR images uses the idea of only matching information that is highly predictable from each other. This usually corresponds to the features closest to the output as the two images are supposed to produce identical segmentation. Therefore, the features computed from the last two layers of CT and pMR segmentation networks are matched by minimizing the squared difference or the L2 loss between them. This is expressed as below.

$$L_{seg} = \mathbb{E}_{x_c \sim X_{CT}} [-\log P(S_{MR}(G_{CT} \rightarrow MR(x_c))) - \log P(S_{CT}(x_c))] + \|\phi_{CT}(x_c) - \phi_{MR}(G_{CT} \rightarrow MR(x_c))\|_F^2, \quad (3)$$

where S_{CT} , S_{MR} are the segmentation networks trained using the CT and pMR images, ϕ_{CT} , ϕ_{MR} are the features computed from these networks, and $G_{CT \rightarrow MR}$ is the cross-modality network used to compute the pMR image, and F stands for Frobenius norm.

The total loss computed from the cross-modality adaptation and the segmentation networks is expressed as:

$$\text{Loss} = L_{adv} + \lambda_{cyc} L_{cyc} + \lambda_{cx} L_{CX} + \lambda_{seg} L_{seg}, \quad (4)$$

where λ_{cyc} , λ_{cx} and λ_{seg} are the weighting coefficients for each loss. During training, we alternatively update the cross-domain adaptation network and the segmentation network with the following gradients, $-\Delta_{\theta_G}(L_{adv} + \lambda_{cyc} L_{cyc} + \lambda_c L_c + \lambda_{cx} L_{cx})$, $-\Delta_{\theta_D}(L_{adv})$ and $-\Delta_{\theta_{seg}} L_{seg}$.

More concretely, the segmentation network is fixed when updating the cross-modality translation and vice versa in each iteration.

2.3 Segmentation architecture:

We implemented the U-net[13] and dense fully convolutional networks (dense-FCN) [14] to evaluate the feasibility of combining hallucinated MR for improving CT segmentation accuracy. These networks are briefly described below.

1. **U-net** was modified using batch normalization after each convolution filter in order to standardize the features computed at the different layers.
2. **Fully Convolutional DenseNets** (Dense-FCN) that is based on [14], uses dense feature maps computed using a sequence of dense feature blocks and concatenated with feature maps from previous computations through residual connections. Specifically, a dense feature block is produced by iterative summation of previous feature maps within that block. As features computed from all image resolutions starting from the image resolution to the lowest resolution are iteratively concatenated, features at all levels are utilized. This in turn facilitates an implicit dense supervision to stabilize training.

2.4 Implementation and training

All networks were implemented using the Pytorch[15] library and trained end to end on Tesla V100 with 16 GB memory and a batch size of 2. The ADAM algorithm [16] with an initial learning rate of $1e-4$ was used during training. The segmentation networks were trained with a learning rate of $2e-4$. We set $\lambda_{adv}=10$, $\lambda_{cx}=1$, $\lambda_{cyc}=1$ and $\lambda_{seg}=5$. For the contextual loss, we use the convolution filters after the Con7, Conv8 and Conv9 due to memory limitations.

3 Datasets and evaluation

We used patients obtained from three different cohorts consisting of (a) the Cancer Imaging Archive (TCIA)[17] with non-small cell lung cancers (NSCLC) [18] consisting of 377 patients (training), (b) 81 longitudinal T2-weighted MR scans (scanned on Philips 3T Ingenia) from 21 patients treated with radiation therapy, and (training) (c) 637 contrast-enhanced tumors treated with immunotherapy at our institution for validation (N=304) and

testing (N=333) such that different sets of patients were used for validation and testing. Early stopping was used during the training to prevent overfitting and the best model selected using validation set was used for testing. Identical CT datasets were used in both CT only and CMEDL approach for equitable comparisons. Expert segmentations were available on all scans.

The segmentation accuracies were evaluated using Dice similarity coefficient (DSC) and Hausdorff distance at 95th percentile (HD95) as recommended in [19]. In addition, we computed the detection rate for the tumors where tumors with at least 50% DSC overlap with expert segmentations were considered as detected.

4 Results

4.1 Tumor detection rate

Our method achieved the most accurate detection using both U-net and Dense-FCN methods for validation and test sets. In comparison the CT-only method resulted in much lower detection rates for both networks (Table 1).

4.2 Segmentation accuracies

The CMEDL approach resulted in more accurate segmentations than CT-only segmentations (see Table 1). In addition, both of the U-net and denseFCN networks trained using CMEDL approach were significantly more accurate than CT only segmentations when evaluated with both DSC ($P < 0.001$) and HD95 ($P < 0.001$) metrics. Figure 2 shows the box plots for the validation and test sets using the two metrics and the two networks. P-values computed using paired Wilcoxon two-sided tests are also shown.

4.3 Visual comparisons

Figure 3 shows visual segmentation results produced by the different networks for representative cases when trained using CT-only and with the CMEDL approach. As seen, in both networks, the CMEDL method closely follows the expert-segmentation that is missed using CT-only networks. Figure 4 shows the feature map activations produced using U-net CT only and with Unet CMEDL. As seen, the feature activations are minimal when using CT-only but shows a clear preferential boundary activation when incorporating the MR information. Figure 4(b) also shows a pseudo MR produced from a CT (Figure 4(a)).

