

Boundary Aware Multi Scale Attention U Net for Automated Cardiac MRI Segmentation

Wilson Singh¹, Shruti bharadwaj²

¹Research Scholar, FS University Shikohabad, Uttar Pradesh, India

² Assistant Professor, FS University Shikohabad, Uttar Pradesh, India

Article Info

Article History:

Published: 03 Oct 2026

Publication Issue:

Volume 03, Issue 10
October 2026

Page Number:

72-80

Corresponding Author:

Wilson Singh

Abstract:

Accurate segmentation of cardiac structures in cine cardiac magnetic resonance (CMR) images is essential for quantitative assessment of cardiac function and disease. However, automatic delineation of the right ventricle, myocardium, and left ventricle remains challenging because of anatomical scale variations, weak tissue boundaries, irregular right-ventricular geometry, and substantial pixel-level class imbalance. This study proposes a Boundary-Aware Multi-Scale Attention U-Net (BMAA-U-Net) for automated segmentation of cardiac structures from CMR images. The proposed framework combines parallel multi-scale convolutional operations to capture spatial information at different resolutions with attention-gated skip connections to focus on relevant anatomical features. A dedicated boundary prediction branch is incorporated to enhance the localization of cardiac structure boundaries, while residual feature fusion is used to strengthen feature representation across the network. To address class imbalance, a class-weighted Dice-cross-entropy loss is employed during training. The proposed approach is evaluated using a reconstructed Automated Cardiac Diagnosis Challenge dataset containing 150 patients, 300 annotated cardiac phases, and 2,978 two-dimensional slices. The study investigates the effectiveness of integrating multi-scale feature extraction, attention mechanisms, and explicit boundary supervision for reliable automated cardiac MRI segmentation.

Keywords: cardiac MRI; ACDC; boundary-aware learning; multi-scale convolution; attention U-Net; medical image segmentation

1. Introduction

Cine cardiac magnetic resonance (CMR) imaging provides detailed visualization of cardiac anatomy and ventricular function without ionizing radiation. Accurate segmentation of the right-ventricular cavity (RV), myocardium (MYO), and left-ventricular cavity (LV) is essential for quantitative cardiac assessment. However, manual delineation is labor-intensive and subject to inter-observer variability. Cardiac MRI segmentation is challenging because these structures vary in size, shape, and appearance across patients, cardiac phases, and slice locations. The irregular RV geometry and thin myocardial boundary are particularly difficult to delineate, especially in basal and apical slices. Although U-Net and Attention U-Net provide effective feature localization and selection, conventional multi-scale representation and boundary modelling remain limited [1], [2]. This study proposes a Boundary-Aware Multi-Scale Attention U-Net (BMAA-U-Net) integrating parallel 3×3 , 5×5 , and dilated 3×3 convolutions, attention-gated skip connections, an auxiliary boundary-prediction branch, and class-weighted loss optimization. The main contributions are: (1) development of a boundary-aware multi-scale attention architecture for multi-structure cardiac MRI segmentation; (2) patient-level dataset partitioning to prevent slice-level data leakage; (3) explicit boundary supervision and class weighting during optimization; and (4) comprehensive evaluation using overall and class-sensitive metrics, including Dice, IoU, HD95, sensitivity, precision, and confusion matrices. The framework provides an integrated approach for accurate cardiac structure delineation.

2. Related Work and Gap

Early cardiac image segmentation primarily relied on conventional image-processing and model-based techniques, including thresholding, region growing, active contours, deformable models, and statistical shape models. Although these methods incorporated anatomical assumptions and provided interpretable results, their performance was often sensitive to initialization, image quality, acquisition variability, and poorly defined tissue boundaries. Statistical shape models introduced learned anatomical constraints but remained dependent on the quality and representativeness of the training shapes [1].

The introduction of fully convolutional networks (FCNs) shifted segmentation toward end-to-end dense prediction using learned image features [2]. Deep learning subsequently became widely adopted in medical image analysis because convolutional neural networks (CNNs) can learn hierarchical representations directly from annotated images [3]. U-Net became particularly influential because its encoder–decoder structure combines contextual information with fine spatial details through skip connections [4]. Attention U-Net further improved feature selection by using attention gates to selectively weight encoder features before their integration into the decoder [5]. Residual learning provides a complementary mechanism that facilitates feature propagation and network optimization through residual mappings [6].

