Echo-SCAR: Temporal Feature Learning for Myocardial Scar Detection from Routine Echocardiography

Mansoor Shehzad[1], Arion Yu[2], Eman Nayaz[3], Mohammad Alkhaleefah[3], Sadeer Al-Kindi[3,1], Guha Balakrishnan[1], and Ashutosh Sabharwal[1]

[1] Rice University, Houston TX 77005, USA
[2] Texas A&M University, College Station TX 77843, USA
[3] Houston Methodist Hospital, Houston TX 77030, USA
mansoorshehzad@rice.edu

Abstract. Myocardial scar detection is a clinically important task with implications for prognosis, arrhythmic risk, and therapeutic decision making. The noninvasive reference standard for scar identification is cardiac magnetic resonance (CMR) with late gadolinium enhancement (LGE), but its use is constrained by cost, limited availability, and logistical barriers. In this study, we develop a CMR-supervised multi-stage automated framework to predict segment-level myocardial scar from routine echocardiography. Using retrospective patient data from Houston Methodist Hospital, we analyzed echocardiogram scans from 1,921 patients with paired CMR data within one year. Myocardial regions were segmented from apical two-, three-, and four-chamber views, covering the full 17-segment AHA model, using view-specific nnU-Net architectures trained on manual annotations across the cardiac cycle. We then compared two feature representations for per-segment scar detection: (1) temporal radiomic features extracted from segmented myocardium and (2) features derived from the encoder bottleneck of trained nnU-Net models. Using class-weighted random forest classifiers, encoder features achieved a mean AUC of 0.822 across 18 view–segment combinations, significantly outperforming radiomics (mean AUC 0.793; $p < 0.05$ in 8 of 18 segments). Combining both feature sets yielded the highest AUC (reaching up to 0.858), although this improvement did not generalize consistently across views. These findings demonstrate that routine echocardiography, enhanced through CMR-supervised feature learning, can serve as a scalable and effective tool for myocardial scar screening.

Keywords: Echocardiography · Radiomics · Myocardial Scar · Feature Extraction · Tissue Characterization

1 Introduction

Myocardial scarring, arising from ischemic or non-ischemic injury, is a critical determinant of patient prognosis and arrhythmic risk [12, 14], and can inform

decisions regarding revascularization, device implantation, and medical management [5]. Cardiac magnetic resonance (CMR) with late gadolinium enhancement (LGE) is the reference standard for non-invasive scar detection, offering high spatial resolution and direct tissue characterization [12]. However, CMR remains limited by high cost ($1,000–$3,000 per study), restricted availability, lengthy acquisition protocols, and specialized staff training requirements, making CMR impractical as a population-level screening tool [13].

Echocardiography is the most widely performed cardiac imaging modality, providing real-time functional assessment at the point of care [17]. Standard echocardiographic parameters, e.g., LV ejection fraction, wall motion score index, and global longitudinal strain, provide valuable functional information but have limited sensitivity for detecting myocardial scar. LV ejection fraction demonstrates poor discriminative ability for segmental scar (AUC $\approx$ 0.54) [10], global longitudinal strain achieves moderate accuracy (AUC $\approx$ 0.70–0.74) for detecting replacement fibrosis [15], and wall motion score index, while correlated with infarct size [6], lacks the granularity for per-segment tissue characterization. A fundamental gap exists between the functional information available from echocardiography and the tissue-level characterization provided by CMR.

Recent advances demonstrate that radiomics features derived from echocardiographic images contain sub-visual prognostic and structural information [10, 4], useful, for example, for detecting changes associated with acute myocardial infarction (MI) [10] and predicting major adverse cardiovascular events (MACE). Critically, these studies demonstrate that echocardiographic radiomics can capture sub-visual tissue changes—suggesting that the same approach may be extensible to the detection of chronic myocardial scar. However, neither acute MI detection nor MACE prediction directly addresses per-segment classification of established replacement fibrosis, and no prior work has systematically attempted this task across diverse etiologies at the segment level. Small sample sizes, limited view coverage, and a lack of analysis on the benefits of radiomics versus learned features to capture subtle tissue change further constrain these studies.

