

Boosting Medical Image Synthesis via Registration-guided Consistency and Disentanglement Learning

Chuanpu Li^{1,2*}, Zeli Chen^{1,2,3*}, Yiwen Zhang^{1,2}, Liming Zhong^{1,2}, and Wei Yang^{1,2}(✉)

¹ School of Biomedical Engineering, Southern Medical University, Guangzhou, China
weiyanggm@gmail.com

² Guangdong Provincial Key Laboratory of Medical Image Processing, Southern Medical University, Guangzhou, China

³ DAMO Academy, Alibaba Group

Abstract. Medical image synthesis remains challenging due to misalignment noise during training. Existing methods have attempted to address this challenge by incorporating a registration-guided module. However, these methods tend to overlook the task-specific constraints on the synthetic and registration modules, which may cause the synthetic module to still generate spatially aligned images with misaligned target images during training, regardless of the registration module’s function. Therefore, this paper proposes registration-guided consistency and incorporates disentanglement learning for medical image synthesis. The proposed registration-guided consistency architecture fosters task-specificity within the synthetic and registration modules by applying identical deformation fields before and after synthesis, while enforcing output consistency through an alignment loss. Moreover, the synthetic module is designed to possess the capability of disentangling anatomical structures and specific styles across various modalities. An anatomy consistency loss is introduced to further compel the synthetic module to preserve geometrical integrity within latent spaces. Experiments conducted on both an in-house abdominal CECT-CT dataset and a publicly available pelvic MR-CT dataset have demonstrated the superiority of the proposed method. The code is available at: <https://github.com/pupuchuan/RegConDIS>.

Keywords: Medical image synthesis · Registration-guided synthesis · Disentanglement learning

1 Introduction

Medical image synthesis can offer advantages by predicting a target imaging modality in situations where acquisition is impractical due to time and cost constraints [24, 25]. For example, MR-to-CT synthesis is a clinically significant solution for MRI-only radiation therapy treatment planning [5, 26], and

* Chuanpu Li and Zeli Chen contributed equally to this article.

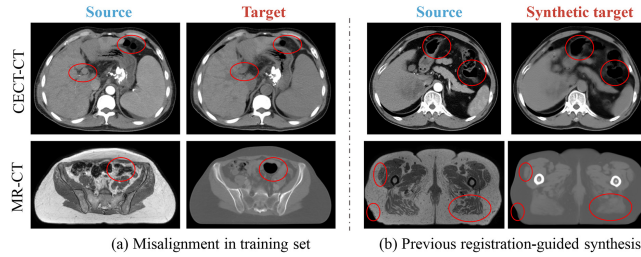


Fig. 1. The misaligned challenges in medical image synthesis. (a) Misaligned anatomical structures result in inaccurate mapping in supervised paradigms. (b) The previous registration-guided synthesis [13] still retains misalignment-induced noise and suffers from blurred boundaries due to the averaging effect.

CECT-to-CT synthesis is beneficial for reducing the registration errors between contrast-enhanced CT (CECT) and CT, the additional exposure, and treatment cost [16, 21].

Existing medical image synthesis methods can be categorized into supervised and unsupervised learning paradigms. In the supervised setting, source and target modalities are registered prior to training [4, 16, 26]. However, as shown in Fig. 1(a), the standard image pre-processing steps cannot entirely eliminate misaligned errors, particularly in anatomical areas like the abdomen or pelvis. This misalignment can result in supervised methods learning inaccurate mappings [4, 6, 28, 29]. Unsupervised methods alleviate the necessity of paired images. Unfortunately, these methods cannot fully utilize retrospective images from the same patient, and the performance of unsupervised synthesis methods is inferior to supervised ones [5, 17, 27]. Therefore, better image quality may be obtained by trying to solve the misalignment problem in the supervised paradigm.

