

A Composite Alignment-Aware Framework for Myocardial Lesion Segmentation in Multi-sequence CMR Images

Yifan Gao^{1,2}, Shaohao Rui^{2,3,5}, Haoyang Su^{2,4,5}, Jinyi Xiang⁶, Lianming Wu⁶,
Xiaosong Wang^{2,5}

¹ School of Biomedical Engineering (Suzhou), Division of Life Science and Medicine, University of Science and Technology of China, Hefei, China

² Shanghai Innovation Institute, Shanghai, China

³ Shanghai Jiao Tong University, Shanghai, China

⁴ Fudan University, Shanghai, China

⁵ Shanghai Artificial Intelligence Laboratory, Shanghai, China

⁶ Department of Radiology, Renji Hospital, Shanghai Jiao Tong University, Shanghai, China

240107100014@sii.edu.cn, wangxiaosong@pjlab.org.cn

Abstract. Accurate segmentation of myocardial lesions from multi-sequence cardiac magnetic resonance imaging is essential for cardiac disease diagnosis and treatment planning. However, achieving optimal feature correspondence is challenging due to intensity variations across modalities and spatial misalignment caused by inconsistent slice acquisition protocols. We propose CAA-Seg, a composite alignment-aware framework that addresses these challenges through a two-stage approach. First, we introduce a selective slice alignment method that dynamically identifies and aligns anatomically corresponding slice pairs while excluding mismatched sections, ensuring reliable spatial correspondence between sequences. Second, we develop a hierarchical alignment network that processes multi-sequence features at different semantic levels, i.e., local deformation correction modules address geometric variations in low-level features, while global semantic fusion blocks enable semantic fusion at high levels where intensity discrepancies diminish. We validate our method on a large-scale dataset comprising 397 patients. Experimental results show that our proposed CAA-Seg achieves superior performance on most evaluation metrics, with particularly strong results in myocardial infarction segmentation, representing a substantial 5.54% improvement over state-of-the-art approaches. The code is available at <https://github.com/yifangao112/CAA-Seg>.

Keywords: Multi-sequence CMR · Myocardial lesion segmentation · Alignment.

 Corresponding author

1 Introduction

Cardiovascular diseases remain a leading global health burden, with myocardial pathologies such as infarction and edema requiring precise quantification for effective clinical management [1, 2]. While manual segmentation of these pathological regions is time-consuming and subject to significant inter-observer variability [3, 4], automatic segmentation presents substantial challenges due to the small size and subtle appearance of infarction regions. Multi-sequence cardiac magnetic resonance (CMR) imaging has emerged as a cornerstone for comprehensive myocardial assessment, where late gadolinium enhancement (LGE) delineates lesion with high contrast [5], while complementary T1 and T2 mapping sequences provide additional pathological indicators for both infarction and edema [6, 7].

Leveraging these imaging sequences offers promising opportunities to enhance segmentation performance by combining distinct pathological markers from each modality [8, 9]. However, the integration of these complementary sequences faces two significant technical barriers, as shown in Fig. 1. First, the substantial difference in slice sampling between sequences creates severe anatomical mismatches. LGE typically employs 6-8 slices to capture localized pathologies, whereas T1/T2 mapping sequences use only 2-3 slices for rapid quantification. When these slices are acquired at different anatomical positions, direct registration becomes anatomically inconsistent and clinically unreliable (Fig. 1b.2). Second, attempting to resolve this mismatch through aggressive resampling introduces significant artifacts. Resampling T1/T2 mapping from 3 to 8 slices forces interpolation across large gaps, creating artificial structural continuity and distorting the very pathological features we aim to segment (Fig. 1b.1). Thus, the challenge lies in establishing both inter-sequence semantic correspondence and intra-sequence spatial consistency across these heterogeneously sampled data streams.