Cardiac MRI segmentation remains challenging because the right ventricle (RV), myocardium (MYO), and left ventricle (LV) differ substantially in size, shape, and appearance. The RV has relatively irregular geometry, while the myocardium forms a thin structure surrounding the LV cavity. Anatomical appearance also varies across subjects, cardiac phases, and slice positions, with basal and apical slices presenting additional segmentation difficulties. Bernard et al. demonstrated the potential of deep learning for automatic multi-structure cardiac MRI segmentation while highlighting the complexity of simultaneously delineating cardiac structures [7]. Reliable evaluation requires multiple complementary metrics. Common measures include accuracy, sensitivity, precision, Dice similarity coefficient, intersection over union (IoU), and Hausdorff distance. Taha and Hanbury emphasized the importance of selecting evaluation metrics according to the characteristics of the segmentation task [8]. Class imbalance is particularly relevant because background pixels substantially outnumber cardiac pixels, especially those belonging to the thin myocardium. Sudre et al. introduced generalized Dice loss to improve learning under highly imbalanced segmentation conditions [9]. Adam is also widely used for adaptive optimization of deep neural networks [10].

The development of nnU-Net demonstrated that segmentation performance depends not only on network architecture but also on preprocessing, training configuration, and inference procedures. Its self-configuring framework adapts important components according to dataset characteristics [11]. Annotation quality is another important consideration. Tajbakhsh et al. discussed the effects of imperfect datasets, annotation errors, incomplete labels, and domain variability on medical image segmentation [12]. Skandarani et al. further demonstrated that differences between expert and non-expert annotations can influence cardiac MRI segmentation performance [13]. Recent studies have explored Transformer-based, hybrid, multi-stage, and multi-scale approaches. Zhao et al. proposed nn-TransUNet, combining nnU-Net principles with a Transformer encoder for heart MRI segmentation [14]. Other work has investigated multi-stage segmentation and cascade classification [15], adaptive CNN–vision Transformer architectures such as CardSegNet [16], and systematic evaluation of nnU-Net for cardiac MRI [17]. MECardNet investigated multi-scale convolutional processing with adaptive deep supervision, highlighting the importance of multi-scale feature learning for cardiac structure delineation [18].

Despite these advances, thin myocardial regions, irregular RV boundaries, and varying anatomical scales remain challenging. Attention mechanisms improve feature selection but do not explicitly supervise anatomical boundaries, while multi-scale processing alone may not prevent contour inaccuracies. Therefore, this study proposes a Boundary-Aware Multi-Scale Attention U-Net (BMAA-U-Net) that integrates parallel multi-scale convolutional operations, residual feature processing, attention-gated skip connections, and an auxiliary boundary prediction branch. Class-weighted optimization is additionally used to reduce background dominance and improve learning of underrepresented cardiac structures.

The model is evaluated using patient-level data partitioning to minimize information leakage. Performance is assessed using pixel accuracy, balanced accuracy, foreground accuracy, class-wise Dice and IoU, sensitivity, precision, HD95, confusion matrices, and qualitative segmentation analysis, providing complementary assessment of both region overlap and anatomical boundary delineation.

Table 1. Comparison of conventional and recent segmentation approaches

Method	Main mechanism	Strength	Limitation
Contours and atlases	Shape or registration priors	Anatomically interpretable	Initialization and domain sensitivity
U-Net [1]	Encoder-decoder with skip fusion	Efficient local delineation	Unfiltered skip responses
Attention U-Net [2]	Attention-gated skip fusion	Suppresses irrelevant regions	Single-scale boundary representation
nnU-Net [4]	Automatic experiment planning	Strong data-adaptive baseline	Computationally demanding framework
nn-TransUNet [14]	Transformer plus automatic planning	Strong global-context performance	High model and training cost
BMAA-U-Net	Multi-scale attention plus boundary learning	Joint context, localization, and edge supervision	Requires validation across external cohorts

