



Adversarial Uni- and Multi-modal Stream Networks for Multimodal Image Registration

Zhe Xu^{1,2}, Jie Luo^{2,3}, Jiangpeng Yan¹, Ritvik Pulya², Xiu Li¹,
William Wells III², and Jayender Jagadeesan²(✉)

¹ Shenzhen International Graduate School, Tsinghua University, Shenzhen, China

² Brigham and Women's Hospital, Harvard Medical School, Boston, USA
jayender@bwh.harvard.edu

³ Graduate School of Frontier Sciences, The University of Tokyo,
Tokyo, Japan

Abstract. Deformable image registration between Computed Tomography (CT) images and Magnetic Resonance (MR) imaging is essential for many image-guided therapies. In this paper, we propose a novel translation-based unsupervised deformable image registration method. Distinct from other translation-based methods that attempt to convert the multimodal problem (e.g., CT-to-MR) into a unimodal problem (e.g., MR-to-MR) via image-to-image translation, our method leverages the deformation fields estimated from both: (i) the translated MR image and (ii) the original CT image in a dual-stream fashion, and automatically learns how to fuse them to achieve better registration performance. The multimodal registration network can be effectively trained by computationally efficient similarity metrics without any ground-truth deformation. Our method has been evaluated on two clinical datasets and demonstrates promising results compared to state-of-the-art traditional and learning-based methods.

Keywords: Multimodal registration · Generative adversarial network · Unsupervised learning

1 Introduction

Deformable multimodal image registration has become essential for many procedures in image-guided therapies, e.g., preoperative planning, intervention, and diagnosis. Due to substantial improvement in computational efficiency over traditional iterative registration approaches, learning-based registration approaches are becoming more prominent in time-intensive applications.

Related Work. Many learning-based registration approaches adopt fully supervised or semi-supervised strategies. Their networks are trained with ground-truth

deformation fields or segmentation masks [5, 12, 13, 16, 19], and may struggle with limited or imperfect data labeling. A number of unsupervised registration approaches have been proposed to overcome this problem by training unlabeled data to minimize traditional similarity metrics, e.g., mean squared intensity differences [4, 11, 15, 17, 21, 26]. However, the performances of these methods are inherently limited by the choice of similarity metrics. Given the limited selection of multimodal similarity metrics, unsupervised registration approaches may have difficulties outperforming traditional multimodal registration methods as they both essentially optimize the same cost functions. A recent trend for multimodal image registration takes advantage of the latent feature disentanglement [18] and image-to-image translation [6, 20, 23]. Specifically, translation-based approaches use Generative Adversarial Network (GAN) to translate images from one modality into the other modality, thus are able to convert the difficult multimodal registration into a simpler unimodal task. However, being a challenging topic by itself, image translation may inevitably produce artificial anatomical features that can further interfere with the registration process.

In this work, we propose a novel translation-based fully unsupervised multimodal image registration approach. In the context of Computed Tomography (CT) image to Magnetic Resonance (MR) image registration, previous translation-based approaches would translate a CT image into an MR-like image (tMR), and use tMR-to-MR registration to estimate the final deformation field ϕ . In our approach, the network estimates two deformation fields, namely ϕ_s of tMR-to-MR and ϕ_o of CT-to-MR, in a dual-stream fashion. The addition of the original ϕ_o enables the network to implicitly regularize ϕ_s to mitigate certain image translation problems, e.g., artificial features. The network further automatically learns how to fuse ϕ_s and ϕ_o towards achieving the best registration accuracy.

Contributions and advantages of our work can be summarized as follows:

1. Our method leverages the deformation fields estimated from the original multimodal stream and synthetic unimodal stream to overcome the shortcomings of translation-based registration;
2. We improve the fidelity of organ boundaries in the translated MR by adding two extra constraints in the image-to-image translation model Cycle-GAN.

We evaluate our method on two clinically acquired datasets. It outperforms state-of-the-art traditional, unsupervised and translation-based registration approaches.

2 Methods

In this work, we propose a general learning framework for robustly registering CT images to MR images in a fully unsupervised manner.

