Lesion Mask-based Simultaneous Synthesis of Anatomic and Molecular MR Images using a GAN

Pengfei Guo[1], Puyang Wang[2], Jinyuan Zhou[3],
Vishal M. Patel[1,2], Shanshan Jiang[3]

[1] Department of Computer Science, Johns Hopkins University, MD, USA
[2] Department of Electrical and Computer Engineering, Johns Hopkins University, MD, USA
[3] Department of Radiology, Johns Hopkins University, Baltimore, MD, USA
sjiang21@jhmi.edu

Abstract. Data-driven automatic approaches have demonstrated their great potential in resolving various clinical diagnostic dilemmas for patients with malignant gliomas in neuro-oncology with the help of conventional and advanced molecular MR images. However, the lack of sufficient annotated MRI data has vastly impeded the development of such automatic methods. Conventional data augmentation approaches, including flipping, scaling, rotation, and distortion are not capable of generating data with diverse image content. In this paper, we propose a method, called synthesis of anatomic and molecular MR images network (SAMR), which can simultaneously synthesize data from arbitrary manipulated lesion information on multiple anatomic and molecular MRI sequences, including T1-weighted (T_1w), gadolinium enhanced T_1w (Gd-T_1w), T2-weighted (T_2w), fluid-attenuated inversion recovery ($FLAIR$), and amide proton transfer-weighted (APTw). The proposed framework consists of a stretch-out up-sampling module, a brain atlas encoder, a segmentation consistency module, and multi-scale label-wise discriminators. Extensive experiments on real clinical data demonstrate that the proposed model can perform significantly better than the state-of-the-art synthesis methods.

Keywords: MRI · Multi-modality Synthesis · GAN

1 Introduction

Malignant gliomas, such as glioblastoma (GBM), remain one of the most aggressive forms of primary brain tumor in adults. The median survival of patients with glioblastomas is only 12 to 15 months with aggressive treatment [15]. For the clinical management in patients who finish surgery and chemoradiation, the treatment responsiveness assessment is relied on the pathological evaluations [16]. In recent years, deep convolutional neural network (CNN) based medical image analysis methods have shown to produce significant improvements

over the conventional methods [2, 6]. However, a large amount of data with rich diversity is required for training effective CNNs models, which is usually unavailable for medical image analysis. Furthermore, lesion annotations and image prepossessing (e.g. co-registration) are labor-intensive, time-consuming and expensive, since expert radiologists are required to label and verify the data. While deploying conventional data augmentations, such as rotation, flipping, random cropping, and distortion, during training partly mitigates such issues, the performance of CNN models still suffer from the limited diversity of the dataset [18]. In this paper, we propose a generative network which can simultaneously synthesize meaningful high quality T_1w, Gd-T_1w, T_2w, $FLAIR$, and APTw MRI sequences from input lesion mask. In particular, APTw is a novel molecular MRI technique, which yields a reliable marker for treatment responsiveness assessment for patients with post-treatment malignant gliomas [8, 17].

Recently Goodfellow et al. [5] proposed generative adversarial network (GAN) which has been shown to synthesize photo-realistic images. Isola et al. [7] and Wang et al. [14] applied GAN under the conditional settings and achieved impressive results on image-to-image translation tasks. When considering the generative models for MRI synthesis alone, several methods have been proposed in the literature. Nguyen et al. [13] and Chartsias et al. [3] proposed CNN-based architectures integrating intensity features from images to synthesize cross-modality MR images. However, their inputs are existing MRI modalities and the diversity of the synthesized images is limited by the training images. Cordier et al. [4] used a generative model for multi-modal MR images with brain tumors from a single label map. However, the input label map contains detailed brain anatomy and the method is not capable of producing manipulated outputs. Shin et al. [11] adopted pix2pix [7] to transfer brain anatomy and lesion segmentation maps to multi-modal MR images with brain tumors. However, it requires to train an extra segmentation network that provides white matter, gray matter, and cerebrospinal fluid (CSF) masks as partial input of synthesis network. Moreover, it is only demonstrated to synthesize anatomical MRI sequences. In this paper, a novel generative model is proposed that can take arbitrarily manipulated lesion mask as input facilitated by brain atlas generated from training data to simultaneously synthesize a diverse set of anatomical and molecular MR images.

