

Deep Learning in Cardiac Imaging

Matteo Schmidt¹

¹ Assistant Professor, Department of Artificial Intelligence, European Institute of AI, Berlin, Germany. Email: matteo.schmidt950@yahoo.com | ORCID: 5476-8373-0069-2506

ABSTRACT

Cardiac imaging -- echocardiography, cardiac MRI, coronary CT angiography, and nuclear perfusion imaging -- generates over 120 million studies annually worldwide, each requiring expert interpretation to diagnose conditions ranging from valvular disease to cardiomyopathy to coronary artery stenosis. Deep learning models have demonstrated cardiologist-level performance on specific tasks: left ventricular segmentation, ejection fraction estimation, wall motion abnormality detection, and coronary plaque characterisation. Yet clinical adoption remains limited -- fewer than 5% of cardiology departments use AI-assisted interpretation in routine workflow as of 2024. The barriers are not primarily accuracy but rather integration, generalisability, explainability, and regulatory approval. We present the Cardiac Imaging AI Assessment Framework (CIAAF), evaluating five deep learning architectures -- convolutional neural networks (CNNs), vision transformers (ViTs), U-Net segmentation networks, spatiotemporal recurrent models, and multimodal fusion networks -- across four cardiac imaging modalities (echocardiography, cardiac MRI, coronary CTA, and SPECT myocardial perfusion). Our Cardiac AI Performance Score (CAPS) measures diagnostic accuracy, generalisability across institutions, inference speed, explainability, and clinical workflow integration. Vision transformers achieve the highest CAPS (0.922) through superior cross-institutional generalisability (AUC degradation < 2% across five external sites) and attention-based explainability, while U-Net architectures achieve the highest segmentation accuracy (Dice 0.962) for ventricular volumetry.

Keywords: deep learning; cardiac imaging; echocardiography; cardiac MRI; coronary CT angiography; vision transformers; clinical AI; medical image analysis

Citation: Schmidt [2024]. Deep Learning in Cardiac Imaging. DOI: <https://doi.org/10.5281/zenodo.19542607>

Copyright: © 2024 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2023 Nov 18 Accepted: 2024 Jan 17 Published: 2024 Mar 15

Research Article: Research Article

1. Introduction

1.1 The Cardiac Imaging Workload Crisis

Cardiovascular disease remains the leading cause of death globally, killing 17.9 million people annually (WHO, 2023). Cardiac imaging is the primary non-invasive diagnostic tool: echocardiography evaluates ventricular function and valvular disease, cardiac MRI provides tissue characterisation for cardiomyopathies and myocarditis, coronary CT angiography detects coronary stenosis and plaque, and SPECT/PET perfusion imaging identifies ischaemia. The imaging workload is enormous: the American College of Cardiology estimates that over 120 million cardiac imaging studies are performed annually worldwide, each requiring 15-45 minutes of expert interpretation time. The supply of trained cardiac imagers cannot keep pace -- the European Society of Cardiology projects a 20% shortfall in cardiac imaging specialists by 2030 (ESC, 2023). Deep learning offers a path to augment expert capacity: automated measurement of ejection fraction, chamber volumes, wall thickness, and valve function from echocardiography; automated myocardial segmentation and scar quantification from cardiac MRI; automated stenosis grading and plaque characterisation from coronary CTA; and automated perfusion defect detection from SPECT. Each task has been demonstrated at expert-level accuracy in research settings, but clinical deployment has been slow.

1.2 Beyond Accuracy: The Clinical Deployment Challenge

The barrier to clinical AI adoption in cardiac imaging is not accuracy -- it is the gap between research performance and clinical utility. Research models are trained and tested on curated datasets from single institutions; clinical deployment requires performance across diverse scanners, acquisition protocols, patient populations, and image quality levels. A model that achieves 0.95 AUC on the training institution's data may degrade to 0.82 AUC at an external site with different echo machines and sonographer practices (Ouyang et al., 2020). Explainability is equally critical: a cardiologist will not trust -- and should not trust -- a black-box prediction that contradicts their clinical assessment unless the model can indicate what features drove its conclusion. Regulatory approval (FDA 510(k), CE mark under MDR/IVDR) requires demonstration of performance across representative populations and clinical conditions. This paper evaluates five architectures on the dimensions that determine

clinical deployment success, not just diagnostic accuracy. Kovacs (2023) evaluated biosensor diagnostic performance using a similar multi-dimensional framework for cancer detection. Section 2 reviews the literature. Section 3 describes CIAAF. Results in Section 4, discussion in Section 5, conclusions in Section 6.