Recent methods impose the registration-guided module to mitigate the misalignment challenge [7, 12, 13, 20, 27–29]. For example, Kong et al. proposed Reg-GAN to train the generator using an additional registration network to fit the misaligned noise distribution adaptively for T1-to-T2 MR image synthesis [13]. Zhong et al. proposed quartet attention aware closed-loop learning (QACL) framework for MR-to-CT synthesis via simultaneous registration [27]. Xin et al. employ four aligners after synthesis module, with cycle-based consistency across them [23]. However, these methods neglect the task-specific constraints on synthetic and registration modules. Specifically, the synthetic module should ideally preserve geometric properties, while the registration module should effectively learn geometric transformations from images suffering from misalignment. In medical image synthesis tasks, the design of the synthesis module is more complex than that of the registration module. *The absence of task-specific constraints may cause the complex synthesis module to directly generate spatially misaligned images during training, regardless of the function of the registration block* [1]. In the testing phase, images synthesized solely using the synthetic network may

still retain misalignment-induced noise and suffer from blurred boundaries due to the averaging effect. (As shown in Fig. 1(b)).

Some methods also integrate disentanglement learning to address the challenge of misalignment [14, 18]. These methods disentangle the anatomical structures and misalignment into distinct features in the latent spaces and synthesize target images using only anatomical features. However, *different styles between modalities may also be falsely interpreted as motions, resulting in implausible deformations* [19, 22]. Therefore, distinguishing the anatomy and the styles from different modalities and constraining anatomical features to be consistent are also beneficial for alleviating the misalignment problem.

In this study, we boost medical image synthesis via registration-guided consistency and disentanglement learning. Unlike prior methods which neglect the task-specific constraints in both synthetic and registration modules, our proposed approach integrates a registration-guided consistency architecture alongside an alignment loss, aiming at minimizing the difference between the images synthesized before and after applying the same deformation field. Through that, the synthetic network is encouraged to eliminate the influence of misalignment-induced noise, under the presumption of consistent output across diverse deformed inputs. Furthermore, we propose an anatomy consistency disentanglement synthetic module (ACDS) to decompose anatomical and style features. The disentangled style features prove advantageous for the registration network in discerning authentic deformations [19]. For the disentangled anatomical features, we introduce an anatomy consistency loss to further mitigate the misalignment issue within the latent space. Experimental validations conducted on both an in-house abdominal CECT-CT dataset and a publicly available pelvic MR-CT dataset demonstrate our method underscores the efficacy relative to several synthetic approaches.

2 Method

Let $\mathcal{S} \subset \mathbb{R}^{D_S \times H_S \times W_S}$ and $\mathcal{T} \subset \mathbb{R}^{D_T \times H_T \times W_T}$ represent domains of source image and target image, where D, H, W are the slice numbers, height, and weight. We denote $\mathbf{S} = \{I_s \in \mathcal{S}\}_{s=1}^N$ and $\mathbf{T} = \{I_t \in \mathcal{T}\}_{t=1}^N$ as paired image sets. Medical image synthesis aims to learn a synthetic network G that maps the source image I_s to its target image $O_t = G(I_s)$. However, the inevitable misalignment between \mathbf{S} and \mathbf{T} often results in inaccurate mapping. To address this challenge, we propose a registration-guided consistency disentanglement learning method.

2.1 Registration-guided Consistency

A key challenge of registration-guided medical synthesis lies in delineating the respective roles of the synthetic network and registration network. Specifically, it's crucial to avoid situations where the synthetic network ignores the function of the registration network and directly synthesizes spatially aligned images during training. To achieve this, we assume that the synthetic and registration processes

