

A Unified Hyper-GAN Model for Unpaired Multi-contrast MR Image Translation

*****@*****

Abstract. Cross-contrast image translation is an important task for completing missing contrasts in clinical diagnosis. However, most existing methods learn separate translator for each pair of contrasts, which is inefficient due to many possible contrast pairs in real scenarios. In this work, we propose a unified Hyper-GAN model for effectively and efficiently translating between different contrast pairs. Hyper-GAN consists of a pair of hyper-encoder and hyper-decoder to first map from the source contrast to a common feature space, and then further map to the target contrast image. To facilitate the translation between different contrast pairs, contrast-modulators are designed to tune the hyper-encoder and hyper-decoder adaptive to different contrasts. We also design a common space loss to enforce that multi-contrast images of a subject share a common feature space, implicitly modeling the shared underlying anatomical structures. Experiments on two datasets of IXI and BraTS 2019 show that our Hyper-GAN achieves state-of-the-art results in both accuracy and efficiency, e.g., improving more than 1.47 and 1.15 dB in PSNR on two datasets with less than half the amount of parameters.

Keywords: Multi-contrast MR · Unpaired image translation · Unified hyper-GAN.

1 Introduction

Magnetic resonance (MR) imaging has been widely utilized in clinical diagnosis, as it has a range of imaging contrasts and largely increases the diversity of diagnostic information. However, due to practical limits, e.g., long scan time [22], image corruption [26], etc., it is often hard to collect all multi-contrast MR images of one subject. To solve this problem, a large variety of synthesis methods [4, 6, 12, 13, 17, 18, 20, 23, 24] try to synthesize missing contrast from the available contrast, and most of them are one-to-one cross-contrast synthesis methods, i.e., one model is trained for each specific pair of contrasts. For example, Dar et al. [6] proposed a conditional generative adversarial network to translate between T1w and T2w images. However, it is impractical to train each network for each pair of contrasts due to a range of commonly used MR contrasts in real scenarios. In this work, we focus on tackling the multi-contrast MR image translation problem by a

more efficient and effective way in an unpaired training setting, i.e., the training multi-contrast MR images are not required to be captured from same subjects.

There are already several unpaired multi-contrast image translation methods in literature, e.g., [2, 5, 14, 21], in recent years. Firstly, CycleGAN [28], as a one-to-one cross-contrast synthesis method, could be extended to learn multi-contrast mappings and require $N \times (N - 1)$ generators for N contrasts, which is impractical in real scenarios. Furthermore, ComboGAN [2] and Domain-Bank [14] decouple generator into encoder/decoder and reduce the requirement to N contrast-specific encoders/decoders for N contrasts, while the parameter size and training time of them scale linearly with the contrast number. In addition, StarGAN [5] and SUGAN [21] share a generator and discriminator for all contrasts, and rely on a contrast indicator to specify the desired output contrast. However, since the contrast indicator is simply concatenated with input image, it might be insufficient to control the translation process [1, 15]. In summary, existing methods either depend on multiple encoders/decoders, or insufficiently control/module the generator to be adaptive to contrast.

In this work, we aim to design a unified deep network for unpaired multi-contrast MR image translation in an efficient and effective way. We design shared encoder and decoder to translate between different contrast pairs based on a common feature space constraint, and the encoding and decoding processes are respectively modulated by the source and target contrast codes. This is inspired by MR imaging [3] that multi-contrast MR images are determined by human intrinsic tissue and scanner imaging parameters. The common feature space implicitly models the intrinsic tissue parameters, and the scanner imaging parameters are encoded to modulate the encoder and decoder in our network design.

Specifically, we first define contrast-specific information as one-hot code indicating MR contrast, and construct two contrast modulators as hyper-network [11] to respectively modulate the encoder and decoder to be adaptive to different MR contrasts. To enforce the common feature space shared by different contrasts, we further design a common space loss to enforce extracted deep features from different contrasts within a common feature space, implicitly modeling the shared underlying anatomical structures, besides traditional adversarial [10] and cycle-consistency [28] losses. This unified multi-contrast MR image translation model, dubbed *Hyper-GAN*, can effectively and efficiently translate between different contrast pairs using a single network. Experiments on two multi-contrast brain MR datasets of IXI and BraTS 2019 show that Hyper-GAN achieves start-of-the-art results in both accuracy and efficiency, e.g., improving more than 1.47 and 1.15 dB in PSNR on two datasets with less than half the amount of parameters.