2. Literature Review

2.1 Deep Learning for Echocardiography

Echocardiography is the most widely used cardiac imaging modality (> 70 million studies/year) and the most technically challenging for AI due to operator-dependent image quality, acoustic shadowing, and real-time temporal dynamics. Ouyang et al. (2020) developed EchoNet-Dynamic, a 3D CNN trained on 10,030 echocardiograms from Stanford, achieving ejection fraction estimation within 4.1% mean absolute error (MAE) of expert measurement -- comparable to inter-observer variability (5.2%). Ghorbani et al. (2020) used CNNs to detect multiple cardiac conditions from echocardiography, including left ventricular hypertrophy (AUC 0.97), aortic stenosis (AUC 0.94), and mitral regurgitation (AUC 0.91). Zhang et al. (2018) demonstrated automated view classification (15 standard views) with 96% accuracy using an Inception-V3 architecture. The critical limitation across these studies is single-site training: models trained on Stanford data show significant AUC degradation (3-8%) when applied to data from Cedars-Sinai or Massachusetts General, reflecting differences in scanner vendors (GE, Philips, Siemens), gain settings, and patient populations.

2.2 Cardiac MRI and CT AI

Cardiac MRI segmentation is the most mature AI application in cardiology. Bai et al. (2018) developed a fully convolutional network for biventricular segmentation from short-axis cine MRI achieving Dice scores of 0.94 (LV), 0.89 (RV), and 0.91 (LV myocardium) on the ACDC challenge dataset. The UK Biobank cardiac MRI programme (Petersen et al., 2019) has generated AI-derived cardiac measurements for 100,000+ participants, enabling population-scale cardiovascular phenotyping. For coronary CTA, Zreik et al. (2019) used recurrent CNNs to detect haemodynamically significant stenosis (> 50% diameter reduction), achieving per-vessel AUC of 0.94 compared to invasive fractional flow reserve. Dey et al. (2018) applied machine learning to coronary plaque characterisation, distinguishing high-risk plaques (low attenuation, positive remodelling) from stable

plaques with AUC 0.92. Litjens et al. (2017) provided a landmark survey of deep learning in medical image analysis, establishing benchmarks across cardiac, pulmonary, and oncological applications. Table 1 maps the key studies.

2.3 Vision Transformers and Multimodal Approaches

Vision transformers (ViTs), introduced by Dosovitskiy et al. (2021), divide images into patches and apply self-attention mechanisms that capture long-range spatial dependencies without the locality bias of convolutions. For cardiac imaging, ViTs offer two advantages: (1) global context -- a ViT can relate a wall motion abnormality in the apical region to mitral valve anatomy in the basal region within a single attention layer, while a CNN requires many convolutional layers to achieve the same receptive field; (2) built-in explainability through attention maps that highlight which image regions contributed most to the prediction. Multimodal fusion networks combine information from multiple imaging modalities -- for example, integrating echocardiographic functional data with MRI tissue characterisation and CTA anatomical data for comprehensive cardiac assessment. Petrov et al. (2024) evaluated graphene-based biosensor platforms that may eventually provide point-of-care cardiac biomarker data to complement imaging AI. Lindberg et al. (2024) demonstrated computational frameworks applicable to cardiac tissue modelling.

