

5/2/21

DeU-Net: Deformable U-Net for 3D Cardiac MRI Video Segmentation

Shunjie Dong¹, Jinlong Zhao¹, Maojun Zhang¹, Zhengxue Shi¹, Jianing Deng¹,
Yiyu Shi², Mei Tian¹, and Cheng Zhuo¹

¹ Zhejiang University, Hangzhou, China

{sj_dong,zhaojl,zhmj,sjwo,dengjn,meitian,czhuo}@zju.edu.cn

² University of Notre Dame, Notre Dame, USA

yshi4@nd.edu

Abstract. Automatic segmentation of cardiac magnetic resonance imaging (MRI) facilitates efficient and accurate volume measurement in clinical applications. However, due to anisotropic resolution and ambiguous border (e.g., right ventricular endocardium), existing methods suffer from the degradation of accuracy and robustness in 3D cardiac MRI video segmentation. In this paper, we propose a novel *Deformable U-Net* (DeU-Net) to fully exploit spatio-temporal information from 3D cardiac MRI video, including a Temporal Deformable Aggregation Module (TDAM) and a Deformable Global Position Attention (DGPA) network. First, the TDAM takes a cardiac MRI video clip as input with temporal information extracted by an offset prediction network. Then we fuse extracted temporal information via a temporal aggregation deformable convolution to produce fused feature maps. Furthermore, to aggregate meaningful features, we devise the DGPA network by employing deformable attention U-Net, which can encode a wider range of multi-dimensional contextual information into global and local features. Experimental results show that our DeU-Net achieves the state-of-the-art performance on commonly used evaluation metrics, especially for cardiac marginal information (ASSD and HD).

1 Introduction

Magnetic Resonance Imaging (MRI) is widely used by cardiologists as the golden modality for cardiac assessment [10]. The segmentation of kinetic MR images along the short axis is complicated but essential to precise morphological and pathological analysis, diagnosis, and surgical planning. In particular, one has to delineate left ventricular endocardium (LV), myocardium (MYO), and right ventricular endocardium (RV) to calculate the volume of the cavities in cardiac MRI video, including end-diastolic (ED) and end-systolic (ES) phases [7].

Recent studies [4, 6, 8, 11, 14, 15] proposed deep learning based approaches to learn robust contextual and semantic features, achieving the state-of-the-art segmentation performance. However, automatic and accurate 3D cardiac MRI video segmentation still remains very challenging due to significant variations in the

DeU-Net: Deformable U-Net for 3D Cardiac MRI Video Segmentation

Shunjie Dong¹, Jinlong Zhao¹, Maojun Zhang¹, Zhengxue Shi¹, Jianing Deng¹,
Yiyu Shi², Mei Tian¹, and Cheng Zhuo¹(✉)

¹ Zhejiang University, Hangzhou, China

{sj_dong,zhaojl,zhmj,sjwo,dengjn,meitian,czhuo}@zju.edu.cn

² University of Notre Dame, Notre Dame, USA
yshi4@nd.edu

Abstract. Automatic segmentation of cardiac magnetic resonance imaging (MRI) facilitates efficient and accurate volume measurement in clinical applications. However, due to anisotropic resolution and ambiguous border (e.g., right ventricular endocardium), existing methods suffer from the degradation of accuracy and robustness in 3D cardiac MRI video segmentation. In this paper, we propose a novel *Deformable U-Net* (DeU-Net) to fully exploit spatio-temporal information from 3D cardiac MRI video, including a Temporal Deformable Aggregation Module (TDAM) and a Deformable Global Position Attention (DGPA) network. First, the TDAM takes a cardiac MRI video clip as input with temporal information extracted by an offset prediction network. Then we fuse extracted temporal information via a temporal aggregation deformable convolution to produce fused feature maps. Furthermore, to aggregate meaningful features, we devise the DGPA network by employing deformable attention U-Net, which can encode a wider range of multi-dimensional contextual information into global and local features. Experimental results show that our DeU-Net achieves the state-of-the-art performance on commonly used evaluation metrics, especially for cardiac marginal information (ASSD and HD).