Existing methods struggle to address these challenges effectively. Conventional approaches predominantly focus on single-modality analysis [10–12], particularly LGE, sacrificing complementary pathological information from other sequences. While recent multi-modal fusion attempts [13–16] demonstrate promise on small datasets (25 patients) [17] with pre-aligned inputs, such idealized conditions rarely reflect the misalignment prevalent in routine clinical acquisitions. This discrepancy forces networks to process anatomically inconsistent features, particularly in mid-ventricular regions where interpolation artifacts dominate.

To address these challenges, we propose CAA-Seg, a composite alignment-aware framework for myocardial lesion segmentation in multi-sequence CMR images. The framework consists of two sequential stages: selective slice alignment (SSA) and hierarchical feature fusion. First, we propose a selective slice alignment method that handles *resampling distortion* and *anatomical mismatch* through dynamic slice matching. Instead of forcing alignment between all slices, SSA employs sliding-window optimization to identify anatomically corresponding slice pairs while excluding mismatched regions. This generates aligned multi-sequence regions for lesion fusion and single-sequence regions for standard seg-




Fig. 1. (a) Visualization of misalignment challenges in multi-sequence CMR. (b.1) **Re-sampling distortion:** Aggressive resampling of T1/T2 mapping (3 to 8 slices) introduces interpolation artifacts and structural distortion. (b.2) **Anatomical mismatch:** LGE and T1/T2 mapping slices were acquired at different anatomical positions, causing misalignment. These challenges motivate our selective slice alignment strategy that preserves only reliable correspondences while avoiding forced matching of incompatible slices. Our slice-level alignment significantly reduces resampling-induced misregistration. Additionally, we use only duplicated LGE slices to avoid anatomical mismatch.

mentation, which are processed through task-specific channels. Second, we develop a hierarchical alignment network (HA-Net) that addresses multi-sequence feature fusion at different semantic levels. At low-level features where local geometric variations dominate, we employ deformable convolutions and pixel-wise modulation to correct spatial misalignments. High-level features, which encode pathological patterns with reduced intensity variations, are aligned through cascaded cross-attention modules to capture semantic relationships between LGE and mapping sequences. A task-aware controller at the bottleneck level adaptively modulates features based on input characteristics, ensuring effective fusion of multi-sequence information.

Our comprehensive validation on a large-scale CMR dataset demonstrates significant improvements over state-of-the-art methods. CAA-Seg achieves 78.09% Dice score for myocardium segmentation and 65.49% for myocardial edema delineation. Notably, for myocardial infarction segmentation, our method achieves 51.11% Dice score, outperforming the second-best approach by a substantial margin of 5.54%. The main technical contributions are:

1. We propose a novel framework for multi-sequence CMR segmentation that effectively integrates heterogeneous cardiac imaging sequences through composite alignment.
2. We propose a selective slice alignment approach that establishes reliable inter-sequence correspondence through dynamic slice matching and mutual information optimization.
3. We develop a hierarchical alignment network that combines local deformation correction and global semantic fusion for effective multi-sequence feature fusion.

2 Method

We propose a two-stage framework for myocardial lesion segmentation in multi-sequence CMR images. The first stage employs selective slice alignment to establish anatomical correspondence between LGE and mapping sequences through dynamic slice matching. The aligned sequences then serve as input to a hierarchical alignment network (HA-Net), which performs multi-level feature fusion while addressing residual spatial and intensity variations for accurate lesion segmentation.

2.1 Selective Slice Alignment

Multi-sequence CMR acquisition results in varying slice counts and positions across sequences, making direct registration between LGE and mapping sequences challenging. Rather than forcing alignment between all slices, which can introduce artifacts and misleading information, we propose the SSA approach that dynamically identifies and aligns only the most reliable slice pairs.