In this work, we present the largest echocardiographic study for myocardial scar detection to date and systematically evaluate which feature representation best captures scar-related tissue changes from echocardiographic video. We first curated retrospective echo–CMR data from Houston Methodist Hospital, consisting of 1,921 patients with per-segment hyperenhancement labels spanning the full AHA 17-segment model. Using this data, we trained a CMR-supervised framework that predicts tissue characterization from routine echocardiography. We then performed a rigorous head-to-head comparison of two complementary feature representations: temporal radiomics and encoder bottleneck embeddings from an nnU-Net [9] model trained for myocardial boundary segmentation. Finally, we performed evaluations of view-level scar screening performance, finding that the segmentation-tuned features consistently outperform radiomics, but that fusing both feature representations does have benefits on a view-dependent basis. Our study is the first to systematically compare feature representations for segment-level scar detection from echocardiography at this scale.

2 Paired Echocardiography–CMR Dataset

We collected retrospective patient data from Houston Methodist Hospital, consisting of 1,921 patients who underwent both transthoracic echocardiography and CMR with LGE within 1 year. Apical four-chamber (A4C), three-chamber (A3C), and two-chamber (A2C) views were identified using automated view-classification models (EchoCV [19], EchoPrime [18]). Inclusion criteria required availability of at least one classified apical view with paired CMR LGE imaging. In total, this resulted in 5,870 A4C, 2,159 A3C, and 1,972 A2C scans. The study was approved by the institutional review board with a waiver of informed consent.

We assessed myocardial hyperenhancement from the CMR LGE images using the AHA 17-segment model [3]. Hyperenhancement was scored on the ordinal transmural extent scale (0–5) from clinical reads; we binarized scar-present (≥ 1) versus absent (0), providing per-segment binary ground truth labels. Each view visualizes a specific subset of AHA segments: A4C covers the septal and lateral walls (segments 3, 6, 9, 12, 14, 16), A3C the anteroseptal and inferolateral walls (segments 2, 5, 8, 11, 13, 15), and A2C the anterior and inferior walls (segments 1, 4, 7, 10, 13, 15). We excluded Segment 17 (apical cap) because reliable feature extraction from this region is compromised by near-field clutter artifact and apical foreshortening inherent to transthoracic imaging, yielding 18 unique view–segment combinations.

3 Methods

We investigate two complementary feature representations for CMR-supervised myocardial scar detection from routine echocardiography: temporal radiomics and encoder bottleneck embeddings extracted from segmentation-trained nnU-Net models. The core hypothesis motivating the encoder approach is that a model trained to delineate myocardial boundaries must implicitly encode tissue-level acoustic properties—information that may be predictive of replacement fibrosis. To enable a controlled comparison, we apply identical temporal summarization to both feature types. Figure 1 provides an overview of the full pipeline, from echocardiographic video input through segmentation, feature extraction, temporal summarization, and per-segment random forest classification against CMR ground truth.

3.1 Automated Myocardial Segmentation

Three view-specific nnU-Net [9] models segmented the left ventricular myocardial wall, each trained on manually annotated echocardiographic frames. Each view had an 8-stage PlainConvUNet encoder with a 512-dimensional bottleneck.

We validated segmentation quality on held-out annotated frames, achieving Dice coefficients of 0.96 (A4C), 0.94 (A3C), and 0.95 (A2C). We mapped segmented regions to the AHA 17-segment model following standard anatomical conventions [3].

