SHINE: An Entropy-Guided Digital Twin Framework for 3D Myocardial Scar Reconstruction from Sparse CMR

Dexiang Zong[1], Shuang Leng[1], Xiaodan Zhao[1], Ru-San Tan[1], Lei Chen[4], Jun Cheng[5], Xulei Yang[5], Ching-Hui Sia[6], Mark Chan[6], Derek Hausenloy[1], and Liang Zhong[1,2,3⋆]

[1] National Heart Centre Singapore, Singapore
[2] Duke-NUS Medical School, Singapore
[3] Department of Biomedical Engineering, National University of Singapore, Singapore
[4] Tongji University, Shanghai, China
[5] Agency for Science, Technology and Research (A*STAR), Singapore
[6] National University Health System, Singapore

Abstract. Accurate reconstruction of the arrhythmogenic myocardial substrate—specifically the conducting isthmuses essential for ventricular tachycardia (VT) ablation—depends on capturing fine-scale scar topology from sparse, anisotropic Late Gadolinium Enhancement (LGE) MRI. Given the pronounced anisotropy (e.g., $1.2 \times 1.2 \times 8.0$ mm) and limited slice sampling, volumetric reconstruction constitutes a highly ill-posed inverse problem, often leading to aliasing artifacts and discontinuous "staircase" geometries under discrete interpolation schemes. To address this challenge, we propose substrate reconstruction as a continuous function learning task defined on a patient-specific anatomical manifold. We introduce SHINE (Scar Heterogeneity via Implicit Neural Encoding), a framework that constructs a mathematically consistent "Digital Twin" of the diseased myocardium. The pipeline integrates automated biventricular segmentation with a Statistical Shape Model (SSM) to establish a continuous anatomical prior. To reconstruct detailed scar morphology, we employ a Bayesian SIREN architecture with sinusoidal activation functions, enabling the representation of high-frequency geometric structures and mitigating the spectral bias persistent in conventional Neural Fields. Importantly, we introduce an Entropy-Weighted Dirichlet Energy loss, which uses epistemic uncertainty to adaptively enforce smoothness in data-sparse regions while preserving high-fidelity details in data-rich regions. Evaluated on a clinical VT cohort of 40 patients, SHINE significantly outperformed discrete interpolation and weakly supervised benchmarks, achieving a Dice Similarity Coefficient of 0.93 and a volumetric error of 3.1%. Notably, our method effectively addresses the Nyquist sampling gap, enabling recovery of sub-voxel conducting channels within a unified, topology-preserving 3D framework to support precision electrophysiology planning.

⋆ Corresponding author: zhong.liang@nhcs.com.sg

Keywords: Myocardial Scar Reconstruction · Physics-Informed Neural Fields · Entropy-Guided Regularization · Anatomical Digital Twin · Sparse-to-Dense · Uncertainty Quantification.

1 Introduction

Sudden Cardiac Death (SCD) secondary to ventricular tachyarrhythmias remains a leading cause of mortality. While the field of medical electronics has evolved significantly since its infancy [1], and quantitative image analysis has become a cornerstone of clinical practice [2], identifying the specific arrhythmogenic substrate remains challenging. This substrate is predominantly characterized by the Heterogeneous Zone (HZ), where dense scar interdigitates with viable myocytes to form complex conducting pathways. Identifying narrow "conducting isthmuses" within this substrate is paramount, as these channels sustain reentrant circuits and constitute critical targets for catheter ablation [3].