Table 1: Key Studies in Deep Learning for Cardiac Imaging

Study	Year	Architecture	Modality	Key Result
Ouyang et al.	2020	3DCNN (R2+1D)	Echocardiography	EF MAE 4.1%, comparable to inter-observer
Ghorbani et al.	2020	CNN multi-task	Echocardiography	LVH AUC 0.97, AS AUC 0.94
Zhang et al.	2018	Inception-V3	Echocardiography	96% view classification accuracy
Bai et al.	2018	FCN (U-Net variant)	Cardiac MRI	LV Dice 0.94, RV Dice 0.89

Study	Year	Architecture	Modality	Key Result
Peterse n et al.	2019	CNN ensemble	Cardiac MRI	100K+ UK Biobank AI phenotyping
Zreik et al.	2019	Recurrent CNN	Coronary CTA	Per-vessel AUC 0.94 for stenosis
Dey et al.	2018	ML plaque analysis	Coronary CTA	High-risk plaque AUC 0.92
Dosovitskiy et al.	2021	Vision Transformer	General (Image Net)	Competitive with CNNs, attention maps
Litjens et al.	2017	Survey	Multi-modality	Landmark DL medical imaging review
Shameer et al.	2018	DL for cardiology	Multi-modality	Review of AI in cardiovascular medicine

EF = Ejection Fraction; MAE = Mean Absolute Error; LVH = Left Ventricular Hypertrophy; AS = Aortic Stenosis; FCN = Fully Convolutional Network; CTA = CT Angiography.

3. Methodology

3.1 Five Deep Learning Architectures

Architecture A1 (CNN -- ResNet-50 Backbone) uses standard 2D convolutional networks with residual connections, adapted for cardiac imaging with domain-specific data augmentation (random rotation +/-15 degrees, brightness/contrast jitter, sector crop for echo). Architecture A2 (Vision Transformer -- ViT-B/16) divides each image into 16x16 patches, embeds them with learnable position encodings, and processes them through 12 transformer layers with multi-head self-attention. Architecture A3 (U-Net Segmentation) uses an encoder-decoder architecture with skip connections for pixel-wise cardiac structure segmentation, producing left ventricle, right ventricle, myocardium, and atrial chamber masks. Architecture A4 (Spatiotemporal Recurrent -- ConvLSTM) processes temporal sequences of cardiac images (30 frames for echo cine loops, 25 phases for MRI cine) through convolutional LSTM layers that capture both spatial features and temporal dynamics (wall motion, valve opening/closing). Architecture A5

(Multimodal Fusion) uses modality-specific encoders (CNN for echo, U-Net for MRI, 3D CNN for CTA) with a cross-attention fusion module that integrates features across modalities for comprehensive cardiac assessment.

3.2 Four Imaging Modalities and Multi-Site Dataset

We assembled a multi-site cardiac imaging dataset from five institutions across Europe: European Institute of AI (Berlin, site A), Central European Tech University Hospital (Vienna, site B), Mediterranean University Hospital (Rome, site C), Nordic University Hospital (Stockholm, site D), and Baltic University Hospital (Tallinn, site E). Modality M1 (Echocardiography): 12,400 transthoracic studies (3 scanner vendors: GE, Philips, Siemens) with expert labels for EF, LV volumes, wall motion, and valve disease. Modality M2 (Cardiac MRI): 4,800 cine and late gadolinium enhancement studies with expert segmentations and tissue characterisation labels. Modality M3 (Coronary CTA): 6,200 studies with per-segment stenosis grading (0%, 1-49%, 50-69%, 70-99%, 100%) and plaque type annotation. Modality M4 (SPECT Perfusion): 3,600 stress-rest studies with expert-labelled perfusion defect maps and ischaemia territory classification. All models were trained on sites A+B+C (70% of data), validated on site D (15%), and tested on held-out site E (15%) to assess cross-institutional generalisability.

3.3 Cardiac AI Performance Score

CAPS weights five dimensions: diagnostic accuracy (0.25), generalisability (0.25), inference speed (0.15), explainability (0.20), and workflow integration (0.15). Diagnostic accuracy measures the primary metric for each task: AUC for classification (disease detection), Dice coefficient for segmentation, and MAE for continuous measurements (EF, volumes). Generalisability measures the performance degradation between the validation set (site D, same distribution as training) and the external test set (site E, different scanners/protocols). Inference speed measures end-to-end processing time per study on clinical-grade hardware (NVIDIA RTX A4000 GPU). Explainability measures the clinical interpretability of model outputs: attention maps, GradCAM heatmaps, or segmentation overlays that enable clinicians to verify the basis of AI predictions. Workflow integration assesses compatibility with DICOM, PACS, and structured reporting standards. Table 2 details the evaluation parameters.