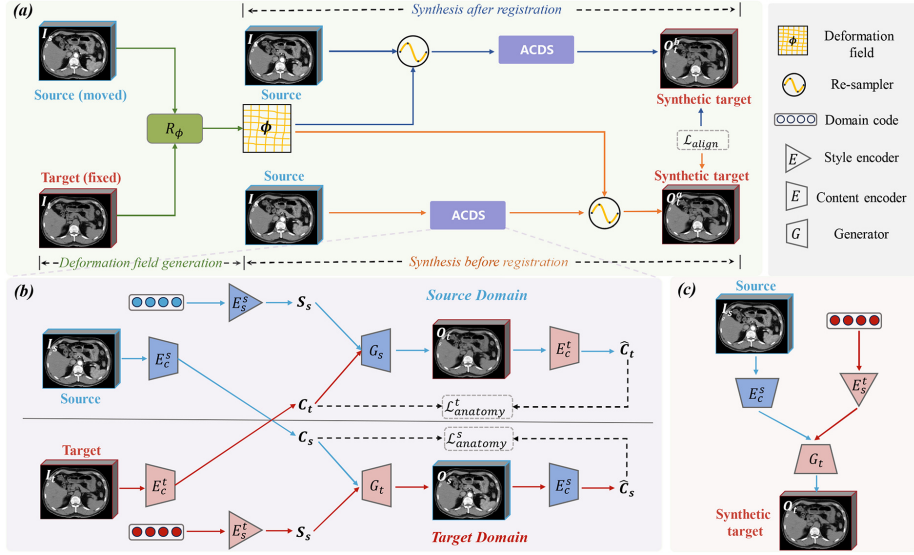


Fig. 2. Overview of our proposed methods. (a) The registration-guided consistency architecture. (b) The anatomy consistency disentanglement synthetic (ACDS) module. (c) Testing phase for medical image synthesis.

are commutative and propose a registration-guided consistency structure (As shown in Fig. 2(a)).

Registration-guided Module. The registration module is composed of a deformation generator R_ϕ and a re-sampler R_S . R_ϕ is designed to take I_s and I_t as input, and produce a deformation field ϕ describing the non-rigid alignment from I_s to I_t :

$$\phi = R_\phi(I_s, I_t). \quad (1)$$

R_S receives the deformation field ϕ and applies it to an image desired to be moved. We have adopted the architectural framework detailed in [3] as the foundation of our registration module, and incorporated its smoothing loss to preserve the smooth displacement of the deformable field:

$$\mathcal{L}_{smooth}(\phi) = \mathbb{E}_{\mathcal{S}, \mathcal{T}}[\|\nabla \phi\|^2]. \quad (2)$$

Synthesis Before Registration. Synthesis before registration first applies the synthetic module G on I_s , followed by the utilization of the re-sampler layer R_S on the synthesized image. This process yields the final output, which can be formally expressed as:

$$O_t^b = R_S(G(I_s), \phi). \quad (3)$$

Synthesis After Registration. In contrast, synthesis after registration first applies the re-sampler layer R_S on I_s , followed by the utilization of the synthetic module G on the warped image. This process yields the other final output, which can be formally expressed as:

$$O_t^a = G(R_S(I_s, \phi)). \quad (4)$$

Alignment Loss. In Eq.3 and Eq.4, the deformation field ϕ used in the re-sampler R_S is both given by the $R_\phi(I_s, I_t)$ and designed to be the same. The sole difference is that the synthetic module G accepts geometrically different images as inputs (I_s and $R_S(I_s, \phi)$). Thus, if we assume that the synthetic and registration processes are commutative, which expect the same output from synthesis before and after registration, the synthetic module G is required to be geometry preserving and finally alleviate the misalignment issue. This assumption can be easily realized by a simple alignment loss, which can be formulated as:

$$\mathcal{L}_{\text{align}}(G, R) = \|O_t^b - I_t\|_1 + \|O_t^a - I_t\|_1. \quad (5)$$

2.2 Anatomy Consistency Disentanglement Synthetic Module

Given that variations in styles across modalities may be falsely interpreted as motions by the registration network [14], it is beneficial for the synthetic module G to disentangle the specific style for each modality. Moreover, disentangling anatomical structures and enforcing their consistency across identical content images is important in mitigating misaligned noise. Therefore, we propose an anatomy consistency disentanglement synthetic module.

Synthetic process. As shown in Fig. 2(b), the ACDS module is composed of two content encoders $\{E_c^s, E_c^t\}$, two style encoders $\{E_s^s, E_s^t\}$, and two generators $\{G_s, G_t\}$. The encoders in both domains respectively encode an input image to an anatomical-invariant space and a modality-specific style space, and the generators respectively combine the anatomical space features and style features to synthesize images. This can be formulated as:

$$O_t = G_t(c_s, s_t) = G_t(E_c^s(I_s), E_s^t(d_t)), \quad (6)$$

$$O_s = G_s(c_t, s_s) = G_s(E_c^t(I_t), E_s^s(d_s)), \quad (7)$$

where O_t and O_s denote the synthetic images, d_t and d_s denote the learned domain code, which is empirically set as an 8-bit vector. The content encoder consists of 3 Conv-IN-ReLU blocks with kernel sizes of 3, strides of 2, channels of 32, 64, 128, followed by 4 residual blocks. The style encoder consists of 3 MLP layers.