2 Method

As shown in Fig. 1, our *Hyper-GAN* utilizes a shared pair of *hyper-encoder* and *hyper-decoder* to translate between different contrast pairs. The source and target contrasts are each represented by a one-hot code, dubbed *contrast code*, with the value of 1 representing the corresponding contrast in a given list of multiple




Fig. 1. Hyper-GAN utilizes shared hyper-encoder and hyper-decoder to translate between different contrast pairs, and the encoding and decoding processes are respectively modulated by the source and target contrast codes. Multi-contrast images of a subject are constrained to be within a common feature space after encoding.

contrasts. These codes are further utilized to adaptively tune the encoding and decoding processes in our hyper-encoder and hyper-decoder. Specifically, our hyper-encoder/hyper-decoder is respectively an encoder/decoder with parameters modulated by a *contrast modulator* with contrast code as input, in order to make the hyper-encoder/hyper-decoder adaptive to different contrasts. The extracted deep features by hyper-encoder are constrained to be within a common feature space shared by different contrast images of a subject. In this way, our Hyper-GAN is an encoder-decoder network adaptive to different source and target contrast pairs. We next introduce the details.

2.1 Network Architecture

As shown in Fig. 2(a) and (b), hyper-encoder E and hyper-decoder G respectively consist of two subnets, i.e., a backbone encoder/decoder and a contrast modulator. The encoder extracts deep features from a source contrast image, while the decoder estimates the target contrast image from the extracted features. The encoder and decoder are respectively paired with a contrast modulator, achieving contrast-adaptive tuning of parameters of encoder and decoder based on the source and target contrast codes. With the contrast code as input, the contrast modulator tunes the parameters using the following two different strategies, i.e., filter scaling or conditional instance normalization.

Filter scaling (FS) [1]. For each convolutional filter f in encoder/decoder, the modulator produces a corresponding scalar α based on the contrast code and modifies this filter as $f' \triangleq \alpha * f$, where $*$ is scalar multiplication operation.

Conditional instance normalization (CIN) [8]. For each instance normalization (IN) layer in encoder/decoder, the modulator estimates its affine parameters γ' and β' based on contrast code, and then IN layer becomes $y = \gamma' \frac{x - \mu(x)}{\sigma(x)} + \beta'$, where $\mu(x)$ and $\sigma(x)$ are mean and standard deviation of the input features x .

Network architecture. Encoder and decoder are respectively taken as the first five and the remaining four residual blocks of the generator in [28], and contrast modulators are set to multilayer perceptrons. Please refer to Fig. 2(a) and

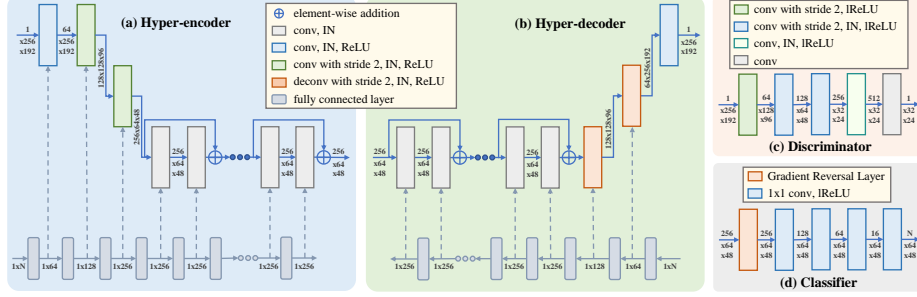


Fig. 2. The illustration of the detailed network architecture in Hyper-GAN.

(b) for the detailed definitions of hyper-encoder and hyper-decoder. As shown in Fig. 2(c), discriminators, as in [28], are fully convolutional networks [19] with five 4×4 convolutional layers to classify whether 70×70 overlapping image patches are real or synthetic. The contrast-classifier is composed of a gradient reversal layer [9] and four 1×1 convolutional layers to classify which contrast the features extracted by hyper-encoder belong to and outputs contrast probabilities.

