

# Robust Deep Learning for Myocardial Scar Segmentation in Cardiac MRI with Noisy Labels

Aida Moafi<sup>1,2</sup>, Danial Moafi<sup>3</sup>, Evgeny M. Mirkes<sup>4</sup>, Gerry P. McCann<sup>1,2</sup>, Abbas S. Alatrany<sup>1,2</sup>, Jayanth Ranjit Arnold<sup>1,2\*</sup>, and Mostafa Mehdipour Ghazi<sup>5\*</sup>

<sup>1</sup> Department of Cardiovascular Sciences, University of Leicester, Leicester, UK

<sup>2</sup> NIHR Leicester Biomedical Research Centre, Glenfield Hospital, Leicester, UK

<sup>3</sup> Department of Information Engineering and Mathematics, University of Siena, Italy

<sup>4</sup> Department of Mathematics, University of Leicester, Leicester, UK

<sup>5</sup> Department of Computer Science, University of Copenhagen, Copenhagen, DK

**Abstract.** The accurate segmentation of myocardial scars from cardiac MRI is essential for clinical assessment and treatment planning. In this study, we propose a robust deep-learning pipeline for fully automated myocardial scar detection and segmentation by fine-tuning state-of-the-art models. The method explicitly addresses challenges of label noise from semi-automatic annotations, data heterogeneity, and class imbalance through the use of Kullback-Leibler loss and extensive data augmentation. We evaluate the model’s performance on both acute and chronic cases and demonstrate its ability to produce accurate and smooth segmentations despite noisy labels. In particular, our approach outperforms state-of-the-art models like nnU-Net and shows strong generalizability in an out-of-distribution test set, highlighting its robustness across various imaging conditions and clinical tasks. These results establish a reliable foundation for automated myocardial scar quantification and support the broader clinical adoption of deep learning in cardiac imaging. The codes are available at: <https://github.com/Danialmoa/YoloSAM>.

**Keywords:** Deep learning · segmentation · myocardial scar · noisy label.

## 1 Introduction

Cardiovascular disease remains a global health concern, responsible for nearly 32% of all deaths worldwide [1]. Among its manifestations, myocardial infarction (MI) is a major cause of morbidity and mortality, often leaving patients with permanent myocardial scarring. Scar extent is a critical prognostic factor, with even a 5% increase in scar size raising mortality risk by 19% [2]. The scar burden also impacts revascularization outcomes and predicts the likelihood of ventricular arrhythmias and heart failure [3,4]. Moreover, scar transmural extent is essential for assessing recovery potential, as regions with less than 50% transmural scar are more likely to regain function [5,6,7]. Given its clinical importance, accurate scar quantification is crucial for personalized treatment and risk stratification.

---

\* Co-senior corresponding authors: jra14@leicester.ac.uk and ghazi@di.ku.dk

Late gadolinium enhancement cardiovascular magnetic resonance imaging (LGE-CMR) is the gold standard for noninvasive scar assessment, offering high-resolution, contrast-enhanced visualization of infarcted tissue. However, current segmentation methods, including manual contouring and semi-automatic thresholding, are time-consuming, labor-intensive, and prone to inter-observer variability [8,9,10]. Manual annotation can take up to 25 minutes per scan [8], limiting scalability in clinical practice. Semi-automatic methods like the signal threshold-to-reference mean (STRM) and full width at half-maximum (FWHM) techniques segment scar tissue based on intensity thresholds but are sensitive to variations in signal-to-noise ratio, region-of-interest selection, and infarct signal characteristics [11,12,13]. These limitations often lead to inconsistent scar delineation.

Deep learning offers promising solutions for automated scar segmentation [14]. Approaches based on 2D and 3D U-Net architectures [15] have demonstrated feasibility, as seen in the EMIDEC Challenge [16], but challenges related to ground-truth variability and imaging protocol differences hinder clinical adoption [17]. A systematic review reported Dice similarity scores up to 0.63 for these methods [18], while manual annotation itself shows limited consistency, with intra- and inter-observer Dice indices of 0.76 and 0.69, respectively [19]. However, variability in LGE sequences and signal intensity further complicates segmentation. Magnitude and phase-sensitive inversion recovery (PSIR) sequences differ in spatial resolution and contrast properties, with PSIR offering better contrast but introducing inconsistencies in absolute intensity values [20,21]. Differences in scanner vendors, acquisition parameters, and post-processing techniques add to the heterogeneity, complicating model generalization. Moreover, the small size and sporadic appearance of scars relative to the overall image can lead to model overfitting on background pixels. The quality of ground-truth labels significantly affects performance, as models often train on manual or semi-automatic annotations with inherent biases. Threshold-based methods like STRM and FWHM, while improving reproducibility, may not fully capture scar extent in cases of diffuse fibrosis, microvascular obstruction (MVO), or signal heterogeneity [22]. These challenges underscore the need for robust, generalizable deep learning models for automated myocardial scar segmentation.

