CardioMesh: A Deep Learning Framework for Cardiac Anatomy and Border Zone Segmentation in LGE-MRI

Background: Late gadolinium enhancement cardiac magnetic resonance (LGE-MRI) is the gold standard for non-invasive characterization of myocardial scar, enabling high-resolution visualization of fibrosis and supporting diagnosis, risk stratification, and treatment planning. Increasing attention has focused on the heterogeneous border zone (BZ) between healthy and scarred myocardium, whose accurate delineation is essential for computational modeling of post-infarction arrhythmias and patient risk assessment.

Aim: We propose CardioMesh, a deep learning framework specifically designed for accurate and spatially consistent segmentation of the BZ from LGE-MRI images.

Methods: CardioMesh follows a two-phase training paradigm combining anatomical pretraining and task-specific fine-tuning. The architecture integrates MANet and UNet++ decoders with an encoder ecosystem based on ResNet-50, DenseNet-121, and ConvNeXtV2. The model operates on 2.5D volumetric inputs, where each sample consists of a target slice and its adjacent context, enabling local 3D information capture while preserving computational efficiency. A physics-informed augmentation pipeline is introduced, modelling MRI-specific artefacts including Rician noise, motion blur, and BZ-focused contrast perturbations to improve robustness to acquisition variability. Training is guided by deep supervision and a composite loss function designed to address class imbalance and boundary precision.

Results: Phase 1 domain adaptation over ACDC and M&Ms-2 datasets yields a new definitive SOTA of DSC=0.923 and HD95=0.66 mm with the MANet+ConvNeXtV2 combination — surpassing all published benchmarks. Phase 2 BZ fine-tuning on the Gold Standard cohort achieves a training DSC of 0.759 (ConvNeXtV2) and a leading test-set HD95 of 6.98 mm. We introduce a Clinical Balanced Score combining DSC and HD95 penalty, a Medical Augmentation Pipeline tailored to LGE-MRI physics (Rician noise, BZ-specific contrast augmentation), and a per-patient 3D volumetric evaluation framework.

Conclusion: The ConvNeXtV2 encoder sets a new definitive SOTA on ACDC anatomy (DSC=0.923, HD95=0.66 mm), confirming that large-kernel pure-CNN architectures with Transformer-inspired design principles outperform residual and self-attention baselines for cardiac MRI.

Table 2. Gold Standard — Phase 2 BZ Results (3D Volumetric, mm).

Split	Architecture	DSC ↑	HD95↓	Notes
TRAIN	**ConvNeXtV2 (MANet)**	**0.759**	**1.62**	**Max BZ assimilation**
TRAIN	UNet++ (R50)	0.716	3.83	Hist. benchmark
VAL	MANet (R50)	0.401	10.49	Best Val DSC
VAL	ConvNeXtV2 (MANet)	0.372	15.20	Stable (diagnostic)
TEST	MANet (R50)	0.391	7.93	Best Test DSC
TEST	**ConvNeXtV2 (MANet)**	**0.335**	**6.98**	**Best spatial fidelity**