Disentanglement Learning. Inspired by [8, 15], the synthetic module realizes disentanglement via generative adversarial loss, self-reconstruction loss, and cycle-consistency loss. We define them as follows:

$$\begin{aligned} \mathcal{L}_{\text{adv}} = & \mathbb{E}_{\mathcal{T}}[\log D_t(I_t)] + \mathbb{E}_{\mathcal{S}}[1 - \log D_t(O_t)] \\ & + \mathbb{E}_{\mathcal{S}}[\log D_s(I_s)] + \mathbb{E}_{\mathcal{T}}[1 - \log D_s(O_s)], \end{aligned} \quad (8)$$

$$\mathcal{L}_{\text{self}} = \|G_s(E_c^s(I_s), E_s^s(d_s)) - I_s\|_1 + \|G_t(E_c^t(I_t), E_s^t(d_t)) - I_t\|_1, \quad (9)$$

$$\mathcal{L}_{\text{cycle}} = \|G_s(E_c^t(O_t), E_s^s(d_s)) - I_s\|_1 + \|G_t(E_c^s(O_s), E_s^t(d_t)) - I_t\|_1, \quad (10)$$

where D_t and D_s represent the discriminators. If the encoders can successfully disentangle the anatomical content from images and the generators uphold geometric fidelity, anatomical content features should be identical. Therefore, we propose an anatomy consistency loss:

$$\mathcal{L}_{\text{anatomy}} = \|E_c^t(O_t) - E_c^s(I_s)\|_1 + \|E_c^s(O_s) - E_c^t(I_t)\|_1. \quad (11)$$

The implementation of the anatomy consistency loss necessitates the ACDS module to further learn and retain anatomical structure information in the latent space. Then, the overall loss function is:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{adv}} + \mathcal{L}_{\text{self}} + \mathcal{L}_{\text{cycle}} + \lambda_1 \mathcal{L}_{\text{anatomy}} + \lambda_2 \mathcal{L}_{\text{smooth}} + \lambda_3 \mathcal{L}_{\text{align}}, \quad (12)$$

where λ is the hyperparameter. As shown in Fig. 2(c), after training, only E_c^s , E_s^t and G_t are required to translate the source image into target image in the testing phase.

3 Experiments

3.1 Experimental Setup.

Datasets. We evaluated our methods on an in-house abdominal CECT-CT dataset and a publicly available pelvic MR-CT derived from the SynthRAD2023 Grand Challenge dataset [9]. The CECT-CT dataset comprised 278 patients diagnosed with liver cancer. Preprocessing of images involved sequential implementation of affine spatial normalization [11] and nonlinear registration techniques [2]. For evaluation, we manually selected 15 well-aligned subjects for the validation set and 40 well-aligned subjects for the test set. The pelvic MR-CT dataset consists of 180 patients. Similarly, the validation set comprised 6 aligned subjects, while the test set included 30 appropriately aligned subjects.

Implementation Details Our method was implemented with PyTorch, using 2 NVIDIA RTX3090 GPUs. The hyperparameter λ_1 , λ_2 , λ_3 were set as 0.5, 10 and 20. For the anatomy consistency disentanglement synthetic module, we used a similar structure as [8] and changed it into 3D. We employed Adam optimizer

Table 1. Quantitative comparison of different methods on the CECT-to-CT synthesis and MR-to-CT synthesis. All p-values from the paired samples t-test are below 0.05.

Methods	CECT-to-CT Synthesis			MR-to-CT Synthesis		
	MAE (HU) ↓	PSNR (dB) ↑	SSIM (%) ↑	MAE (HU) ↓	PSNR (dB) ↑	SSIM (%) ↑
Pix2pix [10]	21.70 ± 1.40	34.89 ± 1.43	87.78 ± 3.02	50.64 ± 4.47	30.61 ± 0.85	82.82 ± 2.14
CycleGAN [30]	22.51 ± 1.78	34.33 ± 1.12	88.14 ± 2.78	82.24 ± 11.42	27.45 ± 0.79	70.79 ± 6.51
UNIT [15]	22.78 ± 1.80	34.04 ± 1.06	88.10 ± 2.78	84.49 ± 11.95	27.80 ± 0.88	70.44 ± 6.22
RegGAN [13]	21.29 ± 1.86	35.10 ± 1.27	88.15 ± 2.84	50.22 ± 3.60	30.44 ± 0.53	83.16 ± 1.92
ResViT [4]	20.79 ± 2.09	35.14 ± 1.41	88.20 ± 2.89	49.90 ± 3.43	30.72 ± 0.60	82.91 ± 2.07
Diffusion [6]	20.33 ± 2.10	35.42 ± 1.52	88.71 ± 2.86	49.65 ± 4.00	30.47 ± 0.72	83.38 ± 2.00
Ours	16.38 ± 1.83	36.52 ± 1.66	89.22 ± 2.83	47.47 ± 4.59	30.86 ± 0.88	83.70 ± 2.47

