



# Robust Multimodal Brain Tumor Segmentation via Feature Disentanglement and Gated Fusion

Cheng Chen<sup>1(✉)</sup>, Qi Dou<sup>2</sup>, Yueming Jin<sup>1</sup>, Hao Chen<sup>1,3</sup>, Jing Qin<sup>4</sup>,  
and Pheng-Ann Heng<sup>1,5</sup>

<sup>1</sup> Department of Computer Science and Engineering, The Chinese University of Hong Kong, Sha Tin, Hong Kong

[cchen@cse.cuhk.edu.hk](mailto:cchen@cse.cuhk.edu.hk)

<sup>2</sup> Department of Computing, Imperial College London, London, UK

<sup>3</sup> Insight Medical Technology Co. Ltd., Shenzhen, China

<sup>4</sup> Centre for Smart Health, School of Nursing, The Hong Kong Polytechnic University, Kowloon, Hong Kong

<sup>5</sup> Guangdong Provincial Key Laboratory of Computer Vision and Virtual Reality Technology, SIAT, Chinese Academy of Sciences, Shenzhen, China

**Abstract.** Accurate medical image segmentation commonly requires effective learning of the complementary information from multimodal data. However, in clinical practice, we often encounter the problem of missing imaging modalities. We tackle this challenge and propose a novel multimodal segmentation framework which is robust to the absence of imaging modalities. Our network uses feature disentanglement to decompose the input modalities into the modality-specific appearance code, which uniquely sticks to each modality, and the modality-invariant content code, which absorbs multimodal information for the segmentation task. With enhanced modality-invariance, the disentangled content code from each modality is fused into a shared representation which gains robustness to missing data. The fusion is achieved via a learning-based strategy to gate the contribution of different modalities at different locations. We validate our method on the important yet challenging multimodal brain tumor segmentation task with the BRATS challenge dataset. With competitive performance to the state-of-the-art approaches for full modality, our method achieves outstanding robustness under various missing modality(ies) situations, significantly exceeding the state-of-the-art method by over 16% in average for Dice on whole tumor segmentation.

## 1 Introduction

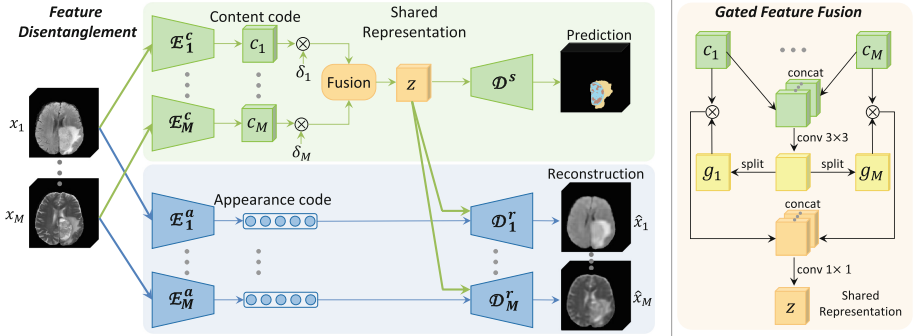
Accurate segmentation of brain tumor is of critical importance for quantitative assessment of tumor progression and preoperative treatment planning. The measurement of tumor-induced tissue changes relies on complementary biological

information provided in multiple Magnetic Resonance Imaging (MRI) modalities, i.e., FLAIR, T1, T1 contrast-enhanced (T1c), and T2. Joint learning from these multimodal images greatly helps to improve segmentation accuracy. A plentiful of multimodal methods have been developed for automated brain tumor segmentation, by either concatenating multiple MRI modalities as inputs [8, 18], or fusing higher-level features from each modality in latent space [5, 14]. However, availability of a full set of the desired modalities is not always guaranteed in real-world scenarios, due to various scanning protocols and diverse patient conditions. In this regard, robustness to one or more missing modality(ies) during inference is essential for a widely-applicable multimodal learning method.

