

Dynamic Position Transformation and Boundary Refinement Network for Left Atrial Segmentation

Fangqiang Xu¹, Wenxuan Tu^{*2}, Fan Feng¹, Malitha Gunawardhana¹,
Jiayuan Yang¹, Yun Gu¹, and Jichao Zhao^{*1}

¹ Auckland Bioengineering Institute, University of Auckland
johnnyxu158@gmail.com, j.zhao@auckland.ac.nz

² School of Computer Science and Technology, Hainan University
wenxuantu@163.com

Abstract. Left atrial (LA) segmentation is a crucial technique for diagnosing irregular heartbeat (i.e., atrial fibrillation). Most current methods for LA segmentation strictly assume that the input data is acquired using object-oriented center cropping, while this assumption may not always hold in practice due to the high cost of manual object annotation. Random cropping is a straightforward data pre-processing approach. However, it 1) introduces significant irregularities and incompleteness in the input data and 2) disrupts the coherence and continuity of object boundary regions. To tackle these issues, we propose a novel **D**ynamic **P**osition transformation and **B**oundary refinement **N**etwork (DPBNet). The core idea is to dynamically adjust the relative position of irregular targets to construct their contextual relationships and prioritize difficult boundary pixels to enhance foreground-background distinction. Specifically, we design a shuffle-then-reorder attention module to adjust the position of disrupted objects in the latent space using dynamic generation ratios, such that the vital dependencies among these random cropping targets could be well captured and preserved. Moreover, to improve the accuracy of boundary localization, we introduce a dual fine-grained boundary loss with scenario-adaptive weights to handle the ambiguity of the dual boundary at a fine-grained level, promoting the clarity and continuity of the obtained results. Extensive experimental results on benchmark datasets have demonstrated that DPBNet consistently outperforms existing state-of-the-art methods.

Keywords: Shuffle-then-Reorder Attention · Dual Fine-Grained Boundary Loss · Left Atrial Segmentation.

1 Introduction

Atrial fibrillation (AF), a prevalent arrhythmia, is primarily treated with left atrial (LA) catheter ablation when pharmacological approaches prove insufficient [1,26]. Considering the variability in patients' heart anatomies, minor errors

* Corresponding Author.

in atrial structure analysis can significantly impact disease diagnosis. Therefore, accurate segmentation of the left atrium is both crucial and challenging.

According to the learning paradigm, existing left atrial segmentation methods can be roughly categorized into one-stage and two-stage models [16,19,20,21]. Two-stage models ensemble two decoupled networks to generate segmentation results by detecting region of interest (ROI), while one-stage model is known for its simplicity and effectiveness, by directly predicting results. The key prerequisite for the success of current one-stage methods lies in the assumption that all samples have a centrally complete foreground target [7,11,15,17]. However, such an assumption may not always hold in practical scenarios considering the difficulty in collecting definite central data, influenced by diverse equipment and variable atrial structures in patients [3,27]. For instance, medical experts must spend a substantial amount of time cropping to extract centralized left atrial data from different, which contradicts the goal of developing a fully automated model to reduce human effort and time.




Fig. 1. Input comparison showcases center cropping in (a), (b), random cropping in (c), (d), with all cropping applied to the original image (e). Red areas represent ground truth, while orange arrows highlight irregular object and discontinuous boundary in random cropping outcomes.

Randomly cropping a fixed-size input from original data is a straightforward processing strategy [2]. However, it has two significant drawbacks that impact model performance. As seen in Fig. 1, **1) Random cropped inputs exhibit greater irregularity and incompleteness compared to central inputs.** This irregularity disrupts original positional relationships, posing challenges for most existing convolutional neural networks (CNNs) [14,12,13] to accurately capture contextual dependencies. Especially in cases involving irregular small targets with unknown range relationships, existing networks often struggle to capture effective dependencies among these irregular target regions. As a consequence, this leads to decreased model confidence and prediction accuracy. **2) Random cropping disrupts the coherence and continuity of boundary regions.** Current loss functions typically treat all pixels equally, neglecting the importance of boundary optimization. Despite advancements in boundary weighting techniques, accurately discerning fine-grained distinctions at dual (foreground-background) boundaries remains challenging. This challenge is par-

ticularly prominent when using random cropping, as it can lead to ambiguity or overflow in the boundary regions. In light of this, we find a random-cropped input framework has been an urgent yet under-developed requirement for LA segmentation.