While Late Gadolinium Enhancement (LGE) MRI offers a unique non-invasive method to assess myocardial fibrosis, its clinical utility is severely impeded by acquisition constraints. To balance breath-hold duration and signal-to-noise ratio (SNR), clinical protocols typically adopt highly anisotropic voxel dimensions (e.g., $1.2 \times 1.2 \times 8.0$ mm). This sparse through-plane resolution leaves approximately 80% of the myocardial volume unobserved [4], rendering the reconstruction of a continuous, topology-preserving 3D structure a mathematically ill-posed inverse problem. Traditional interpolation methods fail to address this fundamental limitation; for instance, Log-Odds interpolation [5] assumes local linearity, resulting in "staircase artifacts" that artificially sever the subtle conducting pathways required for reentry. Similarly, standard deep learning approaches [6, 7] utilizing Convolutional Neural Networks (CNNs) suffer from spectral bias [8]. These networks preferentially learn low-frequency components, effectively smoothing out the sharp, high-frequency details of narrow isthmuses. Recent advances in Implicit Neural Representations (INRs), popularized by Neural Radiance Fields [9], have shown promise in medical imaging [10, 11], yet applying them to sparse LGE-MRI requires careful regularization to avoid hallucinations. Crucially, existing deterministic models lack the capacity to quantify predictive reliability in unobserved gaps. Without uncertainty quantification, these "black-box" models risk synthesizing unrealistic anatomical structures in "blind zones," rendering the reconstruction clinically unsafe for precision ablation planning [16].

To resolve these challenges, we propose SHINE (Scar Heterogeneity via Implicit Neural Encoding), an Entropy-Guided Digital Twin framework. Unlike standard Implicit Neural Representations (Vanilla INRs) that act as deterministic 'black boxes' and are prone to generating non-physical 'hallucinations' in unobserved regions, SHINE introduces a fundamentally robust algorithmic paradigm. By seamlessly integrating patient-specific physical constraints (via anatomical manifolds) with a novel uncertainty-aware loss function, our framework inherently quantifies epistemic uncertainty. This enables the model to adaptively enforce topological continuity in data-sparse 'blind zones' while rigorously

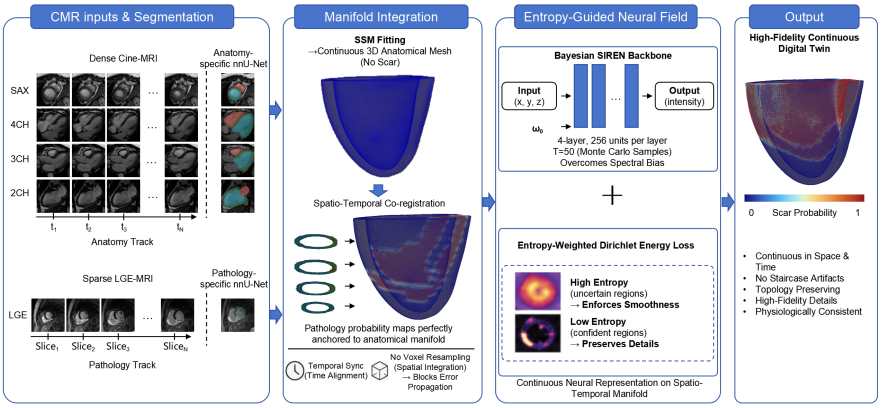

Fig. 1. Overview of SHINE, which integrates Cine-derived anatomical manifolds, sparse LGE pathology, co-registration, and entropy-guided Bayesian SIREN reconstruction into a resolution-independent cardiac digital twin.

preserving high-frequency substrate details where observations are available. We reformulate substrate reconstruction not as discrete voxel in-painting, but as a continuous function learning problem on a patient-specific anatomical manifold. Transcending discrete grids, SHINE constructs a mathematically consistent Digital Twin of the pathological myocardium that is resolution-independent and topologically accurate [12]. Recent studies have further demonstrated the clinical prognostic value of such digital twin models in complex cardiovascular conditions [13]. Our specific contributions are threefold:

1. **Entropy-Guided Physics-Informed Neural Field:** We introduce a reconstruction paradigm that integrates anatomical priors with INRs [14]. By treating the myocardium as a continuous manifold, our method recovers sub-voxel conducting channels, effectively bridging the Nyquist sampling gap.
2. **Bayesian SIREN Backbone:** We employ a Bayesian SIREN architecture [15]. By leveraging periodic activations within a variational framework [16], we enable the representation of high-frequency geometric structures to overcome spectral bias while providing a rigorous measure of predictive uncertainty.
3. **Entropy-Weighted Dirichlet Energy Loss:** We propose a novel regularization mechanism that leverages the model's epistemic uncertainty. This loss function adaptively enforces physical smoothness (Dirichlet energy) in high-uncertainty gaps while preserving high-fidelity topological details in data-rich regions.