To summarize, the following are our key contributions: **1.** A novel conditional GAN-based model is proposed to synthesize meaningful high quality multi-modal anatomic and molecular MR images with controllable lesion information. **2.** Multi-scale label-wise discriminators are developed to provide specific supervision on the region of interest (ROI). **3.** Extensive experiments are conducted and comparisons are performed against several recent state-of-the-art image synthesis approaches. Furthermore, an ablation study is conducted to demonstrate the improvements obtained by various components of the proposed method.

2   Methodology

Fig. 1 gives an overview of the proposed framework. Incorporating multi-scale label-wise discriminators and shape consistency-based optimization, the gener-
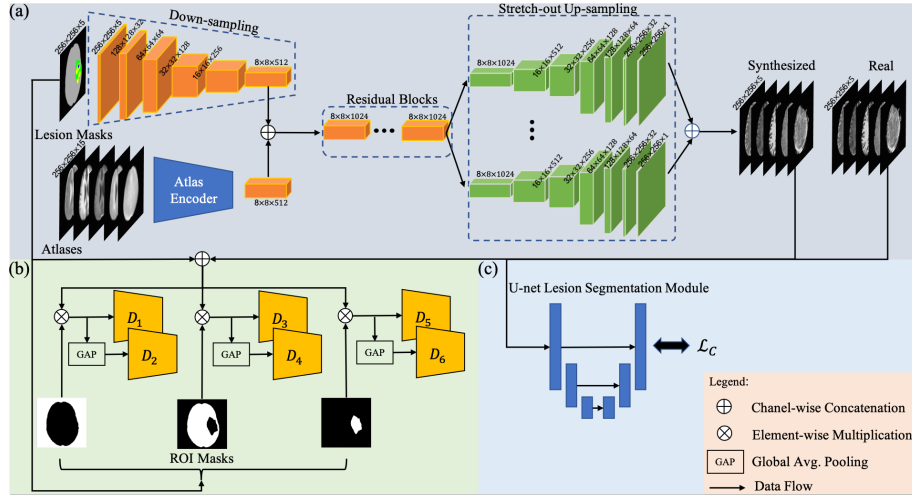
Fig. 1. An overview of the proposed network. (a) Generator network. (b) Multi-scale Label-wise discriminators. Global averaging pooling is used to create the factor of 2 down-sampling input. (c) U-net based lesion segmentation module. We denote lesion shape consistency loss as $\mathcal{L}_C$.

ator aims to produce meaningful high-quality anatomical and molecular MR images with diverse and controllable lesion information. In what follows, we describe different parts of the network in detail.

Multi-modal MRI Sequence Generation. Our generator architecture is inspired by the models proposed by Johnson et al. [9] and Wang et al. [14]. The generator network, consists of four components (see Fig. 1(a): a down-sampling module, an atlas encoder, a set of residual blocks, and a stretch-out up-sampling module. A lesion segmentation map of size $256 \times 256 \times 5$, containing 5 labels: background, normal brain, edema, cavity caused by surgery, and tumor, is passed through the down-sampling module to get a latent feature map. The corresponding multi-model atlas of size $256 \times 256 \times 15$ (details of atlas generation are provided in Section 3) is passed through an atlas encoder to get another latent feature map. Then, the two latent feature maps are concatenated and are passed through residual blocks and stretch-out up-sampling module to synthesize multi-model MRI slices of size $256 \times 256 \times 5$.

The down-sampling module consists of a fully-convolutional module with 6 layers. We set the kernel size and stride equal to 7 and 1, respectively for the first layer. For down-sampling, instead of using maximum-pooling, the stride of other 5 layers is set equal to 2. Rectified Linear Unit (ReLu) activation and batch normalization are sequentially added after each layer. The atlas encoder has the same network architecture but the number of channels in the first convolutional layer is modified to match the input size of the multi-model atlas input. The depth of the network is increased by a set of residual blocks, which is proposed to learn better transformation functions and representations through a deeper perception [18]. The stretch-out up-sampling module contains 5 similar sub-

modules designed to utilize the same latent representations from residual blocks and perform customized synthesis for each sequence. Each sub-module contains one residual learning block and a symmetric architecture with a down-sampling module. All convolutional layers are replaced by transposed convolutional layers for up-sampling. The synthesized multi-model MR images are produced from each sub-model.

