



Unsupervised Multi-modal Style Transfer for Cardiac MR Segmentation

Chen Chen¹(✉), Cheng Ouyang¹(✉), Giacomo Tarroni¹, Jo Schlemper¹, Huaqi Qiu¹, Wenjia Bai^{2,3}, and Daniel Rueckert¹

¹ Biomedical Image Analysis Group, Imperial College London, London, UK
{chen.chen15,c.ouyang}@imperial.ac.uk

² Data Science Institute, Imperial College London, London, UK

³ Department of Medicine, Imperial College London, London, UK

Abstract. In this work, we present a fully automatic method to segment cardiac structures from late-gadolinium enhanced (LGE) images without using labelled LGE data for training, but instead by transferring the anatomical knowledge and features learned on annotated balanced steady-state free precession (bSSFP) images, which are easier to acquire. Our framework mainly consists of two neural networks: a multi-modal image translation network for style transfer and a cascaded segmentation network for image segmentation. The multi-modal image translation network generates realistic and diverse synthetic LGE images conditioned on a single annotated bSSFP image, forming a synthetic LGE training set. This set is then utilized to fine-tune the segmentation network pre-trained on labelled bSSFP images, achieving the goal of unsupervised LGE image segmentation. In particular, the proposed cascaded segmentation network is able to produce accurate segmentation by taking both shape prior and image appearance into account, achieving an average Dice score of 0.92 for the left ventricle, 0.83 for the myocardium, and 0.88 for the right ventricle on the test set.

1 Introduction

Cardiac segmentation from late-gadolinium enhanced (LGE) cardiac magnetic resonance (CMR) which highlights myocardial infarcted tissue is of great clinical importance, enabling quantitative measurements useful for treatment planning and patient management. To this end, the segmentation of the myocardium is an important first step for myocardial infarction analysis.

Since manual segmentation is tedious and likely to suffer from inter-observer variability, it is of great interest to develop an accurate automated segmentation method. However, this is a challenging task due to the fact that (1) the infarcted myocardium presents an enhanced and heterogeneous intensity distribution different from the normal myocardium region and (2) the border between infarcted myocardium and blood pool appears blurry and ambiguous [1]. While

C. Chen and C. Ouyang—Equal contribution.

the borders of the myocardium can be difficult to delineate on LGE images, they are clear and easy to identify on the balanced steady-state free precession (bSSFP) CMR images, which have high signal-to-noise ratio and whose contrast is less sensitive to pathology (see red arrows in Fig. 1(a)). Conventional methods [2, 3] use the segmentation result from the bSSFP CMR of the same patient as prior knowledge to assist the segmentation on LGE CMR images. These methods generally require accurate registration between the bSSFP and LGE images, which can be challenging as the imaging field-of-view (FOV), image contrast and resolution between the two acquisitions can vary significantly [1, 4]. Figure 1(b) visualizes the discrepancy between the intensity distributions of the two imaging modalities in the cardiac structures (specifically, left ventricle (LV), myocardium (MYO), and right ventricle (RV)).




Fig. 1. The differences of image appearance (a) and intensity distributions (b) in the cardiac region (the union of LV, MYO, RV) between LGE images and bSSFP images (Color figure online)

Most recently, a deep neural network-based method has been proposed to segment the three cardiac structures directly from LGE images [5], reporting superior performance. However, this supervised segmentation method requires a large amount of labelled LGE data. Because of the heterogeneous intensity distribution of the myocardium in LGE images and the scarcity of experienced image analysts, it is difficult to perform accurate manual segmentations on LGE images and collect a large training set, compared to that on bSSFP images.