2 Methodology

We formulate 3D myocardial substrate reconstruction as a Bayesian inverse problem on a patient-specific Riemannian manifold. The proposed framework,

SHINE, decouples the geometric reconstruction (Anatomical Digital Twin) from the pathological function learning (Entropy-Guided Neural Field).

2.1 Problem Formulation

Let $\Omega \subset \mathbb{R}^3$ denote the volumetric domain of the left ventricle (LV). We model the myocardial substrate as a continuous scalar field $f : \Omega \to [0, 1]$, where $f(x)$ represents the latent scar probability at a spatial coordinate $x \in \Omega$. Clinical LGE-MRI provides a set of sparse, anisotropic observations $D_{obs} = \{(x_i, y_i)\}_{i=1}^{N}$ sampled exclusively on imaging planes $P_{slices} \subset \Omega$. Our objective is to estimate the posterior distribution $p(f|D_{obs})$.

2.2 The Anatomical Digital Twin: Geometric Prior Construction

To constrain the ill-posed nature of reconstructing f, we impose a patient-specific geometric prior by constructing an anatomical manifold M.

Geometric Manifold Fitting. Let $I = \{I_{sax}, I_{2ch}, I_{3ch}, I_{4ch}\}$ denote multi-view Cine MRI sequences. We utilize a pre-trained nnU-Net [17] to generate dense 2D segmentations. To achieve a continuous representation, we fit a Statistical Shape Model (SSM) [18] to these contours: $S = \bar{S} + \sum_{k=1}^{K} \alpha_k p_k$, where $\bar{S}$ is the population mean shape and p_k are principal modes. Coefficients are optimized via Chamfer distance minimization.

Spatio-Temporal Co-registration. Consistency is maintained via rigorous alignment and multi-view integration [19]:

- Temporal Synchronization: Cardiac phases are synchronized using DICOM metadata.
- Spatial Integration: Physical coordinates rigidly map sparse LGE slices into the 3D Cine-derived manifold without voxel resampling.

2.3 Bayesian SIREN for Pathological Substrate Representation

We represent the scar probability field f as an Implicit Neural Representation (INR) parameterized by weights θ. Unlike discrete voxel arrays, this continuous formulation enables resolution-independent reconstruction on the anatomical manifold.

Input Standardization. To mitigate variability, we utilize softmax probability maps generated by a pathology-specific nnU-Net [17] as input $y_i \in [0, 1]$.

Bayesian SIREN Backbone. To overcome spectral bias, we employ a Bayesian SIREN backbone [15]. The activation for the l-th layer is: $h^{(l)}(x) = \sin(\omega_0(W^{(l)}x + b^{(l)}))$. To quantify predictive reliability, we treat θ as random variables with a variational posterior $q_\phi(\theta)$ approximated via MC Dropout [16]. Predictive mean $\hat{y}(x)$, predictive entropy $H(x)$, and normalized entropy $\tilde{H}(x)$ are estimated through T stochastic forward passes.

2.4 Entropy-Weighted Dirichlet Energy Loss

To prevent non-physical "hallucinations" in unobserved gaps, we introduce an adaptive objective function that leverages predictive uncertainty:

$$L_{total} = L_{BCE}(\hat{y}, y_{obs}) + \lambda \cdot L_{entropy\_reg} \tag{1}$$

where

$$L_{entropy\_reg} = \int_{\Omega_{myo}} \tilde{H}(x) \cdot \|\nabla_x \hat{y}(x)\|_2^2 \, dx \tag{2}$$

Theoretical Mechanism: In data-rich regions ($\tilde{H} \to 0$), the regularization vanishes. In sparse gaps ($\tilde{H} \to 1$), the loss enforces a harmonic constraint ($\Delta f = 0$), forcing smooth transitions.

