DisQ

Disentangling Quantitative MRI Mapping of the Heart

Yang, Changchun; Zhao, Yidong; Huang, Lu; Xia, Liming; Tao, Qian

DOI
10.1007/978-3-031-16446-0_28

Publication date
2022

Document Version
Final published version

Published in
Medical Image Computing and Computer Assisted Intervention – MICCAI 2022 - 25th International Conference, Proceedings

Citation (APA)
Yang, C., Zhao, Y., Huang, L., Xia, L., & Tao, Q. (2022). DisQ: Disentangling Quantitative MRI Mapping of the Heart. In L. Wang, Q. Dou, P. T. Fletcher, S. Speidel, & S. Li (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022 - 25th International Conference, Proceedings* (pp. 291-300). (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Vol. 13436 LNCS). Springer. https://doi.org/10.1007/978-3-031-16446-0_28

DisQ: Disentangling Quantitative MRI Mapping of the Heart

Changchun Yang[1], Yidong Zhao[1], Lu Huang[2], Liming Xia[2], and Qian Tao[1(✉)]

[1] Department of Imaging Physics, Delft University of Technology, Delft, The Netherlands
q.tao@tudelft.nl

[2] Tongji Medical College, Huazhong University of Science and Technology, Wuhan, China

Abstract. Quantitative MRI (qMRI) of the heart has become an important clinical tool for examining myocardial tissue properties. Because heart is a moving object, it is usually imaged with electrocardiogram and respiratory gating during acquisition, to "freeze" its motion. In reality, gating is more-often-than-not imperfect given the heart rate variability and nonideal breath-hold. qMRI of the heart, consequently, is characteristic of varying image contrast as well as residual motion, the latter compromising the quality of quantitative mapping. Motion correction is an important step prior to parametric mapping, however, a long-standing difficulty for registering the dynamic sequence is that the contrast across frames varies wildly: depending on the acquisition scheme some frames can have extremely poor contrast, which fails both traditional optimization-based and modern learning-based registration methods. In this work, we propose a novel framework named *DisQ*, which Disentangles Quantitative mapping sequences into the latent space of *contrast* and *anatomy*, fully unsupervised. The disentangled latent spaces serve for the purpose of generating a series of images with identical contrast, which enables easy and accurate registration of all frames. We applied our DisQ method to the modified Look-Locker inversion recovery (MOLLI) sequence, and demonstrated improved performance of T_1 mapping. In addition, we showed the possibility of generating a dynamic series of baseline images with exactly the same shape, *strictly* registered and perfectly "frozen". Our proposed DisQ methodology readily extends to other types of cardiac qMRI such as T_2 mapping and perfusion.

Keywords: Quantitative magnetic resonance imaging · T_1 mapping · Unsupervised disentangled representation · Motion correction

1 Introduction

Quantitative magnetic resonance imaging (qMRI) has become an important clinical tool for noninvasive evaluation of tissue integrity [23]. In qMRI, quantitative information of tissue is derived from a dynamic sequence of baseline images

L. Wang et al. (Eds.): MICCAI 2022, LNCS 13436, pp. 291–300, 2022.
https://doi.org/10.1007/978-3-031-16446-0_28

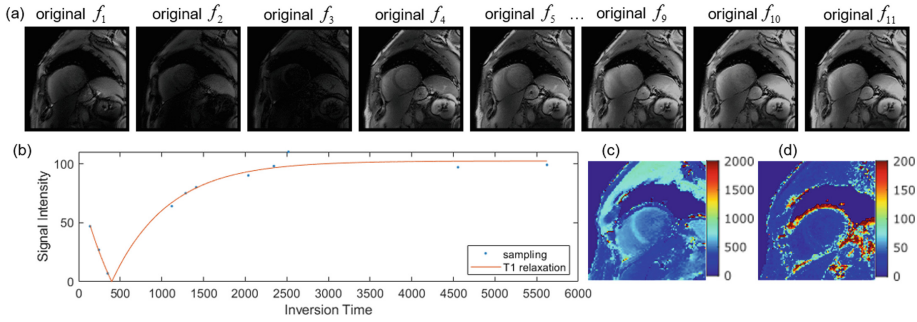

Fig. 1. (a) An example of MOLLI T_1 mapping sequence with 11 baseline images (denoted by f). (b) The 3-parameter signal model for T_1 fitting. (c) The computed T_1 map. Colorbar in the range of 0–2000 ms. (d) The corresponding standard deviation (SD) error map of T_1 mapping with colorbar 0–200 ms.