In this paper, we present a fully automatic framework that addresses the above mentioned issues by training a segmentation model without using manual annotations on LGE images. This is achieved by transferring the anatomical knowledge and features learned on annotated bSSFP images, which are easier to acquire. Our framework mainly consists of two neural networks:

- A multi-modal image translation network: this network is used for translating annotated bSSFP images into LGE images through style transfer. Of note, the network is trained in an unsupervised fashion where the training bSSFP

- images and LGE images are **unpaired**. In addition, unlike common one-to-one translation networks, this network allows the generation of **multiple** synthetic LGE images conditioned on a single bSSFP image;
- A cascaded segmentation network for LGE images consisting of two U-net [6] models (Cascaded U-net): this network is first trained using the labelled bSSFP images and then fine-tuned using the synthetic LGE data generated by the image translation network.

The main contributions of our work are the following: (1) we employ a translation network that can generate **realistic** and **diverse** synthetic LGE images given a single bSSFP image. This network enables generative model-based data augmentation for unsupervised domain adaptation, which not only closes the domain gap between the two modalities, but also improves the generalization properties of the following segmentation network by increasing data variety; (2) we demonstrate that the proposed two-stage cascaded network, which takes both anatomical **shape** information and image **appearance** information into account, produces accurate segmentation on LGE images, greatly outperforming baseline methods; (3) the proposed framework can be easily extended to other unsupervised cross-modality domain adaptation applications where labels of one modality are not available.

2 Methodology

The proposed method aims at learning an LGE image segmentation model using labelled bSSFP $\{(\mathbf{x}_b, \mathbf{y}_b)\}$ and unlabelled LGE $\{\mathbf{x}_l\}$ only. Specifically, the proposed method is a two-stage framework. In the first stage, an unsupervised **image translation** network is trained to translate each bSSFP image \mathbf{x}_b into multiple instances of LGE-like images, noted as $\{\mathbf{x}_{bl}\}$. In the second stage, these LGE-stylized bSSFP images are used together with their original labels $\{(\mathbf{x}_{bl}, \mathbf{y}_b)\}$ to adapt an **image segmentation** network pre-trained on labelled bSSFP images to segment LGE images.

2.1 Image Translation

We employ the state-of-the-art multi-modal unsupervised image-to-image translation network (MUNIT) [7] as our multi-modal image translator. Let $\{\mathbf{x}_l\}$ and $\{\mathbf{x}_b\}$ denote unpaired images from the two different imaging modalities (domains): LGE and bSSFP, given an image drawn from one domain as input, the network is able to change the appearance (i.e. image style) of the image to that of the other domain while preserving the underlying anatomical structure [8]. This is achieved by learning disentangled image representations.

As shown in Fig. 2, each image \mathbf{x} is disentangled into (a) a domain-invariant content code \mathbf{c} : $\mathbf{c} = E^c(\mathbf{x})$ and (b) a domain-specific style code \mathbf{s} : $\mathbf{s} = E^s(\mathbf{x})$ using the content encoder E^c and the style encoder E^s relative to its domain where the content code captures the anatomical structure and the style code

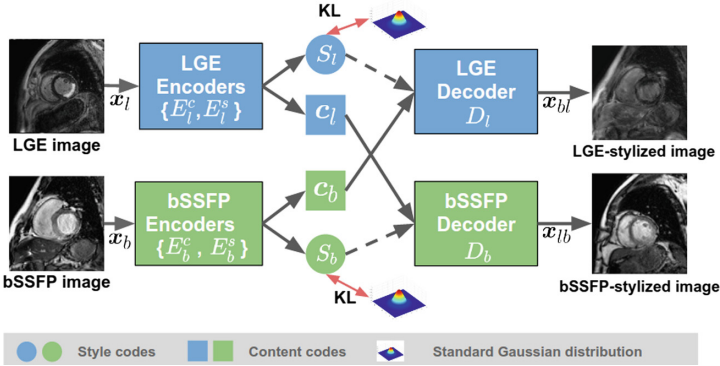


Fig. 2. Overview of the multi-modal image translation network. The network employs the structure of MUNIT [7], which consists of two encoder-decoder pairs for the two domains: bSSFP and LGE, respectively.

carries the information for rendering the structure which is determined by the imaging modality. The image-to-image translation from one domain to the other is achieved by swapping latent codes in two domains. For example, translating a bSSFP image \mathbf{x}_b to be stylized as LGE, is achieved by feeding the content code \mathbf{c}_b for the bSSFP image and the style code \mathbf{s}_l into the LGE decoder D_l : $\mathbf{x}_{bl} = D_l(\mathbf{c}_b, \mathbf{s}_l)$.