2.2 Training Loss

For a N -contrast translation task with contrasts $\{m_k\}_{k=1}^N$ and one-hot contrast codes $\{c_k \in \{0, 1\}^N\}_{k=1}^N$, our Hyper-GAN consists of a pair of hyper-encoder E and hyper-decoder G , a contrast-classifier C and contrast-specific discriminators $\{D_k\}_{k=1}^N$. Note that C and $\{D_k\}_{k=1}^N$ are only utilized in training.

Our Hyper-GAN is trained on unpaired multi-contrast MR images, i.e., different contrasts are acquired from different subjects without requiring accurate registration. When training the network, two random training images I_{m_i}, I_{m_j} in distinct contrasts $m_i, m_j (i \neq j)$ are fed into Hyper-GAN as the source and target contrasts for updating network parameters. For simplicity, we next present the training loss for the translation from contrast m_i to m_j .

Adversarial loss [10]. The hyper-encoder E and hyper-decoder G are required to generate a synthetic image $G(E(I_{m_i}, c_i), c_j)$ close to a real image, while the discriminator D_j is to distinguish between this synthetic image $G(E(I_{m_i}, c_i), c_j)$ and a real image I_{m_j} . The adversarial loss is defined as

$$\mathcal{L}_{adv}(E, G, D_j) = D_j(G(E(I_{m_i}, c_i), c_j))^2 + (1 - D_j(I_{m_j}))^2, \quad (1)$$

where $E(\cdot, c_i)$ and $G(\cdot, c_j)$ respectively denote the outputs of hyper-encoder E and hyper-decoder G tuned by contrast codes c_i and c_j . $G(E(I_{m_i}, c_i), c_j)$ represents the synthetic m_j contrast image translated from an m_i contrast input.

Cycle-consistency loss [28]. A cycle-consistency loss for E and G is to force the reconstructed image $G(E(G(E(I_{m_i}, c_i), c_j), c_i))$ (mapping from contrast m_i to m_j and back) to be identical to the input I_{m_i} , which is written as

$$\mathcal{L}_{cyc}(E, G) = \|G(E(G(E(I_{m_i}, c_i), c_j), c_i)) - I_{m_i}\|_1. \quad (2)$$

Common space loss. The hyper-encoder is adaptive to the contrast of input image, and we expect that the extracted features by hyper-encoder should be in a common space shared by multi-contrast images for each subject, implicitly modeling the intrinsic tissue parameters of the subject. As paired training multi-contrast images are not given, we design the following common space constraints.

(i) Identical loss. As the images I_{m_i} and $G(E(I_{m_i}, c_i), c_j)$ are of the same subject but different contrasts, an identical loss is designed to enforce their extracted features by hyper-encoder E to be identical, which is written as

$$\mathcal{L}_{id}(E, G) = \|E(G(E(I_{m_i}, c_i), c_j), c_j) - E(I_{m_i}, c_i)\|_1 . \quad (3)$$

(ii) Reconstruction loss. A reconstruction loss constrains that image I_{m_i} could reconstruct itself by hyper-encoder and hyper-decoder, which is defined as

$$\mathcal{L}_{rec}(E, G) = \|G(E(I_{m_i}, c_i), c_i) - I_{m_i}\|_1 . \quad (4)$$

(iii) Contrast-classification loss. Contrast-classification loss is to force classifier C to predict the contrast of extracted features by hyper-encoder E . We adversarially train the classifier to make deep features of multiple contrasts extracted by E to be same distributed, i.e., within a common feature space, using a gradient reversal layer [9] in C , which flips gradient sign during backpropagation to force extracted deep features unable to be classified by C . This loss is defined as

$$\begin{aligned} \mathcal{L}_{cla}(E, G, C) = & L_{CE}(C(E(G(E(I_{m_i}, c_i), c_j), c_j)), c_j) \\ & + L_{CE}(C(E(I_{m_i}, c_i)), c_i) , \end{aligned} \quad (5)$$

where L_{CE} computes the cross entropy between estimated contrast probability by C and real contrast code. Then the common space loss is defined as

$$\mathcal{L}_{com}(E, G, C) = \lambda_{id}\mathcal{L}_{id} + \lambda_{rec}\mathcal{L}_{rec} + \lambda_{cla}\mathcal{L}_{cla} . \quad (6)$$

We heuristically set λ_{id} , λ_{rec} and λ_{cla} to 0.5, 10 and 0.2 to make each term in a similar range of loss values as adversarial loss.