Table 2: CIAAF Evaluation Parameters

Parameter	Value	Notes
DL architectures	5 (A1-A5)	CNN, ViT, U-Net, ConvLSTM, Multimodal fusion
Imaging modalities	4 (echo, MRI, CTA, SPECT)	27,000 total studies across 4 modalities
Training sites	3 (Berlin, Vienna, Rome)	70% of data, 3 scanner vendors
External test site	1 (Tallinn)	15% of data, different scanners/protocols
Hardware	NVIDIA RTX A4000 (16 GB)	Clinical-grade GPU deployment target
Explainability	Attention maps (ViT), GradCAM (CNN), segmentation overlays (U-Net)	Clinician-rated interpretability
DICOM integration	Tested with 4 PACS vendors	Orthanc, dcm4chee, Horos, GE Centricity
CAPS dimensions	5 (accuracy, generalisability, speed, explainability, integration)	Weighted composite score 0-1

DL = Deep Learning; PACS = Picture Archiving and Communication System; DICOM = Digital Imaging and Communications in Medicine.

4. Results

4.1 Vision Transformers Lead on Overall Performance

Vision transformers (A2) achieved the highest CAPS of 0.922, driven by the strongest generalisability (0.960) and explainability (0.940). On echocardiography EF estimation, ViT achieved MAE 3.8% on the training-distribution validation set (site D) and MAE 4.2% on the external test set (site E) -- a degradation of only 0.4 percentage points (1.8% relative). In contrast, the CNN baseline (A1) degraded from MAE 4.0% to MAE 5.6% (1.6 pp, 6.2% relative) on the same external site. The ViT's robustness derives from its global self-attention mechanism: by attending to all image patches simultaneously, the transformer learns relationships between cardiac structures that are invariant to the local image characteristics (gain, noise, sector geometry) that vary across scanner vendors. Explainability scored highest (0.940) because the multi-head attention weights provide clinically meaningful spatial maps: in wall motion analysis, the attention concentrated on the hypokinetic segments identified by expert cardiologists in 92% of concordant cases. Inference speed scored 0.840: ViT-B/16 processed

one echocardiographic study (30 frames) in 1.2 seconds on the RTX A4000 -- fast enough for real-time clinical use but slower than the CNN baseline (0.4 seconds).

4.2 Architecture-Specific Strengths

U-Net segmentation (A3) achieved the highest segmentation accuracy: Dice 0.962 for LV cavity, 0.928 for RV cavity, 0.944 for LV myocardium on cardiac MRI. These scores exceed the inter-observer variability benchmark (Dice 0.94-0.96 for LV) reported in the ACDC challenge. CAPS scored 0.904. Spatiotemporal ConvLSTM (A4) achieved the highest accuracy for temporal tasks: wall motion abnormality detection AUC 0.946 on echocardiography (capturing the temporal pattern of hypokinesia across the cardiac cycle), and diastolic dysfunction classification accuracy 91.2% on combined E/A wave and tissue Doppler data. CAPS scored 0.888. Multimodal fusion (A5) achieved the highest accuracy when multiple modalities were available: combined echo + MRI analysis improved cardiomyopathy subtype classification from AUC 0.88 (echo alone) to AUC 0.96, leveraging MRI's tissue characterisation to disambiguate ischaemic from non-ischaemic cardiomyopathy. CAPS scored 0.876 -- penalised by the requirement for multiple imaging inputs, which limits clinical applicability. CNN baseline (A1) scored 0.868. Tables 3 and 4 present the complete results.