Of note, during training, each style encoder is trained to embed images into a latent space that matches the standard Gaussian distribution $\mathcal{N}(0, I)$, minimizing the Kullback-Leibler (KL) divergence between the two. This allows to generate an arbitrary number of synthetic LGE images given a single bSSFP image during inference, by repeatedly sampling the style code from the prior distribution $\mathcal{N}(0, I)$. Of note, although this prior distribution is unimodal, the distribution of translated images in the output space is multi-modal thanks to the nonlinearity of the decoder [7]. We apply this translation network to translate annotated bSSFP images, resulting in a synthetic labelled LGE dataset, which will then be used to finetune a segmentation network. For more details about training the translation network, readers are referred to the original work by Huang *et al.* [7].

2.2 Image Segmentation

Let \mathbf{x}_l be an observed LGE image, the aim of the segmentation task is to estimate label maps \mathbf{y}_l having observed \mathbf{x}_l by modeling the posterior $p(\mathbf{y}_l|\mathbf{x}_l)$. Inspired by curriculum learning [9] and transfer learning, we first train a segmentation network using annotated bSSFP images (source domain; easy examples) and then fine-tune it to segment LGE images (target domain; hard examples). Since labelled LGE images $\{(\mathbf{x}_l, \mathbf{y}_l)\}$ are not available for finetuning, we use a synthetic dataset $\mathcal{X}_{bl} : \{(\mathbf{x}_{bl}, \mathbf{y}_b)\}_{1..N}$ generated by the aforementioned multi-modal image

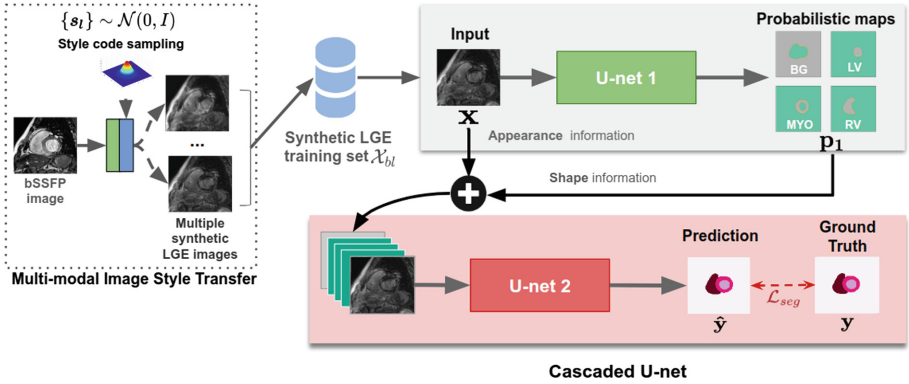


Fig. 3. Overview of the two-stage cascaded segmentation network. The architecture of each U-net is the same as the one of the vanilla U-net, except for two main differences: (1) batch normalization is applied after each convolutional layer; (2) a dropout layer (dropout rate = 0.1) is applied after each concatenation operation in the network’s expanding path to encourage model generalizability. Of note, in this diagram, we simplify the training procedure by omitting the pre-training procedure using labelled bSSFP images.

translator. Ideally, the posterior modelled by the network $p(\mathbf{y}_b|\mathbf{x}_{bl})$ matches $p(\mathbf{y}_l|\mathbf{x}_l)$ when image space and label space are shared. For simplicity, we use \mathbf{x} and \mathbf{y} to denote an image and its corresponding label map from the synthetic dataset in the following paragraphs.