First, given a moving CT image and a fixed MR image, our improved CycleGAN module translates the CT image into an MR-like image. Then, our dual-stream subnetworks, UNet_o and UNet_s, estimate two deformation fields ϕ_o and

ϕ_s respectively, and the final deformation field is fused via a proposed fusion module. Finally, the moving CT image is warped via Spatial Transformation Network (STN) [14], while the entire registration network aims to maximize the similarity between the moved and the fixed images. The pipeline of our method is shown in Fig. 1.

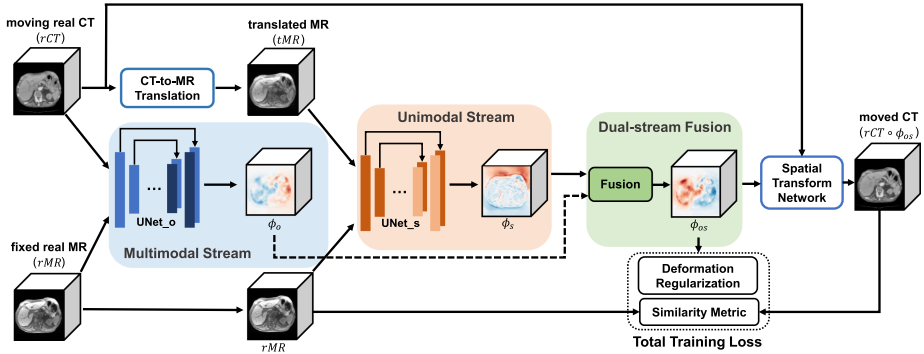


Fig. 1. Illustration of the proposed method. The entire unsupervised network is mainly guided by the image similarity between $rCT \circ \phi_{os}$ and rMR .

2.1 Image-to-Image Translation with Unpaired Data

The CT-to-MR translation step consists of an improved Cycle-GAN with additional structural and identical constraints. As a state-of-the-art image-to-image translation model, Cycle-GAN [28] can be trained without pairwise aligned CT and MR datasets of the same patient. Thus, Cycle-GAN is widely used in medical image translation [1, 9, 25].

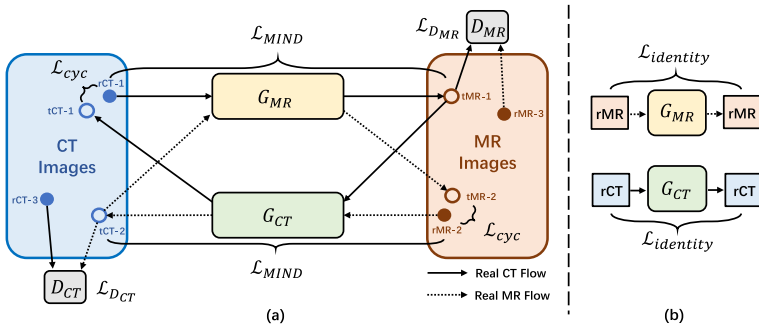


Fig. 2. Schematic illustration of Cycle-GAN with strict constraints. (a) The workflow of the forward and backward translation; (b) The workflow of identity loss.

Our Cycle-GAN model is illustrated in Fig. 2. The model consists of two generators G_{MR} and G_{CT} , which can provide CT-to-MR and MR-to-CT translation respectively. Besides, it has two discriminators D_{CT} and D_{MR} . D_{CT} is used to distinguish between translated CT (tCT) and real CT (rCT), and D_{MR} is for translated MR (tMR) and real MR (rMR). The training loss of original Cycle-GAN only adopts two types of items: adversarial loss given by two discriminators ($\mathcal{L}_{D_{CT}}$ and $\mathcal{L}_{D_{MR}}$) and cycle-consistency loss \mathcal{L}_{cyc} to prevent generators from generating images that are not related to the inputs (refer to [28] for details).

However, training a Cycle-GAN on medical images is difficult since the cycle-consistency loss is not enough to enforce structural similarity between translated images and real images (as shown in the red box in Fig. 3(b)). Therefore, we introduce two additional losses, structure-consistency loss \mathcal{L}_{MIND} and identity loss $\mathcal{L}_{identity}$, to constrain the training of Cycle-GAN.