1 Introduction

Magnetic Resonance Imaging (MRI) is widely used by cardiologists as the golden modality for cardiac assessment [9]. The segmentation of kinetic MR images along the short axis is complicated but essential to precise morphological and pathological analysis, diagnosis, and surgical planning. In particular, one has to delineate left ventricular endocardium (LV), myocardium (MYO), and right ventricular endocardium (RV) to calculate the volume of the cavities in cardiac MRI video, including end-diastolic (ED) and end-systolic (ES) phases [7].

Recent studies [8,6,13,14,10,4] proposed deep learning based approaches to learn robust contextual and semantic features, achieving the state-of-the-art segmentation performance. However, automatic and accurate 3D cardiac MRI video segmentation still remains very challenging due to significant variations in the

subjects, ambiguous borders, inhomogeneous intensity and artifacts, especially for RV. There are two key issues that needs to be resolved: (1) In cardiac MRI video, the RV segmentation performance is influenced by complicated shape and inhomogeneous intensity, especially the partial volume effect close to the free wall. (2) Subtle structures (e.g., MYO) have ambiguous borders and different orientations in different anatomical planes of MRI video, causing segmentation inaccuracy. Thus, it is highly desired to have precise and robust cardiac MRI video segmentation.

In this paper, we propose a new *Deformable U-Net* (DeU-Net) to address the aforementioned issues by fully exploiting the spatio-temporal information from 3D cardiac MRI video and aggregating temporal information to boost segmentation performance. The DeU-Net consists of two parts: Temporal Deformable Aggregation Module (TDAM) and a Deformable Global Position Attention (DGPA) network. To address the partial volume effect of RV in [13,14], the TDAM utilizes the spatio-temporal information of the MRI video clip to produce fused feature maps by a temporal aggregation deformable convolution. To handle the issue of subtle structures in [6], the U-Net based DGPA network jointly encodes a wider range of multi-dimensional contextual information into global and local features, guaranteeing clear and continuous borders of every segmentation map. The experimental results quantitatively and qualitatively show that our proposal achieves the state-of-the-art performance on commonly used metrics, especially for cardiac marginal information (ASSD and HD).

2 Method

The architecture of DeU-Net is plotted in Fig. 1, including a Temporal Deformable Aggregation Module (TDAM) and a Deformable Global Position Attention (DGPA) network. The proposed TDAM consists of two phases: a temporal deformable convolution and an offset prediction network based on U-Net to predict deformable offsets. The fused features produced by TDAM are fed into DGPA for the final segmentation results. The DGPA network which also employs U-Net as backbone introduces deformable convolution for encoders and utilizes the deformable attention block to augment the spatial sampling locations.

2.1 Temporal Deformable Aggregation Module

Many existing methods design very complicated neural networks to achieve performance gain. However, most approaches ignore the spatio-temporal information of 3D MRI video, and treat each frame as a separate object, thereby causing performance degradation. Moreover, in the process of data sampling, various semantic details of the video clips may get lost due to fast variation of cardiac borders and regular convolution, inevitably distorting video local details and pixel-wise connections between the frames. Thus, we propose a Temporal Deformable Aggregation Module (TDAM) to adaptively extract temporal information (motion field) for image interpretation. The proposed TDAM takes a target

