Learning to Synthesize 7 T MRI from 3 T MRI with Few Data by Deformable Augmentation

Jie Wei[1,2,3], Yongsheng Pan[1,2,3], Yong Xia[1,2(✉)], and Dinggang Shen[3,4(✉)]

[1] National Engineering Laboratory for Integrated Aero-Space-Ground-Ocean Big Data Application Technology, School of Computer Science and Engineering, Northwestern Polytechnical University, Xi'an 710072, Shaanxi, China
yxia@nwpu.edu.cn
[2] Research and Development Institute of Northwestern Polytechnical University in Shenzhen, Shenzhen 518057, China
[3] School of Biomedical Engineering, ShanghaiTech University, Shanghai, China
dgshen@shanghaitech.edu.cn
[4] Shanghai United Imaging Intelligence Co., Ltd., Shanghai, China

Abstract. High-quality magnetic resonance imaging (MRI), which is generally acquired by ultra-high field (7-Tesla, 7 T) MRI scanners, may lead to improved performance for brain disease diagnosis, such as Alzheimer's disease (AD). However, 7 T MRI has not been widely used due to higher cost and longer scanning time. To overcome this, we proposed to utilize the generative adversarial networks (GAN)-based techniques to synthesize the 7 T scans from 3 T scans, for which, the most challenge is that we do not have enough data to learn a reliable mapping from 3 T to 7 T. To address this, we further proposed the Unlimited Data Augmentation (UDA) strategy to increase the learning samples via the deformable registration, which can produce enough paired 3 T and 7 T MR images to learning this mapping. Based on this mapping, we synthesize a 7 T MR scan for each subject in Alzheimer's Disease Neuroimaging Initiative (ADNI), and conduct some experiments to evaluate their effect in two tasks of AD diagnosis, including AD identification and mild cognitive impairment (MCI) conversion prediction. Experimental results demonstrate that our UDA strategy is effective to learn a reliable mapping to high-quality MR images, and the synthetic 7 T scans are possible to increase the performance of AD diagnosis.

1 Introduction

Alzheimer's disease (AD) is a neurodegenerative disorder that could accelerate the patient's cognitive loss and lead to dementia [10]. Previous studies have verified that magnetic resonance imaging (MRI) is a relevant important technique in

This work is mainly completed under the collaboration of J. Wei and Y. Pan. J. Wei and Y. Pan contribute equally.

C. Lian et al. (Eds.): MLMI 2021, LNCS 12966, pp. 70–79, 2021.
https://doi.org/10.1007/978-3-030-87589-3_8

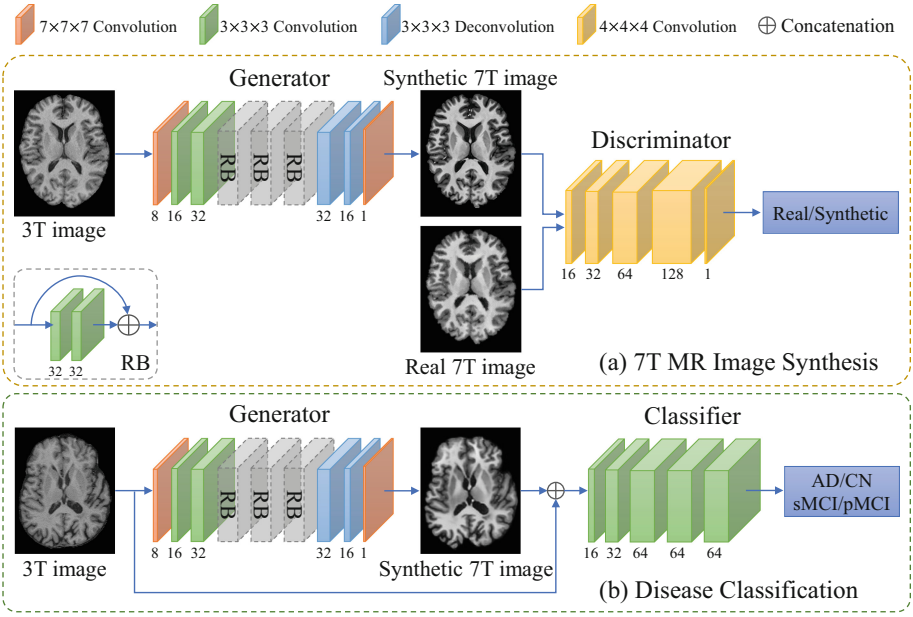

Fig. 1. Framework of our proposed method consisting of two stages, i.e., (a) 7 T MR image synthesis, and (b) disease classification.