MIND (Modality Independent Neighbourhood Descriptor) [8] is a feature that describes the local structure around each voxel. Thus, we minimize the difference in MIND between translated images $G_{CT}(I_{rMR})$ or $G_{MR}(I_{rCT})$ and real images I_{rMR} or I_{rCT} to enforce the structural similarity. We define \mathcal{L}_{MIND} as follows:

$$L_{MIND}(G_{CT}, G_{MR}) = \frac{1}{N_{MR}|R|} \sum_x \|M(G_{CT}(I_{rMR})) - M(I_{rMR})\|_1 + \frac{1}{N_{CT}|R|} \sum_x \|M(G_{MR}(I_{rCT})) - M(I_{rCT})\|_1 \quad (1)$$

where M represents MIND features, N_{MR} and N_{CT} denote the number of voxels in I_{rMR} and I_{rCT} , and R is a non-local region around voxel x .




Fig. 3. CT-to-MR translation examples of original Cycle-GAN and proposed Cycle-GAN tested for (a) pig ex-vivo kidney dataset and (b) abdomen dataset.

The identity loss (as shown in Fig. 2(b)) is included to prevent images already in the expected domain from being incorrectly translated to the other domain. We define it as:

$$\mathcal{L}_{identity} = \|G_{MR}(I_{MR}) - I_{MR}\|_1 + \|G_{CT}(I_{CT}) - I_{CT}\|_1 \quad (2)$$

Finally, the total loss \mathcal{L} of our proposed Cycle-GAN is defined as:

$$\mathcal{L} = \mathcal{L}_{D_{MR}} + \mathcal{L}_{D_{CT}} + \lambda_{cyc}\mathcal{L}_{cyc} + \lambda_{identity}\mathcal{L}_{identity} + \lambda_{MIND}\mathcal{L}_{MIND} \quad (3)$$

where λ_{cyc} , $\lambda_{identity}$ and λ_{MIND} denotes the relative importance of each term.

2.2 Dual-Stream Multimodal Image Registration Network

As shown in Fig. 3, although our improved Cycle-GAN can better translate CT images into MR-like images, the CT-to-MR translation is still challenging for translating “simple” CT images to “complex” MR images. Most image-to-image translation methods will inevitably generate unrealistic soft-tissue details, resulting in some mismatch problems. Therefore, the registration methods that simply convert multimodal to unimodal registration via image translation algorithm are not reliable.

In order to address this problem, we propose a dual-stream network to fully use the information of the moving, fixed and translated images as shown in Fig. 1. In particular, we can use effective similarity metrics to train our multimodal registration model without any ground-truth deformation.

Network Details. As shown in Fig. 1, our dual-stream network is comprised of four parts: multimodal stream subnetwork, unimodal stream subnetwork, deformation field fusion, and Spatial Transformation Network.

In **Multimodal Stream** subnetwork, original CT(rCT) and MR(rMR) are represented as the moving and fixed images, which allows the model to propagate original information to counteract mismatch problems in translated MR(tMR).

Through image translation, we obtain the translated MR(tMR) with similar appearance to the fixed MR(rMR). Then, in **Unimodal Stream**, tMR and rMR are used as moving and fixed images respectively. This stream can effectively propagate more texture information, and constrain the final deformation field to suppress unrealistic voxel drifts from the multimodal stream.

During the network training, the two streams constrain each other, while they are also cooperating to optimize the entire network. Thus, our novel dual-stream design allows us to benefit from both original image information and homogeneous structural information in the translated images.




Fig. 4. Detailed architecture of UNet-based subnetwork. The encoder uses convolution with stride of 2 to reduce spatial resolution, while the decoder uses 3D upsampling layers to restore the spatial resolution.

Specifically, UNet_o and UNet_s adopt the same UNet architecture used in VoxelMorph [4] (shown in Fig. 4). The only difference is that UNet_o is with *multimodal inputs* but UNet_s is with *unimodal inputs*. Each UNet takes a

single 2-channel 3D image formed by concatenating I_m and I_f as input, and outputs a volume of deformation field with 3 channels.

After Uni- and Multi-model Stream networks, we obtain two deformation fields, ϕ_o (for rCT and rMR) and ϕ_s (for tMR and rMR). We stack ϕ_o and ϕ_s , and apply a 3D convolution with size of $3 \times 3 \times 3$ to estimate the final deformation field ϕ_{os} , which is a 3D volume with the same shape of ϕ_o and ϕ_s .