Given a series of 2D slices, let $I_f^k \in \mathbb{R}^{H \times W}$ denote the k -th fixed LGE slice and $I_m^j \in \mathbb{R}^{H \times W}$ represent the j -th moving slice from either T1m or T2m sequence. We formulate the slice-wise registration as an optimization problem:

$$\phi_k^* = \arg \min_{\phi_k} \underbrace{\mathcal{L}_{MMI}(I_f^k, I_m^j \circ \phi_k)}_{\text{Multi-sequence similarity}} + \lambda \underbrace{\mathcal{R}(\phi_k)}_{\text{Regularization}}, \quad k \in [1, K] \quad (1)$$

where K is the total number of slices. To ensure robust multi-modal alignment, we adopt Mattes mutual information (MMI) as the similarity metric:

$$\mathcal{L}_{MMI}(I_f^k, I_m^j \circ \phi_k) = - \sum_{x,y} p_k(x,y) \log \left(\frac{p_k(x,y)}{p_f^k(x)p_m^j(y)} \right) \quad (2)$$

where $p_k(x,y)$ represents the joint intensity distribution. Our transformation model accommodates both global and local deformations through a hierarchical composition:

$$\phi_k = \phi_{k,rigid} \circ \phi_{k,affine} \circ \phi_{k,diff} \quad (3)$$

The key innovation of our approach lies in its selective slice matching strategy. Instead of a fixed window, we employ an adaptive search to identify the most plausible anatomical correspondence for each slice. First, we compute a similarity score for each potential slice pair:

$$S(k,j) = MMI(I_{LGE}^k, \phi_k(I_{T1m}^j)) + MMI(I_{LGE}^k, \phi_k(I_{T2m}^j)) \quad (4)$$

For the k -th LGE slice, we then identify the best-matching moving slice, j_k^* , by searching within a dynamically defined range. This range is constrained by




Fig. 2. Overview of the proposed HA-Net. The framework consists of dual encoders for multi-sequence CMR processing, with specialized alignment modules at different feature levels. Local deformation correction modules address spatial misalignment in low-level features, while cascaded cross-attention blocks enable semantic fusion at high-level features. A task-aware controller at the bottleneck adaptively modulates features based on different input images.

the position of the previously matched slice (j_{k-1}) to maintain sequential order and ensure enough subsequent slices are available for future matches:

$$j_k^* = \arg \min_{j \in [j_{k-1}, N-M+k]} S(k, j) \quad (5)$$

where j_{k-1} is the matching position of the previous slice, N is the total number of moving image slices, and M is the remaining number of slices from the current position k to the end of the LGE sequence. This ensures that:

$$\begin{cases} j_k^* > j_{k-1} & \text{maintain sequence order} \\ j_k^* \leq N - M + k & \text{reserve sufficient slices} \end{cases} \quad (6)$$

This selective mechanism naturally handles through-plane motion while maintaining anatomical consistency within each 2D plane. The final registered volumes are reconstructed only from the reliably matched and transformed slice pairs:

$$V_{reg} = \{I_m^{j_k^*} \circ \phi_k^*\}_{k=1}^K \quad (7)$$

This method enables our framework to focus on high-quality alignments while avoiding error propagation from forced matching of incompatible slices.

2.2 Hierarchical Alignment Network

While SSA provides reliable inter-sequence alignment, minor local deformations and intensity variations may persist at the feature level. The low-level features primarily encode local tissue textures where geometric misalignments affect structural details, while high-level features capture semantic patterns where

intensity variations become less significant. Therefore, we further propose a hierarchical alignment network (HA-Net) that progressively refines and fuses multi-sequence features at different semantic levels. Our network adopts a residual UNet backbone architecture with specialized fusion modules at different scales.

Let $F_m^l \in \mathbb{R}^{H_l \times W_l \times C_l}$ denote features output from encoder at level $l \in \{stem, low, high, bn\}$ for modality $m \in \{LGE, T1m, T2m\}$, where H_l , W_l , and C_l represent spatial dimensions and channels. The network extracts initial features through the stem layer without cross-modal fusion, preserving modality-specific characteristics.