Multi-scale Label-wise Discriminators. In order to efficiently achieve large receptive field in discriminators, we adopt multi-scale PatchGAN discriminators [7], which have identical network architectures but take multi-scale inputs [14]. Conventional discriminators operate on the combination of images and conditional information to distinguish between real and synthesized images. However, optimizing generator to produce realistic images in each ROI cannot be guaranteed by discriminating on holistic images. To address this issue, we propose label-wise discriminators. Based on the radiographic features, original lesion segmentation masks are reorganized into 3 ROIs, including background, normal brain, and lesion. Specifically, the input of each discriminator is the ROI-masked combination of lesion segmentation maps and images. Since proposed discriminators are in a multi-scale setting, for each ROI there are 2 discriminators that operate on original and a down-sampled scales (factor of 2). Thus, there are in total 6 discriminators for 3 ROIs and we refer to these set of discriminators as $\mathbb{D} = \{D_1, D_2, D_3, D_4, D_5, D_6\}$. In particular, $\{D_1, D_2\}, \{D_3, D_4\}$, and $\{D_5, D_6\}$ operate on original and down-sampled versions of background, normal brain, and lesion, respectively. An overview of the proposed discriminators is shown in Fig. 1(b). The objective function for a specific discriminator $\mathcal{L}_{GAN}(G, D_k)$ is as follows:

$$\mathcal{L}_{GAN}(G, D_k) = \mathbb{E}_{(\hat{x}, \hat{y})}[\log D_k(\hat{x}, \hat{y})] + \mathbb{E}_{\hat{x}}[\log(1 - D_k(\hat{x}, \hat{G}(x)))], \qquad (1)$$

where x and y are paired original lesion segmentation masks and real multi-model MR images, respectively. Here, $\hat{x} \triangleq c_k \odot x$, $\hat{y} \triangleq c_k \odot y$, and $\hat{G}(x) \triangleq c_k \odot G(x)$, where $\odot$ denotes element-wise multiplication and c_k corresponds to the ROI mask. For simplicity, we omit the down-sampling operation in this equation.

Multi-task Optimization. A multi-task loss is designed to train the generator and discriminators in an adversarial setting. Instead of only using the conventional adversarial loss $\mathcal{L}_{GAN}$, we also adopt a feature matching loss $\mathcal{L}_{FM}$ [14] to stabilize training, which optimizes generator to match these intermediate representations from the real and the synthesized images in multiple layers of the discriminators. For a specific discriminator, $\mathcal{L}_{FM}(G, D_k)$ is defined as follows:

$$\mathcal{L}_{FM}(G, D_k) = \sum_{i}^{T} \frac{1}{N_i} \left[\|D_k^{(i)}(\hat{x}, \hat{y}) - D_k^{(i)}(\hat{x}, \hat{G}(x)\|_2^2 \right], \qquad (2)$$

where $D_k^{(i)}$ denotes the ith layer of the discriminator D_k, T is the total number of layers in D_k and N_i is the number of elements in the ith layer. If we perform lesion segmentation on images, it is worth to note that there is a consistent relation between the prediction and the real one serving as input for the generator. Lesion labels are usually occluded with each other and brain anatomic structure, which causes ambiguity for synthesizing realistic MR images. To tackle this problem, we

propose a lesion shape consistency loss $\mathcal{L}_C$ by adding a U-net [10] segmentation module (Fig. 1(c)), which regularizes the generator to obey this consistency relation. We adopt Generalized Dice Loss (GDL) [12] to measure the difference between predicted and real segmentation maps and is defined as follows:

$$GDL(R, S) = 1 - \frac{2 \sum_i^N r_i s_i}{\sum_i^N r_i + \sum_i^N s_i}, \tag{3}$$

where R denotes the ground truth and S is the segmentation result. r_i and s_i represent the ground truth and predicted probabilistic maps at each pixel i, respectively. N is the total number of pixels. The lesion shape consistency loss $\mathcal{L}_C$ is then defined as follows:

$$\mathcal{L}_C(x, U(G(x)), U(y)) = GDL(x, U(y)) + GDL(x, U(G(x))), \tag{4}$$

where $U(y)$ and $U(G(x))$ represent the predicted probabilistic maps by taking y and $G(x)$ as inputs in the segmentation module, respectively. The proposed final multi-task loss function for the generator is defined as:

$$\mathcal{L} = \sum_{k=1}^{6} \mathcal{L}_{GAN}(G, D_k) + \lambda_1 \sum_{k=1}^{6} \mathcal{L}_{FM}(G, D_k) + \lambda_2 \mathcal{L}_C(x, U(G(x)), U(y)), \tag{5}$$

where λ_1 and λ_2 two parameters that control the importance of each loss.

3   Experiments and Evaluations

Data Acquisition and Implementation Details. 90 postsurgical patients were involved in this study. MRI scans were acquired by a 3T human MRI scanner (Achieva; Philips Medical Systems). Anatomic sequences of size $512 \times 512 \times 150$ voxels and Molecular APTw sequence of size $256 \times 256 \times 15$ voxels were collected. Detailed imaging parameters and preprocessing pipeline can be found in supplementary material. After preprocessing, the final volume size of each sequence is $256 \times 256 \times 15$. Expert neuroradiologist manually annotated five labels (i.e. background, normal brain, edema, cavity and tumor) for each patient. Then, a multivariate template construction tool [1] was used to create the group average for each sequence (atlas). 1350 instances with the size of $256 \times 256 \times 5$ were extracted from the volumetric data, where the 5 corresponds to five MR sequences. For every instance, the one corresponding atlas slice and two adjacent (in axial direction) atlas slices were extracted to provide human brain anatomy prior. We split these instances randomly into 1080 (80%) for training and 270 (20%) for testing on the patient level. Data collection and processing are approved by the Institutional Review Board.

The synthesis model was trained based on the final objective function Eq. (5) using the Adam optimizer [1]. λ_1 and λ_2 in Eq. (5) were set equal to 5 and 1, respectively. Hyperparameters are set as follows: constant learning rate of 2 $\times 10^{-4}$ for the first 250 epochs then linearly decaying to 0; 500 maximum epochs; batch size of 8. For evaluating the effectiveness of the synthesized MRI sequences on data augmentation, we leveraged U-net [10] to train lesion segmentation models. U-net [10] was trained by the Adam optimizer [1]. Hyperparameters are set as follows: constant learning rate of 2 $\times 10^{-4}$ for the first 100 epochs then linearly decaying to 0; 200 maximum epochs; batch size of 16. In the segmentation