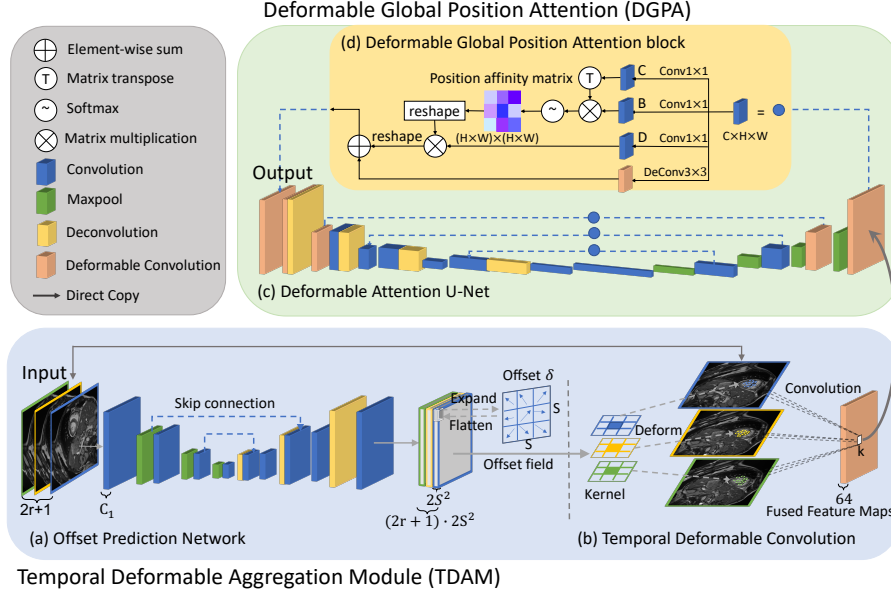


Fig. 1. The architecture of DeU-Net for 3D cardiac MRI video segmentation. Given a video clip ($2r + 1$ concatenated frames) as input, an offset prediction network is designed for deformable offset. Temporal deformable aggregation convolution exploits the offset field to fuse temporal information. The fused feature maps are used by a deformable attention U-Net to enhance segmentation performance. Herein, temporal radius $r = 1$ and deformable kernel size $S = 3$.

frame along with its neighboring reference frames as inputs to jointly predict an offset field. Then, the enhanced contextual information can be fused into the target frame by a temporal aggregation deformable convolution.

We denote the dimension of an input 3D MRI video as $[x, y, z, t]$, where $x - y$ plane is the short-axis plane, z is the short axis, and t is the temporal dimension. Considering the large inter-slice gap in MRI cardiac images along z axis, we select images within the dimension $[x, y, t]$ where stronger correlation may exist. Specially, we define $C_{t_0} \in \mathbb{R}^{H \times W}$ as the target frame in a 3D MRI video at time t_0 , where H and W are the height and width of the input feature map. In order to leverage temporal information, we take the preceding and succeeding r frames as the reference to improve the quality of the target frame C_{t_0} . For a 3D cardiac MRI video clip $\{C_{t_0-r}, \dots, C_{t_0}, \dots, C_{t_0+r}\}$, the conventional temporal fusion scheme (i.e., Early Fusion [5]) can be formulated as a multichannel convolution directly applied on the target frames as:

$$\mathcal{F}(\mathbf{k}) = \sum_{t=t_0-r}^{t_0+r} \sum_{s=1}^{S^2} K_{t,s} \cdot C_t(\mathbf{k} + \mathbf{k}_s), \quad (1)$$

where $\mathcal{F}(\cdot)$ represents the quality-enhanced feature map. S is the size of convolution kernel and $K_t \in \mathbb{R}^{S^2}$ denotes the kernel for t -th channel. \mathbf{k} represents an

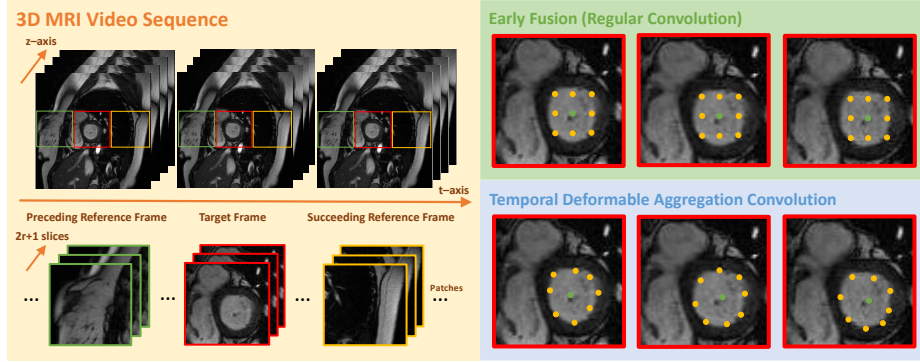


Fig. 2. Comparison between the fixed receptive field in Early Fusion [5] and the adaptive receptive field in TDAM. Each canonical clip is extracted using the samples from the same position in the $x - y$ plane. The yellow points denote the sampling positions of 3×3 convolution window centered at green points. The temporal change of cardiac borders are highlighted, corresponding to the relevant context in the 3D MRI video. The patches are the images collected at different timestamps with the same position.