Based on the aforementioned observations, we propose a novel **D**ynamic **P**osition transformation and **B**oundary refinement **N**etwork (DPBNet) for Left Atrial Segmentation. The key idea is to dynamically adjust the relative position of irregular targets to construct their contextual relationships and focus more on hard boundary pixels to enhance foreground-background distinction. Specifically, we design a Shuffle-then-Reorder Attention Module (SRAM) that utilizes adaptive shuffle ratios to dynamically adjust the positions of feature maps. By incorporating the shuffle-then-reorder operation, the network is encouraged to effectively identify critical information for irregular and incomplete targets, regardless of positional constraints. Moreover, we introduce a dual fine-grained boundary (DFB) loss to enhance boundary precision by applying scenario-specific weights at boundary regions. Unlike straightforward weighting towards the foreground edge [15, 18], the DFB Loss distinguishes the fine-grained optimisation difficulties of dual (fore and background) boundary points, effectively optimizing complex boundary regions even with discontinuity inputs. By incorporating these techniques, DPBNet is capable of adapting to various random inputs and achieving a more precise focus on boundary details. This enables the network to generate more reliable intermediate features, ultimately leading to improved performance.

The main contributions of this work are summarized as follows: **1)** To the best of our knowledge, this is the first attempt to investigate a random cropping input one-stage framework, which is more practical for Left Atrial Segmentation. **2)** A novel Shuffle-then-Reorder Attention Module (SRAM) is proposed to obtain dynamic position dependencies for random cropping inputs. Moreover, we introduce a dual fine-grained boundary (DFB) loss to enable more precise optimization of ambiguity regions by assigning different weights to boundary conditions. **3)** Extensive experimental results have demonstrated that our method can achieve state-of-the-art performance compared to existing methods.

2 Methods

Based on the aforesaid analysis, we are confronted with a challenging yet under-explored task: utilizing irregular and discontinuous input to generate accurate segmentation results. An intuitive solution is to manually modify data sequences (e.g., shifting) and magnify foreground-background boundary differences. To achieve this, we propose a Dynamic Position transformation and Boundary refinement Network (DPBNet) with a Shuffle-then-Reorder Attention Module (SRAM) and the Dual Fine-grained Boundary Loss (DFB Loss), as shown in Fig. 2. The following sections will provide a detailed explanation of the SRAM, followed by a comprehensive exploration of the DFB Loss function.




Fig. 2. The pipeline of our Dynamic Position transformation and Boundary refinement framework utilizes VNet as its backbone architecture. Our learning loss consists of a cross-entropy loss and a dual fine-grained boundary loss on the DFB map.

2.1 Shuffle-then-Reorder Attention module

Shuffle-then-Reorder Operation. Inspired by the shift operation of the Swin Transformer [8,5,4], we devise a shuffle-then-reorder operation with generated shuffle ratios. The key idea is to dynamically adjust the positions of feature maps to capture crucial information that may be easily overlooked due to position separation. As illustrated in Fig. 2, we define the shuffle operation on $F \in \mathbb{R}^{C \times H \times W \times D}$ to generate the shuffled feature map $F_s \in \mathbb{R}^{C \times H \times W \times D}$, by adjusting feature map positions, which can be formalized as:

$$F_s[:, i_s^h, i_s^w, i_s^d] = F[:, \kappa(i^h), \kappa(i^w), \kappa(i^d)] \quad (1)$$

where i^x and i_s^x denote positions before and after shuffling along dimension x , while the symbols C, H, W, and D denote the number of channels, height, width, and depth within a feature map, respectively. Consequently, the shuffled indices i_s^x for each dimension $x \in \{h, w, d\}$ are calculated using the following formulation:

$$\kappa(i^x) = ((i^x - 1) \bmod g^x \times r^x) + \left\lfloor \frac{i^x - 1}{g^x} \right\rfloor + 1 \quad (2)$$

where $\bmod$, r^x and g^x refer to the modulo operation, shuffle ratio and the number of groups in each dimension, respectively, satisfying their product equals the dimension size (e.g., $H = r^h \times g^h$). Similarly, the reorder operation is applied inversely to F_s , utilizing these shuffled indices to restore F to its original spatial position. This ensures consistency in the feature representations.

It is worth noting that the shuffle and reorder processes can be implemented using a simple matrix transpose operation. By utilizing these operations, the

model dynamically adjusts the positions of irregular objects within the latent space. This allows for establishing appropriate distance relationships for such targets without being confined to predetermined input positions.

Shuffle-then-Reorder Attention Module. Accordingly, following the shuffle-then-reorder principle, we design a spatial attention-style approach to capture semantically related object dependencies from local and global aspects. Firstly, we apply max and average pooling operations along the channel and spatial dimensions to generate two feature descriptors, namely F_{MAP}^c and F_{MAP}^s . To obtain dynamic shuffle ratios r^h, r^w and r^d , we perform a standard linear layer on F_{MAP}^c :

$$r_h, r_w, r_d = W(F_{MAP}^c), \quad (3)$$

where $W(\cdot)$ and $\{r^h, r^w, r^d\}$ refer to linear layer and shuffle ratios, respectively. Next, to break the positional constraints of irregular and incomplete objects, we perform shuffle and convolution operations on F_{MAP}^s by the generated shuffle ratios. Finally, we employ a reorder operation with the Sigmoid function to obtain an attention map, which is then combined with the original feature F_{org} through element-wise multiplication and a residual connection. This process aims to refine the feature maps F_{ref} . The whole process can be represented as:

$$F_{ref} = (\Theta(\Phi(F_{MAP}^s)) + 1) \otimes F_{org} \quad (4)$$

where $\Phi(\cdot)$ denotes CNN-based shuffle operation, while $\Theta(\cdot)$ denotes reorder operation with Sigmoid function, respectively. In this way, we present a novel plug-and-play attention mechanism that can be seamlessly integrated into existing architectures to effectively capture both short and long-range relationships. Additionally, by incorporating adaptive geometric transformations on the input data, our model mitigates the challenges of fragmentation and lost connectivity among objects caused by random cropping. This adaptation significantly enhances mutual localization between objects, thereby improving overall performance.

2.2 Dual Fine-grained Boundary Loss

To overcome the limitations of boundary discontinuity, we propose a novel Dual Fine-grained Boundary Loss (DFB Loss). The key idea is to assign different weights to the boundaries of foreground and background regions. This allows us to refine the details of both types of boundaries at a fine-grained level. Initially, we define a point as an interior point if all pixels within its k -neighborhood exhibit uniformity. Conversely, we classify a point as a boundary point if there exists any opposite pixel within this range. Subsequently, the weight of each point is determined based on the number of opposing labels in its k -neighborhood, and interior points are uniformly assigned a weight of 1. This is defined as a dual fine-grained boundary map (DFB Map), and can be represented as follows:

$$w_i = \begin{cases} k^3 - f^3 d(g_i) + 1 & \text{if } i = 1; \\ f^3 d(g_i) + 1 & \text{if } i = 0. \end{cases} \quad (5)$$

where k^3 represents the maximum possible count of foreground points within the k -nearest neighbors. Moreover, we employ a 3D-CNNs to accurately determine the actual number of foreground points (with a value of 1) within the k -nearest neighbors of a central point. As illustrated in Fig. 2, when the central point is classified as a foreground point, we determine the count of surrounding background points by subtracting the output of the 3D-CNNs from k^3 . Ultimately, the DFB loss can be formulated as follows:

$$\ell_{dfb} = 1 - 2 \times \frac{\sum_i^N w_i \cdot p_i \cdot g_i + \epsilon}{\sum_i^N w_i \cdot p_i + w_i \cdot g_i + \epsilon} \quad (6)$$

where N represents the total number of points, p_i denotes the predicted probability for point i , g_i represents the ground truth for point i , and ϵ is a small constant introduced to avoid division by zero. Consequently, the gradient of DFB Loss can be written as:

$$\frac{\partial \ell_{dfb}}{\partial p_i} = -\frac{2Ng_i w_i}{Nw_i(g_i + p_i)} + \frac{Nw_i(2Ng_i p_i w_i) + 2\epsilon}{(Nw_i(g_i + p_i) + \epsilon)^2} \quad (7)$$

The gradient of the DFB loss signifies that the model assigns varying priorities to different types of boundaries during training. In other words, the model places greater emphasis on optimizing the more challenging boundary points, ensuring that these points undergo distinct refinement and receive higher priority in the training process. By enhancing the contrast between background and foreground at the boundaries, the proposed model achieves sharper boundary results, effectively mitigating the discontinuity issue caused by random cropping. To ensure smoother optimization, we further integrate the Cross-Entropy Loss (CE Loss) to optimize the model:

$$\ell_{total} = \ell_{ce} + \ell_{dfb} \quad (8)$$

Here, ℓ_{dfb} represents the DFB Loss, ℓ_{ce} denotes the CE Loss, and ℓ_{total} represents the final combined loss.

3 Experiments and results

Dataset. To evaluate the effectiveness of the proposed DPBNet, we adopt the most commonly used LA dataset from the 2018 Atrial Segmentation Challenge [22]. Specifically, the LA dataset consists of 100 3D late gadolinium-enhanced magnetic resonance images, fully annotated for the left atrial cavity, with an isotropic resolution of $0.625 \times 0.625 \times 0.625$ mm. Additionally, the LA dataset presents images in two distinct resolutions, namely $576 \times 576 \times 88$ and

$640 \times 640 \times 88$. In our DPBNet, we maintain the original resolution of the images throughout both the training and testing phases. Subsequently, the dataset was split into 80 training and 20 testing scans, with a 5-fold cross-validation approach for evaluation.

Implementation Details. For training the networks, we employ the Adam optimizer, with a momentum of 0.9 and a weight decay of $1e-4$, aligning with common practices. The input images of dimensions $256 \times 256 \times 80$ are obtained through random cropping from the original image data. The training process is conducted in a total of 15000 iterations, initiated with a learning rate of $1e-4$. To evaluate and compare the performance of DPBNet against other methods, we employ four widely utilized metrics: Dice score, Jaccard index, 95% Hausdorff Distance (HD95), and Average Symmetric Surface Distance (ASSD). For tests, a sliding window ($256 \times 256 \times 80$, stride: $160 \times 160 \times 4$) extracts patches from original inputs, with results averaged for final predictions [6,24,17]. All experiments are conducted by PyTorch 2.1 (Python 3.8) on a single NVIDIA Tesla A100 GPU and an Intel Xeon Platinum 8168 CPU @ 2.70GHz.

3.1 Comparison Results

To validate the effectiveness of the proposed method, we conduct several experiments with seven state-of-the-art methods on the LA database, which is shown in Table 1 and Fig. 3. It can be seen from Fig. 3. that our model generates more accurate and clear predictions compared to state-of-the-art approaches, with DPBNet displaying enhanced boundary details detection as evidenced by yellow areas. Moreover, DPBNet achieves state-of-the-art (SOTA) performance without dependence on any pre-processing or post-processing, directly mitigating challenges from random cropping.

Table 1. The comparison across seven existing methods on the LA database highlights top results in bold. Notably, 'C-C', 'R-R' denotes center or random cropping for both input and output, while 'R-C' represents random input with center output. R-DPBNet and C-DPBNet denote testing with random or center data, using same parameters.