Local Deformation Correction: At low-level layers, we align features through deformable convolution and pixel-wise modulation. For each modality pair, the offset field Δp and aligned features $F_{aligned}$ are computed as:

$$\begin{aligned}\Delta p &= f_{offset}(F_{LGE}^{low}) \\ F_{aligned} &= \text{DeformConv}(F_{T1m/T2m}^{low}, \Delta p)\end{aligned}\tag{8}$$

The features are then modulated through learned parameters:

$$\begin{aligned}[\gamma, \beta] &= f_{mod}(\text{GlobalPool}([F_{LGE}^{low}, F_{aligned}])) \\ F_{out}^{low} &= \gamma \cdot F_{LGE}^{low} + \beta\end{aligned}\tag{9}$$

where f_{mod} maps the global context to modulation parameters γ and β for pixel-wise feature adjustment.

Global Semantic Fusion: High-level features undergo cascaded cross-attention for semantic alignment. The process consists of two sequential attention blocks:

$$\begin{aligned}F_1^{high} &= \text{Attn}(Q = F_{LGE}^{high}, K, V = F_{T1m}^{high}) \\ F_{out}^{high} &= \text{Attn}(Q = F_1^{high}, K, V = F_{T2m}^{high})\end{aligned}\tag{10}$$

where the attention operation follows the standard scaled dot-product mechanism with multi-head projection.

Task-Aware Controller: Inspired by [18], we incorporate a task-aware controller design that adaptively processes features based on different inputs. At the bottleneck, we integrate task information through a learned prompt P_{task} :

$$F_{out}^{bn} = \text{MLP}(\text{LayerNorm}([F_{LGE}^{bn}; P_{task}]))\tag{11}$$

where $[\cdot; \cdot]$ represents channel-wise concatenation. This controller adapts the feature processing based on the input characteristics from different modalities.

3 Experiments and Results

Dataset and Experimental Setup: We evaluate our method on an IRB-approved, large-scale in-house dataset comprising 397 patients with confirmed myocardial pathologies, randomly split into training (278 cases), validation (39 cases), and testing (80 cases) sets. Among these, all cases included annotations

Table 1. Quantitative evaluation results of different methods in myocardial lesion segmentation dataset. The best and second best results are highlighted in bold and underlined respectively. Myo: myocardium, ME: myocardial edema, MI: myocardial infarction.

Methods	Myo		ME		MI		Overall	
	Dice (%)	HD95 (pixel)	Dice (%)	HD95 (pixel)	Dice (%)	HD95 (pixel)	Dice (%)	HD95 (pixel)
nnU-Net [19]	76.46	10.45	62.68	9.52	45.30	<u>12.78</u>	61.48	<u>10.92</u>
UNet++ [20]	77.51	<u>9.84</u>	65.03	9.77	44.26	14.03	62.27	11.21
TransUNet [21]	77.05	9.87	63.00	11.08	43.25	13.21	61.10	11.39
Swin UNETR [22]	75.23	11.31	61.39	11.21	35.31	14.94	57.30	12.49
UTNet [23]	76.99	10.25	62.16	10.15	42.06	14.23	60.40	11.54
UMamba [24]	<u>77.58</u>	10.56	64.93	9.23 [~]	<u>45.57</u>	13.11	<u>62.69</u>	10.97
MFU-Net [13]	74.84	10.89	59.22	12.53	36.68	14.34	56.91	12.59
AWSNet [14]	76.50	10.71	63.06	9.72	40.00	14.71	59.85	11.71
MyoPS-Net [15]	77.20	9.88	63.90	9.55	44.15	14.56	61.75	11.33
CAA-Seg (ours)	78.09 [~]	9.00 [‡]	65.49 [‡]	<u>9.25</u>	51.11 [‡]	11.43 [‡]	64.89 [‡]	9.89 [‡]

~: p-value $\leq$ 0.05, ‡: p-value $\leq$ 0.05, in comparison between the best and second-best results.

for myocardium (Myo) and myocardial edema (ME), while 208 cases presented with myocardial infarction (MI).