training, all the synthesized data was produced by randomly manipulated lesion masks. For evaluation, we always keep 20% of data unseen for both of the synthesis and segmentation models.
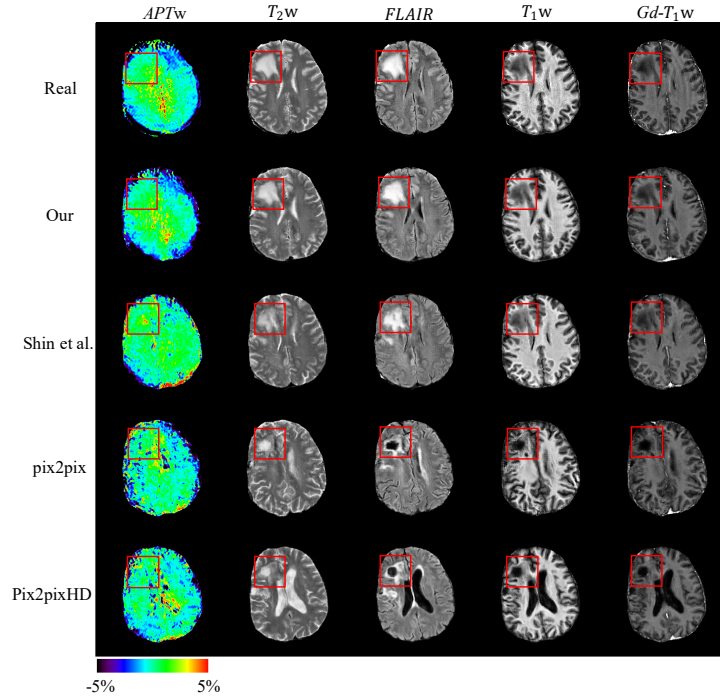
Fig. 2. Qualitative comparison of different methods. The same lesion mask is used to synthesize images from different methods. Red boxes indicate the lesion region.
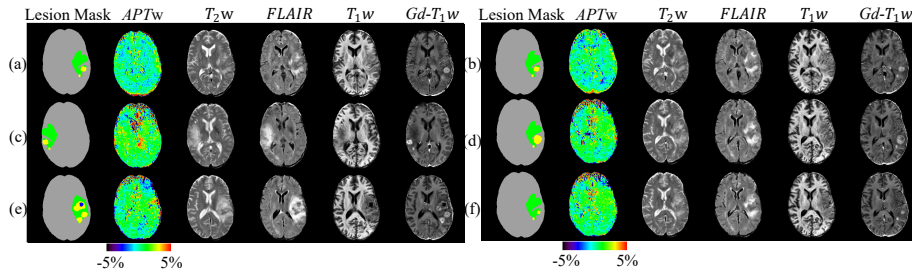
Fig. 3. Examples of lesion mask manipulations. (a) Real images. (b) Synthesized images from the original mask. (c) Synthesized images by mirroring lesion. (d) Synthesized images by increasing tumor size to 100%. (e) Synthesized images by replacing lesion from another patient. (f) Synthesized images by shrinking tumor size to 50%. In lesion masks, gray, green, yellow, and blue represent normal brain, edema, tumor, and cavity caused by surgery, respectively.

MRI Synthesis Results. We evaluate the performance of different synthesis methods by qualitative comparison and human perceptual study. We compare the performance of our method with the following recent state-of-the-art synthesis methods: pix2pix [7], pix2pixHD [14], and Shin et al. [11]. Fig. 2 presents the qualitative comparison of the synthesized multi-model MRI sequences from four different methods. It can be observed that pix2pix [7] and pix2pixHD [14] fail to synthesize realistic looking human brain MR images. There is either an unreasonable brain ventricle (see the last row of Fig. 2) or wrong radiographic features in the lesion region (see the fourth row of Fig. 2). With the help of probability maps of white matter, gray matter and CSF, Shin et al. [11] can produce realistic brain anatomic structures for anatomic MRI sequences. However, there is an obvious disparity between the synthesized and real *APT*w sequence in both normal brain and lesion region. The boundary of the synthezied lesion is also blurred and uncertain (see red boxes in the third row of Fig. 2). The proposed method produces more accurate radiographic features of lesions and more diverse anatomic structure based on the human anatomy prior provided by atlas.

Human Perceptual Study. To verify the pathological information of the synthesized images, we conduct the human perceptual study by an expert neuro-radiologist and the corresponding preference rates are reported in Table 1. It is clear that the images generated by our method are more preferred by an expert neuroradiologist than others showing the practicality of our synthesis method.

Table 1. Preference rates corresponding to the human perceptual study.

Real	Our	Shin et al. [11]	pix2pixHD [14]	pix2pix [7]
100%	**72.1**%	65.6%	39.3%	16.4%

Data Augmentation Results. To further evaluate the quality of the synthezied images, we perform data augmentation by using the synthesized images in training and then perform lesion segmentation. Evaluation metrics in BraTS [2] challenge (i.e. Dice score, Hausdorff distance (95%), Sensitivity, and Specificity) are used to measure the performance of different methods. The data augmentation by synthesis is evaluated by the improvement for lesion segmentation models. We arbitrarily control lesion information to synthesize different number of data for augmentation. The detail of mechanism for manipulating lesion mask can be found in supplementary material. To simulate the piratical usage of data augmentation, we conduct experiments in the manner of utilizing all real data. In each experiment, we vary the percentage of synthezied data to observe the contribution for data augmentation. Table 2 shows the calculated segmentation performance. Comparing with the baseline experiment that only uses real data, the synthesized data from pix2pix [7] and pix2pixHD [14] degrade the segmentation performance. The performance of Shin et al. [11] improves when synthesized data is used for segmentation, but the proposed method outperforms other methods by a large margin. Fig.3 demonstrates the robustness of the proposed model under different lesion mask manipulations (e.g. changing the size of tumor, moving lesion location, and even reassembling lesion information between patients). As can be seen from this figure, our method is robust to various lesion mask

manipulations.