AD diagnosis, e.g., identifying AD individuals and predicting the progress of mild cognitive impairment (MCI) [16]. Generally, higher-quality MR scans, which are recognized with a higher signal-to-noise ratio (SNR) and can be acquired with higher magnetic field strength, are possible to result in a more accurate diagnosis of AD. Comparing to 1.5-Tesla (1.5 T) and 3 T MRI, the ultra-high field (7 T) MRI has evident clinical advantages, including higher SNR, better contrast, higher sensitivity to provide detail of anatomy and pathology of the brain, and easier to detect subtle changes [9,11]. However, in addition to producing quality images, 7 T MRI also brings higher costs and longer scanning time.

Currently, the most widely used MRI is 1.5 T and 3 T scanners, of which a lot of scans have been collected, such as Alzheimer's Disease Neuroimaging Initiative (ADNI) [6]. The low quality may lead to an upper bound of analysis. If these scans can be converted to 7 T MR scans, it is possible to boost the diagnosis performance. This produces a specific scenario of image-to-image translation, for which various techniques have been studied, particularly those techniques based on generative adversarial networks (GAN) [3,5,7,19]. Accordingly, it is potential to learn a mapping from 1.5 T or 3 T MR images to 7 T images based on various GANs, e.g., cycle-consistency GAN (cycGAN) [3], pixel2pixel GAN (p2pGAN) [5], L1GAN [2], feature matching GAN (FMGAN) [17], or sense-consistency GAN (SGAN) [14] with various constraints. As 3 T MRI synthesized from 1.5T MRI can help to improve the performance of AD status prediction [18], synthesizing 7 T MRI from 1.5 T/3 T MRI may further promote the performance.

However, learning image-to-image mappings generally requires abundant data collection. It makes the task of synthesizing 7 T MR scans very difficult because 7 T MRI scanners have not been widely applied and there is no public dataset with accessible 7 T MR scans. Even a previous study [15] collected 15 in-house subjects with both 3 T and 7 T MR scans, it is still not enough to learn a reliable mapping. As diagnosis models generally align images based on image appearances by using image registration methods in their pre-processing pipeline, conventional data augmentation such as up-down flip, left-right flip may not be applicable to brain images.

Based on previous studies, we found that the most variability among different images is the regional characteristic, to address which, a deformable augmentation may be effective to enlarge the learning dataset. Accordingly, we propose an effective strategy to do deformable augmentation, so that to increase training samples to learning reliable mappings from 3 T scans to 7 T scans. We register the 3 T MR scan of each subject with paired 3 T and 7 T scans to a target 3 T MRI in the ADNI dataset via deformable registration and apply the deformable field to the corresponding healthy 7 T MR scan. In this way, we can obtain enough subjects with paired 3 T and 7 T MR images.

In this paper, we first design various GAN-based models to learning the mapping from 3 T to 7 T after doing deformable augmentation. Experimental results demonstrate that all of them are promoted by our deformable augmentation strategy, where SGAN achieves the best synthesis performance. Then, we use the mapping learned by SGAN to synthesis a 7 T scan for each 3 T/1.5 T scan in the ADNI dataset, i.e., each subject in ADNI now has a pair of original 3 T/1.5 T scan and a synthetic 7 T scan. Finally, each pair of the original 3 T/1.5 T image and the synthetic 7 T image are combined as a two-channel 3D image and fed into classifiers. Experimental results demonstrate that our synthetic 7 T scans are effective to improve the performance in AD diagnosis.