Total loss. The total training loss of Hyper-GAN is summation of above training losses over all training image pairs in distinct contrasts, which is defined as

$$\mathcal{L}(E, G, C, D_j) = \mathcal{L}_{adv} + \lambda_{cyc}\mathcal{L}_{cyc} + \mathcal{L}_{com} , \quad (7)$$

where λ_{cyc} is set to 10 as per [28]. To optimize \mathcal{L} , the networks are divided into two groups, i.e., $\{D_j\}_{j=1}^N$ and $\{E, G, C\}$, which are alternately updated, and the networks are totally optimized in 100 epochs. As in [28], the learning rate is set to 0.0002 with a batch size of 1. Our source code will be released on GitHub.

3 Experiments

3.1 Data Sets

IXI dataset¹. We utilize all 319 subjects from Guy’s Hospital, and randomly split them into 150, 5 and 164 subjects for training, validation and testing. Each

¹ <https://brain-development.org/ixi-dataset/>

Table 1. Accuracies of different methods for arbitrary cross-contrast MR image translation on IXI dataset, which are averaged over test set and all 6 translation tasks. Each cell is formatted as “mean (standard deviation)”.

Method	Ablated version	MAE	PSNR	SSIM
CycleGAN		0.0175 (0.0045)	27.18 (2.17)	0.857 (0.055)
StarGAN		0.0174 (0.0052)	27.77 (2.81)	0.856 (0.068)
SUGAN		0.0182 (0.0055)	27.30 (2.88)	0.845 (0.068)
ComboGAN		0.0163 (0.0040)	28.17 (2.34)	0.876 (0.045)
Ours (CIN)	$\mathcal{L}_{adv+cyc}$	0.0145 (0.0039)	28.89 (2.45)	0.897 (0.044)
	$\mathcal{L}_{adv+cyc+id}$	0.0141 (0.0037)	29.12 (2.46)	0.906 (0.040)
	$\mathcal{L}_{adv+cyc+rec}$	0.0144 (0.0043)	29.09 (2.76)	0.900 (0.046)
	$\mathcal{L}_{adv+cyc+cla}$	0.0145 (0.0037)	28.86 (2.32)	0.902 (0.038)
	$\mathcal{L}_{adv+cyc+id+cla}$	0.0138 (0.0037)	29.36 (2.54)	0.909 (0.037)
	$\mathcal{L}_{adv+cyc+id+rec+cla}$	0.0138 (0.0043)	29.45 (3.03)	0.910 (0.042)
Ours (FS)	$\mathcal{L}_{adv+cyc}$	0.0154 (0.0032)	28.28 (1.74)	0.897 (0.030)
	$\mathcal{L}_{adv+cyc+id}$	0.0141 (0.0040)	29.24 (2.73)	0.904 (0.041)
	$\mathcal{L}_{adv+cyc+rec}$	0.0140 (0.0047)	29.39 (3.22)	0.908 (0.045)
	$\mathcal{L}_{adv+cyc+cla}$	0.0141 (0.0036)	29.04 (2.47)	0.903 (0.041)
	$\mathcal{L}_{adv+cyc+id+cla}$	0.0135 (0.0033)	29.49 (2.27)	0.910 (0.034)
	$\mathcal{L}_{adv+cyc+id+rec+cla}$	0.0133 (0.0038)	29.64 (2.70)	0.910 (0.040)

subject contains three contrasts (T1w, T2w and PDw), and only one of three contrasts per subject is used for training to generate unpaired data.

BraTS 2019 dataset². We use all 150 subjects from CBICA institution, and split them into 100, 5 and 45 subjects for training, validation and testing. Each subject contains four contrasts (T1w, T1Gd, T2w, FLAIR), and only one of four contrasts per subject is used for training. All volumes of both datasets are N4 corrected and peak normalized, and the intensities are linearly scaled to [0, 1].