Methods	Setting	Metrics			
		Dice(%)	Jaccard(%)	HD95(mm)	ASSD(mm)
VNet [10] (MICCAI'16)	C-C	$91.85_{\pm 0.31}$	$85.01_{\pm 0.54}$	$2.92_{\pm 0.32}$	$0.96_{\pm 0.11}$
Yang et al. [23] (MICCAI'18)	C-C	92.24	85.64	-	1.49
SEGANet [9] (MICCAI'20)	C-C	$91.0_{\pm 0.2}$	$84.0_{\pm 0.3}$	-	$1.00_{\pm 0.21}$
Zhao et al. [28] (ICPR'21)	C-C	$91.79_{\pm 1.06}$	-	$2.87_{\pm 0.67}$	-
LANet [15] (TMI'22)	C-C	$92.0_{\pm 0.2}$	$86.0_{\pm 0.3}$	$2.88_{\pm 0.56}$	$0.86_{\pm 0.24}$
UMSMLNet [7] (MP'22)	C-C	$92.02_{\pm 0.29}$	$85.28_{\pm 0.49}$	$2.84_{\pm 0.15}$	$0.89_{\pm 0.02}$
CANet [25] (ESA'23)	C-C	91.24	83.96	5.71	1.57
R-DPBNet (Ours)	R-R	$92.50_{\pm 0.25}$	$86.10_{\pm 0.44}$	$2.83_{\pm 0.16}$	$0.94_{\pm 0.07}$
C-DPBNet (Ours)	R-C	$92.57_{\pm 0.33}$	$86.24_{\pm 0.57}$	$2.74_{\pm 0.27}$	$0.85_{\pm 0.04}$




Fig. 3. Visualization results by various methods and DPBNet, showcasing prediction (blue) and ground truth (red). First row shows 2D prediction-ground truth overlap, and second features 3D distance maps from prediction-ground truth differences.

Quantitative analysis in Table 1 demonstrates that our R-DPBNet, utilizing random cropping, surpasses competing methods on most metrics, notably achieving a Dice score of 92.50%. Advantages in HD95 and ASSD for our R-DPBNet are limited by its comprehensive analysis ($640 \times 640 \times 88$) through random cropping, as opposed to the narrower center cropping ($256 \times 256 \times 88$) approach. After employing identical center cropping configurations, as in C-DPBNet, the proposed method in this paper easily achieves the best performance across all metrics. Summarily, our approach breaks impractical assumptions and addresses random cropping challenges, setting new benchmarks for excellence.

Table 2. Ablation studies of our proposed DPBNet on the LA database.

Methods	Setting	Metrics			
		Dice(% $\uparrow$)	Jaccard(% $\uparrow$)	HD95(mm $\downarrow$)	ASSD(mm $\downarrow$)
Vanilla VNet [10]	R-R	90.77	83.37	4.94	1.37
+ SRAM (k=3)	R-R	91.81	84.95	2.96	1.44
+ SRAM (k=5)	R-R	92.21	85.61	3.11	1.06
+ SRAM (k=7)	R-R	92.16	85.54	6.83	1.25
+ Edge Loss [15]	R-R	91.67	84.72	4.60	1.46
+ DFB Loss	R-R	92.29	85.74	3.89	0.90
+ All (k=5)	R-R	92.33	85.83	2.74	1.00

3.2 Ablation Study

To validate the contribution of the SRAM and DFB Loss, we conduct an ablation study on the LA dataset, with united framework VNet [10] as the backbone. As shown in Table 2, all the setting of SRAM generates better results than the

baseline, while kernel size 5 gains better performance than other settings. This implies that dynamic adjustments in positioning can enhance the capability of mutual target localization, generating better performance. Moreover, DFB Loss shows significant improvements over both the baseline and Edge Loss, which adopts single and uniform boundary weights. This improvement stems from DFB Loss enhancing boundary refinement and foreground-background contrast.