2 Method

2.1 Dataset

There are two datasets used in this study. The first is an in-house dataset with 15 subjects, each of which has a pair of 3 T and 7 T T1-weighted MR scans. All 3 T and 7 T scans were acquired with Siemens Magnetom Trio 3 T and 7 T MRI scanners, respectively. Specifically, each 3 T scan consists of 224 coronal slices acquired with the 3D magnetization-prepared rapid gradient-echo (MP-RAGE) sequence. The imaging parameters of 3D MP-RAGE sequence were: repetition time (TR) = 1900 ms, echo time (TE) = 2.16 ms, inversion time (TI) = 900 ms, flip angle (FA) = 9°, and voxel size = $1.0 \times 1.0 \times 1.0$ mm^3. Each 7 T scan consists of 191 sagittal slices acquired with the 3D MP2-RAGE sequence. The imaging parameters of 3D MP2-RAGE sequence were: TR = 6000 ms, TE = 2.95 ms, TI = 800/2700 ms, FA =4°/4°, and voxel size = $0.65 \times 0.65 \times 0.65$ mm^3. Each 3T scan was linearly aligned to the MNI standard space with the voxel size of $1.0 \times 1.0 \times 1.0$ mm^3. The corresponding 7 T scan was then aligned

to the 3 T scan in the MNI space by using FLIRT. After bias field correction, skull removal, and intensity normalization, the intensity values of the 3 T and 7 T images were rescaled to $[-1, 1]$.

The second dataset is the Alzheimer's Disease Neuroimaging Initiative database (ADNI) [6], where we use the baseline T1 scans of its two subsets, i.e., ADNI-1 and ADNI-2. Subjects in ADNI can be divide into (1) AD patients, (2) CN patients, (3) pMCI patients that would progress to AD within 36 months after baseline, and (4) sMCI patients that would not progress to AD. After removing subjects that exist in both ADNI-1 and ADNI-2 from ADNI-2, there are 200 AD, 231 CN, 171 pMCI, and 150 sMCI subjects in ADNI-1, while there are 165 AD, 209 CN, 89 pMCI, and 256 sMCI subjects in ADNI-2. We apply a similar pre-processing pipeline on MR scans in ADNI. After pre-processing, these scans also have the voxel size of $1.0 \times 1.0 \times 1.0$ mm^3.

2.2 Framework

We propose a two-stage deep learning framework to promote the image quality of T1 scans and boost the performance of AD diagnosis as illustrated in Fig. 1. In the *first* stage, we learn a mapping from 3 T scans to 7 T scans via GAN-based techniques, which is used to promote the quality of scans in ADNI. Especially, to address the small training data problem, we propose the Unlimited Data Augmentation (UDA) with deformable registration. In the *second* stage, based on the synthesized 7 T T1 scans, we develop a deep learning method for AD diagnosis, by learning 3T and 7 T features automatically in a data-driven manner. To the best of our knowledge, this is the first attempt to promote AD diagnosis by mapping common-used 3 T MRI to high-quality 7 T MRI.

2.3 Unlimited Data Augmentation

As we only have a few paired 3 T and 7 T scans, which is not enough to learn a reliable generative model to synthesize 7 T scans, we propose the Unlimited Data Augmentation (UDA) to produce more paired 3 T and 7 T MR scans via the deformable registration [1]. The motivation is that the deformable registration can find a non-linear transformation (i.e., deformation field) to establish anatomical correspondences between each two images, and each group of images still keeps the spatial relationships after being applied with the same deformation field. Based on this, we can produce hundreds of paired training samples with potentially large and complex variety.

As our goal is to find a mapping function from 3 T to 7 T MR scans for the ADNI subjects, we randomly select several scans from ADNI as the fixed templates. We first perform deformable registration to each 3 T MR scan of the in-house dataset and each selected template to obtain a deformation field, and then apply the deformation field to this pair of 3 T and 7 T in-house scans. For example, if we select 30 scans from ADNI as templates, then we can produce 15×30 samples in total, which may be enough to train a reliable mapping.

2.4 Synthesizing 7 T MRI with GAN

In the *first* stage of our framework, we attempt to learn a mapping from 3 T to 7 T MR scans, which is conducted by the GAN-based techniques. We follow the structures in [13], to create our GAN structures, which contains a generator and a discriminator.