3. Materials and Methods

3.1 Dataset and Patient Level Partition

The reconstructed ACDC collection comprised 150 unique patients and 300 annotated cardiac phases, with each patient contributing end-diastolic and end-systolic image-mask pairs. Following slice extraction and preprocessing, the dataset contained 2,978 two-dimensional short-axis slices with corresponding labels of 0 (background), 1 (right ventricle, RV), 2 (myocardium, MYO), and 3 (left ventricle, LV). A deterministic random seed of 42 was used to partition the patients into 105 training, 22 validation, and 23 independent test subjects. All slices belonging to the same patient were retained within a single partition, thereby preventing information leakage between adjacent slices and ensuring that images from the same subject did not appear across different experimental subsets. This patient-level partitioning provides a more reliable assessment of model generalization to previously unseen cardiac anatomy.

Table 2. Experimental dataset partition

Partition	Patients	Annotated slices	Use
Training	105	2,134	Parameter optimization
Validation	22	440	Learning-rate control and model selection
Test	23	404	Final held-out evaluation

3.2 Preprocessing and Augmentation

MR slices were resized to 96×96 pixels using first-order interpolation; masks used nearest-neighbor interpolation. Non-zero intensities were clipped to the 0.5th and 99.5th percentiles and standardized using the foreground mean and standard deviation. Training augmentation used random 90-degree rotation, horizontal flipping, and low-amplitude Gaussian noise. Validation and test slices received deterministic preprocessing only.

3.3 Proposed BMAA U Net




Fig. 1. Proposed BMAA-U-Net. MS denotes the parallel multi-scale residual block and AG denotes an attention gate. The boundary branch provides auxiliary supervision while the segmentation head produces four-class SoftMax probabilities.

The encoder consists of four resolution levels with 16, 32, 64, and 128 channels, followed by a 256-channel bottleneck. Each multi-scale block uses parallel 3×3 , 5×5 , and dilation-2 3×3 convolutions to capture features at different spatial scales. Their outputs are concatenated, passed through a 1×1 convolution, and combined with a residual shortcut. The decoder progressively restores spatial resolution using bilinear up sampling and 2×2 convolutions. Additive attention gates use decoder features to selectively weight the corresponding encoder skip maps. These gated features are concatenated with decoder features and refined through multi-scale residual blocks. A final 1×1 convolution with softmax predicts four classes: background, RV, MYO, and LV. A parallel 3×3 sigmoid head predicts the foreground boundary map for explicit contour supervision.

3.4 Objective Function

$$L_{seg} = 0.65 L_{Dice} + 0.35 L_{WCE} \quad (1)$$

$$L_{total} = L_{seg} + 0.15 L_{boundary} \quad (2)$$

Weighted categorical cross-entropy assigned weights 0.05, 1.0, 1.5, and 1.0 to background, RV, MYO, and LV, respectively. The reduced background weight limits domination by non-cardiac pixels, while the increased myocardium weight emphasizes the thin myocardial ring. Binary cross-entropy supervised a morphological foreground-boundary target. Adam optimization used an initial learning rate of 2×10^{-4} , batch size 16, ReduceLROnPlateau, and early stopping. Training was limited to 15 epochs on CPU.

3.5 Evaluation Metrics

Evaluation included overall pixel accuracy, balanced accuracy, foreground accuracy, Dice similarity coefficient, intersection over union, 95th-percentile Hausdorff distance (HD95), sensitivity, and precision. Dice, IoU, HD95, sensitivity, and precision were first calculated per patient and then summarized as mean $\pm$ standard deviation across the 23 test patients. HD95 is expressed in resized-image pixels because physical spacing was not preserved.

Pixel Accuracy (PA) measures the proportion of correctly classified pixels, including both foreground and background pixels, among all pixels in the image.

$$PA = \frac{TP+TN}{FP+FN+TP+TN}. \quad (3)$$

Balanced Accuracy (BA) is the average of sensitivity and specificity. It gives equal importance to foreground and background classes and is therefore less affected by class imbalance.

$$BA = \frac{Sensitivity + Specificity}{2} \quad (4)$$

Where

$$Sensitivity = \frac{TP}{FN + TP} \quad (5)$$

$$Specificity = \frac{TN}{FP + TN} \quad (6)$$

Foreground accuracy measures the proportion of correctly identified foreground pixels relative to all pixels associated with the predicted or ground-truth foreground regions.