4 Conclusion

In this paper, we address the impractical assumption of current LA segmentation methods on object-oriented center cropping due to high manual annotation costs. Specifically, we propose a new Dynamic Position transformation and Boundary refinement Network (DPBNet) by designing a shuffle-then-reorder attention module and a dual fine-grained boundary loss. The core strategy dynamically enhances object interactions and refines foreground-background details. Comprehensive experiments demonstrate that the proposed DPBNet achieves SOTA results on the LA dataset, to the best of our knowledge.

Acknowledgments. This work was supported in part by the Health Research Council of New Zealand (#21355), the Postgraduate Research Student Support (PReSS) funding in University of Auckland, and Postgraduate Scientific Research Innovation Project in Hunan Province under Grant CX20220076.

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

References

1. Burstein, B., Nattel, S.: Atrial fibrosis: mechanisms and clinical relevance in atrial fibrillation. *Journal of the American College of Cardiology* **51**(8), 802–809 (2008)
2. Chen, C., Bai, W., Rueckert, D.: Multi-task learning for left atrial segmentation on ge-mri. In: *Statistical Atlases and Computational Models of the Heart*. pp. 292–301. Springer (2019)
3. Hansen, B.J., Zhao, J., Csepe, T.A., Moore, B.T., Li, N., Jayne, L.A., Kalyanasundaram, A., Lim, P., Bratasz, A., Powell, K.A., et al.: Atrial fibrillation driven by micro-anatomic intramural re-entry revealed by simultaneous sub-epicardial and sub-endocardial optical mapping in explanted human hearts. *European heart journal* **36**, 2390–2401 (2015)
4. Jiang, B., Xu, F., Tu, W., Yang, C.: Channel-wise attention in 3d convolutional networks for violence detection. In: *2019 International Conference on Intelligent Computing and its Emerging Applications (ICEA)*. pp. 59–64. IEEE (2019)
5. Jiang, B., Xu, F., Xia, J., Yang, C., Huang, W., Huang, Y.: Stacked multi-scale attention network for image colorization. In: *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 2225–2229. IEEE (2022)

6. Li, S., Zhang, C., He, X.: Shape-aware semi-supervised 3d semantic segmentation for medical images. In: *Medical Image Computing and Computer Assisted Intervention*. pp. 552–561. Springer (2020)
7. Liu, Y., Wang, W., Luo, G., Wang, K., Liang, D., Li, S.: Uncertainty-guided symmetric multilevel supervision network for 3d left atrium segmentation in late gadolinium-enhanced mri. *Medical Physics* **49**(7), 4554–4565 (2022)
8. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 10012–10022 (2021)
9. Lourenço, A., Kerfoot, E., Dibblin, C., Alskaf, E., Anjari, M., Bharath, A.A., King, A.P., Chubb, H., Correia, T.M., Varela, M.: Left atrial ejection fraction estimation using seganet for fully automated segmentation of cine mri. In: *Statistical Atlases and Computational Models of the Heart*. pp. 137–145. Springer (2021)
10. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: *2016 fourth international conference on 3D vision (3DV)*. pp. 565–571. IEEE (2016)
11. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention*. pp. 234–241. Springer (2015)
12. Tu, W., Guan, R., Zhou, S., Ma, C., Peng, X., Cai, Z., Liu, Z., Cheng, J., Liu, X.: Attribute-missing graph clustering network. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. pp. 15392–15401 (2024)
13. Tu, W., Xiao, B., Liu, X., Zhou, S., Cai, Z., Cheng, J.: Revisiting initializing then refining: An incomplete and missing graph imputation network. *IEEE Transactions on Neural Networks and Learning Systems* pp. 1–14 (2024)
14. Tu, W., Zhou, S., Liu, X., Ge, C., Cai, Z., Liu, Y.: Hierarchically contrastive hard sample mining for graph self-supervised pretraining. *IEEE Transactions on Neural Networks and Learning Systems* pp. 1–14 (2023)
15. Uslu, F., Varela, M., Boniface, G., Mahenthiran, T., Chubb, H., Bharath, A.A.: La-net: A multi-task deep network for the segmentation of the left atrium. *IEEE transactions on medical imaging* **41**(2), 456–464 (2021)
16. Vesal, S., Ravikumar, N., Maier, A.: Dilated convolutions in neural networks for left atrial segmentation in 3d gadolinium enhanced-mri. In: *Statistical Atlases and Computational Models of the Heart*. pp. 319–328. Springer (2019)
17. Wu, Y., Xu, M., Ge, Z., Cai, J., Zhang, L.: Semi-supervised left atrium segmentation with mutual consistency training. In: *Medical Image Computing and Computer Assisted Intervention*. pp. 297–306. Springer (2021)
18. Xia, J., Tan, G., Xiao, Y., Xu, F., Leung, C.S.: Edge-aware multi-scale progressive colorization. In: *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 1655–1659. IEEE (2021)
19. Xia, Q., Yao, Y., Hu, Z., Hao, A.: Automatic 3d atrial segmentation from ge-mris using volumetric fully convolutional networks. In: *Statistical Atlases and Computational Models of the Heart*. pp. 211–220. Springer (2019)
20. Xiong, Z., Fedorov, V.V., Fu, X., Cheng, E., Macleod, R., Zhao, J.: Fully automatic left atrium segmentation from late gadolinium enhanced magnetic resonance imaging using a dual fully convolutional neural network. *IEEE transactions on medical imaging* **38**(2), 515–524 (2018)
21. Xiong, Z., Nalar, A., Jamart, K., Stiles, M.K., Fedorov, V.V., Zhao, J.: Fully automatic 3d bi-atria segmentation from late gadolinium-enhanced mris using double