2.5 Implementation Details

To ensure reproducibility, our framework utilizes a 4-layer Bayesian SIREN backbone with 256 hidden units per layer. The predictive mean and epistemic uncertainty are estimated through $T = 50$ stochastic forward passes. The entropy regularization weight is set to $\lambda = 5 \times 10^{-4}$. The source code and pre-trained models are publicly available at https://github.com/FluxVisionAI/SHINE.

3 Results and Discussion

3.1 Comparative Analysis and the Fidelity Paradox

We evaluated SHINE on a clinical cohort of 40 patients originating from two centers with different MRI scanners (1.5T Siemens Avanto vs. 3.0T Biograph mMR), providing evidence of robustness across acquisition settings. All acquisitions were analyzed at a nominal $1.2 \times 1.2 \times 8.0$ mm sparse LGE resolution. We employed a dual-track validation protocol: (i) a Geometric Track using high-resolution Cine-MRI as an anatomical reference, and (ii) a Pathological Track via Leave-One-Slice-Out (LOSO) cross-validation. Specifically, the LOSO protocol evaluated SHINE against 455 physical slices covering extreme infarct sizes (4.7%–56.1%), validating the framework's reliability across diverse pathological conditions. Although absolute 3D ground truth data, such as 3D single-breathhold LGE-CMR, are not available in the current study, our data provides relatively dense spatial sampling (8-16 slices/patient). To address this limitation, we utilize the LOSO cross-validation protocol as a rigorous surrogate ground truth. By withholding true physical LGE slices (completely isolated from optimization), we establish a reliable baseline. Evaluated against these hidden slices, SHINE achieves high spatial overlap (DSC 0.93) and strong intensity fidelity (MSE 4×10^{-5}), supporting the anatomical plausibility of the reconstructed fine-scale structures rather than non-physical hallucinations. As shown in Table 1, SHINE significantly outperforms baselines, achieving a DSC of 0.93 and an

Table 1. Performance benchmark against state-of-the-art reconstruction methods.

Method	Dice Score (DSC) ↑	MSE (Fidelity) ↓	Vol. Error ↓
Log-Odds [5]	0.71	8.2×10^{-2}	23.5%
SR-CNN [7]	0.88	4.5×10^{-2}	10.8%
Weakly Supervised DL [6]	0.91	1.9×10^{-2}	7.1%
Vanilla INR [15]	0.91	2.0×10^{-3}	15.2%
SHINE (Proposed)	**0.93**	$\mathbf{4.0 \times 10^{-5}}$	**3.1%**

intensity fidelity (MSE) of 4×10^{-5}—three orders of magnitude superior to Log-Odds interpolation. This performance substantiates our validation: the negligible MSE reflects high pathological fidelity (Track 2), while the 3.1% volumetric error confirms precise anatomical anchoring (Track 1). Qualitative results (Fig. 2, 3) further demonstrate that SHINE maintains topological continuity across 8.0 mm gaps, resolving the 'Fidelity Paradox' where Vanilla INRs exhibit high DSC but catastrophic 15.2% volumetric errors due to fragmented hallucinations.

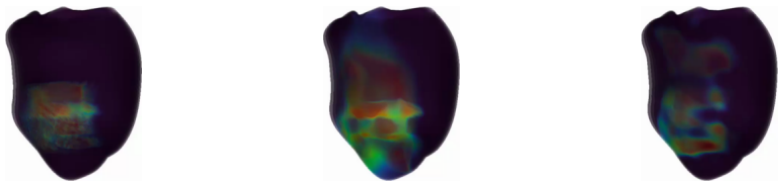

(A) Log-Odds (B) DL Baseline (C) SHINE (Proposed)

Fig. 2. Qualitative 3D reconstruction benchmarking. SHINE recovers sub-voxel topological continuity, whereas Log-Odds and DL baselines show staircase artifacts or blurred channels.