A typical solution is to synthesize the missing modality(ies) with available ones [15]. Such method requires to build a specific model for each modality from all possible combinations of available modalities, which is complicated. Alternatively, Havaei et al. [6] propose hetero-modal image segmentation (HeMIS), which fuses multimodal information by computing statistics (i.e., mean and variance) across individual features. This method is easily scalable to various data missing situations, as the fusion in latent space adapts to any number of modalities. Furthermore, Chartsias et al. [2] and Van Tulder et al. [16] enhance the modality-invariance of latent representations by minimizing the L1 or L2 distance of features from different modalities. However, different MRI modalities vary in intensity distributions with modality-specific appearance, thus simply encouraging the features from different modalities to be close under L1- or L2-Norm may not achieve optimal representations with desired modality-invariance. Instead, a concurrent work [12] uses adversarial learning to ensure the model generate similar features under missing modalities as in a full modality situation.

To effectively extract modality-invariant representations conveying essential content of tumor, learning to cancel out the modality-specific information may be helpful. This can be achieved using feature disentanglement by decomposing inputs to latent space of interpretable factors [7, 13]. In medical imaging, disentangling representations has recently demonstrated effectiveness for liver lesion classification [1], myocardial segmentation [3] and multimodal deformable registration [11]. However, these works are for uni-modal or bi-modal data. To our best knowledge, the potential of feature disentanglement for robust multimodal segmentation at arbitrary modality number has not been tapped yet.

We propose a novel multimodal learning framework with feature disentanglement and gated feature fusion, which is robust to missing modalities. Our network disentangles multimodal features into modality-specific appearance code and modality-invariant content code. The content code of each modality is fused into a shared representation containing discriminative information for segmentation task. To enhance its modality-invariance, the shared representation is required to reconstruct any modality given corresponding appearance code, even in the absence of some modality(ies). Furthermore, we employ a novel gated feature fusion strategy to automatically learn weight maps and gate the contribution from different modalities at different locations. We validate our proposed method on the task of multimodal brain tumor segmentation with BRATS challenge [10].



**Fig. 1.** Proposed multimodal segmentation framework. **Left:** Feature disentanglement for multimodal learning. **Right:** Detailed structure of the gated feature fusion module.

With competitive performance to the state-of-the-art methods for full modality, our method is highly robust to various missing modality(ies) situations.

## 2 Method

An overview of our proposed multimodal segmentation framework is in Fig. 1. We first introduce feature disentanglement to encode multimodal inputs to modality-specific appearance code and modality-invariant content code. Next, we present a learning-based gating strategy for integrating the complementary disentangled content code from individual modality to a more expressive fused representation, and detailed learning process and network architecture are described in the end.

### 2.1 Feature Disentanglement for Robust Multimodal Learning

We denote the multimodal images by  $\{x_1, \dots, x_M\}$ , where  $M = 4$  in our brain tumor segmentation task. Each modality  $x_i$  is input to its own appearance encoder  $E_i^a$  and content encoder  $E_i^c$ , and we correspondingly obtain its disentangled appearance code  $a_i = E_i^a(x_i)$  and content code  $c_i = E_i^c(x_i)$ . For the appearance code, we follow the common practice in [9] and set it as a 8-bit vector assuming its prior distribution to be a centered isotropic Gaussian  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ . The Kullback-Leibler (KL) divergence is computed to encourage the estimated distribution of  $p(a_i)$  to be as close to the normal distribution. In this way, we obtain the loss of  $\mathcal{L}_{\text{KL}} = \sum_{i=1}^M \mathbb{E}[D_{\text{KL}}(p(a_i) || \mathcal{N}(\mathbf{0}, \mathbf{I}))]$  for training the appearance encoders  $\{E_i^c\}$ .