with a poly decay strategy, a batch size of 2, and a maximum of 200 epochs to dynamically adjust the learning rate from 2×10^{-4} . To evaluate the efficacy of the proposed model, we employed three standard evaluation metrics: the mean absolute error (MAE), the structural similarity index measurement (SSIM), and the peak signal-to-noise ratio (PSNR).

3.2 Results

Comparisons with State-of-the-Art. We compared our method with several synthesis methods including supervised methods such as Pix2pix [10], ResViT [4], diffusion model [6], unsupervised methods such as CycleGAN [30], disentanglement methods such as UNIT [15], and registration-guided method such as RegGAN [13]. To ensure a fair comparison, we converted all these methods to 3D models. As illustrated in Table 1, our proposed method demonstrates superior performance, exhibiting a MAE of 16.38 ± 1.83 HU, a PSNR of 36.52 ± 1.66 dB and an SSIM of 89.22 ± 2.83 % for CECT-to-CT synthesis, and an MAE of 47.47 ± 4.59 HU, a PSNR of 30.86 ± 0.88 dB and a SSIM of 83.70 ± 2.47 % for MR-to-CT synthesis. Fig. 3 and Fig. 4 show the visual results. The pink boxes show the structural details where misalignment are often encountered during training. The supervised and previous registration-guided methods still retain misalignment noise and have blurred boundaries due to the averaging effect. In contrast, our methods preserve anatomical consistency and clear boundaries in the regions where misalignment is often encountered during training.

Ablation Study. An ablation study was undertaken to assess the efficacy of the registration-guided consistency and ACDS module. Table 2 shows the results obtained from the validation set for both CECT-to-CT and MR-to-CT tasks. The third row in Table 2 demonstrates that the registration-guided consistency design yields a reduction of 2-3 HU in the MAE compared with synthesis before registration in both tasks. The results in the fourth row also indicate the integration of the ACDS module facilitates an enhancement in the quality of synthesis.




Fig. 3. Visual comparison of synthesized results on CECT-CT dataset.

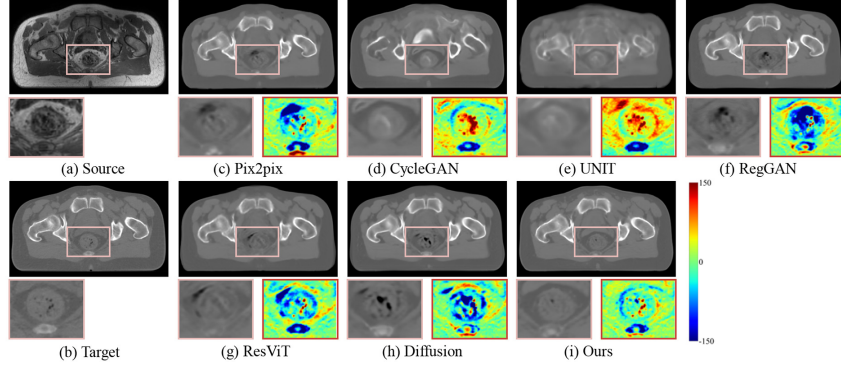


Fig. 4. Visual comparison of synthesized results on MR-CT dataset.

Table 2. Ablation study of critical components of the proposed method. AFT represents synthesis after registration branch, BEF represents synthesis after registration branch, and ACDS represents anatomy consistency disentanglement synthetic module.