arbitrary spatial position and \mathbf{k}_s indicates the regular sampling offsets. However, noisy content can be easily introduced due to cardiac temporal change in video. Inspired by Multi-Frame Quality Enhancement [12] for Video Quality Enhancement (VQE), we design TDAM to augment the regular sampling offset with extra learnable offset $\delta_{(t,\mathbf{k})} \in \mathbb{R}^{2S^2}$ for the potential spatio-temporal correlations:

$$\mathbf{k}_s \leftarrow \mathbf{k}_s + \delta_{(t,\mathbf{k}),s}. \quad (2)$$

Note that the deformable offset $\delta_{(t,\mathbf{k})}$ is designed for each convolution window centered at a spatio-temporal position (t, \mathbf{k}) , as shown in Fig. 1. We propose to take the whole clip into consideration and jointly predict all the deformable offsets with a U-Net based network [8]. Maxpool and deconvolutional layers are used for downsampling and upsampling, respectively. Convolutional layer with stride 1, zero paddings are designed to retain the feature size. With such a scheme, the spatial deformations with temporal dynamics in 3D MRI video clips can be simultaneously modeled.

Compared with the existing VQE approaches which achieve explicit motion compensation before fusion to alleviate the negative effects of temporal motion, TDAM implicitly focuses on cardiac border cues with position-specific sampling. As shown in Fig. 2, adjacent deformable convolution windows can independently sample the contents, achieving higher flexibility and performance boost.

2.2 Deformable Global Position Attention.

Regular convolution is restricted by the kernel size as well as the fixed geometric structures, which results in limited performance in modeling geometric transformations. In practice, it is very difficult to reduce false-positive predictions due

to unclear borders between the cardiac instances. To address these issues, we propose a Deformable Global Position Attention (DGPA) network to capture a sufficiently large receptive field and semantic global contextual information. DGPA augments the spatial sampling locations in the modules with additional offsets, which is designed to model complex geometric transformations. Thus, long-range contextual information can be collected, which helps obtain more discriminative cardiac border for pixel-level prediction.

As shown in Fig. 1, the fused local features $I \in \mathbb{R}^{N \times H \times W}$ are regarded as input to the DGPA block, where N represents the number of input channels, H and W indicate the height and width of the input features, respectively. We first feed the input features I with a 3×3 deformable convolution layer to capture cardiac geometric information. The formulation is as below:

$$O = I \otimes \mathcal{K}_{l, \delta'}, \quad (3)$$

where $O \in \mathbb{R}^{N \times H \times W}$ the feature map, \mathcal{K} is deformable convolution kernel, l is the kernel size and δ' is the deformable offset. The input feature map is reshaped to three new feature maps $B, C, D \in \mathbb{R}^{N \times M}$, where M denotes the number of pixels ($M = H \times W$). In order to exploit the high-level features of cardiac borders, a dot-product is conducted between B and the transpose of C . Then the result is applied into a softmax layer to calculate the attention map $P \in \mathbb{R}^{N \times N}$:

$$p_{ji} = \frac{\exp(B_i \cdot C_j^T)}{\sum_{i=1}^N \exp(B_i \cdot C_j^T)}, \quad (4)$$

where p_{ji} represents the i^{th} pixel's impact on the j^{th} pixel. The more similar feature representations of the two pixels indicate a stronger correlation between them. Then we perform a matrix multiplication between the transpose of P and D to reshape the result to $\mathbb{R}^{N \times H \times W}$. Finally, an element-wise sum operation is applied with the feature map O from deformable block to obtain the output features $\mathcal{Z} \in \mathbb{R}^{N \times H \times W}$ as below:

$$\mathcal{Z}_j = \alpha \sum_{i=1}^N (p_{ji} \cdot D_i) + O_j, \quad (5)$$

where α is the scale parameter belonging to the position affinity matrix. Each element in \mathcal{Z} is a weighted sum of the features globally and selectively aggregates input features I . Long-range dependencies of the feature map are calculated to improve intra-class compact and semantic consistency.