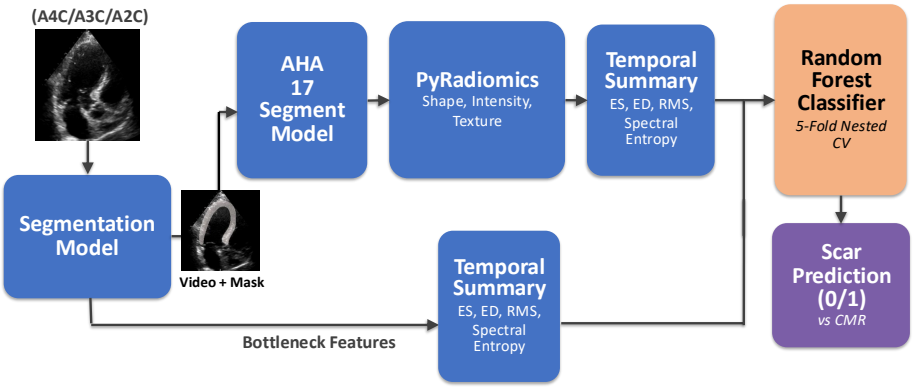

Fig. 1: **Overview of the proposed framework for myocardial scar prediction.** We segment echocardiographic videos using view-specific nnU-Net models. We extract two feature representations from the segmented myocardium: radiomics (PyRadiomics) and encoder bottleneck features (global average pooled). We apply identical temporal summaries to both, enabling direct comparison. Random forest classifiers predict per-segment hyperenhancement, validated against CMR ground truth.

3.2 Temporal Feature Extraction

Radiomic Features. For each patient, we extracted radiomic features from segmented myocardial regions mapped to the AHA 17-segment model using PyRadiomics across all frames of the cardiac cycle [8]. Features included shape descriptors (area, perimeter, sphericity, elongation), first-order intensity statistics (mean, standard deviation, skewness, kurtosis, energy, entropy), and texture descriptors from gray-level co-occurrence (GLCM), run-length (GLRLM), dependence (GLDM), size-zone (GLSZM), and neighboring gray-tone difference (NGTDM) matrices. Each frame yielded 103 features per segment.

Encoder Bottleneck Features. We repurposed the encoder branch of each trained nnU-Net as a feature extractor. For each echocardiographic frame, we passed the image through the encoder, and reduced the final bottleneck feature map ($1 \times 512 \times h \times w$) to a 512-dimensional vector via global average pooling. This approach leverages cardiac-specific representations learned during segmentation training without any scar-specific supervision—the encoder never encountered hyperenhancement labels. Each frame yielded 512 encoder features.

Temporal Summary Statistics. For each feature (radiomic or encoder channel), we computed four temporal summaries from the per-frame time series:

– **ES/ED values**: Feature at end-systolic and end-diastolic frames, capturing extremes of the cardiac cycle.

- **Root mean square (RMS)**: RMS $= \sqrt{\frac{1}{T}\sum_{t=1}^{T} x_t^2}$, quantifying overall feature magnitude across the cycle.
- **Spectral entropy**: $H_s = -\sum_k P(f_k)\log_2 P(f_k)$, where $P(f_k)$ is the normalized power spectral density, capturing the regularity of temporal variation.

This analysis yielded 412 radiomic temporal summaries per segment (103 PyRadiomics features $\times$ 4 temporal summaries) and 2,048 encoder temporal summaries (512 channels $\times$ 4 summaries).

4 Experiments and Results

4.1 Model Implementation and Experimental Protocol

We trained class-weighted Random Forest (RF) classifiers [2] for per-segment hyperenhancement prediction under three feature configurations: (1) radiomics, (2) encoder features (Enc), and (3) early fusion (Radiomics + Enc). We applied class weighting to mitigate significant segment-level class imbalance (10.5% hyperenhancement prevalence). For high-dimensional encoder and fusion features, we used principal component analysis (PCA) to reduce dimensionality to 50 components prior to classification. We preprocessed radiomics features with collinearity filtering and univariate feature selection (SelectKBest).

To prevent data leakage, we used patient-stratified nested cross-validation throughout: an outer 5-fold `StratifiedGroupKFold` (grouped by patient MRN) for performance estimation, with an inner 3-fold `StratifiedGroupKFold` inside `GridSearchCV` for hyperparameter selection within each outer training fold. This ensures all videos from a given patient appear exclusively in either the training or evaluation set at each fold. We report results on the outer folds.