Next, for the content code  $\{c_i\}$  which is expected to gain modality-invariance after evaporating the stylized appearances of different image modalities, we fuse them into an integrated representation  $z = \mathcal{F}(\{c_i\})$  expressing essential semantic contents of the tumor.  $\mathcal{F}$  is an automatically learned fusion strategy for which we will elaborate in Sect. 2.2. From the perspective of successful disentanglement, the obtained content representation  $z$  should enable to be re-rendered into the

original image given any appearance code of a certain modality. To encourage such reconstruction capability, we develop the pseudo-cycle-consistency loss by introducing a set of modality-specific decoders  $\{D_i^r\}$ , as follows:

$$\mathcal{L}_{\text{rec}} = \sum_{i=1}^M \|D_i^r(z, a_i) - x_i\|_1, \text{ where } z = \mathcal{F}(\delta_i c_1, \dots, \delta_M c_M), \quad (1)$$

where we employ the L1-Norm to alleviate the generated images getting blurred. With a Bernoulli indicator  $\delta_i$ , we are aiming to grant the content representation  $z$  extra robustness to missing data, i.e., still producing a high-quality reconstructed  $\hat{x}_i = D_i^r(z, a_i)$  even in the absence of  $c_i$  when fusing content codes. We perform the modality dropout in latent space by randomly setting  $\delta_i$  to 0, which turns off the content code  $c_i$  in current learning iteration.

After the disentanglement procedure which cancels out the effect of modality-specific appearance features while aggregates the complementary content information from arbitrary combination of multimodal data, we can perform accurate and robust brain tumor segmentation. We build a segmentation decoder  $\hat{Y} = D^s(z)$  which learns discriminative pattern based on our derived representative and robust  $z$ . We jointly use Dice loss and weighted cross-entropy loss to handle the unbalanced object sizes in multi-class segmentation:

$$\mathcal{L}_{\text{seg}} = \text{Dice}(\hat{Y}, Y) + H(\hat{Y}, Y) = - \sum_{k \in K} \left( \frac{2 \sum_{j=1}^N y_j^k \hat{y}_j^k}{\sum_{j=1}^N y_j^k y_j^k + \hat{y}_j^k \hat{y}_j^k + \epsilon} + \sum_{j=1}^N w^k \cdot y_j^k \log q_j^k \right), \quad (2)$$

where  $y_j^k$ ,  $q_j^k$ ,  $\hat{y}_j^k$  respectively denote the ground truth, probability prediction and one-hot output of voxel  $j$  for class  $k$ . Directly combing the two types of segmentation loss works well in practice, without need of particularly tuning a balancing weight between them. The  $\epsilon = e^{-7}$  is a constant for numerical stabilization, and  $w^k$  is calculated online in a batch, treating class imbalance in cross-entropy loss.

## 2.2 Multimodal Content Fusion with Learned Gating

Effectively fusing the complementary information from various modalities is crucial in a multimodal learning framework. This also holds for our scenario, though we disentangle the content code and enforce robustness to missing data. In fact, feature fusion plays a more important role in unusual inference situations such as unavailability of some modality(ies). If not considered carefully, the fused representation would otherwise be infected by the noisy information from empty-input channel(s), then the model's performance is inevitably degraded. Existing approaches tackle this using average [6] or max operation [2]. However, the average operation makes each modality contribute equally, which may disregard highly informative features from a certain modality. On the contrary, the max operation only retains the largest response, neglecting information from all the others.

Instead of hard-coding a fusion operation, we propose to automatically learn the mapping function to integrate multimodal features. The contribution weights

of a modality are not necessarily identical across locations, as a modality contains different amount of information for areas of each class. For example, T1c shows clear structure of enhancing tumor, but not for the area of edema. In this regard, we dynamically learn a weight map to gate the scale of information from each content code  $c_i$ , with a flexibility of voxel-by-voxel. Then, the gated content from individual modalities are fused to form the integrated representation.