acquired with modulated MR imaging parameters. Based on the underlying physics, quantification of tissue properties is obtained by fitting a parametric signal model, under the assumption that the series of baseline images are aligned anatomically. However, this assumption is often violated in cardiac qMRI where the object is constantly moving. Even with careful electrocardiogram and respiratory gating, the baseline images often contain residual motion, which compromises the quality of quantitative mapping and undermines the value of qMRI.

Quantification of myocardial T_1 relaxation time is among the most important applications of qMRI in current radiology practice [6]. A widely used MRI sequence is the modified Look-Locker inversion (MOLLI) recovery [13], normally with 11 baseline images, governed by the following 3-parameter function:

$$s\left(t_{\text{inv}}\right) = A - B \cdot \exp\left(-\frac{t_{\text{inv}}}{T_1^*}\right) \tag{1}$$

where s is the signal intensity at t_{inv}, the inversion time during acquisition (11 in total), and A, B and T_1^* are the three parameters. The true T_1 is calculated as $T_1 = (\frac{B}{A} - 1) \cdot T_1^*$. Figure 1 illustrates an example of MOLLI T_1 baseline images (a) and parametric mapping (b-d). In this example, we can appreciate the dynamic change in baseline images and poor myocardium-blood contrast in some of them, e.g. the 3rd image in (a), as well as the residual motion in (b).

To realize accurate quantitative mapping, motion correction by image registration is an important step prior to parametric mapping. Popular registration methods include traditional optimization-based methods and modern learning-based methods. Xue *et al.* [22] proposed to use synthetic image estimation for myocardial motion correction, iteratively improved mapping accuracy. PCA-based method was proposed at [7,21] for groupwise registration. Learning-based methods explode [1,16,17] with the potential of deep learning, can be divided into two categories: supervised [18] and unsupervised (VoxelMorph [1]).

A fundamental difficulty for registering the dynamic sequence is that the contrast across frames varies wildly: depending on acquisition scheme some frames

can have extremely poor contrast (e.g. near the signal nulling point), which can fail both traditional optimization-based and modern learning-based registration methods. In this work, we propose a novel solution to this problem by first addressing the issue of contrast, inspired by the recent success of unsupervised disentangled representation learning in computer vision [4,9,12] and medical imaging [5,14,17,24]. Our rationale is as follows: according to the underlying physics, an MR image can be modeled as an function of anatomical tissue property and acquisition parameters. Therefore, when appropriately formulated, cardiac qMRI images can be disentangled to their *anatomical representation* and *contrast representation*. With such disentanglement, we may unify baseline images either in terms of contrast (for easy image registration), or anatomy (for direct quantitative mapping).

For the problem to be well-posed, existing methods for anatomy and shape disentanglement in medical imaging normally requires the dataset to share at least one common factor, i.e. multiple contrast of the same anatomy, or same contrast of different anatomy. As such, most work focused on brain MRI as the same anatomy requirement can be easily satisfied. However, for a moving object, cardiac qMRI is characteristic of varying image contrast as well as residual motion. In this work, we propose a framework named "Disentangling Quantitative MRI" (DisQ) to decompose the dynamic cardiac images under the condition of simultaneous anatomy and contrast change. We validated the method on MOLLI T_1 mapping, the most popular qMRI application of heart, but the methodology can be extended to other quantitative sequences. Our contributions include:

- This is among the first work to address cardiac qMRI analysis from an anatomy-contrast disentanglement perspective;
- We propose a novel network architecture and a number effective bootstrapping strategies, dedicate to cardiac qMRI (characteristic of simultaneous contrast and anatomy change), evaluate on the clinical T_1 mapping data;
- We demonstrate the possibility of generating *strictly* registered baseline images for cardiac qMRI, beyond any existing registration methods.