To evaluate the dissimilarity between moved and fixed images, we integrate spatial transformation network (STN) [14] to warp the moving image using ϕ_{os} . The loss function consists of two components as shown in Eq. (4).

$$\mathcal{L}_{total}(I_{rMR}, I_{rCT}, \phi_{os}) = \mathcal{L}_{sim}(I_{rMR}, I_{rCT} \circ \phi_{os}) + \lambda \mathcal{L}_{smooth}(\phi_{os}) \quad (4)$$

where λ is a regularization weight. The first loss \mathcal{L}_{sim} is similarity loss, which is to penalize the differences in appearance between fixed and moved images. Here we adopt SSIM [22] for experiments. Suggested by [4], deformation regularization \mathcal{L}_{smooth} adopts a L2-norm of the gradients of the final deformation field ϕ_{os} .

3 Experiments and Results

Dataset and Preprocessing. We focus on the application of abdominal CT-to-MR registration. We evaluated our method on two proprietary datasets since there is no designated public repository.

- 1) *Pig Ex-vivo Kidney CT-MR Dataset.* This dataset contains 18 pairs of CT and MRI kidney scans from pigs. All kidneys are manually segmented by experts. After preprocessing the data, e.g., resampling and affine spatial normalization, we cropped the data to $144 \times 80 \times 256$ with 1 mm isotropic voxels and arbitrarily divided it into two groups for training (15 cases) and testing (3 cases).
- 2) *Abdomen (ABD) CT-MR Dataset.* This 50-patient dataset of CT-MR scans was collected from a local hospital and annotated with anatomical landmarks. All data were preprocessed into $176 \times 176 \times 128$ with the same resolution (1 mm³) and were randomly divided into two groups for training (45 cases) and testing (5 cases).

Implementation. We trained our model using the following settings: (1) The Cycle-GAN for CT-MR translation network is based on the existing implementation [27] with changes as discussed in Sect. 2.1. (2) The Uni- and Multi-modal stream registration networks were implemented using Keras with the Tensorflow backend and trained on an NVIDIA Titan X (Pascal) GPU.

3.1 Results for CT-to-MR Translation

We extracted 1792 and 5248 slices from the transverse planes of the Pig kidney and ABD dataset respectively to train the image translation network. Parameters λ_{cyc} , $\lambda_{identity}$ and λ_{MIND} were set to 10, 5, and 5 for training.

Table 1. Quantitative results for image translation.

Pig Kidney	Method	PSNR	SSIM	ABD	Method	PSNR	SSIM
	Cycle-GAN	32.07	0.9025		Cycle-GAN	22.95	0.7367
	Ours	32.74	0.9532		Ours	23.55	0.7455

Since our registration method is for 3D volumes, we apply the pre-trained CT-to-MR generator to translate moving CT images into MR-like images slice-by-slice and concatenate 2D slices into 3D volumes. The qualitative results are visualized in Fig. 3. In addition, to quantitatively evaluate the translation performance, we apply our registration method to obtain aligned CT-MR pairs and utilize SSIM [22] and PSNR [10] to judge the quality of translated MR (shown in Table 1). In our experiment, our method predicts better MR-like images on both datasets.

3.2 Registration Results

Affine registration is used as the baseline method. For traditional method, only mutual information (MI) based **SyN** [2] is compared since it is the only metric (available in ANTs [3]) for multimodal registration. In addition to SyN, we implemented the following learning-based methods: 1) **VM_MIND** and **VM_SSIM** which extends VoxelMorph with similarity metrics MIND [8] and SSIM [22]. 2) **M2U** which is a typical translation-based registration method. It generates tMR from CT and converts the multimodal problem to tMR-to-MR registration. It’s noteworthy that the parameters of all methods are optimized to the best results on both datasets.

Two examples of the registration results are visualized in Fig. 5, where the red and yellow contours represent the ground truth and registered organ boundaries respectively. As shown in Fig. 5, the organ boundaries aligned by the traditional SyN method have a considerable amount of disagreement. Among all learning-based methods, our method has the most visually appealing boundary alignment for both cases. **VM_SSIM** performed significantly worse for the kidney. **VM_MIND** achieved accurate registration for the kidney, but its result for the ABD case is significantly worse. Meanwhile, **M2U** suffers from artificial features in the image translation, which leads to an inaccurate registration result.