Specifically, the disentangled content codes  $\{c_1, \dots, c_M\}$  from each modality are concatenated, and then input to a convolutional layer with an output channel of  $M$ . With sigmoid activation, we obtain the gating weight matrix  $G$ , which can be split into  $M$  separate maps of  $\{g_1, \dots, g_M\}$ , one for each modality. Next, we re-weight the content code as  $\tilde{c}_i = c_i \cdot g_i$  with element-wise multiplication. These outputs  $\{\tilde{c}_i\}$  are concatenated and forwarded to a bottleneck  $1 \times 1$  convolution, followed by Leaky ReLU activation. As shown in Eq. (1), we randomly set some content code(s) to be 0 with  $\{\delta_i\}$  during training, to enhance model’s robustness to missing data. Overall, we obtain the fused content code  $z = \mathcal{F}(\delta_i c_1, \dots, \delta_M c_M)$ , which has the same feature map size and channel as individual code  $c_i$ .

It is worth noting that our learning-based gating strategy is general for multimodal feature fusion, which is superior to existing average or max hard-coding way, by properly aggregating complementary content with data-dependent weight. In our framework, we jointly use it with the disentanglement procedure, and form an accurate and robust end-to-end multimodal learning method.

### 2.3 Learning Process and Network Architecture

The entire framework is learned with the overall objective function as:

$$\mathcal{L}_{\text{total}}(D^s, \{E_i^c, E_i^a, D_i^r\}) = \mathcal{L}_{\text{seg}} + \lambda \mathcal{L}_{\text{rec}} + \beta \mathcal{L}_{\text{KL}}, \quad (3)$$

where the  $\lambda, \beta$  are trade-off parameters weighting the importance of each component, which are both empirically set as 0.1 in our experiments. We utilize an Adam optimizer with an initial learning rate of  $1e^{-4}$ , and progressively multiply it by  $(1 - \text{epoch}/\text{max\_epoch})^{0.9}$  during training. The intensive components with our model only allow us to set the batch size as 1 using one Nvidia Xp GPU.

Our encoders  $\{E_i^c\}$  and decoder  $D^s$  for segmentation task adopt 3D U-Net [4] architecture, except that one independent encoder is used for each input modality. In each downsampling stage, the content features of individual modality are fused via the learning-based gating strategy with 0.5 probability of  $\delta_i$  to be zero for dropping  $c_i$ . Each fused features are then skip-connected to the corresponding upsampling stage. Each  $E_i^c$  consists of 4 residual blocks with instance normalization and Leaky ReLU activation. Between each block, the image dimensions are progressively reduced by 2 and the feature channels are doubled. All convolutions use kernel size of  $3 \times 3 \times 3$  and the initial channel number is 16. The  $D^s$  also has 4 residual blocks which is similar to  $E_i^c$ , except that the feature map size is upsampled by 2 with channel number halved after each stage. For image reconstruction, we generally follow the practice in [7]. Specifically, each  $E_i^a$  consists of 5 convolutional layers followed by a global average pooling and

a fully-connected layer to obtain the appearance code. Each  $D_i^r$  uses 4 residual blocks followed by 4 upsampling and convolutional layers to produce  $\hat{x}_i$ .

### 3 Experiments


**Dataset and Preprocessing.** We validate our proposed method with the 2015 Brain Tumor Segmentation Challenge (BRATS) dataset [10]. The training set consists of 274 cases with ground truth being provided. The test set contains 110 cases with reference labels being held by the organizers and the evaluation

**Table 1.** Comparison of brain tumor segmentation performance on BRATS 2015 test set. The values are obtained by submitting our results to the online evaluation system.

Methods	Dice (%)			Precision (%)			Sensitivity (%)		
	Complete	Core	Enhancing	Complete	Core	Enhancing	Complete	Core	Enhancing
Kamnitsas et al. [8]	84	67	63	82	<b>85</b>	<b>64</b>	<b>89</b>	62	66
Zhao et al. [17]	82	<b>72</b>	62	84	78	60	83	<b>73</b>	69
OM-Net [18]	<b>86</b>	71	<b>64</b>	<b>86</b>	83	61	88	68	<b>72</b>
Ours	84	<b>72</b>	<b>64</b>	84	80	<b>64</b>	<b>89</b>	69	68

**Table 2.** Robustness comparison of our method against HeMIS [6] and the imputation MLP [6] on the test split of BRATS 2015 training set. The Dice score is presented for every combination case of modalities being available (✓) or missing (-).