The segmentation network is a two-stage cascaded network which consists of two U-nets [6], see Fig. 3. Specifically, given an image \mathbf{x} as input, the first U-net (U-net 1) aims at predicting four-class pixel-wise probabilistic maps $\mathbf{p}_1 = f_{\text{U-net}}^1(\mathbf{x}; \theta)$ for the three cardiac structures (i.e. LV, MYO, RV) and the background class (BG). Inspired by the auto-context architecture [10], we combine these learned probabilistic maps \mathbf{p}_1 from the first network with the raw image \mathbf{x} to form a 5-channel input to train the second U-net (U-net 2) for fine-grained segmentation: $\mathbf{p}_2 = f_{\text{U-net}}^2(\mathbf{x}, \mathbf{p}_1; \phi)$. By combining the appearance information from the image \mathbf{x} with the shape prior information from the initial segmentation \mathbf{p}_1 as input, the cascaded network has the potential to produce more precise and robust segmentations even in the presence of unclear boundaries for the different cardiac structures.

To train the network, we use a **composite segmentation loss** function \mathcal{L}_{seg} which consists of two loss terms: $\mathcal{L}_{seg} = \mathcal{L}_{wce} + \lambda \mathcal{L}_{edge}$. The first term \mathcal{L}_{wce} is a weighted cross entropy loss: $\mathcal{L}_{wce} = -\sum_m \omega^m \mathbf{y}^m \log(\mathbf{p}^m)$ where w^m denotes the weight for class m and \mathbf{p}^m is the corresponding predicted probability map. We set the weight for myocardium ω^{MYO} to be higher than the weights for the other three classes to address class imbalance problem since there is a lower percentage of pixels that corresponds to the myocardium class in CMR images. The second term \mathcal{L}_{edge} is an edge-based loss which penalizes the disagreement on

the contours of the cardiac structures. Specifically, we apply two 2D 3×3 Sobel filters [11] S_k ($k=1, 2$) to the soft prediction maps \mathbf{p} as well as the one-hot heatmaps \mathbf{y} of the ground truth to extract edge information along horizontal and vertical directions. The edge loss is then computed by calculating the l_2 distance between the predicted edge maps and the ground truth edge maps: $\mathcal{L}_{edge} = \sum_{m, m \neq BG} \sum_{k=1,2} \|f_{S_k}(\mathbf{p}^m) - f_{S_k}(\mathbf{y}^m)\|_2$, where $f_{S_k}(\mathbf{p}^m)$ is the edge map extracted by applying the sobel filter S_k to the predicted probabilistic map \mathbf{p}^m for foreground class m .

By using the edge loss together with the weighted cross entropy for optimization, the network is encouraged to focus more on the contours of the three structures and the myocardium, which are usually more difficult to delineate. In our experiments, we set $\lambda = 0.5$ to balance the contribution of the two losses.

2.3 Post-processing

At inference time, each slice from a previously unseen LGE stack is fed to the cascaded network to get the probabilistic maps for the four classes. Dense conditional random field (CRF) [12] is then applied to refine the 2D predicted segmentation mask slice by slice. After that, 3D morphological dilation and erosion operations are applied to the whole segmentation stack to further improve the global smoothness. In particular, we perform the operations in a hierarchical order: first we apply them to the binary map covering all the three structures, then to the MYO and the LV labels, separately.

3 Experiments and Results

3.1 Data

The framework was trained and evaluated on the Multi-sequence Cardiac MR Segmentation Challenge (MS-CMRSeg 2019) dataset¹. We used a subset of 40 bSSFP and 40 LGE images to train the image translation network. Then, we created a synthetic dataset by applying the learned translation network to 30 labelled bSSFP images. Specifically, for each bSSFP image, we randomly sampled the style code from $\mathcal{N}(0, I)$ five times, resulting in a set of 150 synthetic LGE images in total. This synthetic dataset and the original 30 bSSFP images with corresponding labels formed the training set for the segmentation network. Exemplar results of these synthetic LGE images are provided in the supplemental material. For validation, we used a subset of 5 annotated LGE images provided by the challenge organizers.