The generator is an encoder-decoder network, where the encoder consists of three 3D convolutional layers with strides of {1, 2, and 2} and channels of {8, 16, and 32}, respectively, while the decoder consists of two 3D deconvolutional layers (with 32 and 16 channels, respectively) and a 3D convolutional layer with 1 channel. Between the encoder and decoder, we also insert 3 residual blocks (RB), each of which contains two convolutional layers with a shortcut connection. The kernel sizes of the first and last convolutional layers are set to $7 \times 7 \times 7$, while which of the other layers are set to $3 \times 3 \times 3$. All convolutional/deconvolutional layers except the last one are followed by the batch normalization and rectified linear unit (ReLu) activation while the last layer uses the "tanh" activation.

The discriminator is a fully convolutional network with five 3D convolutional layers with the kernel size of $4 \times 4 \times 4$. The channels of these layers are 16, 32, 64, 128, and 1, respectively, where the strides are 2, 2, 2, 2, and 1, respectively. The former four layers are followed by batch normalization and ReLu activation.

During the training phase, the generator attempts to create the mapping to approximate the difference between 3 T and 7 T MR images while the discriminator outputs a binary indicator to distinguish the original and the synthetic 7 T MR images. Besides, with different constraints, the generator and discriminator can form different GANs, such as the basic GAN with only the adversarial loss [3], the cycle-consistency GAN (cycGAN) with the adversarial loss and cycle-consistency loss [19], the pixel-to-pixel GAN (p2pGAN) with the adversarial loss and the mean absolute error (MAE) loss [5], the 3D encoder-decoder with only MAE loss (L1GAN) [2], the feature matching GAN (FMGAN) with the adversarial loss and the feature matching loss [17] and the sense-consistency GAN (SGAN) with the MAE loss and the feature matching loss [14]. Specifically, to adopt the potential location shift, we apply patch-wise training by randomly sampling patches of size $192 \times 192 \times 192$ from the whole volume at the center of each brain with a maximum random shift of $4 \times 4 \times 4$ as the inputs. It should be noted that all these patches can cover the whole brain regions of all subjects. When synthesizing a target 7 T scan during the test phase, we randomly select 50 patches from each 3 T scan, feed them into the generator to obtain 50 synthetic patches, and calculate the average of these synthetic patches as the synthetic 7 T MR scan.

2.5 Disease Classification

After synthesizing 7 T scans for all subjects in ADNI, we further involve these 7 T scans in the classification tasks. Therefore, at the *second* stage, we built a classification model using both original 3 T and synthetic 7 T MR images to identify AD subjects (AD vs. NC classification) and predict MCI progression (pMCI vs. sMCI classification).

Table 1. Results (% except PSNR) of image synthesis achieved by different methods for 3 T scans of subjects in in-house data.

Method	MAE			PSNR			SSIM		
	#1	#10	#20	#1	#10	#20	#1	#10	#20
GAN	49.33	21.14	19.81	16.73	22.37	23.27	39.70	51.11	59.86
cycGAN	9.53	6.81	5.96	27.10	28.89	29.96	74.01	79.84	82.42
p2pGAN	7.14	5.80	5.25	29.03	29.96	30.39	78.74	83.48	84.58
L1GAN	7.18	5.12	4.89	29.35	30.72	31.02	82.32	84.58	85.87
FMGAN	7.00	5.81	5.77	29.18	30.11	30.62	81.47	84.19	84.47
SGAN	**6.55**	**4.81**	**4.53**	**29.79**	**30.91**	**31.40**	**83.90**	**85.02**	**86.29**
	#30	#40	#50	#30	#40	#50	#30	#40	#50
SGAN	4.15	4.08	4.11	31.86	31.96	32.03	87.45	87.51	87.69

The classification model employs the same structure as the disease-image specific neural network (DSNN) [12], which consists of 5 convolutional layers with channels of 16, 32, 64, 64, and 64, respectively. Each layer has the kernel size of $3 \times 3 \times 3$ and stride of 1 and is followed by batch normalization and ReLu activation. Each of the first four convolutional layers is followed by a max-pooling layer to downsample the output with the stride of 2. The last layer is followed by a spatial cosine kernel to utilize the spatial information. While using both original 3 T and synthetic 7 T MR scans as inputs, each pair of 3 T and 7 T images are concatenated along the channel dimension to form a two-channel volume. We also use the same patch-wise training for the DSNN with the patch size of $160 \times 176 \times 160$ and a maximum random shift of $2 \times 2 \times 2$. While testing, we only select the patches at the center of the brain from the 3 T and 7 T scans.