We develop a deep-learning pipeline for fully automated segmentation of myocardial scars by fine-tuning state-of-the-art models [23,24]. Our contributions are six-fold: (1) We propose an end-to-end pipeline achieving high-accuracy segmentation, producing smoother results despite being trained on noisy labels. (2) We introduce the Kullback-Leibler (KL) loss to mitigate the impact of imperfect manual or semi-automatic annotations. (3) We highlight the limitations of overlap-based metrics for evaluating small lesions and propose perimeter similarity as a more reliable size-based metric. (4) We use data augmentation and detection models to guide the segmentation process, addressing class imbalance and label noise. (5) We validate our approach on a heterogeneous, multi-cohort dataset, demonstrating its robustness across diverse imaging conditions. (6) We evaluate segmentation accuracy on an out-of-distribution test set, including MI scars, demonstrating generalizability to unseen data with a different task.

## 2 Methods

### 2.1 Data Acquisition and Annotation

This study leverages a multi-cohort dataset comprising four independent studies on chronic reperfused ischemic scars, summarized in Table 1. The datasets were collected from multiple centers, encompassing diverse MRI acquisition protocols with variations in imaging sequences (e.g., magnitude and phase), scanner manufacturers (e.g., Siemens and Philips), magnetic field strengths (e.g., 1.5T and 3T), and spatial resolutions (e.g., from  $1 \times 1 \times 6 \text{ mm}^3$  to  $2 \times 2 \times 10 \text{ mm}^3$ ). These variations introduce heterogeneity in image quality, intensity, and contrast, posing challenges for consistent segmentation performance. The PSIR images were used when standard LGE images showed inconsistencies, as they are less dependent on inversion time selection and operator adjustments during acquisition. This approach enhanced the reliability of myocardial structure identification.

A total of 348 patients were selected, with LGE scar segmentation performed manually using commercially available Circle CVI software. To ensure high-quality annotations, all segmentations were reviewed by two expert cardiac MRI specialists. For reproducibility, scar delineation was performed using the widely used FWHM method. As a threshold-based technique, the FWHM method has inherent limitations, including sensitivity to intensity variations and partial volume effects, affecting segmentation accuracy. The labeling process involved three key steps: (1) Myocardial contouring, where the epicardial and endocardial borders were manually traced to define the myocardial region of interest. (2) Scar core identification, in which an ROI was placed within the region of highest signal intensity in the myocardium, corresponding to the scar core. (3) Refinement and exclusion, where a final revision step was conducted to exclude potential imaging artifacts and ensure consistency across cases.

### 2.2 Preprocessing and Augmentation

The annotated slices were preprocessed, accounting for DICOM pixel spacing, and scaled to  $256 \times 256$  pixels. We implemented different augmentation tech-

Table 1: Summary of the multi-cohort dataset [25] used in the study.

	# Images			# Subjects			Age		Sex
	Train	Valid	Test	Train	Valid	Test	Mean	$\pm$ SD	
Cohort 1	450	74	91	49	9	11	59.44	$\pm 9.85$	25%
Cohort 2	493	82	88	59	12	13	45.41	$\pm 21.20$	65%
Cohort 3	304	55	63	38	8	9			
Cohort 4	471	54	41	49	8	6			
Cohort 5	81	32	18	12	4	2			
Cohort 6	93	22	34	10	3	4			
Cohort 7	129	27	26	15	3	3			
Cohort 8	117	37	49	12	4	5			

niques to enhance the model’s robustness to anatomical variability and imaging artifacts. In addition to standard augmentations with spatial transformations like flipping, rotation, scaling, and deformation, we integrated medical image-specific techniques to simulate real-world variations in MRIs. These augmentations included gamma correction, adaptive histogram equalization, brightness adjustment, and Rician noise addition to mimic MRI-specific noise characteristics. Bounding box jittering was also introduced to simulate variability in manual annotations and enhance detection robustness.