SHINE (Proposed) achieves near-perfect intensity fidelity (MSE = 4×10^{-5}), which is vital for clinical simulations where small errors in scar border can significantly impact reentry prediction. This confirms that standard INRs lack the necessary constraints to maintain physical consistency, leading to significant overestimation of scar burden.

The volumetric error paradox in Table 2 further highlights the need for topology-aware evaluation. Although removing entropy regularization yields a lower volumetric error (1.4%), its high hidden-slice MSE (1.9×10^{-3}) and Fig. 4A indicate fragmented salt-and-pepper artifacts caused by cancellation between over- and under-estimation. SHINE instead prioritizes topological fidelity, achieving low MSE (4×10^{-5}) and reliable conducting-channel continuity despite a 3.1% volumetric error.

To address potential upstream error propagation [21, 22], we analyzed the relationship between segmentation/SSM geometric deviations and downstream reconstruction errors. No significant correlation was observed (r=0.20, p=0.34; mean deviation 1.59 ± 0.22 mm), indicating that SHINE's manifold alignment

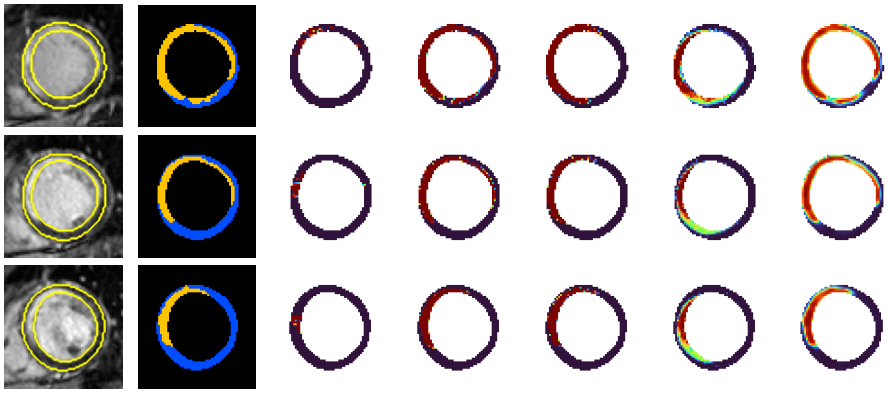

Fig. 3. Qualitative slice-wise benchmarking. Compared with Log-Odds [5], Vanilla INR [15], SR-CNN [7], and Weakly Supervised DL [6], SHINE best preserves sharp boundaries and sub-voxel conducting-channel continuity.

Table 2. Architectural ablation study and component-wise performance analysis.

Method	Backbone	Align.	Ent-Reg.	Dice Score (DSC) ↑	MSE ↓	Vol. Error ↓
SHINE (Proposed)	**SIREN**	✓	✓	**0.93**	$\mathbf{4.0 \times 10^{-5}}$	**3.1%**
(a) w/o Entropy-Reg	SIREN	✓	×	0.92	1.9×10^{-3}	1.4%
(b) w/o Alignment	SIREN	×	✓	0.89	0.6×10^{-3}	2.6%
(c) w/o SIREN	ReLU+PE	✓	✓	0.74	0.6×10^{-3}	14.0%
(d) Naive Baseline	ReLU+PE	×	×	0.86	4.6×10^{-3}	5.4%
(e) Vanilla INR	SIREN	×	×	0.91	2.0×10^{-3}	15.2%

decouples anatomical geometry from pathological texture learning and mitigates propagation of geometric noise into scar topology.

3.2 Ablation Study and Spectral Bias Resolution

We conducted a systematic ablation study to isolate the contribution of each core component. Quantitative results are summarized in Table 2, revealing a critical "Performance Inversion" between Variants (c) and (d). Visual evidence in Fig. 4 corroborates the findings.