4.2 Per-Segment Classification: Feature Space Comparison

Table 1 presents per-segment AUC for all three feature configurations across 18 view–segment combinations. Three key findings emerge from this systematic comparison:

Encoder features consistently outperform radiomics. Across all 18 segments, encoder features achieved a mean AUC of 0.822 compared to 0.793 for radiomics. DeLong tests confirmed significant superiority of encoder features on 8/18 segments ($p < 0.05$), with the largest improvements in A2C segments (Seg 10: 0.730 vs. 0.820, $p = 0.0004$) where texture-based radiomics appear least effective.

The incremental value of fusion over encoder features is inconclusive. Encoder and fusion features achieved comparable overall performance (mean AUC 0.822 vs. 0.830). Fusion yielded the highest AUC on A3C (mean 0.858, significantly outperforming encoder on 6/6 segments), but on A2C, encoder features alone performed best (mean 0.826), with fusion *degrading* performance on two segments (Seg 04: 0.820 $\rightarrow$ 0.790, $p = 0.039$; Seg 07: 0.844 $\rightarrow$

Table 1: Per-segment AUC for hyperenhancement detection by feature type (nested 5-fold CV on development set). Bold: best feature type per segment.

View	Seg	Radiomics AUC	Encoder AUC	Fusion AUC	Best	Rad vs Enc p	Rad vs Fus p
A4C (n=5870)	03	**0.830**	0.819	0.821	Rad	0.29	0.39
	06	0.771	0.782	**0.780**	Enc	0.42	0.51
	09	0.806	0.828	**0.832**	Fus	**0.034**	**0.015**
	12	0.812	0.806	**0.819**	Fus	0.63	0.55
	14	0.809	**0.842**	0.830	Enc	**0.015**	0.09
	16	0.802	0.839	**0.840**	Fus	**0.008**	**0.006**
	Mean	*0.805*	*0.819*	*0.820*			
A3C (n=2159)	02	0.811	0.819	**0.852**	Fus	0.56	**0.004**
	05	0.836	0.825	**0.868**	Fus	0.44	**0.018**
	08	0.809	0.827	**0.869**	Fus	0.21	**<0.001**
	11	0.816	0.835	**0.868**	Fus	0.25	**0.004**
	13	0.733	0.771	**0.815**	Fus	**0.027**	**<0.001**
	15	0.800	0.850	**0.876**	Fus	**0.012**	**<0.001**
	Mean	*0.801*	*0.821*	*0.858*			
A2C (n=1972)	01	0.705	**0.770**	0.756	Enc	0.06	0.11
	04	0.757	**0.820**	0.790	Enc	**0.010**	0.13
	07	0.833	**0.844**	0.808	Enc	0.56	0.14
	10	0.730	**0.820**	0.803	Enc	**0.0004**	**0.002**
	13	0.782	0.834	**0.844**	Fus	**0.027**	**0.004**
	15	0.824	0.867	**0.869**	Fus	**0.041**	**0.018**
	Mean	*0.772*	*0.826*	*0.812*			
Overall Mean		**0.793**	**0.822**	**0.830**			

0.808, $p = 0.019$). On A4C, the approaches were statistically indistinguishable. No consistent pattern emerged across views, likely reflecting issues related to class imbalance in the training data, and cohort size.

Both feature types significantly outperform clinical baselines. Conventional echocardiographic parameters—ejection fraction (AUC $\approx$ 0.54 [10]), global longitudinal strain (AUC $\approx$ 0.70–0.74 [15]), and wall motion score index (AUC $\approx$ 0.62 for ischemic injury [6])—provide substantially lower discriminative ability for per-segment scar detection. Even the weakest segment (Seg 01, A2C: Radiomics = 0.705, Enc = 0.770) exceeds the best reported clinical baseline.

4.3 View-Level Scar Screening

For view-level "any scar" detection, we used encoder features from each video frame to generate summary statistics and trained a separate RF model to predict whether the patient harbored hyperenhancement in any AHA segment visible from that view. Encoder features achieved strong screening performance across

Table 2: View-level "any scar" detection using encoder features with random forest (nested 5-fold CV on development set).