We use seven encoder stages with feature dimensions of {32, 64, 128, 256, 512, 512, 512} channels. The deformable convolution modules use 3×3 kernels with 18 sampling points. Input images are resampled and standardized to 384×384 pixels via crop or pad. The features naturally separate into low-level (64-256 channels) and high-level (512 channels) representations. The cross-attention blocks employ 8 attention heads with a hidden dimension of 256. The task controller generates a 128-dimensional task embedding that is projected to match feature dimensions.

Our comparisons include general medical segmentation frameworks (nnU-Net [19], UNet++ [20], TransUNet [21], Swin UNETR [22], UTNet [23], UMamba [24]) and cardiac-specific approaches (MFU-Net [13], AWSNet [14], MyoPS-Net [15]). General frameworks were evaluated using only LGE images due to performance degradation with multi-sequence input, while cardiac-specific approaches utilized multi-sequence data with their dedicated fusion mechanisms. For fair comparison, all methods follow identical preprocessing and postprocessing protocols as nnU-Net, and are trained for 200 epochs with an initial learning rate of 0.01 using cosine decay scheduler and batch size of 16 on an NVIDIA RTX 4090 GPU. The network was trained using a combination of Dice loss and Cross-Entropy loss, consistent with the standard nnU-Net framework. The results are assessed using Dice coefficient (%) and 95% Hausdorff Distance (HD95, in pixels) across three targets: Myo, ME, and MI. Statistical significance is assessed using the Wilcoxon signed rank test between our method and other approaches.

Results: The experimental results comparing different methods for myocardial pathology segmentation are presented in Table 1. Our CAA-Seg achieves consistent improvements across all metrics, with notable gains in the challenging MI segmentation task where it reaches a Dice score of 51.11% (p<0.05 improvement over the second-best result). The performance advantage is particularly evident in overall metrics, where CAA-Seg obtains a mean Dice of 64.89% and




Fig. 3. Visual comparison of myocardial lesion segmentation results across different methods with myocardium shown in red, myocardial edema in green, and myocardial infarction in blue.

Table 2. Ablation study of CAA-Seg, analyzing the impact of different registration methods, network architectures, and input sequences on segmentation performance.

Settings	Registration			Network		Input		Overall	
	Baseline	MvMM [25]	SSA	nnU-Net	HA-Net	LGE	Multi	Dice (%)	HD95 (pixel)
1	✓			✓		✓		59.63	12.17
2	✓				✓	✓		60.12	11.92
3			✓		✓	✓		62.98	10.83
4			✓		✓	✓	✓	61.51	11.85
5 (proposed)			✓		✓	✓	✓	64.89	9.89

HD of 9.89 pixels, representing statistically significant improvements over previous methods. Interestingly, while UMamba shows competitive performance in Myo and ME segmentation, its effectiveness diminishes in MI detection, highlighting the advantages of our alignment-aware design. In terms of computational efficiency, CAA-Seg remains competitive. The average inference time per case was 0.71s with 1.8GB of GPU memory, compared to nnU-Net (0.52s, 1.1GB) and UMamba (0.62s, 1.3GB). While slightly more intensive due to the explicit alignment, our framework offers substantial performance gains for a modest increase in computational cost.

The qualitative results visualized in Fig. 3 further demonstrate how our method effectively handles complex cases where traditional approaches struggle with multi-sequence alignment and feature fusion. CAA-Seg achieves more accurate target localization, especially in regions with minor pathological changes that are typically challenging to identify.

Ablation Study: We further conduct ablation experiments to validate the effectiveness of each component in our framework. The results in Table 2 demonstrate that our SSA strategy combined with the HA-Net achieves the best performance, significantly outperforming both traditional registration methods and single-modality approaches. Notably, while the multivariate mixture model (MvMM) [25] registration shows competitive results, the integration of SSA with multi-sequence input further improves performance by handling spatial misalignment more effectively.