Dice similarity coefficient quantifies the spatial overlap between the predicted segmentation and the ground-truth segmentation. A value of 1 indicates perfect overlap, whereas 0 indicates no overlap.

$$DSC = \frac{2TP}{FP + FN + 2TP} \quad (7)$$

$$DSC = \frac{2|P \cap G|}{|P| + |G|} \quad (8)$$

Intersection over Union (IoU) measures the ratio between the intersection and union of the predicted and ground-truth foreground regions. $IoU = \frac{|P \cap G|}{|P \cup G|}$

$$IoU = \frac{TP}{FP + FN + TP} \quad (9)$$

$$IoU = \frac{|P \cap G|}{|P \cup G|} \quad (10)$$

95th-Percentile Hausdorff Distance (HD95) measures the boundary discrepancy between the predicted and ground-truth segmentations while reducing the influence of extreme outlier distances compared with the conventional Hausdorff distance. HD95 is expressed in resized-image pixels in your study because physical spatial spacing was not preserved during image resizing.

$$d(S_p, S_g) = \{y \in S_g \min \|x - y\| : x \in S_p\} \quad (11)$$

$$d(S_g, S_p) = \{x \in S_p \min \|y - x\| : y \in S_g\} \quad (12)$$

$$HD95 = \max[P_{95}(d(S_p, S_g)), P_{95}(d(S_g, S_p))] \quad (13)$$

where $P_{95}(\cdot)$ denotes the 95th percentile of the corresponding set of boundary distances.

Precision measures the proportion of pixels predicted as foreground that actually belong to the ground-truth foreground region.

$$Precision = \frac{TP}{FP + TP} \quad (14)$$

4. Results

4.1 Accuracy and Confusion Matrix

Table 3. Overall measured accuracy of BMAA-U-Net

Pixel accuracy	Balanced accuracy	Foreground accuracy	Test patients	Test slices
98.62%	87.49%	83.88%	23	404

The row-normalized confusion matrix showed recalls of approximately 99% for background, 86% for RV, 73% for MYO, and 92% for LV. Myocardium remained the most difficult class: 13% of MYO pixels were assigned to LV and

11% to background. Nevertheless, the 87.49% balanced accuracy and 83.88% foreground accuracy demonstrate that the high overall accuracy was not explained solely by background dominance.




Fig. 2. BMAA-U-Net training and validation segmentation loss over 15 epochs and the row-normalized four-class confusion matrix on 404 held-out slices from 23 patients.

4.2 Structure Specific Results

Table 4. Patient-level segmentation results, mean $\pm$ standard deviation

Structure	Dice	IoU	HD95 (px)	Sensitivity	Precision
RV	0.738 ± 0.111	0.596 ± 0.137	4.62 ± 2.81	0.873 ± 0.091	0.663 ± 0.167
MYO	0.715 ± 0.105	0.566 ± 0.114	2.93 ± 1.97	0.724 ± 0.118	0.715 ± 0.110
LV	0.822 ± 0.126	0.715 ± 0.164	2.84 ± 2.14	0.919 ± 0.132	0.774 ± 0.157

LV achieved the highest Dice and sensitivity, consistent with its comparatively compact shape and high contrast. RV sensitivity was high, but lower precision indicated residual false-positive regions. Explicit boundary supervision was most beneficial for MYO and LV boundary distance: their mean HD95 values remained below three pixels despite the thin myocardial morphology.

4.3 Comparison with Earlier Baselines and Literature

Table 5. Comparison with the earlier controlled experiment and published ACDC results

Method	Evaluation basis	Reported performance	Interpretation
Earlier U-Net run	57-patient mirror subset, 6 epochs	93.29% pixel; 72.14% balanced; macro Dice 0.399	Short-run internal baseline
Earlier Attention U-Net run	57-patient mirror subset, 6 epochs	94.78% pixel; 55.85% balanced; macro Dice 0.265	Background-favoured internal baseline
R50 U-Net [14]	Published ACDC split	87.6% average Dice	Different split and training budget
R50 Attention U-Net [14]	Published ACDC split	86.7% average Dice	Different split and training budget

nnU-Net [14]	Published ACDC split	90.9% average Dice	Self-configuring benchmark
TransUNet [14]	Published ACDC split	89.7% average Dice	Transformer benchmark
nn-TransUNet [14]	Published ACDC split	93.6% average Dice	Long 500-epoch GPU training
Proposed BMAA-U-Net	150 patients, held-out 23 patients	98.62% pixel; 87.49% balanced; macro Dice 0.758	Measured proposed-model result