2 Methodology

2.1 Overall Framework: Disentangling Latent Spaces

A schematic plot of our proposed method is shown in Fig. 2. Let $f_t^s \in F_T^S$ denote the input baseline MOLLI image of t-th inversion time of the s-th subject. As shown in Fig. 2(a), we aim to decompose an image pair $\{f_i^s, f_j^s\}$ of the same subject to their anatomical representations $\{a_i^s = E^A(f_i^s), a_j^s = E^A(f_j^s)\}$ by an anatomical encoder E^A and separate contrast representations $\{c_i^s = E^C(f_i^s), c_j^s = E^C(f_j^s)\}$ by a contrast encoder E^C. The generator G then reconstructs the images from their anatomical and contrast representations. As in prior work [3,14], we optimize the *self-reconstruction* and *cross-reconstruction* losses to learn the disentangled latent spaces. With a pair $\{a_i^s, c_j^s\}$ derived from

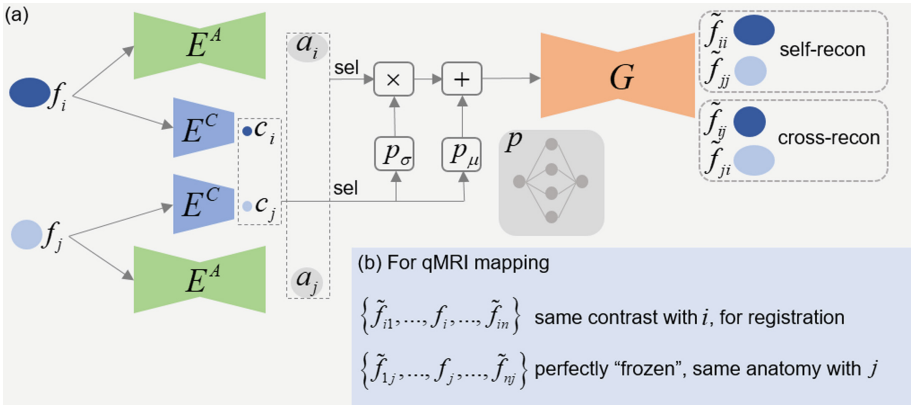

Fig. 2. (a) DisQ: the overview of network architecture to disentangle anatomy and contrast of paired baseline images f_i and f_j. The a and c decomposed from each image will be selected one at a time for reconstruction. See text for Projector p and architecture details. (b) Two ways DisQ can potentially be used in analyzing cardiac qMRI: (1) unify the contrast for motion correction, (2) unify the anatomy for direct quantitative mapping.

images of any two baseline images, G can synthesize an image $\tilde{f}^s_{ji}$, which should be similar to the image f^s_j with contrast c^s_j.

$$L_{\text{self-recon}} = \frac{1}{ST}\sum_{s=1}^{S}\sum_{i=1}^{T}\mathbb{E}_{f^s_i \sim F^S_T}\left\|\tilde{f}^s_{ii} - f^s_i\right\|_1, \tag{2}$$

$$L_{\text{cross-recon}} = \frac{1}{ST(T-1)}\sum_{s=1}^{S}\sum_{i=1}^{T}\sum_{j=1,j\neq i}^{T}\mathbb{E}_{f^s_i, f^s_j \sim F^S_T}\left\|\tilde{f}^s_{ji} - f^s_j\right\|_1, \tag{3}$$

where $\tilde{f}^s_{ji} = G(E^C(f^s_j), E^A(f^s_i))$. Under this generic framework, we present further technical novelties that enable disentanglement of cardiac qMRI.

2.2 Bootstrapping Disentangled Representations

Anatomy Encoder. Our shared anatomical encoder E^A is built from the basic architecture of the U-Net [20], to extract anatomical information a. It is desirable that the extracted a is limited in capacity (with minimal information on contrast), but at the same time captures the anatomy. We therefore design a to be a one-hot encoded multi-channel map through a straight-through Gumbel-softmax (STGS) layer [5,24]. Consequently, the generator G cannot reconstruct images without extra information of contrast since the one-hot encoding strictly restricts the capacity of a.