Modalities				Dice (%)								
				Complete			Core			Enhancing		
F	T1	T1c	T2	Ours	HeMIS	MLP	Ours	HeMIS	MLP	Ours	HeMIS	MLP
-	-	-	✓	<b>85.49</b>	58.48	61.50	<b>58.66</b>	40.18	37.32	<b>37.66</b>	20.31	18.62
-	-	✓	-	<b>71.86</b>	33.46	2.04	<b>72.87</b>	44.55	17.70	<b>70.22</b>	49.93	32.92
-	✓	-	-	<b>68.40</b>	33.22	2.07	<b>50.00</b>	17.42	10.52	<b>22.67</b>	4.67	10.78
✓	-	-	-	<b>83.02</b>	71.26	63.81	<b>46.67</b>	37.45	34.26	<b>28.30</b>	5.57	15.90
-	-	✓	✓	<b>87.53</b>	67.59	64.97	<b>78.46</b>	63.39	49.38	<b>76.82</b>	65.38	60.30
-	✓	✓	-	<b>74.59</b>	45.93	1.99	<b>76.40</b>	55.06	26.55	<b>73.95</b>	62.40	40.93
✓	✓	-	-	<b>87.66</b>	80.28	78.13	<b>60.17</b>	49.52	48.97	<b>35.28</b>	22.26	25.18
-	✓	-	✓	<b>87.87</b>	69.56	66.88	<b>64.88</b>	47.26	43.66	<b>41.05</b>	23.56	26.37
✓	-	-	✓	<b>89.08</b>	82.10	81.35	<b>63.51</b>	53.42	52.41	<b>39.72</b>	23.19	25.01
✓	-	✓	-	<b>88.01</b>	79.80	81.13	<b>78.09</b>	66.12	65.51	<b>76.62</b>	67.12	66.19
✓	✓	✓	-	<b>87.73</b>	80.88	82.19	<b>80.68</b>	69.26	69.34	<b>78.81</b>	71.30	70.93
✓	✓	-	✓	<b>89.07</b>	83.87	80.40	<b>65.99</b>	57.76	53.46	<b>43.04</b>	28.46	28.34
✓	-	✓	✓	<b>89.06</b>	82.78	83.37	<b>79.47</b>	70.62	70.45	<b>78.07</b>	70.52	70.56
-	✓	✓	✓	<b>88.26</b>	70.98	67.85	<b>80.84</b>	66.60	55.40	<b>78.56</b>	67.84	64.81
✓	✓	✓	✓	<b>89.07</b>	83.15	82.43	<b>81.19</b>	72.50	71.46	<b>79.13</b>	75.37	72.08
Average				<b>84.45</b>	68.22	60.01	<b>69.19</b>	54.07	47.09	<b>57.33</b>	43.86	41.93

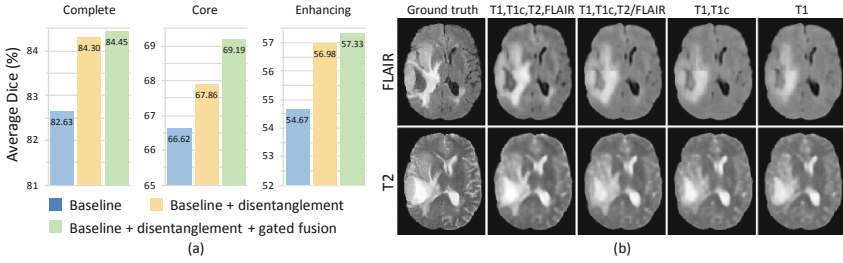


**Fig. 2.** Segmentation results from our method for complete tumor (yellow), tumor core (red), and enhancing tumor (blue). Input modalities at inference are indicated. (Color figure online)

can be obtained via an online system. Each case contains four MRI modalities including FLAIR, T1, T1c, and T2. The task of the challenge is to segment three tumor classes, i.e. complete, core, and enhancing tumor. The dataset is preprocessed with being skull-stripped, co-registered, and resampled to isotropic  $1\text{ mm}^3$  resolution, by the organizer. We further normalized the intensity of each volume to zero mean and unit variance within the brain tissue area. A patch of size  $80 \times 80 \times 80$  was randomly cropped during training as input to the network.