To ensure a robust and unbiased evaluation, we partitioned the datasets at the patient level, separating data from each cohort into training (70%), validation (15%), and test (15%) subsets. This prevented data leakage and ensured the subsets were representative of the different cohorts. Given the inherent class imbalance in scar segmentation, we stratified the subsets based on the total number of scar pixels per patient. This approach maintained the full distribution of scar burden across the subsets, preventing the underrepresentation of critical cases. Each LGE-MRI scan was then processed at the slice level, with individual slices serving as network input. Since many slices in the LGE series did not contain the myocardium, this approach helped mitigate background imbalance and ensured that only relevant regions were considered during model development.

### 2.3 Detection and Segmentation with Noisy Labels

Fig. 1 illustrates the fully automated pipeline developed in this study for detecting and segmenting myocardial scars with noisy ground truth labels. The pipeline employs the You Only Look Once (YOLO) architecture [23] with default configurations for scar bounding box detection, serving as prompts for the Segment Anything Model (SAM) [24]. This approach enables automated scar localization and provides a streamlined segmentation framework, eliminating the need for manual intervention. The detection step is crucial to focus on the tiny and fragmented scar regions. This method effectively integrates anatomical knowledge into deep learning, improving segmentation performance in challenging imaging scenarios and clinical applications where traditional methods struggle.

Despite reproducibility, the FWHM semi-automatic labels are sharp and often misaligned with true anatomical boundaries due to thresholding artifacts and partial volume effects. Hence, we consider them noisy in the sense of structured boundary uncertainty, rather than incorrect. To address the issue with noisy labels, we propose incorporating the KL divergence as an additional loss alongside the Dice loss and cross-entropy. We apply Gaussian smoothing with  $\sigma = 2$  to the target mask to obtain soft labels, capturing uncertainty in transitional boundary regions. The model predictions are mapped to  $[0, 1]$  using a sigmoid activation, followed by a log-softmax to obtain the predicted distribution. The prediction and the soft target scores are then normalized as probabilities, and KL divergence is applied between the probability distributions of predicted and soft target, penalizing overconfident predictions near ambiguous regions. The three losses are combined in a weighted manner, with weighting factors of 0.2, 0.2, and 0.6, respectively, determined by a grid search.




Fig. 1: Overview of the proposed fully automated pipeline for myocardial scar detection and segmentation. The YOLO model is fine-tuned to generate automatic bounding box prompts, guiding the fine-tuned SAM for precise segmentation of the localized scar. KL loss and extensive data augmentation strategies are employed to mitigate data variations and handle noisy label training.

## 2.4 Evaluation Metrics

We evaluate segmentation performance using Hausdorff distance (HD), Dice similarity coefficient (DSC), area similarity (AS), and perimeter similarity (PS). Since DSC is an overlapping metric sensitive to small errors and may not fully capture performance on tiny, noisy (semi-automatic) scar labels, AS and PS provide complementary shape-based assessments. AS measures the pixel count similarity between ground truth ( $T$ ) and prediction ( $Y$ ) as  $1 - ||Y| - |T|| / (|Y| + |T|)$ , while PS applies the same formula to boundary pixel counts.

## 3 Results

**Baseline Comparison.** We first evaluated the nnU-Net [26] for segmentation, with results summarized in Table 2. Next, we designed an experiment using ground-truth bounding boxes as prompts and fine-tuned SAM and modified DeepLabV3+ [27] with a ResNet-50 backbone and additional input layers for handling bounding boxes. This approach significantly improved performance over the baseline nnU-Net, as shown in the table. These findings motivated the automation of SAM with YOLO-based bounding-box detection.

**YOLO-SAM.** Table 3 presents the results of the proposed method based on the fine-tuned YOLO-SAM model. Although this fully automated pipeline intro-

Table 2: Test segmentation accuracy (mean $\pm$ SD) of the different models across all cohorts. Note that Box-SAM and Box-DeepLabV3+ refer to the models fine-tuned using ground-truth bounding boxes as input.