For the same subject, the learned multi-channel anatomical representations a^s_i, a^s_j should share similarity, but are not exactly identical due to the residual

motion of the heart. Instead of enforcing identity of anatomy, we consider the the two anatomies similar, as two weak augmentations of the true "frozen" shape of the subject. We formulate this into an anatomical similarity loss to encourage loosely similar a between the two learned anatomy representations:

$$L_{\text{anatomy}} = 1 - \frac{\langle a_i^s, a_j^s \rangle}{\|a_i^s\|_2 \cdot \|a_j^s\|_2}. \tag{4}$$

which promotes the two anatomy representation a_i^s and a_j^s to be as close as possible, while allowing minor deviations (residual motion). We will present ablation study to validate this loss.

Contrast Encoder. The second latent space is the contrast representations c capturing the contrast information in different baseline images. Given that the underlying T_1 relaxation function (Eq. 1) is simple, the latent space of contrast should be intrinsically low-dimensional. We encode the contrast information into a low-dimensional vector by a shared encoder E^C. We employ an information bottleneck loss [2,19] here to limit the information capacity of c and avoid informative leakage:

$$L_{\text{contrast}} = \left\| \|c\|_2^2 - C \right\|_1, \tag{5}$$

where C is the bottleneck capacity controlling the amount of information in the latent contrast representation. The choice of C will be presented in section Implementation.

Projector. The input for DisQ is two qMRI frames $\{f_i^s, f_j^s\}$ of the same subject, but with different acquisition parameters (in the case of MOLLI at different t_{inv}). Feeding them into the DisQ network, we can obtain $\{a_i^s, a_j^s\}$ and $\{c_i^s, c_j^s\}$ respectively, to represent their anatomies and contrasts. Consequently, by combining a and c in pairs, we can generate four synthetic images. Two of them are self-reconstruction, with c broadcasted to the same height and width as a. A code z is obtained after the broadcasted c being concatenated with the selected a in the channel dimension, and is sent to the generator G for reconstruction. The other two are cross-reconstruction, where we adopt a different concatenated mechanism. As one-hot encoding of STGS tends to have high variance with this gradient estimator, we proposed to reduce the variance of STGS inspired by Rao-Blackwellization [15]. We thereby introduce bias here to counteract the variance of STGS, which is realized by a projector p, expressed by:

$$z_{ji}^s = p_\sigma(c_j^s) \cdot a_i^s + p_\mu(c_j^s), \tag{6}$$

where p_σ and p_μ are two fully connected layers constructing the projector p.

Overall Loss. Our overall loss function is defined as $L_{\text{overall}} = \lambda_1 L_{\text{recon}} + \lambda_2 L_{\text{per}} + \lambda_3 L_{\text{anatomy}} + \lambda_4 L_{\text{contrast}}$, where L_{recon} sums up $L_{\text{self-recon}}$ (Eq. 2) and $L_{\text{cross-recon}}$ (Eq. 3), L_{per} is the perceptual loss introduced in [8]: $\|\text{VGG}(\tilde{f}) - \text{VGG}(f)\|_1$, where f is the original image, $\tilde{f}$ is the reconstructed image.

3 Experiments

3.1 Dataset

In total 102 MOLLI T_1 acquisitions were included in this study. The images were acquired by a 3.0T Ingenia MR-scanner (Philips Healthcare,Best, The Netherlands), in three short-axis slices: apical, mid, and basal. Both native and post-contrast T_1 mapping were performed using the same 3-3-5 scheme provided by the manufacturer. Each data has a dimension of $256 \times 256 \times 3 \times 11$. We randomly split the dataset into 80 for training, 11 for validation, and 11 for testing. The myocardium of left ventricle were manually annotated as the region of interest.

3.2 Implementation

Training. The four hyperparameters for our objective function L_{overall} were set to $\lambda_1 = 2, \lambda_2 = 0.03, \lambda_3 = 0.02, \lambda_4 = 10^{-8}$. The hyperparameter C in Eq. 5 was increased per epoch: $1000 \times e^{0.002i}$. The channel numbers of a and c in DisQ were set to 3 and 2, and our model was trained for 300 epochs by the Adam optimizer with learning rate of 3×10^{-4}. During training, we randomly selected two baseline images from the same MOLLI sequence, but at two different inversion time. Our codes are released at https://github.com/Changchun-Yang/DisQ.