3 Experiments and Results

3.1 Setup

Evaluation Datasets. We evaluate our proposal and competitive approaches on the publicly available data of ACDC MICCAI 2017 Challenge [1] with additional labeling done by experience radiologists [10]. The dataset has right ventricle, myocardium, and left ventricle segmentation frames from MRI videos with labels

provided by experienced radiologists. The collected images following the common clinical SSFP cine-MRI sequence have similar properties as 3D cardiac MRI videos. The MRI sequence, as a series of short-axis slices of end diastolic and end systolic instant, starts from the mitral valves down to the apex of the left ventricle. We resize the exams into 256×256 images, and no additional pre-processing was conducted. The dataset has 150 exams from different patients with 100 for training and 50 for testing.

Implementation Details. The proposed method is implemented based on PyTorch library with reference to MMDetection toolbox [3] for deformable convolution, using a NVIDIA GTX 1080Ti GPU. For the training set, standard data augmentation (i.e., mirror, axial flip or rotation) is further used to exploit training samples better. We use Adam optimizer to update the network parameters. The initial learning rate is set to 2×10^{-4} and a weight decay of 1×10^{-4} . We use a batch size of at least 12. The number of reference frames r from Eq. 1 is set as 1. The training is stopped if the Dice score does not increase by 20 epochs. In our experiments, we perform 5-fold cross-validation.

Comparison Methods. We compare the proposed DeU-Net with several state-of-the-art approaches for 3D Cardiac MRI Video segmentation: (1) k U-Net [2] combined convolutional and recurrent neural networks to exploit the intra-slice and inter-slice contexts, respectively. (2) GridNet [14] incorporated a shape prior whose registration on the input image is learned by the model, learning both high-level and low-level features. (3) Attention U-Net [6] proposed a novel self-attention gating module to learn irrelevant regions in an input image for dense label predictions. For a fair comparison, all these methods are modified and fully trained for 3D cardiac MRI video segmentation.

3.2 Results

For quantitative evaluation, Table 1 details the comparison among U-Net, GridNet, Attention U-Net, k U-Net, and the proposed DeU-Net on average symmetric surface distance (ASSD), Hausdorff distance (HD), and Dice score. In addition to the scores at ED and ES phases, we evaluate the average score of the entire 3D cardiac MRI video for all the three metrics (with more detailed statistics reported in the supplementary materials due to the space limit). We can observe that DeU-Net significantly outperforms all the prior networks on most metrics. It is worth noting that our proposal substantially improves segmentation performance on RV, where the ventricle has complicated shape and intensity inhomogeneities. Note that, the proposed method achieves the best results on ASSD that is lower than GridNet by an average of 0.19 mm. Compared to existing methods, HD is lower for our approach by an average of 7.15 mm, and the Dice score is better by an average of 5%. This is a strong indication that DeU-Net exploits spatio-temporal information and reduces the negative effect of cardiac border ambiguity.

To separate the contributions of TDAM and DGPA, we also evaluate the performance of three variants of DeU-Net: (1) DeU-Net(t), the variant without

Table 1. Average scores of the 3D cardiac MRI video different metrics and approaches.