	DSC	HD	AS	PS
nnU-Net	0.579 $\pm$ 0.338	11.786 $\pm$ 16.004	0.702 $\pm$ 0.345	0.748 $\pm$ 0.341
Box-SAM	0.653 $\pm$ 0.135	7.807 $\pm$ 7.922	0.793 $\pm$ 0.133	0.918 $\pm$ 0.080
Box-DeepLabV3+	0.621 $\pm$ 0.169	8.029 $\pm$ 8.818	0.838 $\pm$ 0.138	0.905 $\pm$ 0.095

Table 3: Test segmentation accuracy (mean $\pm$ SD) of the proposed method (YOLO-SAM) across different cohorts.

	DSC	HD	AS	PS
Cohort 1	0.579 $\pm$ 0.338	11.786 $\pm$ 16.004	0.702 $\pm$ 0.345	0.748 $\pm$ 0.341
Cohort 2	0.636 $\pm$ 0.258	9.362 $\pm$ 8.469	0.814 $\pm$ 0.243	0.865 $\pm$ 0.232
Cohort 3	0.525 $\pm$ 0.280	15.993 $\pm$ 18.012	0.753 $\pm$ 0.301	0.782 $\pm$ 0.299
Cohort 4	0.579 $\pm$ 0.406	8.931 $\pm$ 13.915	0.628 $\pm$ 0.412	0.670 $\pm$ 0.418
Cohort 5	0.596 $\pm$ 0.308	15.741 $\pm$ 15.535	0.779 $\pm$ 0.317	0.805 $\pm$ 0.271
Cohort 6	0.516 $\pm$ 0.373	15.388 $\pm$ 16.646	0.636 $\pm$ 0.373	0.680 $\pm$ 0.373
Cohort 7	0.666 $\pm$ 0.377	13.217 $\pm$ 24.006	0.767 $\pm$ 0.368	0.774 $\pm$ 0.357
Cohort 8	0.628 $\pm$ 0.361	7.599 $\pm$ 11.248	0.734 $\pm$ 0.347	0.778 $\pm$ 0.339
Total	0.601 $\pm$ 0.330	10.728 $\pm$ 14.217	0.753 $\pm$ 0.318	0.797 $\pm$ 0.311

duces a slight performance drop, it eliminates the need for human intervention, unlike the semi-automatic approach that requires manual bounding box input.

Fig. 2 shows two segmented sample images using the proposed deep learning model compared to FWHM-based ground-truth labels. As seen, the automatic segmentation aligns with the ground-truth labels. However, a slight spatial shift in the predicted labels resulted in a very low Dice score of 0.15 due to the small overlap in the tiny scar regions in the first sample. This suggests that while the model captures the general extent of the scar, with a high perimeter similarity of 0.85, it struggles with precise localization. However, a closer inspection of the predicted labels reveals that they are often placed in more anatomically accurate locations. This observation highlights not only the limitations of the noisy, semi-automatically generated ground-truth labels but also the shortcomings of overlap-based metrics like the Dice score in evaluating segmentation accuracy. In contrast, size-based measures provide a more robust evaluation by assessing similarity independent of exact spatial alignment. Furthermore, the automatic predictions produce smoother and more coherent label regions than semi-automatic annotations, which often appear fragmented into small, scattered components.




Fig. 2: Results of the proposed deep learning pipeline for myocardial scar segmentation. Each row shows a sample with three columns: (1) YOLO-based bounding box detection (red), (2) SAM-based automatic segmentation (red), and (3) semi-automatic FWHM-based ground truth (green). A slight spatial shift in the first predictions leads to a low Dice score of 0.15 despite a high perimeter similarity of 0.85, reflecting the model’s accurate scar extent but imprecise ground-truth localization. The second prediction is smooth, achieving a perimeter similarity of 0.98, though the fragmented ground truth results in a lower Dice score of 0.52.

Despite the Dice score of 0.52 for the second sample, influenced by the discontinuity and small size of the ground-truth labels, the perimeter similarity remains high at 0.98, better reflecting the quality of the predicted segmentation.