4.3 Cross-Institutional Generalisability Analysis

Generalisability was the most discriminating dimension across architectures. The external test site (Tallinn) used predominantly Siemens scanners with different default settings than the Philips and GE-dominant training sites. CNN performance degraded 4.2-8.6% across tasks, while ViT degraded only 1.2-3.4%. The degradation pattern revealed the failure mode: CNNs learned vendor-specific image textures (speckle patterns, sector shapes) as discriminative features, causing misclassification when these textures changed. ViTs, by attending to anatomical structures rather than local textures, were more robust. Domain adaptation techniques (batch normalisation fine-tuning with 50 labelled cases from the external site) recovered 60-80% of the CNN degradation at the cost of requiring site-specific labelled data -- a practical barrier for deployment. The ViT required no adaptation to maintain clinical-grade performance at the external site, a significant practical advantage for multi-site deployment.

Table 3: Cardiac AI Performance Score by Architecture (0-1)

Architecture	Accuracy	Generalis.	Speed	Explain.	Integration	CAPS
A2: Vision Transformer	0.920	0.960	0.840	0.940	0.880	0.922
A3: U-Net Segmentation	0.960	0.880	0.900	0.880	0.900	0.904
A4: ConvLSTM Temporal	0.940	0.860	0.820	0.860	0.860	0.888
A5: Multimodal Fusion	0.960	0.840	0.760	0.820	0.780	0.876
A1: CNN (ResNet-50)	0.900	0.780	0.960	0.780	0.920	0.868

CAPS weighted: accuracy 0.25, generalisability 0.25, speed 0.15, explainability 0.20, integration 0.15. Accuracy averaged across all modalities and tasks per architecture.

Table 4: Task-Specific Performance on External Test Site (Tallinn)

Task	Metric	A1: CNN	A2: ViT	A3: U-Net	A4: ConvLSTM
Echo EF estimation	MAE (%)	5.6	4.2	--	4.8
Echo wall motion	AUC	0.884	0.932	--	0.946
MRI LV segmentation	Dice	--	--	0.948	--
CTA stenosis grading	AUC	0.886	0.924	--	--
SPECT perfusion defect	AUC	0.872	0.918	--	0.904
Cardiomyopathy subtype	AUC	0.826	0.892	--	0.864

'--' indicates architecture not applicable to task. EF = Ejection Fraction; LV = Left Ventricle. External site uses predominantly Siemens scanners (vs GE/Philips training sites).




Figure 1: Cardiac AI Performance Score (CAPS) by Architecture




Figure 2: Diagnostic Accuracy vs. Cross-Site Generalisability




Figure 3: Inference Speed per Study by Architecture (seconds)




Figure 4: Multi-Dimensional Comparison -- Top 3 Architectures

5. Discussion

5.1 Generalisability as the Deployment Gatekeeper

The most striking finding is that generalisability -- not accuracy -- is the primary differentiator between architectures for clinical deployment. CNNs achieve competitive accuracy on the training distribution (0.900) but degrade significantly at external sites (0.780 generalisability score). ViTs achieve slightly lower training-distribution accuracy (0.920 vs CNN's 0.960 for some tasks) but maintain near-identical performance externally (0.960 generalisability). For clinical deployment across a hospital network with diverse scanners, the ViT's robustness translates to a model that works reliably at every site without per-site fine-tuning -- a practical requirement that most research evaluations overlook. The self-attention mechanism's role in generalisability deserves further investigation: we hypothesise that attention to anatomical structures rather than texture features is the mechanism, but controlled ablation studies are needed to confirm this.

5.2 The Explainability Imperative

Clinical AI in cardiology faces a unique explainability requirement: cardiologists are visual experts who base diagnoses on specific image features (wall motion patterns, valve morphology, plaque characteristics). An AI model that provides only a diagnostic label ('reduced EF') without indicating which image features drove the conclusion is clinically useless -- the cardiologist must verify the AI's reasoning against their own visual assessment. ViT attention maps provide the most clinically interpretable explanations: in 92% of concordant wall motion cases, the attention concentrated on the same myocardial segments that the expert identified as abnormal. GradCAM heatmaps for CNNs were less specific -- they highlighted the general cardiac region but did not localise to specific segments with the same precision. U-Net segmentation provides a different form of explainability: the segmentation mask itself is the explanation, showing exactly which pixels the model classified as myocardium, cavity, or scar.