3.2 Implementation Details

Image Preprocessing. To deal with the different image size and heterogeneous pixel spacing between different imaging modalities, all images were resampled to

¹ <https://zmiclab.github.io/mscmrseg19/>.

a pixel spacing of $1.25 \text{ mm} \times 1.25 \text{ mm}$ and then cropped to 192×192 pixels, with the heart roughly at the center of each image. This spatial normalization would reduce the computational cost and task complexity in the following training procedure of image translation and segmentation, making the networks focus on the relevant regions. To identify the heart, we trained a localization network based on U-net using the 30 annotated bSSFP images in the training set to produce rough segmentations for the three structures. The localization network employs instance normalization layers which perform style normalization [13], encouraging the network invariance to image style changes (e.g. image contrast). As a result, the network is able to produce coarse masks localizing the heart on all bSSFP images and most LGE images even though it was trained on bSSFP images only. In case that this network might fail to locate the heart on certain LGE slices, we summed the segmentation masks across slices in each volume and then cropped them according to the center of the aggregated mask. After cropping, each image was intensity normalized.

Network Training. (1) For the image translation network, we used the official implementation² of [7]. Network configuration and hyper-parameters were kept the same as in [7] except the input and output images are 2D, single-channel. It was trained for 20k iterations with a batch size of 1. (2) For the segmentation network, we first trained the first U-net with the labelled bSSFP images and then fine-tuned it with synthetic LGE images. This procedure was replicated to train the second U-net with the parameters of the first U-net being fixed. Both networks were optimized using the composite loss \mathcal{L}_{seg} where adam was used for stochastic gradient descent. The learning rate was initially set to 0.001 and was then decreased to 1×10^{-5} for fine-tuning. The weights for BG, LV, MYO, and RV in \mathcal{L}_{wce} were empirically set to 0.2 : 0.25 : 0.3 : 0.25. During training, we applied data augmentation on the fly. Specifically, elastic deformations, random scaling and random rotations as well as gamma augmentation [14] were used. The algorithm was implemented using python and PyTorch and was trained for 1000 epochs in total on an NVIDIA Tesla P40 GPU.

3.3 Results

To evaluate the accuracy of segmentation results, the Dice metric and the average surface distance (ASD) between the automatic segmentation and the corresponding manual segmentation for each volume were calculated.

We compare the proposed method with two baseline methods: (1) a registration-based method and (2) a single U-net. Specifically, for the registration-based method, each LGE segmentation result was obtained by directly registering the corresponding bSSFP labels to the LGE image using MIRTk toolkit³ for ease of comparison. The transformation matrix was learned by applying mutual information-based registration (Rigid+Affine+FFD)

² <https://github.com/NVlabs/MUNIT>.

³ <https://mirtk.github.io/>.

between the two images. For U-net, we trained it with two settings: (a) **U-net**: trained on labelled bSSFP images only; (b) **U-net with fine-tuning (FT)**: trained on labelled bSSFP images and then fine-tuned using the synthetic LGE data, which is the same training procedure of the proposed method. Quantitative and qualitative results are shown in Table 1 and Fig. 4.

While the registration-based method (MIRTK) outperforms the U-net (see row 1 and row 2 in Table 1), it still fails to produce accurate segmentation on the myocardium (see the italic number in row 1), indicating the limitation of this registration-based method. However, by contrast, neural network-based methods (row 3–5) fine-tuned using the *synthetic LGE dataset* significantly improves the segmentation accuracy, increasing the Dice score for MYO by $\sim 15\%$. This improvement demonstrates the learned translation network is capable of generating realistic LGE images while preserving the domain-invariant structural information that is informative to optimize the segmentation network. In particular, compared to U-net (FT), the proposed **Cascaded U-net (FT)** achieves more accurate segmentation performance with improvement in terms of both Dice and ASD (see bold numbers). The model even produces robust segmentation results on the challenging apical and basal slices (please see the last column in Fig. 4). This demonstrates the benefit of integrating the high-level shape knowledge and low-level image appearance to guide the segmentation procedure. In addition, the proposed post-processing further refines the segmentation results through smoothing, reducing the average ASD from 1.37 to 1.26 (see the last row in Table 1).