Evaluation. For every MOLLI data, we choose the t-th baseline image f_t as the reference, then all other frames $i \in \{1...T\}, i \neq t$ along with f_t are fed into the DisQ to get all the reconstructed results. We then generate two new sequence of images with reference to f_t: $\{\tilde{f}_{t1}, ..., f_t, ..., \tilde{f}_{tT}\}$, and $\{\tilde{f}_{1t}, ..., f_t, ..., \tilde{f}_{Tt}\}$. The first sequence share the same contrast with f_t, but retains the anatomy of the original baseline images. This sequence of images (with the same contrast) is then used for residual motion correction. The derived deformation field is then applied to the original baseline image series for a motion-corrected MOLLI. The second sequence keeps the original contrast of baseline images while sharing the same anatomy, hence with cardiac motion perfectly "frozen". This sequence of generated images can be directly used for T1 mapping. Quantitative metrics include the value and standard deviation (SD) error of the T_1 map as in [10]. The unsupervised registration network is adopted from the baseline VoxelMorph [1]. We set t as 5 in our experiments, but our results were not sensitive to its choice.

Comparative and Ablation Study. We evaluated the proposed bootstrapping strategies by comparative and ablation studies. In particular, we performed ablation study for the proposed anatomy loss L_{anatomy} and projector p. As a baseline, we implemented the same network architecture, by substituting L_{anatomy} with the common MAE loss, and removing the projector p. This baseline is denoted as **Dis**. The proposed anatomy loss and projector were then integrated in this baseline model one by one to create ablation models. All models in comparison however carried the contrast loss L_{contrast}, which is important for reasonable disentanglement of contrast. In addition, the T_1 mapping results of the

originally acquired data (with residual motion) was denoted as **Org**. The T_1 mapping results after VoxelMorph registration is denoted as **Morph**. We further implemented the PCA-based groupwise registration [7,21] using the traditional *elastix* toolbox [11] as another method in comparison, denoted as **Groupwise**.

3.3 Results

Quantitative Analysis. We first present and analyze our quantitative results. As shown in Table 1, we calculated the mean and standard deviation (std) of SD error within the myocardium region for all listed methods. We see that the accuracy of fitting is the lower on the uncorrected original MOLLI data, while the mapping results of learning-based **Morph** and optimization-based **Groupwise** were both significantly improved compared with **Org**. When using the disentanglement framework **Dis**, the T1 mapping based on the generated dataset of **Dis** achieved the worst results. This implies sub-optimal disentanglement, i.e., information leakage between a and c in **Dis**. The results improved when adding L_{anatomy} and $\boldsymbol{p}$. Specifically, the former improved mean and the latter std. This confirms that L_{anatomy} guarantees anatomy disentanglement in presence of residual motion, and that the proposed projector p is efficient in reducing variance. Our **DisQ** achieved further improved results in both mean and std. The mean SD error of ours was still slightly higher than **Groupwise**, however latter demanded lengthy optimization.

Table 1. The mean and standard deviation of fitting quantitative T_1 maps. (Unit: *ms*)

Method	Org	Morph	Dis	Dis+L_{anatomy}	Dis+p	Dis+L_{anatomy}+ p (Our DisQ)	Groupwise
Mean	47.9	39.1	57.9	41.2	43.8	**36.6**	**32.2**
Standard deviation	24.6	22.5	26.3	25.9	21.3	**19.9**	21.0

Qualitative Analysis. We select 11 baseline images of one subject from our test MOLLI sequence, and original 11 frames are shown in Fig. 1(a). Then we show the generated cross-reconstructed data using DisQ in Fig. 3, which is unified through two strategies, either in terms of contrast (for easy image registration, Fig. 3(a)), or anatomy (for direct quantitative mapping, Fig. 3(b)). They share the contrast or anatomy from the selected inversion time respectively. We also show the quantitative native and post-contrast T_1 maps and their SD in Fig. 4, it can be seen that compared with **DisQ**, the SD of **Org** is very obvious at the motion boundary, and **Morph** is affected by drastic changes in contrast and may locally produce large errors.