The quantitative results are presented in Table 2. We compare different methods by the Dice score [7] and target registration error (TRE) [24]. We also provide the average run-time for each method. As shown in Table 2, our method consistently outperformed other methods and was able to register a pair of images in less than 2 s (when using GPU).




Fig. 5. Visualization results of our model compared to other methods. Upper: Pig Kidney. Bottom: Abdomen (ABD). The red contours represent the ground truth organ boundary while the yellow contours are the warped contours of segmentation masks. (Color figure online)

Table 2. Quantitative results for Pig Kidney Dataset and Abdomen (ABD) Dataset.

Metric	Organ		Affine	SyN	M2U	VM_SSIM	VM_MIND	Ours
Dice (%)	Pig	Kidney	89.53	89.87	90.21	93.75	96.48	98.57
		ABD	Kidney	80.03	82.36	78.96	82.21	84.58
		Spleen	79.58	80.38	77.76	81.79	83.11	87.01
		Liver	78.74	79.13	78.83	82.05	81.98	83.34
TRE (mm)	ABD	spleen	4.16	4.20	3.76	3.58	3.65	2.47
		Liver	6.55	5.61	5.91	4.72	4.87	3.64
Time(s) GPU/CPU	Pig		-/103	-/121	1.08/20	1.06/20	1.07/20	1.12/21
	ABD		-/108	-/137	1.23/24	1.22/22	1.21/23	1.27/24

3.3 The Effect of Each Deformation Field

In order to validate the effectiveness of the deformation field fusion, we compare ϕ_s , ϕ_o and ϕ_{os} together with warped images (shown in Fig. 6). The qualitative result shows that ϕ_s from the unimodal stream alleviates the voxel drift effect from the multimodal stream. While ϕ_o from the multimodal stream uses the original image textures to maintain the fidelity and reduce artificial features for the generated tMR image. The fused deformation field ϕ_{os} produces better alignment than both streams alone, which demonstrates the effectiveness of the joint learning step.




Fig. 6. Visualizations of the deformation field fusion. (a) moving image; (h) fixed image; (b/d/f) deformation fields; (c/e/g) images warped by (b/d/f), corresponding average Dice scores (%) of all organs are calculated. The contours in red represent ground truth, while yellow shows the warped segmentation mask. (Color figure online)

4 Conclusion

We proposed a fully unsupervised uni- and multi-modal stream network for CT-to-MR registration. Our method leverages both CT-translated-MR and original CT images towards achieving the best registration result. Besides, the registration network can be effectively trained by computationally efficient similarity metrics without any ground-truth deformation. We evaluated the method on two clinical datasets, and it outperformed state-of-the-art methods in terms of accuracy and efficiency.

Acknowledgement. This project was supported by the National Institutes of Health (Grant No. R01EB025964, R01DK119269, and P41EB015898) and the Overseas Cooperation Research Fund of Tsinghua Shenzhen International Graduate School (Grant No. HW2018008).

References

1. Armanious, K., Jiang, C., Abdulatif, S., Küstner, T., Gatidis, S., Yang, B.: Unsupervised medical image translation using cycle-MedGAN. In: 2019 27th European Signal Processing Conference (EUSIPCO), pp. 1–5. IEEE (2019)
2. 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)
3. 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**, 2033–2044 (2011)
4. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Gutttag, J., Dalca, A.V.: An unsupervised learning model for deformable medical image registration. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 9252–9260 (2018)
5. Cao, X., et al.: Deformable image registration based on similarity-steered CNN regression. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) MICCAI 2017. LNCS, vol. 10433, pp. 300–308. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66182-7_35