3.2 Experimental Results

We compare our Hyper-GAN with the state-of-the-art unpaired multi-contrast image translation methods, including StarGAN [5], SUGAN [21] and ComboGAN [2], for arbitrary cross-contrast MR image translation. Our Hyper-GAN includes two versions, i.e., the contrast modulator using filter scaling (FS) or conditional instance normalization (CIN). All the experiments are performed on 2D sagittal slices. Quantitatively, we compute the mean absolute error (MAE), peak signal-to-noise ratio (PSNR), and structural similarity (SSIM) between 3D volumes of ground truth and image translation results. The final accuracy is averaged over the test set and all translation tasks between all possible source/target

² <https://www.med.upenn.edu/cbica/brats2019.html>

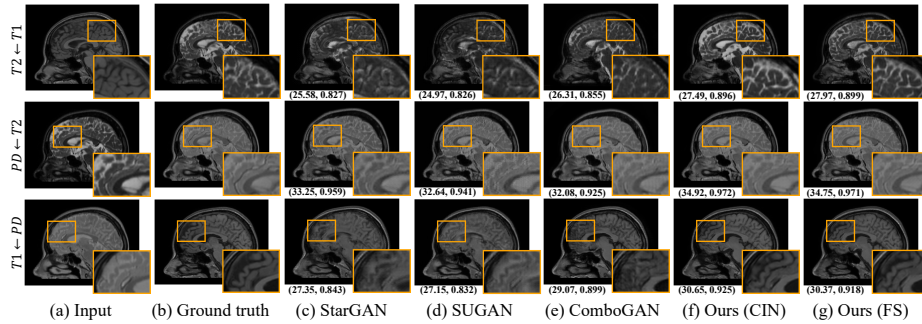


Fig. 3. Visual comparison of cross-contrast MR image translation results on IXI dataset. Top to bottom: T1w-to-T2w, T2w-to-PDw and PDw-to-T1w translation results on a test subject. The values under each sub-image are PSNR and SSIM scores.

contrast pairs, i.e., 6 tasks for IXI and 12 tasks for BraTS 2019. A paired two-sided Wilcoxon signed-rank test is conducted to compare the performance.

Results on IXI dataset. Table 1 reports the results on IXI dataset of three contrasts. It shows that, compared with StarGAN, SUGAN and ComboGAN, our Hyper-GAN (i.e., $\mathcal{L}_{adv} + cyc + id + rec + cla$) achieves significantly better performance in all metrics ($p < .001$) and produces 29.45/29.64 (using CIN/FS respectively) in PSNR, comparing favorably with 27.77 of StarGAN, 27.30 of SUGAN and 28.17 of ComboGAN.

Effectiveness of training losses. In Table 1, we also compare Hyper-GAN trained with different losses defined in Sect. 2.2. Interestingly, even our baseline (i.e., $\mathcal{L}_{adv} + cyc$) performs better than compared methods ($p < .005$) and obtains 28.89/28.28 in PSNR, justifying effectiveness of our network design with hyper-encoder/hyper-decoder using contrast modulators. Starting from this baseline, each extra loss enforcing common feature space constraint improves results in all metrics except that contrast-classification loss in CIN version produces higher SSIM but comparable MAE and PSNR scores. Specifically, the identical and reconstruction losses respectively improve PSNR from 28.89/28.28 of baseline to 29.12/29.24 and 29.09/29.39. In addition, the PSNR scores are further improved to 29.36/29.49 when both identical and contrast-classification losses are utilized and our whole model performs best. This indicates the effectiveness of our defined common space loss implicitly modeling the shared underlying anatomical structures of each subject. Both FS and CIN versions achieve good performance while FS works generally better. Figure 3 shows visual comparison results.