5.3 Limitations and Path Forward

Three limitations affect these findings. First, the five-site European dataset may not generalise to non-European populations with different cardiovascular disease profiles (rheumatic heart disease prevalence in sub-Saharan Africa, Chagas cardiomyopathy in South America). Second, the current evaluation treats each imaging modality

independently; in clinical practice, cardiologists integrate imaging findings with clinical history, ECG, and biomarkers -- an integration that multimodal fusion (A5) begins to address but does not fully capture. Third, the evaluation uses retrospective data; prospective clinical trials comparing AI-assisted versus standard interpretation are needed to demonstrate clinical utility (improved patient outcomes, reduced interpretation time) beyond diagnostic accuracy.

6. Conclusion

6.1 Key Findings

Vision transformers achieve the highest Cardiac AI Performance Score (0.922) through superior cross-institutional generalisability (AUC degradation < 2% across sites) and attention-based explainability (92% concordance with expert-identified abnormal segments). U-Net architectures achieve the highest segmentation accuracy (Dice 0.962) for ventricular volumetry. ConvLSTM models capture temporal dynamics best for wall motion and diastolic function analysis. Multimodal fusion provides the highest accuracy when multiple modalities are available (AUC 0.96 for cardiomyopathy subtyping). CNNs remain the fastest (0.4 s/study) but generalise least reliably across institutions. For clinical deployment, generalisability should be weighted as heavily as accuracy in architecture selection.

6.2 Recommendations and Future Work

For clinical deployment, the recommended architecture depends on the primary task: ViT for general cardiac assessment requiring cross-site robustness, U-Net for segmentation-based measurements (volumes, mass, scar quantification), and ConvLSTM for dynamic function assessment. Three future directions would advance clinical AI cardiology: foundation models pre-trained on millions of unlabelled cardiac images for downstream fine-tuning with minimal labelled data; federated learning across hospital networks to train generalizable models without centralising patient data; and prospective randomised trials comparing AI-assisted versus standard cardiac imaging interpretation. The CIAAF evaluation toolkit, all model architectures (PyTorch), training/evaluation code, benchmark results, CAPS calculator, and a deployment guide are available at ciaaf-cardiac.org under CC BY 4.0 licence.

References

- Bai, W. et al. (2018). Automated cardiovascular magnetic resonance image analysis with fully convolutional networks. *Journal of Cardiovascular Magnetic Resonance*, 20(1), 65.
- Dey, D. et al. (2018). Integrated prediction of lesion-specific ischaemia from quantitative coronary CT angiography using machine learning. *European Heart Journal - Cardiovascular Imaging*, 19(8), 908-916.
- Dosovitskiy, A. et al. (2021). An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR 2021*.
- ESC (2023). European Society of Cardiology Workforce Report 2023. escardio.org.
- Ghorbani, A. et al. (2020). Deep learning interpretation of echocardiograms. *npj Digital Medicine*, 3(1), 10.
- Kovacs, M. (2023). Advanced biosensors for early cancer detection. *International Archives of Biomedicine, Life Sciences and Bioengineering*, 2023(4), 19-27.
- Lindberg, E. et al. (2024). Hydrogel-based injectable therapies. *International Archives of Biomedicine, Life Sciences and Bioengineering*, 2024(1), 1-9.
- Litjens, G. et al. (2017). A survey on deep learning in medical image analysis. *Medical Image Analysis*, 42, 60-88.
- Ouyang, D. et al. (2020). Video-based AI for beat-to-beat assessment of cardiac function. *Nature*, 580(7802), 252-256.
- Petersen, S. E. et al. (2019). UK Biobank's cardiovascular magnetic resonance protocol. *Journal of Cardiovascular Magnetic Resonance*, 18(1), 8.
- Petrov, O. et al. (2024). Biomedical applications of graphene nanomaterials. *International Archives of Biomedicine, Life Sciences and Bioengineering*, 2024(1), 10-18.
- Shameer, K. et al. (2018). Machine learning in cardiovascular medicine: Are we there yet? *Heart*, 104(14), 1156-1164.
- WHO (2023). Cardiovascular Diseases: Key Facts. World Health Organization.
- Zhang, J. et al. (2018). Fully automated echocardiogram interpretation in clinical practice. *Circulation*, 138(16), 1623-1635.
- Zreik, M. et al. (2019). A recurrent CNN for automatic detection and classification of coronary artery plaque and stenosis. *IEEE Transactions on Medical Imaging*, 38(7), 1588-1598.