**Performance of Robust Brain Tumor Segmentation.** We first compare our method with the state-of-the-art methods on the test set of BRATS 2015 for full modalities. Results were obtained directly from the online evaluation system and compared without postprocessing. In Table 1, our method achieves the highest Dice score of the core and enhancing tumor, with the other evaluations highly competitive with the ranking 1st approach OM-Net [18], validating the effectiveness of our segmentation backbone.

We then evaluate the robustness of our method to missing modality inference. The absence of modality  $i$  is implemented by setting  $\delta_i$  to be zero for dropping  $c_i$  at inference. For direct comparison with the HeMIS method and image synthesis method using multilayer perceptron (MLP) [6], we used the same data split of BRATS training set as in [6] and directly referenced the results from their paper. In Table 2, our method significantly outperforms the HeMIS and imputation MLP methods for all the 15 possible combination situations of unavailable modalities and all the three tumor classes. This demonstrates the outstanding robustness of our multimodal segmentation method. From the results, we can see that FLAIR and T2 modalities are more informative than others for the complete tumor segmentation, and T1c is discriminative for accurate prediction of enhancing tumor. In Fig. 2, we show that with the increase of the number of missing modalities, the segmentation results produced by our robust model just gradually degrade, rather than encountering sudden failure. Even with T1 alone, we can achieve decent segmentation for the complete tumor and tumor core.



**Fig. 3.** (a) Ablation study of key components in our method. (b) Example reconstruction of FLAIR and T2 images for different combinations of input modalities.

**Ablation Study.** We investigate the effectiveness of feature disentanglement and gated fusion, as two key components in our method. We first set up a baseline network which uses average fusion without feature disentanglement. Then we add the feature disentanglement and gated fusion one by one into the baseline network. In Fig. 3(a), we compare the performance of the three networks on the Dice score, averaging over the 15 possible combination situations of input modalities. Both the feature disentanglement and gated fusion bring performance improvement across all the tumor parts, achieving the highest Dice score in most situations (10, 13, and 11 situations out of 15 for the complete tumor, tumor core, and enhancing tumor respectively). Figure 3(b) shows the reconstruction results of FLAIR and T2 image by combining their corresponding appearance code with the shared representation fused from content code of different combination of input modalities. Even when some modality(ies) are missing, our network can still reconstruct the missing modality with the shared representation, indicating that the shared representation  $z$  successfully yields the essential tumor content.

## 4 Conclusion

We propose a novel multimodal segmentation framework which jointly uses the feature disentanglement and gated feature fusion to obtain a modality-invariant and discriminative representation. We validate our method on brain tumor segmentation under both full modalities and various combination situations of missing modalities, achieving new state-of-the-art results on BRATS benchmark. The outstanding robustness to great inference variations can make our method widely-applicable in real-world clinical scenarios.

**Acknowledgments.** This work was supported in part by the National Basic Program of China 973 Program under Grant 2015CB351706, the National Natural Science Foundation of China, under Project No. U1613219, the Research Grants Council of Hong Kong Special Administrative Region, under Project No. CUHK14225616, and the Hong Kong Innovation and Technology Commission, under Project No. ITS/319/17.