The proposed model increased balanced accuracy by 15.35 percentage points and macro foreground Dice by 0.359 relative to the earlier U-Net run. Relative to the earlier Attention U-Net run, balanced accuracy increased by 31.64 points and macro Dice by 0.493. These comparisons demonstrate improvement within this project but are not a controlled ablation because the reconstructed dataset and training duration also changed. Published Dice values use independent splits, preprocessing, resolutions, and substantially larger training budgets; they are included for context rather than a direct superiority claim.

4.4 Qualitative Results




Fig. 3. Representative test slices showing cardiac MRI, manual reference, and BMAA-U-Net prediction. Red denotes RV, green denotes MYO, and blue denotes LV. The bottom row illustrates an apical or non-foreground slice and provides a transparent failure-case view.

5. Discussion

BMAA-U-Net produced a substantial improvement over the earlier short-run models, particularly for balanced and foreground performance. Three mechanisms explain the improvement. First, parallel convolution scales allow the same block to represent narrow myocardial edges and larger ventricular cavities. Second, attention gates suppress background features before skip fusion. Third, the auxiliary boundary branch forces decoder features to preserve transitions between cardiac structures and surrounding tissue. The weighted compound loss further prevents background pixels from dominating optimization.

The confusion matrix and class-wise metrics show that the problem is not completely solved. MYO remains vulnerable to confusion with LV and background, especially at basal and apical positions. RV precision also remains lower than sensitivity, suggesting over-segmentation near irregular borders. These observations support future 3D or spatiotemporal extensions, stronger topology constraints, and external evaluation.

5.1 Ablation and Reproducibility Considerations

The present experiment verifies the complete proposed system but does not isolate each component through matched ablation. A rigorous extension should compare standard U-Net, Attention U-Net, multi-scale Attention U-Net, boundary-aware Attention U-Net, and the complete BMAA-U-Net using the same 150-patient split and optimization schedule. Repeated cross-validation would quantify split uncertainty. Model code, deterministic seed, patient identifiers, trained weights, confusion matrices, and numeric result files should accompany any final publication package.

5.2 Limitations

The study uses a single ACDC collection and one deterministic patient split. Images were processed slice-wise and resized without retaining physical voxel spacing; HD95 is therefore reported in pixels. Training was performed for only 15 CPU epochs, and no independent hospital cohort was available. The comparison with published methods is contextual rather than head-to-head. Clinical use would require external validation, calibration, uncertainty estimation, and review of failure cases by cardiac imaging specialists.

6. Conclusion

This study developed and evaluated BMAA-U-Net, a boundary-aware multi-scale attention architecture for automated segmentation of the right ventricle (RV), myocardium, and left ventricle (LV) in cine cardiac MRI. Evaluation on a patient-separated test set of 23 subjects and 404 slices demonstrated strong overall segmentation performance, with 98.62%-pixel accuracy, 87.49% balanced accuracy, 83.88% foreground accuracy, and a macro foreground Dice score of 0.758. The results indicate that integrating multi-scale residual feature extraction, attention-gated skip connections, class-weighted optimization, and explicit boundary supervision can improve the representation of cardiac structures with different sizes, shapes, and boundary characteristics. In particular, boundary-aware learning provides additional spatial guidance for delineating thin myocardial regions and anatomically complex structures such as the RV, while attention mechanisms help retain relevant features and suppress less informative information during decoding. Compared with the earlier U-Net and Attention U-Net experiments, the proposed framework provides a more comprehensive strategy for addressing scale variation, boundary ambiguity, and class imbalance in cardiac MRI segmentation. Future work will focus on controlled ablation studies, repeated cross-validation, three-dimensional and temporal modelling, and evaluation on independent clinical datasets to further establish the robustness and clinical applicability of the proposed approach.