4 Discussion and Conclusion

In this work, we presented CAA-Seg, a composite alignment-aware framework for myocardial lesion segmentation in multi-sequence CMR images. Our method effectively handles the heterogeneous nature of multi-sequence cardiac imaging data, where misalignments stem from both acquisition protocols and tissue characteristics. The selective slice alignment strategy successfully establishes reliable correspondence between LGE and mapping sequences, while the hierarchical alignment network enables effective feature fusion across modalities with varying intensity distributions and spatial resolutions. Experimental validation on a large-scale dataset of 397 cases demonstrates significant performance improvements across all lesion regions.

We acknowledge that the performance of CAA-Seg may degrade in cases with extreme inter-sequence misalignment beyond the capture range of our alignment module, or for very small and indistinct lesions that are challenging even for expert analysis. Future work could focus on extending the framework to handle more complex clinical scenarios and incorporating temporal information from cine sequences. The composite alignment strategy demonstrated in this work shows promise for application in other multi-sequence medical imaging tasks that require cross-modality feature fusion.

Disclosure of Interests. The authors have no competing interests.

References

1. Muthiah Vaduganathan, George A Mensah, Justine Varieur Turco, Valentin Fuster, and Gregory A Roth. The global burden of cardiovascular diseases and risk: a compass for future health, 2022.
2. Shaohao Rui, Haoyang Su, Jinyi Xiang, Lian-Ming Wu, and Xiaosong Wang. Cardiocot: Hierarchical reasoning for multimodal survival analysis. *arXiv preprint arXiv:2505.19195*, 2025.
3. Yifan Gao, Wei Xia, Dingdu Hu, Wenkui Wang, and Xin Gao. Desam: Decoupled segment anything model for generalizable medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 509–519. Springer, 2024.
4. Yifan Gao, Wei Xia, Wenkui Wang, and Xin Gao. Mba-net: Sam-driven bidirectional aggregation network for ovarian tumor segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 437–447. Springer, 2024.
5. Alexandre Unger, Jérôme Garot, Solenn Toupin, Suzanne Duhamel, Francesca Sanguineti, Thomas Hovasse, Stéphane Champagne, Thierry Untersee, Bernard Chevalier, Mariama Akodad, et al. Prognostic value of cardiac mri late gadolinium enhancement granularity in participants with ischemic cardiomyopathy. *Radiology*, 314(1):e240806, 2025.
6. Xi Liu, Yue Gao, Ying-Kun Guo, Chun-Chao Xia, Rui Shi, Li Jiang, Meng-Ting Shen, Lin-Jun Xie, Wan-Lin Peng, Wen-Lei Qian, et al. Cardiac magnetic resonance t1 mapping for evaluating myocardial fibrosis in patients with type 2 diabetes mellitus: correlation with left ventricular longitudinal diastolic dysfunction. *European Radiology*, 32(11):7647–7656, 2022.