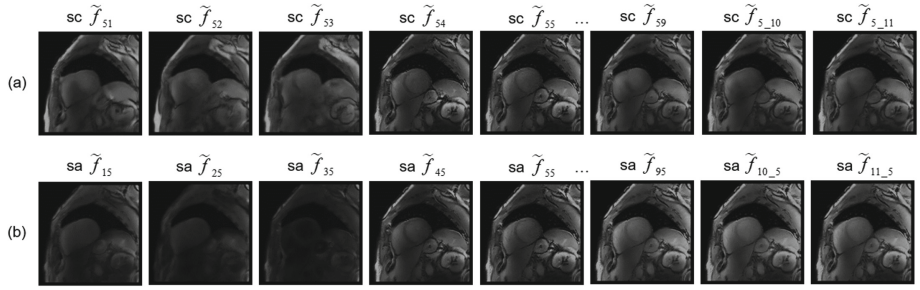

Fig. 3. Utilizing DisQ to analyze cardiac qMRI: (a) all images have the same contrast (*sc*, from the 5th frame), respective anatomies, (b) all images share the same anatomy (*sa*, also from the 5th frame), while preserving their respective contrasts. $\tilde{f}_{ij}$ represents contrast from frame i, anatomy from j.

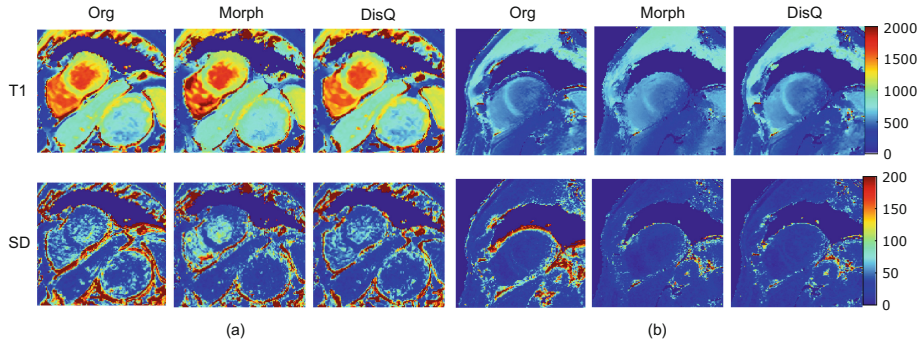

Fig. 4. The resulting quantitative T_1 maps and corresponding SD error maps of (a) native and (b) post-contrast MOLLI sequences. Colorbar in the unit of ms. (Color figure online)

Computational Time. The training time for our disentanglement architecture is ∼10 h on one 3090Ti GPU, and ∼300 ms for inference on a pair of cross reconstructed images. For the registration network, we use the original Voxelmorph, and training time is ∼6 h and evaluation time is ∼400 ms for pairwise registration. For Groupwise registration by Elastix toolbox, the inference time is ∼9000 s. In comparison, our pipeline only takes ∼7 s for disentangling and registering of one MOLLI sequence.

4 Conclusion

In this work, we propose a novel image disentanglement framework DisQ (Disentangling Quantitative MRI) to discompose cardiac qMRI images into their *anatomical representation* and *contrast representation* in the latent space. This is among the first work to address cardiac qMRI analysis from an anatomy-contrast

disentanglement perspective, with effective bootstrapping strategies proposed to tackle simultaneous changes of contrast and anatomy in cardiac qMRI. We applied DisQ to analyze the clinical MOLLI sequences (both native and post-contrast), and demonstrated improved precision for the final cardiac T1 map. Our proposed DisQ methodology is generic, which readily extends to other types of cardiac qMRI such as T_2 mapping and perfusion. Future work is warranted to investigate its generalizability to other qMRI sequences with different underlying physics.

Acknowledgement. The authors gratefully acknowledge TU Delft AI Initiative for financial support.