Table 2. Quantitative results corresponding to image segmentation when the synthesized data is used for data augmentation. For each experiments, the first row reports the percentage of synthesized/real data for training and the number of instances of synthesized/real data in parentheses. Exp.4 reports the results of ablation study.

	Dice Score			Hausdorff95 Distance			Sensitivity			Specificity		
	Edema	Cavity	Tumor	Edema	Cavity	Tumor	Edema	Cavity	Tumor	Edema	Cavity	Tumor
Exp.1: 50% Synthesized+ 50% Real (1080 + 1080)												
pix2pix [7]	0.589	0.459	0.562	13.180	21.003	10.139	0.626	0.419	0.567	0.995	0.998	0.999
pix2pix HD [14]	0.599	0.527	0.571	17.406	8.606	10.369	0.630	0.494	0.570	0.996	0.998	0.999
Shin et al. [11]	0.731	0.688	0.772	7.306	6.290	6.294	0.701	0.662	0.785	0.997	0.999	0.999
our	**0.794**	**0.813**	**0.821**	**6.049**	**1.568**	**2.293**	**0.789**	**0.807**	**0.841**	**0.997**	**0.999**	**0.999**
Exp.2: 25% Synthesized+ 75% Real (540 + 1080)												
pix2pix [7]	0.602	0.502	0.569	10.706	9.431	10.147	0.640	0.463	0.579	0.995	0.999	0.998
pix2pix HD [14]	0.634	0.514	0.663	17.754	9.512	9.061	0.670	0.472	0.671	0.996	0.999	0.999
Shin et al. [11]	0.673	0.643	0.708	14.835	7.798	6.688	0.664	0.602	0.733	0.997	0.999	0.998
our	**0.745**	**0.780**	**0.772**	**8.779**	**6.757**	**4.735**	**0.760**	**0.788**	**0.805**	**0.997**	**0.999**	**0.999**
Exp.3: 0% Synthesized + 100% Real (0 + 1080)												
Baseline	**0.646**	**0.613**	**0.673**	**8.816**	**7.856**	**7.078**	**0.661**	**0.576**	**0.687**	**0.996**	**0.999**	**0.998**
Exp.4: **Ablation Study**												
w/o Stretch-out	0.684	0.713	0.705	6.592	5.059	4.002	0.708	0.699	0.719	0.997	0.999	0.999
w/o Label-wise D	0.753	0.797	0.785	7.844	2.570	2.719	0.735	0.780	0.783	0.998	0.999	0.999
w/o Atlas	0.677	0.697	0.679	13.909	11.481	7.123	0.691	0.689	0.723	0.997	0.999	0.998
w/o $\mathcal{L}_C$	0.728	0.795	0771	8.604	3.024	3.233	0.738	0.777	0.777	0.997	0.999	0.999
Our	**0.794**	**0.813**	**0.821**	**6.049**	**1.568**	**2.293**	**0.789**	**0.807**	**0.841**	**0.998**	**0.999**	**0.999**

Ablation Study. We conduct extensive ablation study to separately evaluate the effectiveness of using stretch-out up-sampling module, label-wise discriminators, atlas, and lesion shape consistency loss $\mathcal{L}_C$ in our method using the same experimental setting as exp.1 in Table 2. The contribution of modules for data augmentation by synthesis is reported in Table 2 exp.4. We find that when atlas is not used in our method, it significantly affects the synthesis quality due to the lack of human brain anatomy prior. Losing the customized reconstruction for each sequence (stretch-out up-sampling module) can also degrade the synthesis quality. Moreover, dropping either $\mathcal{L}_C$ or label-wise discriminators in the training reduces the performance, since the shape consistency loss and the specific supervision on ROIs are not used to optimize the generator to produce more realistic images.