6. Cao, X., Yang, J., Wang, L., Xue, Z., Wang, Q., Shen, D.: Deep learning based inter-modality image registration supervised by intra-modality similarity. In: Shi, Y., Suk, H.-I., Liu, M. (eds.) MLMI 2018. LNCS, vol. 11046, pp. 55–63. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00919-9_7
7. Dice, L.R.: Measures of the amount of ecologic association between species. *Ecology* **26**(3), 297–302 (1945)
8. Heinrich, M.P., et al.: Mind: modality independent neighbourhood descriptor for multi-modal deformable registration. *Med. Image Anal.* **16**(7), 1423–1435 (2012)
9. Hiasa, Y., et al.: Cross-modality image synthesis from unpaired data using CycleGAN. In: Gooya, A., Goksel, O., Oguz, I., Burgos, N. (eds.) SASHIMI 2018. LNCS, vol. 11037, pp. 31–41. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00536-8_4
10. Hore, A., Ziou, D.: Image quality metrics: PSNR vs. SSIM. In: 2010 20th International Conference on Pattern Recognition, pp. 2366–2369. IEEE (2010)
11. Hu, X., Kang, M., Huang, W., Scott, M.R., Wiest, R., Reyes, M.: Dual-stream pyramid registration network. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11765, pp. 382–390. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_43
12. Hu, Y., et al.: Label-driven weakly-supervised learning for multimodal deformable image registration. In: 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018), pp. 1070–1074. IEEE (2018)
13. Hu, Y., et al.: Weakly-supervised convolutional neural networks for multimodal image registration. *Med. Image Anal.* **49**, 1–13 (2018)
14. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. In: Advances in Neural Information Processing Systems, pp. 2017–2025 (2015)
15. Kuang, D., Schmah, T.: FAIM – a ConvNet method for unsupervised 3D medical image registration. In: Suk, H.-I., Liu, M., Yan, P., Lian, C. (eds.) MLMI 2019. LNCS, vol. 11861, pp. 646–654. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32692-0_74
16. Liu, C., Ma, L., Lu, Z., Jin, X., Xu, J.: Multimodal medical image registration via common representations learning and differentiable geometric constraints. *Electron. Lett.* **55**(6), 316–318 (2019)
17. Mahapatra, D., Ge, Z., Sedai, S., Chakravorty, R.: Joint registration and segmentation of Xray images using generative adversarial networks. In: Shi, Y., Suk, H.-I., Liu, M. (eds.) MLMI 2018. LNCS, vol. 11046, pp. 73–80. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00919-9_9
18. 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
19. Sedghi, A., et al.: Semi-supervised image registration using deep learning. In: Medical Imaging 2019: Image-Guided Procedures, Robotic Interventions, and Modeling, vol. 10951, p. 109511G. International Society for Optics and Photonics (2019)
20. Tanner, C., Ozdemir, F., Profanter, R., Vishnevsky, V., Konukoglu, E., Goksel, O.: Generative adversarial networks for mr-ct deformable image registration. arXiv preprint [arXiv:1807.07349](https://arxiv.org/abs/1807.07349) (2018)
21. de Vos, B.D., Berendsen, F.F., Viergever, M.A., Sokooti, H., Staring, M., Išgum, I.: A deep learning framework for unsupervised affine and deformable image registration. *Med. Image Anal.* **52**, 128–143 (2019)

22. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* **13**(4), 600–612 (2004)
23. Wei, D., et al.: Synthesis and inpainting-based MR-CT registration for image-guided thermal ablation of liver tumors. In: Shen, D., et al. (eds.) *MICCAI 2019*. LNCS, vol. 11768, pp. 512–520. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32254-0_57
24. West, J.B., et al.: Comparison and evaluation of retrospective intermodality image registration techniques. In: *Medical Imaging 1996: Image Processing*, vol. 2710, pp. 332–347. International Society for Optics and Photonics (1996)
25. Yang, H., et al.: Unpaired brain MR-to-CT synthesis using a structure-constrained CycleGAN. In: Stoyanov, D., et al. (eds.) *DLMIA/ML-CDS -2018*. LNCS, vol. 11045, pp. 174–182. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00889-5_20
26. Zhao, S., Lau, T., Luo, J., Eric, I., Chang, C., Xu, Y.: Unsupervised 3D end-to-end medical image registration with volume tweening network. *IEEE J. Biomed. Health Inform.* (2019)
27. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Cyclegan (2017). <https://github.com/xhujoy/CycleGAN-tensorflow>
28. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2223–2232 (2017)