Architectural Synergy and Spectral Bias. The results substantiate that the SIREN backbone is the most critical architectural element. Comparing Vanilla INR (0.91) with the Naive Baseline (0.86) reveals a significant performance leap provided by sinusoidal activations. Interestingly, replacing SIREN with a ReLU+PE MLP (Variant c) [8] leads to a substantial collapse in DSC to 0.74, which is notably lower than the Naive Baseline. This performance inversion between (c) and (d) identifies a conflict where low-capacity ReLU kernels struggle to satisfy strict manifold constraints, indicating that SIREN provides a critical inductive

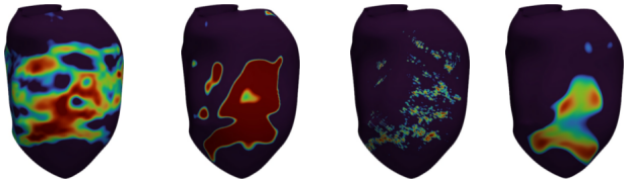

(A) w/o Ent-Reg (B) w/o Align (C) w/o SIREN (D) SHINE

Fig. 4. Qualitative analysis of ablation variants. Removing entropy regularization, alignment, or SIREN introduces hallucinations, geometric distortion, or spectral-bias-induced blurring, respectively.

bias for this task. Furthermore, the excessive Volumetric Error in Variant (e) (15.2%) substantiates the critical role of entropy-guided regularization. Without this adaptive mechanism, the neural field is prone to non-physical oscillations and "hallucinations" in unobserved regions.

3.3 Clinical Safety Maps and Computational Efficiency

Beyond numerical accuracy, SHINE provides essential uncertainty-aware clinical safety maps (Fig. 5). High entropy alerts clinicians to unobserved inter-slice gaps, while low entropy correlates with verified LGE findings. Regarding computational complexity, the patient-specific manifold fitting requires approximately 7.5 minutes per case. Crucially, this is strictly a pre-operative process, which seamlessly enables highly efficient, real-time intraoperative inference taking only 0.54 seconds per patient on an NVIDIA RTX 4090 GPU.

(A) Predictive Uncertainty (B) Blind Zones (C) Clinical Safety Map

Fig. 5. Uncertainty-aware clinical safety maps. High-entropy regions indicate unobserved inter-slice gaps, whereas low-entropy regions correspond to verified LGE findings.

4 Conclusion

This paper introduced SHINE, an entropy-guided Bayesian neural field for 3D myocardial scar reconstruction from sparse clinical CMR. By learning a continuous scar probability function on patient-specific anatomical manifolds, SHINE

mitigates discretization artifacts and spectral bias while quantifying epistemic uncertainty. In a 40-patient VT cohort evaluated with leave-one-slice-out validation, SHINE achieved a DSC of 0.93, a volumetric error of 3.1%, and real-time inference in 0.54 s per case. Future work will validate the framework on a multicenter cohort of over 500 cases across three institutions [20].

Acknowledgments. This study is supported by the National Health Innovation Centre (NHIC-I2D-2502365), the National Medical Research Council Open Fund Individual Research Grant (MOH-001903-01), the SingHealth Duke-NUS AMC–A*STAR Healthcare Translation Partnership (07 FY2023 HTP P2 15-A2), and the SingHealth Duke-NUS Academic Medicine Research Grants AM/TP098/2025 and AM/TP032/2020. L. Zhong is supported by the National Medical Research Council Clinician Innovator Award (MOH-001647-00).