## References

1. Ben-Cohen, A., Mechrez, R., et al.: Improving CNN training using disentanglement for liver lesion classification in CT. arXiv preprint [arXiv:1811.00501](https://arxiv.org/abs/1811.00501) (2018)
2. Chartsias, A., Joyce, T., Giuffrida, M.V., Tsaftaris, S.A.: Multimodal MR synthesis via modality-invariant latent representation. *IEEE TMI* **37**(3), 803–814 (2018)
3. Chartsias, A., et al.: Factorised spatial representation learning: application in semi-supervised myocardial segmentation. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) *MICCAI 2018*. LNCS, vol. 11071, pp. 490–498. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-00934-2\\_55](https://doi.org/10.1007/978-3-030-00934-2_55)
4. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: Ourselin, S., Jostkowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) *MICCAI 2016*. LNCS, vol. 9901, pp. 424–432. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46723-8\\_49](https://doi.org/10.1007/978-3-319-46723-8_49)
5. Fidon, L., et al.: Scalable multimodal convolutional networks for brain tumour segmentation. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) *MICCAI 2017*. LNCS, vol. 10435, pp. 285–293. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-66179-7\\_33](https://doi.org/10.1007/978-3-319-66179-7_33)
6. Havaei, M., Guizard, N., Chapados, N., Bengio, Y.: HeMIS: hetero-modal image segmentation. In: Ourselin, S., Jostkowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) *MICCAI 2016*. LNCS, vol. 9901, pp. 469–477. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46723-8\\_54](https://doi.org/10.1007/978-3-319-46723-8_54)
7. 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). [https://doi.org/10.1007/978-3-030-01219-9\\_11](https://doi.org/10.1007/978-3-030-01219-9_11)
8. Kamnitsas, K., Ledig, C., Newcombe, V.F., et al.: Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *MIA* **36**, 61–78 (2017)
9. 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). [https://doi.org/10.1007/978-3-030-01246-5\\_3](https://doi.org/10.1007/978-3-030-01246-5_3)
10. Menze, B.H., Jakab, A., Bauer, S., et al.: The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE TMI* **34**(10), 1993–2024 (2015)
11. Qin, C., Shi, B., Liao, R., Mansi, T., Rueckert, D., Kamen, A.: Unsupervised deformable registration for multi-modal images via disentangled representations. In: Chung, A.C.S., Gee, J.C., Yushkevich, P.A., Bao, S. (eds.) *IPMI 2019*. LNCS, vol. 11492, pp. 249–261. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-20351-1\\_19](https://doi.org/10.1007/978-3-030-20351-1_19)
12. Shen, Y., Gao, M.: Brain tumor segmentation on MRI with missing modalities. In: Chung, A.C.S., Gee, J.C., Yushkevich, P.A., Bao, S. (eds.) *IPMI 2019*. LNCS, vol. 11492, pp. 417–428. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-20351-1\\_32](https://doi.org/10.1007/978-3-030-20351-1_32)
13. Tsai, Y.H.H., Liang, P.P., Zadeh, A., Morency, L.P., Salakhutdinov, R.: Learning factorized multimodal representations. arXiv preprint [arXiv:1806.06176](https://arxiv.org/abs/1806.06176) (2018)
14. Tseng, K.L., Lin, Y.L., Hsu, W., et al.: Joint sequence learning and cross-modality convolution for 3D biomedical segmentation. In: *CVPR*, pp. 6393–6400 (2017)

15. van Tulder, G., de Bruijne, M.: Why does synthesized data improve multi-sequence classification? In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9349, pp. 531–538. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-24553-9\\_65](https://doi.org/10.1007/978-3-319-24553-9_65)
16. van Tulder, G., deBruijne, M.: Learning cross-modality representations from multi-modal images. *IEEE TMI* **38**(2), 638–648 (2019)
17. Zhao, X., Wu, Y., Song, G., Li, Z., Zhang, Y., Fan, Y.: A deep learning model integrating FCNNs and CRFs for brain tumor segmentation. *MIA* **43**, 98–111 (2018)
18. Zhou, C., Ding, C., Lu, Z., Wang, X., Tao, D.: One-pass multi-task convolutional neural networks for efficient brain tumor segmentation. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11072, pp. 637–645. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-00931-1\\_73](https://doi.org/10.1007/978-3-030-00931-1_73)