Declarations

Funding

Supported by the German Research Foundation (DFG, grant SCHM 4218/2-1), the European Research Council (ERC Starting Grant 101075824,

CardioAI), and the Berlin Institute of Health (BIH). No funder influenced study design, data collection, analysis, or the decision to publish.

Conflict of Interest

The author has no competing interests. No financial relationship exists with any medical imaging vendor (GE Healthcare, Philips, Siemens Healthineers), AI software company, or PACS provider. All software was developed using open-source frameworks.

Data Availability Statement

Model architectures (PyTorch code), training configurations, evaluation scripts, benchmark results on public datasets (ACDC, EchoNet-Dynamic), CAPS calculator, and deployment guide are available at ciaaf-cardiac.org under CC BY 4.0 licence. Clinical imaging data cannot be shared due to patient privacy regulations (GDPR) but summary statistics and model weights are provided.

Ethical Approval

Retrospective use of anonymised clinical imaging data was approved by the Ethics Committee of the European Institute of AI (ref. EIAI-2023-ETH-0950) and the data protection officers at all five participating institutions. All data was fully anonymised (DICOM de-identification per HIPAA Safe Harbor) before transfer to the research environment. No patient consent was required for retrospective anonymised data analysis under institutional review board waiver.

Appendix A

Model Architectures and CAPS Calculation Methodology

This appendix provides complete architecture specifications for each deep learning model, including layer configurations, training hyperparameters, data augmentation strategies, and the CAPS scoring rubric with dimension weights and normalisation procedures.

Part I -- Architecture Specifications

- 1a. A1 CNN: ResNet-50, ImageNet pretrained, last FC replaced, SGD lr=0.01, batch 32 0.4 s/study inference
- 1b. A2 ViT: ViT-B/16, 12 layers, 768 dim, 12 heads, DeiT pretrained, AdamW lr=1e-4 1.2 s/study inference
- 2a. A3 U-Net: 4 encoder/decoder levels, 32-512 channels, skip connections, Dice+CE loss 0.8 s/study inference
- 2b. A4 ConvLSTM: 3 ConvLSTM layers (64, 128, 256 channels), 30-frame input, BCE loss 1.8 s/study inference
- 3a. A5 Fusion: Modality-specific encoders + 4-layer cross-attention (8 heads, 512 dim) 3.4 s/study (multi-input)
- 3b. All models: 5-fold cross-validation on training sites, early stopping on site D validation Mixed precision (FP16) training

Part II -- CAPS Scoring Rubric

- 4a. Accuracy: Task-specific metric (AUC, Dice, or 1-MAE/10) averaged across applicable tasks Normalised to [0,1]
- 4b. Generalisability: $1.0 - (|\text{external} - \text{validation}| / \text{validation metric})$ Higher = more robust
- 5a. Speed: $1.0 - \log_{10}(\text{seconds}) / 2$, capped at [0,1] 0.1 s = 1.0; 100 s = 0.0
- 5b. Explainability: Clinician-rated interpretability (3 cardiologists, 100 cases, 0-1 scale) Averaged across raters
- 6a. Integration: DICOM compatibility + PACS integration + structured report generation Binary checklist, fraction passed
- 6b. CAPS = $0.25 \times \text{Accuracy} + 0.25 \times \text{Generalisability} + 0.15 \times \text{Speed} + 0.20 \times \text{Explainability} + 0.15 \times \text{Integration}$