- convolutional neural networks. In: *Statistical Atlases and Computational Models of the Heart*. pp. 63–71. Springer (2020)
22. Xiong, Z., Xia, Q., Hu, Z., Huang, N., Bian, C., Zheng, Y., Vesal, S., Ravikumar, N., Maier, A., Yang, X., et al.: A global benchmark of algorithms for segmenting the left atrium from late gadolinium-enhanced cardiac magnetic resonance imaging. *Medical image analysis* **67**, 101832 (2021)
 23. Yang, X., Wang, N., Wang, Y., Wang, X., Nezafat, R., Ni, D., Heng, P.A.: Combating uncertainty with novel losses for automatic left atrium segmentation. In: *Statistical Atlases and Computational Models of the Heart*. pp. 246–254. Springer (2019)
 24. Yu, L., Wang, S., Li, X., Fu, C.W., Heng, P.A.: Uncertainty-aware self-ensembling model for semi-supervised 3d left atrium segmentation. In: *Medical Image Computing and Computer Assisted Intervention*. pp. 605–613. Springer (2019)
 25. Zhao, C., Xiang, S., Wang, Y., Cai, Z., Shen, J., Zhou, S., Zhao, D., Su, W., Guo, S., Li, S.: Context-aware network fusing transformer and v-net for semi-supervised segmentation of 3d left atrium. *Expert Systems with Applications* **214**, 119105 (2023)
 26. Zhao, J., Butters, T.D., Zhang, H., Pullan, A.J., LeGrice, I.J., Sands, G.B., Smaill, B.H.: An image-based model of atrial muscular architecture: effects of structural anisotropy on electrical activation. *Circulation: Arrhythmia and Electrophysiology* **5**(2), 361–370 (2012)
 27. Zhao, J., Hansen, B.J., Wang, Y., Csepe, T.A., Sul, L.V., Tang, A., Yuan, Y., Li, N., Bratasz, A., Powell, K.A., et al.: Three-dimensional integrated functional, structural, and computational mapping to define the structural “fingerprints” of heart-specific atrial fibrillation drivers in human heart ex vivo. *Journal of the American Heart Association* **6**(8), e005922 (2017)
 28. Zhao, Z., Puybureau, E., Boutry, N., Géraud, T.: Do not treat boundaries and regions differently: An example on heart left atrial segmentation. In: *2020 25th International Conference on Pattern Recognition (ICPR)*. pp. 7447–7453. IEEE (2021)