4   Conclusion

We proposed an effective generation model for multi-model anatomic and molecular *APT*w MRI sequences. It was shown that the proposed multi-task optimization under adversarial training further improves the synthesis quality in each ROI. The synthesized data can be used for data augmentation, particularly for those images with pathological information of malignant gliomas, to improve the performance of segmentation. Moreover, the proposed approach is an automatic, low-cost solution to produce high quality data with diverse content that can be used for training of data-driven methods.

In our future work, we will generalize data augmentation by synthesis to other tasks, such as classification. Furthermore, the proposed method will be extended to 3D synthesis once better quality molecular MRI data is available for training the models.

References

1. Avants, B.B., Tustison, N.J., Song, G., Cook, P.A., Klein, A., Gee, J.C.: A reproducible evaluation of ants similarity metric performance in brain image registration. Neuroimage **54**(3), 2033–2044 (2011)
2. Bakas, S., Reyes, M., Jakab, A., Bauer, S., Rempfler, M., Crimi, A., Shinohara, R.T., Berger, C., Ha, S.M., Rozycki, M., et al.: Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. arXiv preprint arXiv:1811.02629 (2018)
3. Chartsias, A., Joyce, T., Giuffrida, M.V., Tsaftaris, S.A.: Multimodal mr synthesis via modality-invariant latent representation. IEEE transactions on medical imaging **37**(3), 803–814 (2017)
4. Cordier, N., Delingette, H., Lê, M., Ayache, N.: Extended modality propagation: image synthesis of pathological cases. IEEE transactions on medical imaging **35**(12), 2598–2608 (2016)
5. 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. pp. 2672–2680 (2014)
6. Guo, P., Li, D., Li, X.: Deep oct image compression with convolutional neural networks. Biomedical Optics Express **11**(7), 3543–3554 (2020)
7. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1125–1134 (2017)
8. Jiang, S., Eberhart, C.G., Lim, M., Heo, H.Y., Zhang, Y., Blair, L., Wen, Z., Holdhoff, M., Lin, D., Huang, P., et al.: Identifying recurrent malignant glioma after treatment using amide proton transfer-weighted mr imaging: a validation study with image-guided stereotactic biopsy. Clinical Cancer Research **25**(2), 552–561 (2019)
9. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: European conference on computer vision. pp. 694–711. Springer (2016)
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. pp. 234–241. Springer (2015)
11. Shin, H.C., Tenenholtz, N.A., Rogers, J.K., Schwarz, C.G., Senjem, M.L., Gunter, J.L., Andriole, K.P., Michalski, M.: Medical image synthesis for data augmentation and anonymization using generative adversarial networks. In: International workshop on simulation and synthesis in medical imaging. pp. 1–11. Springer (2018)

12. Sudre, C.H., Li, W., Vercauteren, T., Ourselin, S., Cardoso, M.J.: Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In: Deep learning in medical image analysis and multimodal learning for clinical decision support, pp. 240–248. Springer (2017)

13. Van Nguyen, H., Zhou, K., Vemulapalli, R.: Cross-domain synthesis of medical images using efficient location-sensitive deep network. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 677–684. Springer (2015)

14. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High-resolution image synthesis and semantic manipulation with conditional gans. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2018)

15. Wen, P.Y., Kesari, S.: Malignant gliomas in adults. New England Journal of Medicine **359**(5), 492–507 (2008)

16. Woodworth, G.F., Garzon-Muvdi, T., Ye, X., Blakeley, J.O., Weingart, J.D., Burger, P.C.: Histopathological correlates with survival in reoperated glioblastomas. Journal of neuro-oncology **113**(3), 485–493 (2013)

17. Zhou, J., Payen, J.F., Wilson, D.A., Traystman, R.J., van Zijl, P.C.: Using the amide proton signals of intracellular proteins and peptides to detect ph effects in mri. Nature medicine **9**(8), 1085–1090 (2003)

18. Zhou, Y., He, X., Cui, S., Zhu, F., Liu, L., Shao, L.: High-resolution diabetic retinopathy image synthesis manipulated by grading and lesions. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 505–513. Springer (2019)