2.6 Implement Details

We trained all networks on a platform with an NVIDIA TITAN Xp GPU (32 GB). For the 7 T MR image synthesis stage, we adopted the Adam optimizer and set the batch size to 1, maximum epoch number to 500, and initial learning rate to 0.001, which was divided by 10 after 100 iterations. For the disease classification stage, we adopted the stochastic gradient descent (SGD) optimizer and set the batch size to 6, maximum epoch number to 300, and initial learning rate to 0.001, which was divided by 10 after 50 iterations. And we apply 3-fold cross-validation on 7 T MR image synthesis and train the model three times for disease classification.

3 Experiments and Results

3.1 Evaluation of Image Synthesis

We evaluate the performance of different 7 T MR image synthesis methods with three commonly used image quality metrics, including (1) the mean absolute

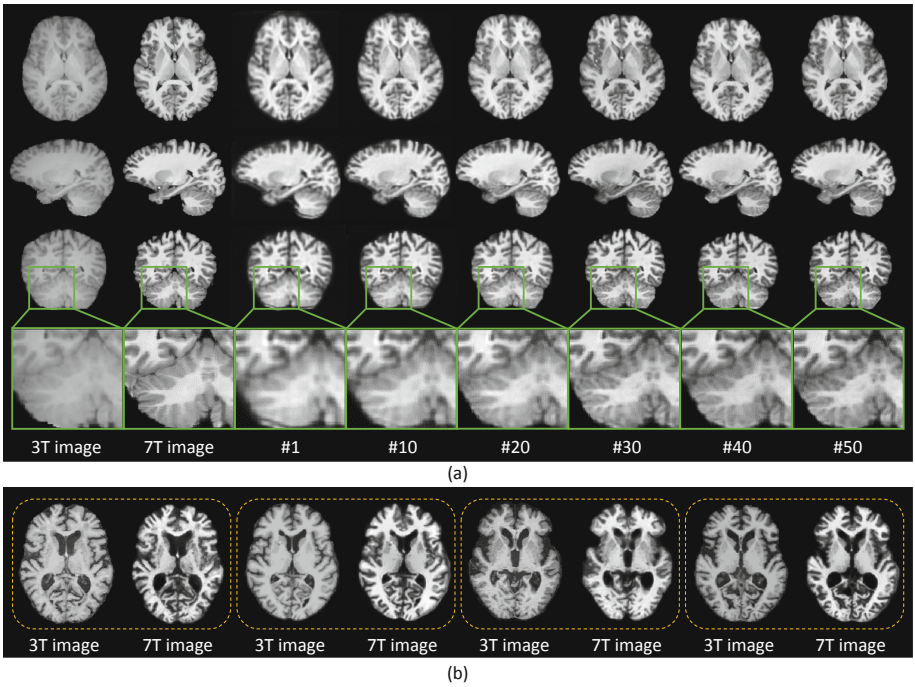

Fig. 2. (a) Comparison of typical synthesis results of SGAN with different expanded factors on one subject from in-house data. (b) The synthesis results of SGAN with the expanded factor of 30 on four subjects from ADNI. The X of $\#X$ is the expanded factor.

error (MAE), (2) the peak signal-to-noise ratio (PSNR), and (3) the structural similarity index measure (SSIM) [4]. Table 1 shows the MAE, PSNR, and SSIM values computed based on the predicted images from different methods, including the basic GAN, cycGAN, p2pGAN, L1GAN, FMGAN, and SGAN, all of which are implemented using the same architecture and details for a fair comparison. The X of $\#X$ is the number of selected fixed templates, which means that the training data was expanded by a factor of X. It can be observed that the performances of all models on all metrics are improved with the increase of the expanded factor. This can be attributed to our effective UDA performed by the deformable registration. Specifically, among these compared models, SGAN achieves the best performance in terms of all three metrics. Moreover, we further evaluate the influence of the expand factor for SGAN by increasing the selected templates. As shown in the last row of Table 1, the benefits from UDA become less and less while the time costs are increasing linearly. Considering the trade-off between performance and efficiency, we think the expanded factor of 30 is the best choice. Therefore, we use SGAN with the expanded factor of 30 to synthesize 7 T MR scans for our second-stage study.