View	n	AUC [95% CI]	Sens.	Spec.
A2C	1,972	0.865 [0.849–0.882]	0.752	0.770
A3C	2,159	0.850 [0.833–0.871]	0.616	0.902
A4C	5,870	0.851 [0.841–0.860]	0.802	0.695

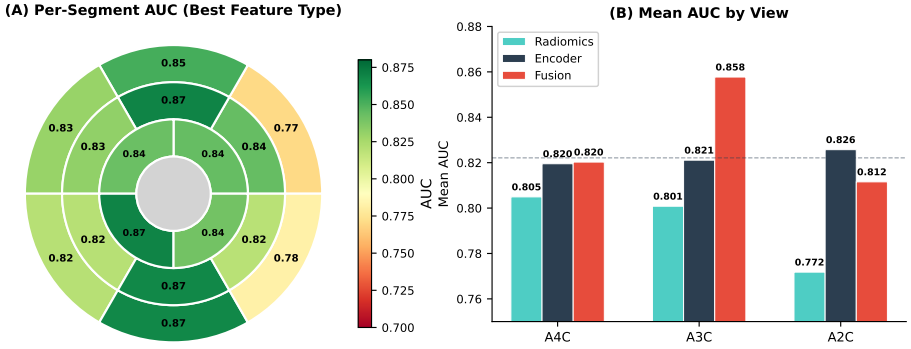

Fig. 2: (A) AHA bullseye plot displaying per-segment AUC using the best-performing feature type for each segment. (B) Mean AUC by view and feature configuration, showing encoder consistently outperforming Radiomics, with fusion providing inconsistent additional benefit.

all views (Table 2). A2C yielded the highest AUC (0.865) with balanced sensitivity and specificity, while A3C achieved the highest specificity (0.902) suitable for a rule-in strategy. All three views exceeded an AUC of 0.85, supporting clinical utility as a CMR triage tool.

5 Discussion

We present the largest echocardiographic study for myocardial scar detection to date, demonstrating that routine echocardiography can be enhanced with CMR-level tissue characterization through temporal feature analysis. Our systematic comparison of radiomics and segmentation-derived encoder features across 18 view–segment combinations reveals three actionable insights for the field.

Segmentation encoders capture scar-related tissue information. The nnU-Net encoders, trained exclusively for myocardial segmentation, produced bottleneck representations that significantly outperformed radiomics for scar detection (mean AUC 0.822 vs. 0.793). This indicates that learning to delineate myocardium also implicitly encodes tissue properties, a finding consistent with the observation that myocardial infarction alters the acoustic properties that

influence both segmentation boundaries and tissue texture [16, 1]. Importantly, this "free" tissue characterization requires no additional scar-specific training.

Clinical utility as a screening tool. View-level AUC of 0.85–0.87 across all three views positions this framework as a viable CMR triage tool. With echocardiography performed as part of routine clinical care, an automated scar screening module could flag patients at elevated risk for myocardial scar, enabling targeted CMR referral. Many echocardiograms are performed for non-specific indications such as dyspnea or atypical chest pain; automated scar screening may be most impactful precisely in these cases, where clinical suspicion is insufficient to prompt direct CMR referral. Both feature approaches substantially exceed the discriminative ability of conventional echocardiographic parameters (EF: AUC $\approx$ 0.54; GLS: AUC $\approx$ 0.70–0.74; WMSI: AUC $\approx$ 0.62), confirming the incremental value of quantitative feature extraction.

Comparison with prior work. While Jamthikar et al. demonstrated the capacity of radiomics to detect acute myocardial infarction (AUC 0.87, 684 patients, 6 centers), our study establishes that standard echocardiography also harbors the structural acoustic signatures necessary to detect chronic replacement fibrosis from diverse etiologies. This result indicates that learned representations can extract sufficient tissue characterization from standard clinical views. We provide the largest systematic head-to-head comparison of radiomics and encoder features for cardiac tissue characterization, with statistical testing at every view–segment combination confirming that segmentation models implicitly learn superior tissue representations.