We also emphasize the efficiency and clinical impact of predicted segmentations in scar detection and classification. To this end, we calculated the scar mass in the semi-automatic ground truth and the predicted segmentation in the test set. The estimated mass of the semi-automatic (ground truth) scars was  $12.510 \pm 8.784$  g, compared to  $11.337 \pm 8.201$  g for the fully-automatic (predicted) scars, indicating a very close estimation. The Wilcoxon rank sum test did not reveal a statistically significant difference between the estimated scar masses from the semi-automatic and predicted labels ( $p = 0.42$ ).

**Acute MI Generalization.** To assess generalization to different clinical domains, we evaluated the trained model in a multicenter acute MI cohort dataset with 50 patients and 406 images. The goal was to compare scar pattern differences between acute and chronic cases and to assess whether the predicted segmentations capture the expected clinical distinctions. Since the spatial characteristics of scars, including size, shape, and continuity, differ between acute and chronic cases, we extracted several morphological characteristics from the predicted scar regions. These features included scar size (number of pixels in the scar area), the number of separate lesions (number of connected components),



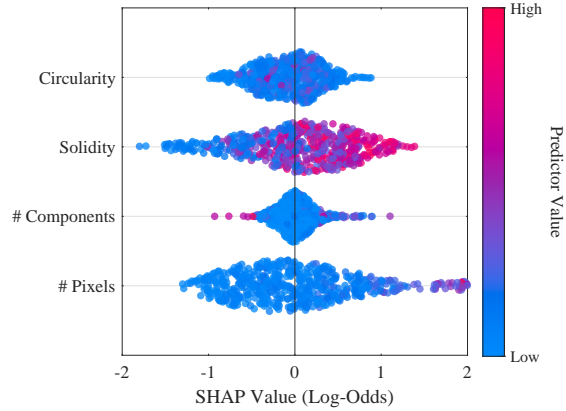


Fig. 3: Shapley summary plot for detecting acute MI cases based on morphological features extracted from the predicted scar segmentations.

compactness (solidity) defined as the ratio of the lesion area to its convex hull area, and elongation (circularity) calculated as  $4\pi \cdot \text{area} / \text{perimeter}^2$ .

To distinguish between acute and chronic scars, we trained an XGBoost classifier using these features extracted from the predicted segmentations. We then applied the Shapley method to explain the importance of each feature in the classification process. Fig. 3 shows the Shapley summary plot for detecting acute cases based on predicted scars, highlighting the greater burden and solidity of the infarct observed in acute cases. A statistical analysis of the extracted features further confirmed their significance in differentiating acute from chronic scars, with  $p$ -values of  $4.8 \times 10^{-7}$ ,  $1.4 \times 10^{-7}$ ,  $1.7 \times 10^{-4}$ , and 0.03 for scar size, solidity, circularity, and the number of connected components, respectively.

Using shape-based analysis of segmented scars offers the potential to develop diagnostic classifiers and predict patient risk and treatment outcomes. Such insights could enable cardiologists to rapidly distinguish between infarct types, which helps with both diagnosis and therapeutic decision-making. In particular, acute infarcts often exhibit regions of edema that may recover over time, whereas chronic infarcts typically represent irreversible damage. Identifying these distinct patterns could guide more effective treatment strategies. This experiment also underscored the advantage of training the model exclusively in chronic cases with more homogeneous and well-defined scar patterns, thus minimizing bias from the fragmented and less distinct appearance often seen in acute scars.

## 4 Conclusion

We developed a robust deep-learning pipeline for fully automated myocardial scar segmentation by fine-tuning state-of-the-art SAM and YOLO models. Our approach tackled key challenges in medical image analysis, including label noise



from semi-automatic annotations, data heterogeneity, and class imbalance, using KL loss and extensive data augmentation strategies. Through a comprehensive evaluation of both acute and chronic cases, the model demonstrated its ability to produce smooth and accurate segmentations despite the presence of noisy labels.

The integration of a detection model to guide segmentation significantly improved accuracy by focusing on the scar region. Furthermore, the combination of diverse training data and systematic data augmentation enabled the model to generalize well to unseen test data with varying anatomical and pathological characteristics, as well as different acquisition parameters, including higher-resolution scans (e.g.,  $0.6 \times 0.6 \times 8$  mm<sup>3</sup>). Finally, the classification experiment using segmented acute MI scars provided a meaningful validation of the model's ability to distinguish between acute and chronic cases, even when trained solely on chronic data. These findings underscore the potential of our approach for broader clinical application in automated cardiac imaging analysis.