In Fig. 2(a), we further visualize an example of our in-house dataset to display the synthetic 7 T scans of SGAN over expanded factors. From top to bottom are, respectively, the central slice in axial, sagittal, coronal, and enlarged coronal views of the original 3 T and 7 T MR images and the synthetic 7 T MR images. It can be seen that the 7 T MR scan has better contrast than the 3 T scan, and the quality of synthetic images is improving along with the expanded factors. Figure 2(b) gives four examples from ADNI to show the original 3 T and synthetic 7 T scans obtained by SGAN with the expanded factor of 30.

Table 2. Diagnosis results (%) achieved by different methods for ADNI.

Method	AUC	ACC	SEN	SPE	F1S	MCC
	AD vs. CN					
DSNN (3 T)	93.02 ± 0.93	85.62 ± 1.04	78.52 ± 1.95	93.16 ± 1.78	85.90 ± 1.37	70.23 ± 2.12
DSNN (7 T)	91.22 ± 0.91	85.56 ± 0.89	77.42 ± 0.65	94.63 ± 1.46	83.82 ± 1.16	73.37 ± 1.67
DM-DSNN	**94.74 ± 0.29**	**87.40 ± 0.83**	**78.64 ± 3.47**	**95.23 ± 1.08**	**86.07 ± 0.96**	**74.82 ± 1.28**
	pMCI vs. sMCI					
DSNN (3 T)	82.60 ± 0.42	77.58 ± 0.60	65.17 ± 1.83	81.90 ± 0.66	72.21 ± 0.77	44.84 ± 1.58
DSNN (7 T)	80.05 ± 0.67	77.39 ± 0.47	57.12 ± 1.15	84.11 ± 0.66	70.61 ± 0.71	41.23 ± 1.41
DM-DSNN	**83.25 ± 0.30**	**78.61 ± 0.25**	**66.97 ± 1.41**	**84.57 ± 0.39**	**73.55 ± 0.22**	**47.10 ± 0.42**

3.2 Evaluation of Disease Classification

We further evaluated the synthetic 7 T MR scans in both tasks of AD identification (AD vs. CN) and MCI conversion prediction (pMCI vs. sMCI). We use the DSNN in [12] as our classification method, and report both the results that using only 3 T or 7 T modality and using both 3T and 7 T modalities for comparison. Six metrics are computed (as shown in Table 2) to evaluate the classification results, including the area under the receiver operating characteristic (AUC), accuracy (ACC), sensitivity (SEN), specificity (SPE), F1-Score (F1S), and Matthews correlation coefficient (MCC) [8]. Subjects from ADNI-1 and ADNI-2 were used for training and testing the models, respectively.

From Table 2, we can find that the classification results of DM-DSNN that use both original 3 T and synthetic 7 T images outperform each of the single-modality DSNN that use only original 3 T or synthetic 7 T images. It implies that our synthetic 7 T MR images could help to improve the classification performance across external datasets. For instance, the DM-DSNN achieves the best AUC values in both AD vs. CN classification (94.74%), and pMCI vs. sMCI classification (83.25%), while DSNN achieves 93.02%/82.60% while using 3 T scans and 91.22%/80.05% while using synthetic 7 T scans. Besides, while using only a single modality, using our synthetic 7 T images perform comparable but slightly less well than using original 3 T images. It reveals that our synthetic 7 T images can not take place of the real 7 T images even it is useful in disease diagnosis. Nevertheless, we achieve better performance than DSNN, which demonstrates that our two-stage framework is effective to boost diagnosis performance with only T1 MR images.