Limitations. This is a single-center study; multi-center validation with vendor harmonization (e.g., ComBat [11, 7]) is an important next step. The retrospective design introduces selection bias toward patients who received both echocardiography and CMR. CMR and echocardiography pairs were matched within a one-year window; scar positives in some cases may reflect pathology that matured between the two studies. The encoder features, while more powerful, are less interpretable than radiomics, though the temporal summary statistics provide some insight into which cardiac phases are most discriminative. The view-dependence of fusion benefit may in part reflect properties of the acquisition view (e.g., image quality, segment geometry, class imbalance) that modulate the complementarity of the two feature spaces; disentangling these factors is a direction for future work. Finally, we used random forest classifiers throughout; end-to-end approaches using scar-supervised fine-tuning of video foundation models (e.g., DINOv2 or masked autoencoders pre-trained on echocardiography) represent natural extensions that may further improve performance.

6 Conclusion

We developed a multi-stage pipeline demonstrating that routine echocardiography can detect CMR-confirmed myocardial scar through temporal feature analysis from nnU-Net-segmented myocardium. Crucially, our systematic evaluation reveals that semantic segmentation models in echocardiography inherently learn

underlying tissue state as a byproduct of boundary detection. Encoder bottleneck features, trained exclusively for spatial segmentation without scar-specific labels, consistently outperform dedicated radiomics for tissue characterization. This feature extraction paradigm, combined with view-level screening performance, offers a highly scalable approach to opportunistic scar screening for targeted CMR referrals and as a risk marker for adverse events.

Acknowledgments

This work was supported by the NSF PATHS-UP Engineering Research Center (ERC) under cooperative agreement EEC-1648451. The authors thank Krish Kabra for many insightful conversations.

Disclosure of Interests

The authors declare no competing interests.

References

1. Bijnens, B., Cikes, M., Butakoff, C., Sitges, M., Crispi, F.: Myocardial motion and deformation: What does it tell us and how does it relate to function? Fetal Diagnosis and Therapy **32**(1–2), 5–16 (2012)
2. Breiman, L.: Random forests. Machine Learning **45**(1), 5–32 (2001)
3. Cerqueira, M.D., Weissman, N.J., Dilsizian, V., Jacobs, A.K., Kaul, S., Laskey, W.K., Pennell, D.J., Rumberger, J.A., Ryan, T., Verani, M.S.: Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart: A statement for healthcare professionals from the cardiac imaging committee of the council on clinical cardiology of the American Heart Association. Circulation **105**(4), 539–542 (2002)
4. Dhamdhere, R., Modanwal, G., Makhlouf, M.H.E., Shafiabadi Hassani, N., Bharadwaj, S., Fu, P., Milioglou, I., Rahman, M., Al-Kindi, S., Madabhushi, A.: STAR-Echo: A novel biomarker for prognosis of MACE in chronic kidney disease patients using spatiotemporal analysis and transformer-based radiomics models. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2023. LNCS, vol. 14225, pp. 284–294. Springer Nature Switzerland (2023). https://doi.org/10.1007/978-3-031-43987-2_28
5. 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. JACC: Cardiovascular Imaging **9**(9), 1046–1055 (2016)
6. Eek, C., Grenne, B., Brunvand, H., Aakhus, S., Endresen, K., Hol, P.K., Smith, H.J., Smiseth, O.A., Edvardsen, T., Skulstad, H.: Strain echocardiography and wall motion score index predicts final infarct size in patients with non-ST-segment-elevation myocardial infarction. Circulation: Cardiovascular Imaging **3**(2), 187–194 (2010)