## Acknowledgments

A.M. is funded by a University of Leicester Future 100 Studentship. M.M.G. is supported by the Pioneer Centre for AI (DNRF grant P1). J.R.A. is supported by a NIHR Clinician Scientist Award (CS-2018-18-ST2-007). G.P.M. is supported by a NIHR Research Professorship Award (RP-2017-08-ST2-007).

## Disclosure of Interests

The authors have no competing interests in the paper.

## References

1. Lindstrom, M., DeCleene, N., Dorsey, H., Fuster, V., Johnson, C.O., LeGrand, K.E., Mensah, G.A., Razo, C., Stark, B., Varieur Turco, J., et al.: Global burden of cardiovascular diseases and risks collaboration, 1990-2021. *Journal of the American College of Cardiology* **80**(25) (2022) 2372–2425
2. Stone, G.W., Selker, H.P., Thiele, H., Patel, M.R., Udelson, J.E., Ohman, E.M., Maehara, A., Eitel, I., Granger, C.B., Jenkins, P.L., et al.: Relationship between infarct size and outcomes following primary PCI: patient-level analysis from 10 randomized trials. *Journal of the American College of Cardiology* **67**(14) (2016)
3. Kancharla, K., Weissman, G., Elagha, A.A., Kancherla, K., Samineni, S., Hill, P.C., Boyce, S., Fuisz, A.R.: Scar quantification by cardiovascular magnetic resonance as an independent predictor of long-term survival in patients with ischemic heart failure treated by coronary artery bypass graft surgery. *Journal of Cardiovascular Magnetic Resonance* **18**(1) (2016) 45
4. Disertori, M., Rigoni, M., Pace, N., Casolo, G., Masè, M., Gonzini, L., Lucci, D., Nollo, G., Ravelli, F.: Myocardial fibrosis assessment by LGE is a powerful predictor of ventricular tachyarrhythmias in ischemic and nonischemic LV dysfunction: a meta-analysis. *JACC: Cardiovascular Imaging* **9**(9) (2016) 1046–1055

5. Choi, K.M., Kim, R.J., Gubernikoff, G., Vargas, J.D., Parker, M., Judd, R.M.: Transmural extent of acute myocardial infarction predicts long-term improvement in contractile function. *Circulation* **104**(10) (2001) 1101–1107
6. Enein, F.A., Allaaboun, S., Khayyat, S., Andijani, M., Alkhuzai, M.M., Aljunied, A.A., Al Adhrai Sr, M., Khayyat, S.F., Andijani, M.A.: Association between myocardial scar burden and left ventricular ejection fraction in ischemic cardiomyopathy. *Cureus* **12**(12) (2020)
7. Holtackers, R.J., Emrich, T., Botnar, R.M., Kooi, M.E., Wildberger, J.E., Kreitner, K.F.: Late gadolinium enhancement cardiac magnetic resonance imaging: from basic concepts to emerging methods. In: *RöFo-Fortschritte auf dem Gebiet der Röntgenstrahlen und der bildgebenden Verfahren*. Volume 194., Georg Thieme Verlag KG (2022) 491–504
8. Mewton, N., Liu, C.Y., Croisille, P., Bluemke, D., Lima, J.A.: Assessment of myocardial fibrosis with cardiovascular magnetic resonance. *Journal of the American College of Cardiology* **57**(8) (2011) 891–903
9. Khan, J.N., Nazir, S.A., Horsfield, M.A., Singh, A., Kanagala, P., Greenwood, J.P., Gershlick, A.H., McCann, G.P.: Comparison of semi-automated methods to quantify infarct size and area at risk by cardiovascular magnetic resonance imaging at 1.5 T and 3.0 T field strengths. *BMC Research Notes* **8** (2015) 1–15
10. Kramer, C.M., Barkhausen, J., Bucciarelli-Ducci, C., Flamm, S.D., Kim, R.J., Nagel, E.: Standardized cardiovascular magnetic resonance imaging (CMR) protocols: 2020 update. *Journal of Cardiovascular Magnetic Resonance* **22**(1) (2020)
11. Klem, I., Heiberg, E., Van Assche, L., Parker, M.A., Kim, H.W., Grizzard, J.D., Arheden, H., Kim, R.J.: Sources of variability in quantification of cardiovascular magnetic resonance infarct size-reproducibility among three core laboratories. *Journal of Cardiovascular Magnetic Resonance* **19**(1) (2016) 62
12. Heiberg, E., Engblom, H., Carlsson, M., Erlinge, D., Atar, D., Aletras, A.H., Arheden, H.: Infarct quantification with cardiovascular magnetic resonance using "standard deviation from remote" is unreliable: validation in multi-centre multi-vendor data. *Journal of Cardiovascular Magnetic Resonance* **24**(1) (2022) 53
13. Flett, A.S., Hasleton, J., Cook, C., Hausenloy, D., Quarta, G., Ariti, C., Muthurangu, V., Moon, J.C.: Evaluation of techniques for the quantification of myocardial scar of differing etiology using cardiac magnetic resonance. *JACC: Cardiovascular Imaging* **4**(2) (2011) 150–156
14. Chen, C., Qin, C., Qiu, H., Tarroni, G., Duan, J., Bai, W., Rueckert, D.: Deep learning for cardiac image segmentation: a review. *Frontiers in Cardiovascular Medicine* **7** (2020) 25
15. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: *18th International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer (2015) 234–241
16. Lalande, A., Chen, Z., Pommier, T., Decourselle, T., Qayyum, A., Salomon, M., Ginjac, D., Skandarani, Y., Boucher, A., Brahim, K., et al.: Deep learning methods for automatic evaluation of delayed enhancement-MRI. the results of the EMIDEC challenge. *Medical Image Analysis* **79** (2022) 102428
17. Hoh, T., Margolis, I., Weine, J., Joyce, T., Manka, R., Weisskopf, M., Cesarovic, N., Fuetterer, M., Kozerke, S.: Impact of late gadolinium enhancement image acquisition resolution on neural network based automatic scar segmentation. *Journal of Cardiovascular Magnetic Resonance* **26**(1) (2024) 101031
18. Jathanna, N., Podlasek, A., Sokol, A., Auer, D., Chen, X., Jamil-Copley, S.: Diagnostic utility of artificial intelligence for left ventricular scar identification using