4 Conclusion

We aim to improve the quality of MRI by learning a mapping from 3T scans to 7 T scans, for which the limited data is the most challenging difficulty. To address this, we propose the Unlimited Data Augmentation (UDA) strategy to increase the training samples by deformable registration. Experiments suggest that this strategy can boost the performance of various GAN-based synthesis techniques, and the performance of the AD diagnosis tasks, i.e., identifying AD subjects and predicting MCI conversion, can be further improved with our synthetic 7 T scans.

Acknowledgment. This work was supported partly by the National Natural Science Foundation of China under Grants 61771397, partly by the CAAI-Huawei MindSpore Open Fund under Grants CAAIXSJLJJ-2020-005B, and partly by the China Postdoctoral Science Foundation under Grants BX2021333.

References

1. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. Med. Image Anal. **12**(1), 26–41 (2008)
2. Cohen, J.P., Luck, M., Honari, S.: Distribution matching losses can hallucinate features in medical image translation. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11070, pp. 529–536. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00928-1_60
3. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S.: Generative adversarial networks. In: Advances in Neural Information Processing Systems (NIPS), pp. 2672–2680 (2014)
4. Horé, A., Ziou, D.: Image quality metrics: PSNR vs. SSIM. In: 20th International Conference on Pattern Recognition (ICPR) (2010)
5. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 5967–5976 (2017)
6. Jack, C.R., Bernstein, M.A., Fox, N.C., Thompson, P., Weiner, M.W.: The Alzheimer's disease neuroimaging initiative (ADNI): MRI methods. J. Magn. Reson. Imaging **27**(4), 685–691 (2010)
7. Jiang, J., et al.: PSIGAN: joint probabilistic segmentation and image distribution matching for unpaired cross-modality adaptation-based MRI segmentation. IEEE Trans. Med. Imaging **39**(12), 4071–4084 (2020)
8. Koyejo, O.O., Natarajan, N., Ravikumar, P.K., Dhillon, I.S.: Consistent binary classification with generalized performance metrics. In: Advances in Neural Information Processing Systems (NIPS), pp. 2744–2752 (2014)
9. Lian, C., Zhang, J., Liu, M., Zong, X., Hung, S.C., Lin, W., Shen, D.: Multi-channel multi-scale fully convolutional network for 3D perivascular spaces segmentation in 7T MR images. Med. Image Anal. **46**, 106–117 (2018)
10. Long, J.M., Holtzman, D.M.: Alzheimer disease: an update on pathobiology and treatment strategies. Cell **179**(2) (2019)
11. Obusez, E.C., et al.: 7T MR of intracranial pathology: preliminary observations and comparisons to 3T and 1.5T. Neuroimage **168**, 459–476 (2018)

12. Pan, Y., Liu, M., Lian, C., Xia, Y., Shen, D.: Disease-image specific generative adversarial network for brain disease diagnosis with incomplete multi-modal neuroimages. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11766, pp. 137–145. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32248-9_16
13. Pan, Y., Liu, M., Lian, C., Xia, Y., Shen, D.: Spatially-constrained fisher representation for brain disease identification with incomplete multi-modal neuroimages. IEEE Trans. Med. Imaging **39**(9), 2965–2975 (2020)
14. Pan, Y., Xia, Y.: Ultimate reconstruction: understand your bones from orthogonal views. In: 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI), pp. 1155–1158 (2021)
15. Qu, L., Zhang, Y., Wang, S., Yap, P.T., Shen, D.: Synthesized 7T MRI from 3T MRI via deep learning in spatial and wavelet domains. Med. Image Anal. **62**, 101663 (2020)
16. Sperling, R., Mormino, E., Johnson, K.: The evolution of preclinical Alzheimer's disease: implications for prevention trials. Neuron **84**(3), 608–622 (2014)
17. 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: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 8798–8807 (2018)
18. Zhou, X., Qiu, S., Joshi, P.S., Xue, C., Kolachalama, V.B.: Enhancing magnetic resonance imaging-driven Alzheimer's disease classification performance using generative adversarial learning. Alzheimer's Res. Ther. **13**(1) (2021)
19. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: IEEE International Conference on Computer Vision (ICCV), pp. 2242–2251 (2017)