Results on BraTS 2019 dataset. Table 2 reports the results on BraTS 2019 dataset of four contrasts. Our Hyper-GAN obtains significantly higher accuracies than compared methods in all metrics ($p < .001$) and achieves 31.66/31.90 (using CIN/FS respectively) in PSNR, comparing favorably with 30.75 of StarGAN, 30.37 of SUGAN and 30.55 of ComboGAN. Figure 4 shows the T1Gd-to-FLAIR image translation results, and our method performs best even with lesions.

Table 2. Accuracies of different methods for arbitrary cross-contrast MR image translation on BraTS 2019 dataset, which are averaged over test set and all 12 translation tasks. Each cell is formatted as “mean (standard deviation)”.

Method	MAE	PSNR	SSIM	# of parameters
StarGAN	0.0083 (0.0029)	30.75 (2.64)	0.905 (0.028)	50.72M
SUGAN	0.0087 (0.0021)	30.37 (1.97)	0.909 (0.013)	82.63M
ComboGAN	0.0081 (0.0025)	30.55 (2.17)	0.913 (0.015)	71.02M
Ours (CIN)	0.0072 (0.0022)	31.66 (2.46)	0.922 (0.018)	23.71M
Ours (FS)	0.0070 (0.0026)	31.90 (2.74)	0.930 (0.018)	22.57M




Fig. 4. Visual comparison of T1Gd-to-FLAIR image translation results on BraTS 2019 dataset. The values under each sub-image are PSNR and SSIM scores.

Comparison with CycleGAN. We also compare our Hyper-GAN with CycleGAN [28] in Table 1. Note that concatenation of backbones of hyper-encoder and hyper-decoder is identical to generator in CycleGAN as discussed in Sect. 2.1. We train individual CycleGAN for each possible source/target contrast pair and average the results over all pairs as final accuracy. Hyper-GAN significantly outperforms CycleGAN ($p < .005$), demonstrating the effectiveness of our network design based on contrast modulator and common feature space constraint.

Computational Efficiency. As shown in Table 2, our parameter size is 2 and 3 times smaller than StarGAN and SUGAN/ComboGAN respectively. Also for BraTS 2019 dataset, CycleGAN requires 169.56M parameters that is 7 times larger than ours. This is because Hyper-GAN utilizes only a single pair of hyper-encoder and hyper-decoder, while the number of encoders/decoders in ComboGAN or CycleGAN scales linearly or quadratically with contrast number.

4 Conclusion

We have proposed a unified GAN for unpaired multi-contrast MR image translation. It can flexibly translate between different contrast pairs using a unified network consisting of hyper-encoder/hyper-decoder and common feature space constraint. This design enables the network to fully investigate the common anatomical structures by common feature space modeling, and contrast-specific imaging by hyper-network design of hyper-encoder and hyper-decoder. It achieves state-of-the-art performance, and outperforms previous GAN-based image translation models that depend on multiple encoders/decoders for different contrast pairs. In the future, we are interested in the extension to multi-contrast, multi-institute setting [7], as well as the combination with segmentation task [16, 25, 27].