References

1. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: VoxelMorph: a learning framework for deformable medical image registration. IEEE Trans. Med. Imaging **38**(8), 1788–1800 (2019)
2. Burgess, C.P., et al.: Understanding disentangling in β-VAE. arXiv preprint arXiv:1804.03599 (2018)
3. Chartsias, A., et al.: Disentangle, align and fuse for multimodal and semi-supervised image segmentation. IEEE Trans. Med. Imaging **40**(3), 781–792 (2020)
4. Denton, E.L., et al.: Unsupervised learning of disentangled representations from video. Adv. Neural Inf. Process. Syst. **30** (2017)
5. Dewey, B.E., et al.: A disentangled latent space for cross-site MRI harmonization. In: Martel, A.L., et al. (eds.) MICCAI 2020, Part VII. LNCS, vol. 12267, pp. 720–729. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59728-3_70
6. Haaf, P., Garg, P., Messroghli, D.R., Broadbent, D.A., Greenwood, J.P., Plein, S.: Cardiac T1 mapping and extracellular volume (ECV) in clinical practice: a comprehensive review. J. Cardiovasc. Magn. Reson. **18**(1), 1–12 (2017)
7. Huizinga, W., et al.: PCA-based groupwise image registration for quantitative MRI. Med. Image Anal. **29**, 65–78 (2016)
8. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016, Part II. LNCS, vol. 9906, pp. 694–711. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46475-6_43
9. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4401–4410 (2019)
10. Kellman, P., Hansen, M.S.: T1-mapping in the heart: accuracy and precision. J. Cardiovasc. Magn. Reson. **16**(1), 1–20 (2014). https://doi.org/10.1186/1532-429X-16-2
11. Klein, S., Staring, M., Murphy, K., Viergever, M.A., Pluim, J.P.: Elastix: a toolbox for intensity-based medical image registration. IEEE Trans. Med. Imaging **29**(1), 196–205 (2009)
12. Locatello, F., et al.: Challenging common assumptions in the unsupervised learning of disentangled representations. In: International Conference on Machine Learning, pp. 4114–4124. PMLR (2019)

13. Messroghli, D.R., Radjenovic, A., Kozerke, S., Higgins, D.M., Sivananthan, M.U., Ridgway, J.P.: Modified look-locker inversion recovery (MOLLI) for high-resolution T1 mapping of the heart. Magn. Reason. Med. Off. J. Int. Soc. Magn. Reason. Med. **52**(1), 141–146 (2004)
14. Ouyang, J., Adeli, E., Pohl, K.M., Zhao, Q., Zaharchuk, G.: Representation disentanglement for multi-modal brain MRI analysis. In: Feragen, A., Sommer, S., Schnabel, J., Nielsen, M. (eds.) IPMI 2021. LNCS, vol. 12729, pp. 321–333. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-78191-0_25
15. Paulus, M.B., Maddison, C.J., Krause, A.: Rao-blackwellizing the straight-through gumbel-softmax gradient estimator. arXiv preprint arXiv:2010.04838 (2020)
16. Qin, C., et al.: Joint learning of motion estimation and segmentation for cardiac MR image sequences. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018, Part II. LNCS, vol. 11071, pp. 472–480. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00934-2_53
17. 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
18. Qiu, H., Qin, C., Le Folgoc, L., Hou, B., Schlemper, J., Rueckert, D.: Deep learning for cardiac motion estimation: supervised vs. unsupervised training. In: Pop, M., et al. (eds.) STACOM 2019. LNCS, vol. 12009, pp. 186–194. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-39074-7_20
19. Ren, X., Yang, T., Wang, Y., Zeng, W.: Rethinking content and style: exploring bias for unsupervised disentanglement. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 1823–1832 (2021)
20. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015, Part III. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
21. Tao, Q., van der Tol, P., Berendsen, F.F., Paiman, E.H., Lamb, H.J., van der Geest, R.J.: Robust motion correction for myocardial t1 and extracellular volume mapping by principle component analysis-based groupwise image registration. J. Magn. Reson. Imaging **47**(5), 1397–1405 (2018)
22. Xue, H., et al.: Motion correction for myocardial t1 mapping using image registration with synthetic image estimation. Magn. Reson. Med. **67**(6), 1644–1655 (2012)
23. van Zijl, P., et al.: Quantitative assessment of blood flow, blood volume and blood oxygenation effects in functional magnetic resonance imaging. Nat. Med. **4**(2), 159–167 (1998)
24. Zuo, L., et al.: Unsupervised MR harmonization by learning disentangled representations using information bottleneck theory. NeuroImage **243**, 118569 (2021)