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

References

1. Lusted, L.B.: Medical electronics. New Engl. J. Med. **252**(14), 580–585 (1955)
2. Giger, M.L., Chan, H.P., Boone, J.: Anniversary Paper: History and status of CAD and quantitative image analysis: The role of Medical Physics and AAPM. Med. Phys. **35**(11), 5799–5820 (2008)
3. Perez-David, E., Arenal, A., Rubio-Guivernau, J.L., et al.: Noninvasive identification of ventricular tachycardia-related conducting channels using contrast-enhanced magnetic resonance imaging in patients with chronic myocardial infarction. J. Am. Coll. Cardiol. **57**(2), 184–194 (2011)
4. Lustig, M., Donoho, D., Pauly, J.M.: Sparse MRI: The application of compressed sensing for rapid MR imaging. Magn. Reson. Med. **58**(6), 1182–1195 (2007)
5. Ukwatta, E., Arevalo, H., Rajchl, M., White, J., Pashakhanloo, F., Prakosa, A., Herzka, D.A., McVeigh, E., Lardo, A.C., Trayanova, N.A., Vadakkumpadan, F.: Image-based reconstruction of three-dimensional myocardial infarct geometry for patient-specific modeling of cardiac electrophysiology. Med. Phys. **42**(8), 4579–4590 (2015)
6. Sen, A., Rohrer, U., Bhagirath, P., et al.: Weakly supervised learning for scar reconstruction in personalized cardiac models: Integrating 2D MRI to 3D anatomical models. Comput. Biol. Med. **198**, 111219 (2025)
7. Schlemper, J., Caballero, J., Hajnal, J.V., et al.: A deep cascade of convolutional neural networks for dynamic MR image reconstruction. IEEE Trans. Med. Imaging **37**(2), 491–503 (2018)
8. Rahaman, N., Baratin, A., Arpit, D., et al.: On the spectral bias of neural networks. In: Proc. ICML, vol. 97, pp. 5301–5310 (2019)
9. Mildenhall, B., Srinivasan, P.P., Tancik, M., et al.: NeRF: Representing scenes as neural radiance fields for view synthesis. In: Proc. ECCV, pp. 405–421 (2020)
10. Zhu, Y., Liu, Y., Zhang, Y., Liang, D.: Implicit neural representation for medical image reconstruction. Phys. Med. Biol. **70**(12) (2025)
11. Shen, C., Zhu, H., Zhou, Y., et al.: Continuous 3D myocardial motion tracking via echocardiography. IEEE Trans. Med. Imaging **43**(12), 4236–4252 (2024)

12. Corral-Acero, J., Mezger, F., Jagoda, P., et al.: The 'digital twin' to enable the vision of precision cardiology. Eur. Heart J. **41**(48), 4556–4564 (2020)
13. Leng, S., Guo, J., Alabed, S., et al.: Digital twin model predicts prognosis in pulmonary arterial hypertension. JACC Cardiovasc. Imaging **19**(4), 565–567 (2026)
14. Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. J. Comput. Phys. **378**, 686–707 (2019)
15. Sitzmann, V., Martel, J., Bergman, A., et al.: Implicit Neural Representations with Periodic Activation Functions. In: Proc. NeurIPS, vol. 33, pp. 7462–7473 (2020)
16. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In: Proc. ICML, pp. 1050–1059 (2016)
17. Isensee, F., Jaeger, P.F., Kohl, S.A.A., et al.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. Nat. Methods **18**(2), 203–211 (2021)
18. Mauger, C.A., Gilbert, K., Lee, A.M., et al.: Right ventricular shape and function: cardiovascular magnetic resonance reference morphology and biventricular risk factor morphometrics in UK Biobank. J. Cardiovasc. Magn. Reson. **21**, 41 (2019)
19. Sun, A., Chen, X., Wan, X., et al.: Deep learning reconstruction enhanced myocardial scar detection on late gadolinium enhancement MRI in patients with ventricular arrhythmias. Meta-Radiology **4**(2), 100204 (2026)
20. Yang, P., Leng, S., Zong, D., et al.: Prognostic value of end-to-end deep learning assessment of myocardial scar and microvascular obstruction on late gadolinium enhancement cardiovascular magnetic resonance. J. Cardiovasc. Magn. Reson. 28(1), 102712 (2026)
21. Li, L., Camps, J., Wang, Z.J., et al.: Toward enabling cardiac digital twins of myocardial infarction using deep computational models for inverse inference. IEEE Trans. Med. Imaging **43**(7), 2466–2478 (2024)
22. Zou, R., Li, Y., Zhou, N., Guan, X., Li, G., Xie, D., Bo, J., Li, S., Wu, M., Deng, J., Fan, X., Luo, C.: ScarElastic: continuous elasticity field modeling for myocardial scar delineation in LGE-CMR. NPJ Digit. Med. **9**(1), 3 (2025)