References

1. Alharbi, Y., Smith, N., Wonka, P.: Latent filter scaling for multimodal unsupervised image-to-image translation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1458–1466 (2019)
2. Anoosheh, A., Agustsson, E., Timofte, R., Van Gool, L.: Combogan: Unrestrained scalability for image domain translation. In: IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPR-W). pp. 783–790 (2018)
3. Bernstein, M.A., King, K.F., Zhou, X.J.: Handbook of MRI pulse sequences. Elsevier (2004)
4. Bui, T.D., Nguyen, M., Le, N., Luu, K.: Flow-based deformation guidance for unpaired multi-contrast mri image-to-image translation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI). pp. 728–737 (2020)
5. Choi, Y., Choi, M., Kim, M., Ha, J.W., Kim, S., Choo, J.: Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8789–8797 (2018)
6. Dar, S.U., Yurt, M., Karacan, L., Erdem, A., Erdem, E., Çukur, T.: Image synthesis in multi-contrast MRI with conditional generative adversarial networks. IEEE Transactions on Medical Imaging **38**(10), 2375–2388 (2019)
7. Dewey, B.E., Zuo, L., Carass, A., He, Y., Liu, Y., Mowry, E.M., Newsome, S., Oh, J., Calabresi, P.A., Prince, J.L.: A disentangled latent space for cross-site MRI harmonization. In: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI). pp. 720–729 (2020)
8. Dumoulin, V., Shlens, J., Kudlur, M.: A learned representation for artistic style. In: International Conference on Learning Representations (ICLR) (2017)
9. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: International Conference on Machine Learning (ICML). pp. 1180–1189 (2015)
10. 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). pp. 2672–2680 (2014)
11. Ha, D., Dai, A.M., Le, Q.V.: Hypernetworks. In: International Conference on Learning Representations (ICLR) (2017)
12. Huang, Y., Shao, L., Frangi, A.F.: DOTE: Dual convolutional filter learning for super-resolution and cross-modality synthesis in MRI. In: International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI). pp. 89–98 (2017)
13. Huang, Y., Zheng, F., Cong, R., Huang, W., Scott, M.R., Shao, L.: MCMT-GAN: Multi-task coherent modality transferable GAN for 3D brain image synthesis. IEEE Transactions on Image Processing **29**, 8187–8198 (2020)
14. Hui, L., Li, X., Chen, J., He, H., Yang, J.: Unsupervised multi-domain image translation with domain-specific encoders/decoders. In: International Conference on Pattern Recognition (ICPR). pp. 2044–2049 (2018)
15. 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. 1125–1134 (2017)
16. Jiang, J., Veeraraghavan, H.: Unified cross-modality feature disentangler for unsupervised multi-domain MRI abdomen organs segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI). pp. 347–358 (2020)

17. Jog, A., Carass, A., Roy, S., Pham, D.L., Prince, J.L.: Random forest regression for magnetic resonance image synthesis. *Medical Image Analysis* **35**, 475–488 (2017)
18. Liu, X., Xing, F., Prince, J.L., Carass, A., Stone, M., Fakhri, G.E., Woo, J.: Dual-cycle constrained bijective VAE-GAN for tagged-to-cine magnetic resonance image synthesis. In: *International Symposium on Biomedical Imaging (ISBI)* (2021)
19. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 3431–3440 (2015)
20. Roy, S., Chou, Y.Y., Jog, A., Butman, J.A., Pham, D.L.: Patch based synthesis of whole head MR images: Application to EPI distortion correction. In: *International Workshop on Simulation and Synthesis in Medical Imaging (SASHIMI)*. pp. 146–156 (2016)
21. Sohail, M., Riaz, M.N., Wu, J., Long, C., Li, S.: Unpaired multi-contrast MR image synthesis using generative adversarial networks. In: *International Workshop on Simulation and Synthesis in Medical Imaging (SASHIMI)*. pp. 22–31 (2019)
22. Vranic, J., Cross, N., Wang, Y., Hippe, D., de Weerd, E., Mossa-Basha, M.: Compressed sensing-sensitivity encoding (CS-SENSE) accelerated brain imaging: reduced scan time without reduced image quality. *American Journal of Neuroradiology* **40**(1), 92–98 (2019)
23. Yu, B., Zhou, L., Wang, L., Shi, Y., Fripp, J., Bourgeat, P.: Ea-GANs: edge-aware generative adversarial networks for cross-modality MR image synthesis. *IEEE Transactions on Medical Imaging* **38**(7), 1750–1762 (2019)
24. Yu, B., Zhou, L., Wang, L., Shi, Y., Fripp, J., Bourgeat, P.: Sample-adaptive GANs: Linking global and local mappings for cross-modality MR image synthesis. *IEEE Transactions on Medical Imaging* **39**(7), 2339–2350 (2020)
25. Yuan, W., Wei, J., Wang, J., Ma, Q., Tasdizen, T.: Unified attentional generative adversarial network for brain tumor segmentation from multimodal unpaired images. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 229–237 (2019)
26. Zaitsev, M., Maclaren, J., Herbst, M.: Motion artifacts in MRI: A complex problem with many partial solutions. *Journal of Magnetic Resonance Imaging* **42**(4), 887–901 (2015)
27. Zhou, T., Canu, S., Vera, P., Ruan, S.: Brain tumor segmentation with missing modalities via latent multi-source correlation representation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 533–541 (2020)
28. 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. 2223–2232 (2017)