- cardiac magnetic resonance imaging—a systematic review. *Cardiovascular Digital Health Journal* **2**(6) (2021) S21–S29
19. Lalande, A., Chen, Z., Decourselle, T., Qayyum, A., Pommier, T., Lorgis, L., de La Rosa, E., Cochet, A., Cottin, Y., Gin hac, D., et al.: EMIDEC: a database usable for the automatic evaluation of myocardial infarction from delayed-enhancement cardiac MRI. *Data* **5**(4) (2020) 89
  20. Stirrat, J., Joncas, S.X., Salerno, M., Drangova, M., White, J.: Influence of phase correction of late gadolinium enhancement images on scar signal quantification in patients with ischemic and non-ischemic cardiomyopathy. *Journal of Cardiovascular Magnetic Resonance* **17** (2015) 1–9
  21. Schulz-Menger, J., Bluemke, D.A., Bremerich, J., Flamm, S.D., Fogel, M.A., Friedrich, M.G., Kim, R.J., von Knobelsdorff-Brenkenhoff, F., Kramer, C.M., Pennell, D.J., et al.: Standardized image interpretation and post-processing in cardiovascular magnetic resonance-2020 update: Society for cardiovascular magnetic resonance (SCMR): Board of trustees task force on standardized post-processing. *Journal of Cardiovascular Magnetic Resonance* **22**(1) (2020) 19
  22. Hennemuth, A., Seeger, A., Friman, O., Miller, S., Klumpp, B., Oeltze, S., Peitgen, H.O.: A comprehensive approach to the analysis of contrast enhanced cardiacMR images. *IEEE Transactions on Medical Imaging* **27**(11) (2008) 1592–1610
  23. Tian, Y., Ye, Q., Doermann, D.: YOLOv12: Attention-centric real-time object detectors. *arXiv preprint arXiv:2502.12524* (2025)
  24. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: *Proceedings of the IEEE/CVF International Conference on Computer Cision*. (2023) 4015–4026
  25. Author, A.: Anonymous title. *Anonymous Journal* (Anonymous Year)
  26. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* **18**(2) (2021) 203–211
  27. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: *Proceedings of the European Conference on Computer Vision*. (2018) 801–818