7. Fortin, J.P., Parker, D., Tunç, B., Watanabe, T., Elliott, M.A., Ruparel, K., Roalf, D.R., Satterthwaite, T.D., Gur, R.C., Gur, R.E., Scheffer, R.T., Verma, R., Shinohara, R.T.: Harmonization of multi-site diffusion tensor imaging data. NeuroImage **161**, 149–170 (2017)
8. van Griethuysen, J.J.M., Fedorov, A., Parmar, C., Hosny, A., Aucoin, N., Narayan, V., Beets-Tan, R.G.H., Fillion-Robin, J.C., Pieper, S., Aerts, H.J.W.L.: Computational radiomics system to decode the radiographic phenotype. Cancer Research **77**(21), e104–e107 (2017)
9. Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation. Nature Methods **18**(2), 203–211 (2021)
10. Jamthikar, A.D., Hathaway, Q.A., Maganti, K., Hamirani, Y., Bokhari, S., Yanamala, N., Sengupta, P.P.: Ultrasonic texture analysis for predicting acute myocardial infarction. JACC: Cardiovascular Imaging **18**(11), 1185–1199 (2025). https://doi.org/10.1016/j.jcmg.2025.06.018
11. Johnson, W.E., Li, C., Rabinovic, A.: Adjusting batch effects in microarray expression data using empirical Bayes methods. Biostatistics **8**(1), 118–127 (2007)
12. Kim, R.J., Wu, E., Rafael, A., Chen, E.L., Parker, M.A., Simonetti, O., Klocke, F.J., Bonow, R.O., Judd, R.M.: The use of contrast-enhanced magnetic resonance imaging to identify reversible myocardial dysfunction. New England Journal of Medicine **343**(20), 1445–1453 (2000)
13. Kramer, C.M., Barkhausen, J., Flamm, S.D., Kim, R.J., Nagel, E.: Standardized cardiovascular magnetic resonance (CMR) protocols 2013 update. Journal of Cardiovascular Magnetic Resonance **15**(1), 91 (2013)
14. Kwong, R.Y., Chan, A.K., Brown, K.A., Chan, C.W., Reynolds, H.G., Tsang, S., Davis, R.B.: Impact of unrecognized myocardial scar detected by cardiac magnetic resonance imaging on event-free survival in patients presenting with signs or symptoms of coronary artery disease. Circulation **113**(23), 2733–2743 (2006)
15. Le, T.T., Huang, W., Singh, G.K., Toh, D.F., Ewe, S.H., Tang, H.C., Loo, G., Bryant, J.A., Ang, B., Tay, E.L.W., et al.: Echocardiographic global longitudinal strain is associated with myocardial fibrosis and predicts outcomes in aortic stenosis. Frontiers in Cardiovascular Medicine **8**, 750016 (2021). https://doi.org/10.3389/fcvm.2021.750016
16. Mimbs, J.W., O'Donnell, M., Bauwens, D., Miller, J.G., Sobel, B.E.: The dependence of ultrasonic attenuation and backscatter on collagen content in dog and rabbit hearts. Circulation Research **47**(1), 49–58 (1980). https://doi.org/10.1161/01.res.47.1.49
17. Ouyang, D., He, B., Ghorbani, A., Yuan, N., Ebinger, J., Langlotz, C.P., Heidenreich, P.A., Harrington, R.A., Liang, D.H., Ashley, E.A., Zou, J.Y.: Video-based AI for beat-to-beat assessment of cardiac function. Nature **580**(7802), 252–256 (2020)
18. Vukadinovic, M., Chiu, I.M., Tang, X., Yuan, N., Chen, T.Y., Cheng, P., Li, D., Cheng, S., He, B., Ouyang, D.: Comprehensive echocardiogram evaluation with view primed vision language AI. Nature **650**(8103), 970–977 (2025). https://doi.org/10.1038/s41586-025-09850-x
19. Zhang, J., Gajjala, S., Agrawal, P., Tison, G.H., Hallock, L.A., Beussink-Nelson, L., Lassen, M.H., Fan, E., Aras, M.A., Jordan, C., Fleischmann, K.E., Melisko, M., Qasim, A., Shah, S.J., Bajcsy, R., Deo, R.C.: Fully automated echocardiogram interpretation in clinical practice. Circulation **138**(16), 1623–1635 (2018)