7. William Warnica, Amna Al-Arnawoot, Aleksandra Stanimirovic, Paaladinesh Thavendiranathan, Rachel M Wald, Mini Pakkal, Gauri Rani Karur, Bernd J Wintersperger, Valeria Rac, and Kate Hanneman. Clinical impact of cardiac mri t1 and t2 parametric mapping in patients with suspected cardiomyopathy. *Radiology*, 305(2):319–326, 2022.
8. Yifan Gao, Yin Dai, Fayu Liu, Weibing Chen, and Lifu Shi. An anatomy-aware framework for automatic segmentation of parotid tumor from multimodal mri. *Computers in Biology and Medicine*, 161:107000, 2023.
9. Yin Dai, Yifan Gao, and Fayu Liu. Transmed: Transformers advance multi-modal medical image classification. *Diagnostics*, 11(8):1384, 2021.
10. Qian Yue, Xinzhe Luo, Qing Ye, Lingchao Xu, and Xiahai Zhuang. Cardiac segmentation from lge mri using deep neural network incorporating shape and spatial priors. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part II 22*, pages 559–567. Springer, 2019.
11. Xinhua Yu, Junxin Chen, Bo Fang, Wei Wang, Li-bo Zhang, and Zhihan Lv. Cardiac lge mri segmentation with cross-modality image augmentation and improved u-net. *IEEE journal of biomedical and health informatics*, 27(2):588–597, 2021.
12. Lei Li, Veronika A Zimmer, Julia A Schnabel, and Xiahai Zhuang. Medical image analysis on left atrial lge mri for atrial fibrillation studies: A review. *Medical image analysis*, 77:102360, 2022.
13. Haochuan Jiang, Chengjia Wang, Agisilaos Chartsias, and Sotirios A Tsaftaris. Max-fusion u-net for multi-modal pathology segmentation with attention and dynamic resampling. In *Myocardial Pathology Segmentation Combining Multi-Sequence Cardiac Magnetic Resonance Images: First Challenge, MyoPS 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 4, 2020, Proceedings 1*, pages 68–81. Springer, 2020.
14. Kai-Ni Wang, Xin Yang, Juzheng Miao, Lei Li, Jing Yao, Ping Zhou, Wufeng Xue, Guang-Quan Zhou, Xiahai Zhuang, and Dong Ni. Awsnet: An auto-weighted supervision attention network for myocardial scar and edema segmentation in multi-sequence cardiac magnetic resonance images. *Medical Image Analysis*, 77:102362, 2022.
15. Junyi Qiu, Lei Li, Sihan Wang, Ke Zhang, Yinyin Chen, Shan Yang, and Xiahai Zhuang. Myops-net: Myocardial pathology segmentation with flexible combination of multi-sequence cmr images. *Medical image analysis*, 84:102694, 2023.
16. Xu Xu, Junxin Chen, Dipanwita Thakur, and Duo Hong. Multi-modal disease segmentation with continual learning and adaptive decision fusion. *Information Fusion*, page 102962, 2025.
17. Lei Li, Fuping Wu, Sihan Wang, Xinzhe Luo, Carlos Martín-Isla, Shuwei Zhai, Jianpeng Zhang, Yanfei Liu, Zhen Zhang, Markus J Ankenbrand, et al. Myops: A benchmark of myocardial pathology segmentation combining three-sequence cardiac magnetic resonance images. *Medical Image Analysis*, 87:102808, 2023.
18. Yiwen Ye, Yutong Xie, Jianpeng Zhang, Ziyang Chen, and Yong Xia. Uniseg: A prompt-driven universal segmentation model as well as a strong representation learner. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 508–518. Springer, 2023.
19. Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods*, 18(2):203–211, 2021.

20. Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. Unet++: A nested u-net architecture for medical image segmentation. In *Deep learning in medical image analysis and multimodal learning for clinical decision support: 4th international workshop, DLMIA 2018*, pages 3–11. Springer, 2018.
21. Jieneng Chen, Jieru Mei, Xianhang Li, Yongyi Lu, Qihang Yu, Qingyue Wei, Xiangde Luo, Yutong Xie, Ehsan Adeli, Yan Wang, et al. Transunet: Rethinking the u-net architecture design for medical image segmentation through the lens of transformers. *Medical Image Analysis*, 97:103280, 2024.
22. Ali Hatamizadeh, Vishwesh Nath, Yucheng Tang, Dong Yang, Holger R Roth, and Daguang Xu. Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. In *International MICCAI brainlesion workshop*, pages 272–284. Springer, 2021.
23. Yunhe Gao, Mu Zhou, and Dimitris N Metaxas. Utnet: a hybrid transformer architecture for medical image segmentation. In *Medical image computing and computer assisted intervention–MICCAI 2021: 24th international conference, Strasbourg, France, September 27–October 1, 2021, proceedings, Part III 24*, pages 61–71. Springer, 2021.
24. Jun Ma, Feifei Li, and Bo Wang. U-mamba: Enhancing long-range dependency for biomedical image segmentation. *arXiv preprint arXiv:2401.04722*, 2024.
25. Xiahai Zhuang. Multivariate mixture model for myocardial segmentation combining multi-source images. *IEEE transactions on pattern analysis and machine intelligence*, 41(12):2933–2946, 2018.