Component			CECT-to-CT Synthesis			MR-to-CT Synthesis		
AFT	BEF	ACDS	MAE (HU) ↓	PSNR (dB) ↑	SSIM (%) ↑	MAE (HU) ↓	PSNR (dB) ↑	SSIM (%) ↑
•	○	○	28.84 ± 2.70	33.49 ± 1.07	85.30 ± 3.06	86.84 ± 11.23	27.61 ± 1.01	69.74 ± 2.30
○	•	○	21.56 ± 1.49	34.65 ± 1.03	87.22 ± 3.11	55.59 ± 3.33	29.65 ± 0.31	81.74 ± 1.59
•	•	○	19.43 ± 1.84	35.30 ± 1.09	87.47 ± 3.12	53.48 ± 6.56	30.12 ± 0.65	81.98 ± 1.52
•	•	•	16.96 ± 1.92	35.64 ± 1.34	88.27 ± 3.16	52.68 ± 5.01	30.18 ± 0.35	82.20 ± 1.40

4 Conclusion

This paper introduces a registration-guided consistency architecture and an ACDS module to address the misalignment issue for medical image synthesis. The proposed registration-guided consistency design customizes the registration and synthetic module to be task-specific and encourages the synthetic module to eliminate the influence of misalignment-induced noise. Additionally, disentanglement learning and an anatomy consistency loss are applied to the ACDS, enhancing its ability to preserve geometric integrity, thus further avoiding the misalignment issue. Experimental results have demonstrated the superiority of the proposed method.

Acknowledgments. This work was supported by the grants from the National Natural Science Foundation of China (No. 82172020, No. 82472050).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Arar, M., Ginger, Y., Danon, D., Bermano, A.H., Cohen-Or, D.: Unsupervised multi-modal image registration via geometry preserving image-to-image translation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13410–13419 (2020)
2. Avants, B.B., Tustison, N., Song, G., et al.: Advanced normalization tools (ants). *Insight j* **2**(365), 1–35 (2009)
3. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging* **38**(8), 1788–1800 (2019)
4. Dalmaz, O., Yurt, M., Çukur, T.: Resvit: Residual vision transformers for multi-modal medical image synthesis. *IEEE Transactions on Medical Imaging* **41**(10), 2598–2614 (2022)
5. Dayarathna, S., Islam, K.T., Uribe, S., Yang, G., Hayat, M., Chen, Z.: Deep learning based synthesis of mri, ct and pet: Review and analysis. *Medical Image Analysis* p. 103046 (2023)
6. Dorjsembe, Z., Pao, H.K., Odonchimed, S., Xiao, F.: Conditional diffusion models for semantic 3d brain mri synthesis. *IEEE Journal of Biomedical and Health Informatics* (2024)
7. Honkamaa, J., Khan, U., Koivukoski, S., Valkonen, M., Latonen, L., Ruusuvuori, P., Marttinen, P.: Deformation equivariant cross-modality image synthesis with paired non-aligned training data. *Medical Image Analysis* **90**, 102940 (2023)
8. Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. In: Proceedings of the European conference on computer vision (ECCV). pp. 172–189 (2018)
9. Huijben, E.M., Terpstra, M.L., Pai, S., Thummerer, A., Koopmans, P., Afonso, M., van Eijnatten, M., Gurney-Champion, O., Chen, Z., Zhang, Y., et al.: Generating synthetic computed tomography for radiotherapy: Synthrad2023 challenge report. *Medical image analysis* **97**, 103276 (2024)

10. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition. pp. 1125–1134 (2017)
11. Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W., Smith, S.M.: *Fsl. Neuroimage* **62**(2), 782–790 (2012)
12. Jung, S., Choi, Y., Al-masni, M.A., Jung, M., Kim, D.H.: Deformation-aware segmentation network robust to motion artifacts for brain tissue segmentation using disentanglement learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 213–222. Springer (2024)
13. Kong, L., Lian, C., Huang, D., Hu, Y., Zhou, Q., et al.: Breaking the dilemma of medical image-to-image translation. *Advances in Neural Information Processing Systems* **34**, 1964–1978 (2021)
14. Lin, Y., Han, H., Zhou, S.K.: Deep non-linear embedding deformation network for cross-modal brain mri synthesis. In: 2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2022)
15. Liu, M.Y., Breuel, T., Kautz, J.: Unsupervised image-to-image translation networks. *Advances in neural information processing systems* **30** (2017)
16. Pang, H., Qi, S., Wu, Y., Wang, M., Li, C., Sun, Y., Qian, W., Tang, G., Xu, J., Liang, Z., et al.: Ncct-cect image synthesizers and their application to pulmonary vessel segmentation. *Computer Methods and Programs in Biomedicine* **231**, 107389 (2023)
17. Phan, V.M.H., Liao, Z., Verjans, J.W., To, M.S.: Structure-preserving synthesis: Maskgan for unpaired mr-ct translation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 56–65. Springer (2023)
18. Qin, C., Shi, B., Liao, R., Mansi, T., Rueckert, D., Kamen, A.: Unsupervised deformable registration for multi-modal images via disentangled representations. In: International Conference on Information Processing in Medical Imaging. pp. 249–261. Springer (2019)
19. Reaungamornrat, S., Sari, H., Catana, C., Kamen, A.: Multimodal image synthesis based on disentanglement representations of anatomical and modality specific features, learned using uncooperative relativistic gan. *Medical image analysis* **80**, 102514 (2022)
20. Su, R., van der Sluijs, M., Cornelissen, S., van Zwam, W., van der Lugt, A., Niessen, W., Ruijters, D., van Walsum, T., Dalca, A.: Angiomoco: Learning-based motion correction in cerebral digital subtraction angiography. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 770–780. Springer (2023)
21. Sumida, I., Magome, T., Kitamori, H., Das, I.J., Yamaguchi, H., Kizaki, H., Aboshi, K., Yamashita, K., Yamada, Y., Seo, Y., et al.: Deep convolutional neural network for reduction of contrast-enhanced region on ct images. *Journal of Radiation Research* **60**(5), 586–594 (2019)
22. Wang, C., Yang, G., Papanastasiou, G., Tsaftaris, S.A., Newby, D.E., Gray, C., Macnaught, G., MacGillivray, T.J.: Dicyc: Gan-based deformation invariant cross-domain information fusion for medical image synthesis. *Information Fusion* **67**, 147–160 (2021)
23. Xin, B., Young, T., Wainwright, C., Blake, T., Lebrat, L., Gaass, T., Benkert, T., Stemmer, A., Coman, D., Dowling, J.: Deformation-aware GAN for medical image synthesis with substantially misaligned pairs. In: Medical Imaging with Deep Learning (2024), <https://openreview.net/forum?id=2iAsI3SZ8C>

24. Zhang, J., Sun, K., Yang, J., Hu, Y., Gu, Y., Cui, Z., Zong, X., Gao, F., Shen, D.: A generalized dual-domain generative framework with hierarchical consistency for medical image reconstruction and synthesis. *Communications Engineering* **2**(1), 72 (2023)
25. Zhang, Y., Peng, C., Wang, Q., Song, D., Li, K., Zhou, S.K.: Unified multi-modal image synthesis for missing modality imputation. *IEEE Transactions on Medical Imaging* (2024)
26. Zhong, L., Chen, Z., Shu, H., Zheng, K., Li, Y., Chen, W., Wu, Y., Ma, J., Feng, Q., Yang, W.: Multi-scale tokens-aware transformer network for multi-region and multi-sequence mr-to-ct synthesis in a single model. *IEEE Transactions on Medical Imaging* (2023)
27. Zhong, L., Chen, Z., Shu, H., Zheng, Y., Zhang, Y., Wu, Y., Feng, Q., Li, Y., Yang, W.: Qacl: Quartet attention aware closed-loop learning for abdominal mr-to-ct synthesis via simultaneous registration. *Medical Image Analysis* **83**, 102692 (2023)
28. Zhong, L., Huang, P., Shu, H., Li, Y., Zhang, Y., Feng, Q., Wu, Y., Yang, W.: United multi-task learning for abdominal contrast-enhanced ct synthesis through joint deformable registration. *Computer Methods and Programs in Biomedicine* **231**, 107391 (2023)
29. Zhou, L., Ni, X., Kong, Y., Zeng, H., Xu, M., Zhou, J., Wang, Q., Liu, C.: Mitigating misalignment in mri-to-ct synthesis for improved synthetic ct generation: an iterative refinement and knowledge distillation approach. *Physics in Medicine & Biology* **68**(24), 245020 (2023)
30. 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)