Table 1. Dice scores and ASD (mm) of the proposed segmentation method (Cascaded U-net) and baseline methods on the validation set. Bold numbers indicate the best scores among the results obtained by those methods before post-processing (PP) whereas italic numbers are those mean Dice scores under 0.700. FT: fine-tuning using the synthetic LGE dataset. N/A means that the ASD value cannot be calculated due to missing predictions for that cardiac structure.

Method	Dice				ASD			
	LV	MYO	RV	AVG ^a	LV	MYO	RV	AVG ^a
MIRTK	0.819	<i>0.665</i>	0.831	0.772	2.56	1.65	2.11	2.11
U-net	<i>0.624</i>	<i>0.441</i>	<i>0.577</i>	<i>0.547</i>	10.03	6.07	N/A	N/A
U-net (FT)	0.874	0.781	0.896	0.850	1.78	1.50	1.28	1.52
Cascaded U-net (FT)	0.895	0.812	0.898	0.868	1.41	1.46	1.23	1.37
Cascaded U-net (FT) + PP	0.897	0.816	0.895	0.869	1.17	1.42	1.18	1.26

^aFor ease of comparison, we calculate the average (AVG) Dice score and the average ASD score over the three structures for each method.

Finally, we applied ensemble learning to improve our model’s performance in the test phase. Specifically, we trained the proposed segmentation network for multiple times, each time regenerating a new synthetic LGE dataset for fine-tuning. We trained four models in total. Our final submission result for each test image was obtained by averaging the probabilistic maps from these models

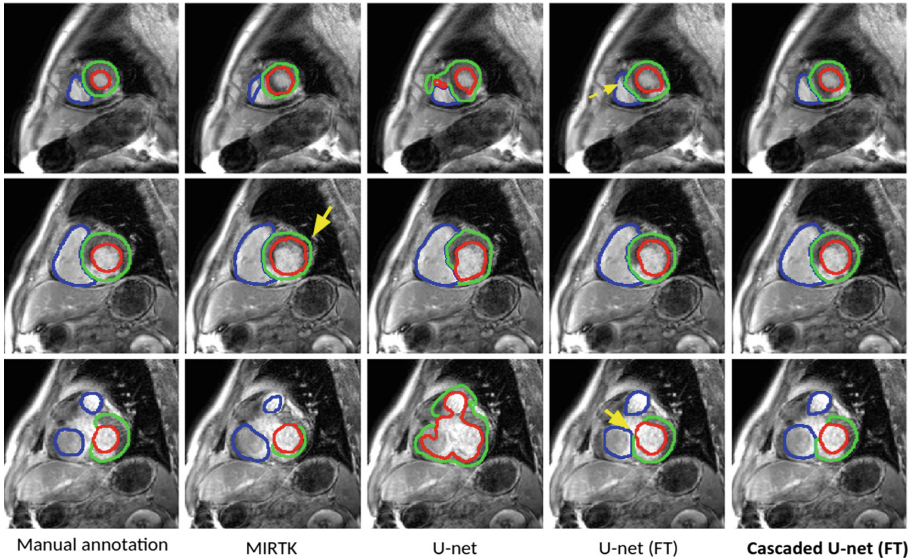


Fig. 4. Segmentation results for the proposed Cascaded U-net and the baseline approaches. Our proposed method (the right-most column) produces more anatomically plausible segmentation results on the images, greatly outperforming the baseline methods, especially in the challenging cases: the apical (the top row) and the basal slices (the bottom row).

and then assigning to each pixel the class with the highest score. In the testing stage of the competition, the method achieves very promising segmentation performance on a relative large test set (40 subjects), with an average Dice score of 0.92 for LV, 0.83 for MYO, and 0.88 for RV; an ASD of 1.66 for LV, 1.76 for MYO, and 2.16 for RV.