5 Discussion

We developed a novel approach for segmenting lung tumors located in areas with low soft-tissue contrast by leveraging learned prior information from more informative MR modality. These cross-modality priors are learned from unrelated patients and are used to hallucinate MRI to inform CT segmentation. Through extensive experiments on two different network architectures, we showed that leveraging a more informative modality (MRI) to inform inference in a less informative modality (CT), improves segmentation. Our work is limited by lack of sufficiently large MR datasets to potentially improve the accuracy of cross-

domain adaptation models. Nevertheless, this is the first approach to our knowledge that used the cross-modality information in a novel way to generate CT segmentation.

6 Conclusions

We introduced a novel approach for segmenting on CT datasets that can leverage more informative MR modality through cross-modality learning. Our approach implemented on two different segmentation architectures shows improved performance over CT-only methods.

Acknowledgements

This work was supported by the MSK Cancer Center support grant/core grant P30 CA008748, and NCI R01 CA198121-03.

References

1. Njeh C: Tumor delineation: The weakest link in the search for accuracy in radiotherapy. *Journal of medical physics/Association of Medical Physicists of India* 33(4) (2008) 136
2. Devic S: MRI simulation for radiotherapy treatment planning. *Medical Physics* 39(11) (2012) 6701–11 [PubMed: 23127064]
3. 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
4. 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
5. Jiang J, Hu YC, Tyagi N, Zhang P, Rimner A, Mageras GS, Deasy JO, Veeraraghavan H: Tumor-aware, adversarial domain adaptation from ct to mri for lung cancer segmentation In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer (2018) 777–785
6. 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
7. Long M, Zhu H, Wang J, Jordan MI: Deep transfer learning with joint adaptation networks In: *Proceedings of the 34th International inConference on Mache Learning-Volume 70, JMLR. org* (2017) 2208–2217
8. Kamnitsas K, Ledig C, Newcombe VF, Simpson JP, Kane AD, Menon DK, Rueckert D, Glocker B: Efficient multi-scale 3d cnn with fully connected crf for accurate brain lesion segmentation. *Medical Image Analysis* 36 (2017) 61–78 [PubMed: 27865153]
9. Vanya VV, Nick P, Martin R, Ioannis L, Eric OA, Andrea GR, Daniel R, Ben G: 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 2018, Lake Tahoe, NV, USA, March 12–15, 2018* (2018) 547–556
10. Cai J, Zhang Z, Cui L, Zheng Y, Yang L: Towards cross-modal organ translation and segmentation: A cycle-and shape-consistent generative adversarial network. *Medical Image Analysis* (2018)
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. Mechrez R, Talmi I, Zelnik-Manor L: The contextual loss for image transformation with non-aligned data In: *Proceedings of the European Conference on Computer Vision (ECCV)*. (2018) 768–783

13. 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
14. Jégou S, Drozdal M, Vazquez D, Romero A, Bengio Y: The one hundred layers tiramisu: Fully convolutional densenets for semantic segmentation In: Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on, IEEE (2017) 1175–1183
15. Paszke A, Gross S, Chintala S, Chanan G, Yang E, DeVito Z, Lin Z, Desmaison A, Antiga L, Lerer A: Automatic differentiation in pytorch. (2017)
16. Kingma DP, Ba J: Adam: A method for stochastic optimization. Proceedings of the 3rd International Conference on Learning Representations (ICLR) (2014)
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. Aerts HJ, Velazquez ER, Leijenaar RT, Parmar C, Grossmann P, Carvalho S, Bussink J, Monshouwer R, Haibe-Kains B, Rietveld D, et al.: Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach. Nature communications 5 (2014) 4006
19. Menze BH, Jakab A, Bauer S, Kalpathy-Cramer J, Farahani K, Kirby J, Burren Y, Porz N, Slotboom J, Wiest R, et al.: The multimodal brain tumor image segmentation benchmark (BRATS). IEEE Transactions on Medical Imaging 34(10) (2015) 1993 [PubMed: 25494501]

Fig.1:

Overview of the comparison of different combinations. x_c is the CT image; x_m is the MRI image; $G_{C \rightarrow M}$ and $G_{M \rightarrow C}$ are the CT and MRI transfer networks; x'_m is the translated MRI image from x_c ; x'_c is the translated MRI image from x_m .

Fig.2:

Box plots comparing CT-only and CMEDL-based networks.

Fig.3:

Representative segmentations produced using CT-only and CMEDL-based segmentations for U-net and DenseFCN networks. The Dice similarity coefficient (DSC) is also shown for each method. Red corresponds to algorithm, green to expert and yellow to overlap between algorithm and expert.

Fig.4:

Feature map activations from the 21 channel of last layer of Unet. (a) the original CT (b) the translated pMRI (c) activation from CT only (d) activation from pMRI (e) activation from CMEDL

Table 1:

Detection and segmentation accuracy using the two networks.

Method	Validation			Test		
	Detection rate	DSC	HD95 mm	Detection rate	DSC	HD95 mm
U-net CT	80%	0.67±0.18	7.44±7.18	79%	0.68±0.17	9.35±7.08
DenseFCN CT	77%	0.70±0.15	7.25±6.71	75%	0.68±0.16	9.34±9.68
U-net CMEDL	85%	0.71±0.15	6.57±7.15	85%	0.72±0.14	8.22±6.89
DenseFCN CMEDL	84%	0.74±0.12	5.89±5.87	84%	0.73±0.12	7.19±8.55