Method	ASSD LV		ASSD MYO		ASSD RV		ASSD
	ED	ES	ED	ES	ED	ES	Average
U-Net	0.34±.09	0.51±.08	0.36±.07	0.41±.05	0.81±.06	1.65±.07	0.78±.06
kU-Net	0.20±.12	0.32±.07	0.34±.03	0.41±.07	0.62±.08	0.58±.06	0.51±.11
GridNet	0.16±.10	0.25±.09	0.24±.07	0.26±.09	0.27±.11	0.49±.08	0.38±.08
Attention U-Net	0.17±.13	0.32±.11	0.24±.09	0.27±.09	0.26±.08	0.56±.13	0.39±.12
ToFlow U-Net	0.12±.07	0.28±.04	0.21±.05	0.15±.04	0.30±.06	0.57±.08	0.27±.10
DeU-Net(t)	0.07±.08	0.23±.09	0.19±.02	0.18±.06	0.22±.04	0.51±.09	0.24±.08
DeU-Net(d)	0.10±.04	0.18±.06	0.17±.03	0.20±.03	0.19±.05	0.49±.02	0.22±.05
DeU-Net	0.04±.01	0.12±.04	0.12±.02	0.12±.00	0.13±.03	0.41±.03	0.19±.04
Method	HD LV		HD MYO		HD RV		HD
	ED	ES	ED	ES	ED	ES	Average
U-Net	6.17±.88	8.29±.62	15.26±.98	17.92±.23	20.51±.54	21.21±.88	19.89±.73
kU-Net	4.59±.32	5.40±.55	7.11±.79	6.45±.64	11.92±.52	14.83±.42	14.38±.70
GridNet	5.96±.42	6.57±.41	8.68±.45	8.99±.84	13.48±.53	16.66±.83	15.06±.25
Attention U-Net	4.39±.76	5.27±.74	7.02±.88	7.35±.90	12.65±.58	10.99±.41	13.95±.21
ToFlow U-Net	3.21±.47	4.32±.63	6.20±.91	6.02±.69	10.39±.81	13.00±.24	11.19±.43
DeU-Net(t)	4.08±.32	4.59±.09	5.91±.33	5.31±.70	11.58±.12	12.23±.13	9.28±.14
DeU-Net(d)	2.76±.36	4.27±.55	4.77±.40	4.22±.43	12.13±.23	11.97±.09	7.69±.24
DeU-Net	2.48±.27	3.25±.30	4.56±.29	4.20±.14	9.88±.17	9.02±.11	6.80±.17
Method	Dice LV		Dice MYO		Dice RV		Dice
	ED	ES	ED	ES	ED	ES	Average
U-Net	0.96±.00	0.90±.01	0.78±.01	0.76±.02	0.88±.02	0.80±.02	0.81±.03
kU-Net	0.96±.00	0.90±.00	0.88±.02	0.89±.03	0.91±.03	0.82±.03	0.83±.01
GridNet	0.96±.01	0.91±.01	0.88±.03	0.90±.03	0.90±.01	0.82±.03	0.85±.02
Attention U-Net	0.96±.01	0.91±.02	0.88±.02	0.90±.01	0.91±.01	0.83±.02	0.84±.03
ToFlow U-Net	0.96±.00	0.91±.01	0.90±.01	0.90±.01	0.92±.01	0.84±.02	0.87±.02
DeU-Net(t)	0.96±.00	0.91±.00	0.89±.00	0.90±.00	0.92±.01	0.83±.01	0.86±.03
DeU-Net(d)	0.96±.01	0.91±.01	0.88±.01	0.91±.01	0.92±.00	0.84±.00	0.88±.02
DeU-Net	0.97±.00	0.92±.00	0.90±.00	0.91±.01	0.93±.00	0.86±.01	0.90±.01

TDAM. (2) DeU-Net(d), the variant without DGPA. (3) ToFlow U-Net, the variant replacing TDAM by Task-oriented Flow (ToFlow) [11]. For the differences among RV, MYO and LV, Table 1 shows that DeU-Net(d) works better than ToFlow U-Net and DeU-Net(t) on both HD and ASSD, indicating that TDAM can fully explore temporal correspondences across multiple frames, especially close to the borders. Moreover, DeU-Net also performs better over all the three variants, validating the necessity of having both TDAM and DGPA in the flow.

Finally, Fig. 3 illustrates a visual comparison of the ground truth, GridNet, Attention U-Net, ToFlow U-Net and DeU-Net in a 3D cardiac MRI video clip. It can be seen that GridNet and Attention U-Net successfully produce accurate results on most slices of each 3D volume, but the shape of the target region is not as accurate as ToFlow U-Net and DeU-Net. Such observations are especially apparent in the rows of 2, 3, and 4 of Fig. 3. Moreover, DeU-Net accurately extracts the borders of the target regions on most of the slices, as shown in the last row of Fig. 3. Note that, the segmentation performance of RV, labeled as blue, is significantly lower than that of MYO and LV due to the irregular shape and the ambiguous borders.