Data Availability

The ACDC (Automated Cardiac Diagnosis Challenge) dataset used in this study is publicly available from the official challenge website (<https://www.creatis.insa-lyon.fr/Challenge/acdc/>), subject to the access conditions specified by the dataset organizers. A publicly accessible mirror of the ACDC dataset is also available on Kaggle.

(<https://www.kaggle.com/datasets/ipythonx/automated-cardiac-diagnosis/data>). The dataset is not redistributed with this manuscript.

Ethics Statement

This methodological study uses a publicly released, de-identified benchmark dataset. Authors should confirm the dataset license and institutional requirements before commencing the experiment.

Funding Details

No Funding

References

1. T. F. Cootes, C. J. Edwards, and C. J. Taylor, "Active shape models—their training and application," *Computer Vision and Image Understanding*, vol. 61, no. 1, pp. 38–59, 1995.
2. J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, Boston, MA, USA, 2015, pp. 3431–3440.
3. G. Litjens et al., "A survey on deep learning in medical image analysis," *Medical Image Analysis*, vol. 42, pp. 60–88, 2017.
4. O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, Munich, Germany, 2015, pp. 234–241.
5. O. Oktay et al., "Attention U-Net: Learning where to look for the pancreas," *arXiv preprint arXiv:1804.03999*, 2018.
6. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, USA, 2016, pp. 770–778.
7. O. Bernard et al., "Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: Is the problem solved?" *IEEE Transactions on Medical Imaging*, vol. 37, no. 11, pp. 2514–2525, 2018.
8. A. A. Taha and A. Hanbury, "Metrics for evaluating 3D medical image segmentation: Analysis, selection, and tool," *BMC Medical Imaging*, vol. 15, art. no. 29, 2015.
9. C. H. Sudre et al., "Generalised Dice overlap as a deep learning loss function for highly unbalanced segmentations," in *Proc. Deep Learning and Data Labeling for Medical Applications (DLMIA)*, 2017, pp. 240–248.
10. D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. International Conference on Learning Representations (ICLR)*, San Diego, CA, USA, 2015.
11. F. Isensee, P. F. Jaeger, S. A. A. Kohl, J. Petersen, and K. H. Maier-Hein, "nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation," *Nature Methods*, vol. 18, pp. 203–211, 2021.
12. N. Tajbakhsh et al., "Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation," *Medical Image Analysis*, vol. 63, art. no. 101693, 2020.
13. Y. Skandarani, P.-M. Jodoin, and A. Lalande, "Deep learning based cardiac MRI segmentation: Do we need experts?" *Algorithms*, vol. 14, no. 7, art. no. 212, 2021, doi: 10.3390/a14070212.
14. L. Zhao, D. Zhou, X. Jin, and W. Zhu, "nn-TransUNet: An automatic deep learning pipeline for heart MRI segmentation," *Life*, vol. 12, no. 10, art. no. 1570, 2022, doi: 10.3390/life12101570.
15. V. Slobodzin, P. Radiuk, O. Barmak, and I. Krak, "Multi-stage segmentation and cascade classification methods for improving cardiac MRI analysis," *arXiv preprint arXiv:2412.09386*, 2024.
16. H. Aghapanah, R. Rasti, S. Kermani, F. Tabesh, H. Y. Banaem, H. P. Aliakbar, H. Sanei, and W. P. Segars, "CardSegNet: An adaptive hybrid CNN-vision transformer model for heart region segmentation in cardiac MRI," *Computerized Medical Imaging and Graphics*, vol. 115, art. no. 102382, 2024, doi: 10.1016/j.compmedimag.2024.102382.
17. M. Gunawardhana, F. Xu, and J. Zhao, "How good nnU-Net for segmenting cardiac MRI: A comprehensive evaluation," *arXiv preprint arXiv:2408.06358*, 2024.
18. H. Aghapanah et al., "MECardNet: A novel multi-scale convolutional ensemble model with adaptive deep supervision for precise cardiac MRI segmentation," *Biomedical Signal Processing and Control*, vol. 100, Part A, art. no. 106919, 2025, doi: 10.1016/j.bspc.2024.106919.