4 Conclusion

In this paper, we showed that synthesizing multi-modal LGE images from labelled bSSFP images to finetune a pre-trained segmentation network shows impressive segmentation performance on LGE images even though the network has not seen *real* labelled LGE images before. We also demonstrated that the proposed segmentation network (Cascaded U-net) outperformed the baseline methods by a significant margin, suggesting the benefit of integrating the high-level shape knowledge and low-level image appearance to guide the segmentation procedure. More importantly, our cascaded segmentation network is independent of the particular architecture of underlying convolutional neural networks. In other words, the basic neural network (U-net) in our work can be replaced with any of the state-of-the-art segmentation network to potentially improve the prediction accuracy and robustness. Moreover, the proposed solution based on

unsupervised multi-modal style transfer is not only limited to the cardiac image segmentation but can be extended to other multi-modal image analysis tasks where the manual annotations of one modality are not available. Future work will focus on the application of the method to the problems such as domain adaptation for multi-modality brain segmentation.

Supplemental Material

See Fig. 5.




Fig. 5. Exemplar synthetic LGE images generated from bSSFP images using the multi-modal image translation network. Given **one** bSSFP image (column 1), the translation network translates the image into **multi-modal** LGE-like images (column 2 to 4). These translated images differ in image brightness and contrast as well as the intensity distribution in the cardiac region, while preserving the same cardiac anatomy. These synthetic images, in together with the annotations on the original bSSFP images (the last column) contribute to the synthetic dataset which is used to fine-tune the proposed segmentation network.

References

1. Zhuang, X.: Multivariate mixture model for cardiac segmentation from multi-sequence MRI. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9901, pp. 581–588. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46723-8_67

2. YingLi, L., et al.: Automatic myocardium segmentation of LGE MRI by deformable models with prior shape data. *JCMR* **15**(1), P14 (2013)
3. Tao, Q., et al.: Automated left ventricle segmentation in late gadolinium-enhanced MRI for objective myocardial scar assessment. *JMRI* **42**(2), 390–399 (2015)
4. Zhuang, X.: Multivariate mixture model for myocardium segmentation combining multi-source images. *PAMI* **41**(12), 2933–2946 (2018). <https://ieeexplore.ieee.org/document/8458220>
5. Yue, Q., Luo, X., Ye, Q., Xu, L., Zhuang, X.: Cardiac segmentation from LGE MRI using deep neural network incorporating shape and spatial priors. In: Shen, D., et al. (eds.) *MICCAI 2019*. LNCS, vol. 11765, pp. 559–567. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_62
6. 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*. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
7. Huang, X., Liu, M.-Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) *ECCV 2018*. LNCS, vol. 11207, pp. 179–196. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01219-9_11
8. 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
9. Bengio, Y., Louradour, J., Collobert, R., Weston, J.: Curriculum learning. In: *Proceedings of the 26th Annual International Conference on Machine Learning, ICML 09*, New York, NY, USA, pp. 41–48. ACM (2009)
10. Tu, Z., Bai, X.: Auto-context and its application to high-level vision tasks and 3D brain image segmentation. *PAMI* **32**, 1744–1757 (2010)
11. Sobel, I., Feldman, G.: A 3x3 isotropic gradient operator for image processing. *Pattern Classif. Scene Anal.* 271–272, January 1973. [https://www.scirp.org/\(S\(351jmbntvnsjt1aadkozje\)\)/reference/ReferencesPapers.aspx?ReferenceID=83629](https://www.scirp.org/(S(351jmbntvnsjt1aadkozje))/reference/ReferencesPapers.aspx?ReferenceID=83629)
12. Krähenbühl, P., et al.: Efficient inference in fully connected CRFs with Gaussian edge potentials. In: *NeurIPS* (2011)
13. Huang, X, et al.: Arbitrary style transfer in real-time with adaptive instance normalization. In: *ICCV* (2017)
14. Chen, C., Bai, W., Rueckert, D.: Multi-task Learning for left atrial segmentation on GE-MRI. In: Pop, M., et al. (eds.) *STACOM 2018*. LNCS, vol. 11395, pp. 292–301. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-12029-0_32