Fig. 3. Visualization of the segmentation results by different methods on the testing data. RV, MYO, and LV are labeled in red, green and blue, respectively.

4 Discussions and Conclusions

In this paper, we propose a *Deformable U-Net* (DeU-Net) to fully exploit spatio-temporal information from 3D cardiac MRI video, including a Temporal Deformable Aggregation Module (TDAM) and a Deformable Global Position Attention (DGPA) network. Based on the temporal correlation across multiple frames, TDAM aggregates temporal information with learnable sampling offsets, and capture sufficient semantic context. To obtain the discriminative and compact features in subtle structures, the DGPA network encodes a wider range of multi-dimensional fused contextual information into global and local features. Experimental results show that our proposal achieves the state-of-the-art performance on commonly used metrics, especially for cardiac marginal information (ASSD and HD). In the future, it would be of interest to apply our proposal to other datasets (such as myocardial contrast echocardiography). Our segmentation method will facilitate the translation of neural networks to clinical practice.

Acknowledgement. This work was supported in part by National Key Research and Development Program Program of China [No. 2018YFE0126300], Key Area Research and Development Program of Guangdong Province [No. 2018B030338001], and Information Technology Center, Zhejiang University.

References

1. Bernard, O., Lalonde, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P.A., Cetin, I., Lekadir, K., Camara, O., Ballester, M.A.G., et al.: Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE transactions on medical imaging* **37**(11), 2514–2525 (2018)
2. Chen, J., Yang, L., Zhang, Y., Alber, M., Chen, D.Z.: Combining fully convolutional and recurrent neural networks for 3d biomedical image segmentation. In: *Advances in neural information processing systems*. pp. 3036–3044 (2016)
3. Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., et al.: Mmdetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155* (2019)
4. Deng, J., Wang, L., Pu, S., Zhuo, C.: Spatio-temporal deformable convolution for compressed video quality enhancement. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. pp. 10696–10703 (2020)
5. Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., Fei-Fei, L.: Large-scale video classification with convolutional neural networks. In: *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. pp. 1725–1732 (2014)
6. Oktay, O., Schlemper, J., Folgoc, L.L., Lee, M., Heinrich, M., Misawa, K., Mori, K., McDonagh, S., Hammerla, N.Y., Kainz, B., et al.: Attention u-net: Learning where to look for the pancreas. *arXiv preprint arXiv:1804.03999* (2018)
7. Peng, P., Lekadir, K., Gooya, A., Shao, L., Petersen, S.E., Frangi, A.F.: A review of heart chamber segmentation for structural and functional analysis using cardiac magnetic resonance imaging. *Magnetic Resonance Materials in Physics, Biology and Medicine* **29**(2), 155–195 (2016)
8. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical image computing and computer-assisted intervention*. pp. 234–241. Springer (2015)
9. Vick III, G.W.: The gold standard for noninvasive imaging in coronary heart disease: magnetic resonance imaging. *Current opinion in cardiology* **24**(6), 567–579 (2009)
10. Wang, T., Xiong, J., Xu, X., Jiang, M., Yuan, H., Huang, M., Zhuang, J., Shi, Y.: Msu-net: Multiscale statistical u-net for real-time 3d cardiac mri video segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 614–622. Springer (2019)
11. Xue, T., Chen, B., Wu, J., Wei, D., Freeman, W.T.: Video enhancement with task-oriented flow. *International Journal of Computer Vision* **127**(8), 1106–1125 (2019)
12. Yang, R., Xu, M., Wang, Z., Li, T.: Multi-frame quality enhancement for compressed video. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 6664–6673 (2018)
13. Zheng, H., Yang, L., Han, J., Zhang, Y., Liang, P., Zhao, Z., Wang, C., Chen, D.Z.: Hfa-net: 3d cardiovascular image segmentation with asymmetrical pooling

- and content-aware fusion. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 759–767. Springer (2019)
14. Zotti, C., Luo, Z., Humbert, O., Lalande, A., Jodoin, P.M.: Gridnet with automatic shape prior registration for automatic mri cardiac segmentation. In: International Workshop on Statistical Atlases and Computational Models of the Heart. pp. 73–81. Springer (2017)