

# Tagged-Informed Prior for Motion Quantification in Cine CMR Using Implicit Neural Representations

**Laura Alvarez-Florez**<sup>1,2,5</sup>

L.ALVAREZFLOREZ@AMSTERDAMUMC.NL

**Soufiane Ben Haddou**<sup>1,2,5</sup>

S.BENHADDOU@AMSTERDAMUMC.NL

**Fleur V. Y. Tjong**<sup>3,5</sup>

F.V.TJONG@AMSTERDAMUMC.NL

**Ivana Išgum**<sup>1,4,5,6</sup>

ISGUM.IVANA@MAYO.EDU

<sup>1</sup> *Department of Biomedical Engineering and Physics, Amsterdam University Medical Center, The Netherlands*

<sup>2</sup> *Informatics Institute, University of Amsterdam, The Netherlands*

<sup>3</sup> *Heart Center, Department of Clinical and Experimental Cardiology, Amsterdam University Medical Center, The Netherlands*

<sup>4</sup> *Department of Radiology and Nuclear Medicine, Amsterdam University Medical Center, The Netherlands*

<sup>5</sup> *Amsterdam Cardiovascular Sciences, Amsterdam, The Netherlands*

<sup>6</sup> *Department of Radiology, Mayo Clinic, Rochester, United States of America*

**Editors:** Under Review for MIDL 2026

## Abstract

Accurate quantification of myocardial motion from cine cardiac magnetic resonance (CMR) is essential for assessing cardiac function. Although tagged CMR provides high-fidelity measurements of myocardial deformation, its longer acquisition time limits routine clinical use, making cine CMR motion estimation the more widely applicable approach. Implicit neural representations (INRs) offer a promising framework for cine-based motion estimation by modelling cardiac motion as a continuous spatiotemporal function. However, they require subject-specific optimisation and are sensitive to initialization, leading to slow convergence. Furthermore, optimisation from random initialization can lead to large number of solutions that may not guarantee biomechanically plausible motion. To address these limitations, we propose a strategy to improve and accelerate INR-based registration of cine CMR by leveraging a population-level prior derived from tagged CMR data. First, we train subject-specific INRs on the tagged cine dataset to encode characteristic myocardial deformation patterns. Second, we aggregate their parameters across subjects to form a tagged-informed population prior. Third, we use this prior initialization to warm-start the optimization of cine INRs. The resulting prior provides a physiologically meaningful starting point for cine-only INR optimisation, reducing the search space and promoting more realistic cardiac motion. We develop and test the method on the UK Biobank. Compared with standard initialization, the proposed prior enables the INR to reach near-optimal performance using only half as many optimisation steps, achieving a 4% improvement in Dice and a 15% reduction in Hausdorff distance. These gains also translate to another test set of 855 subjects from a different institution, encompassing different pathologies, where the prior yields smoother and more physiologically plausible strain curves. To support reproducibility, upon publication and approval, we will publicly release the prior weights and accompanying code.

**Keywords:** Cardiac MRI, Tagged CMR, Implicit Neural Representations, Image Registration, Cardiac Motion Quantification, Prior Initialization

## 1. Introduction

Characterizing myocardial motion is a key component in the assessment of cardiovascular disease. Quantitative descriptors such as ejection fraction and myocardial strain provide insight into global and regional contractility, supporting risk stratification and treatment planning across a range of cardiovascular diseases (Lange and Schuster, 2021; Fudim et al., 2021). Obtaining motion measurements, however, requires precise estimation of myocardial deformation throughout the cardiac cycle, which remains a challenge.

Dedicated sequences such as tagged Cardiac Magnetic Resonance (CMR) provide high-fidelity measurements of regional myocardial deformation and are widely regarded as the reference standard for non-invasive assessment of cardiac motion. However, they require additional scanning time and specialised acquisition protocols, which restrict their use in routine clinical workflows (Ibrahim, 2011). As a result, there is strong interest in deriving motion and strain directly from standard cine CMR routinely acquired in clinic.

Image registration methods have long been used to assess cardiac motion by aligning myocardial anatomy across the cardiac cycle (Zhang et al., 2024). Deep learning (DL) methods build on this principle by learning the registration mapping directly from data, enabling faster inference and robustness to image noise. Most DL-based motion estimation frameworks adopt a registration formulation in which consecutive cine frames are aligned by maximising image similarity (De Vos et al., 2019). Registration methods based on convolutional networks have shown strong performance in this setting (Upendra et al., 2020; Morales et al., 2021; Qin et al., 2020), but they operate on discrete voxel grids, limiting the spatial and temporal continuity of the estimated motion fields. Implicit neural representations (INRs) have recently emerged as an alternative (Wolterink et al., 2022) and have outperformed other registration methods for cardiac motion quantification (López et al., 2023). INRs model motion as a continuous function of spatial and temporal coordinates, fitting a small multilayer perceptron to each subject’s cine sequence. Although this yields highly smooth and differentiable displacement fields, a major practical limitation is that INRs must be optimised from scratch for every new subject, resulting in long inference times that hinder clinical usability. Moreover, image registration remains an ill-posed problem, and optimisation guided mainly by image similarity does not guarantee that the estimated cardiac motion is biomechanically plausible.

Recent work suggests that initialization plays a crucial role in addressing these challenges, both by reducing the per-subject optimisation burden and by guiding the network toward more meaningful solutions (Koneputugodage et al., 2025). Initializing the INR closer to a physiologically plausible deformation, instead of starting from random weights, can accelerate convergence and help avoid unrealistic motion estimates. Our prior work demonstrated that, when estimating motion between consecutive cine frames, warm-starting each INR from the previous time point in the cardiac cycle speeds up convergence and improves registration quality (Alvarez-Florez et al., 2024). In this work, we extend this idea by constraining the INR optimisation space using high-fidelity measurements of myocardial deformation encoded in tagged images, with the aim of promoting more biomechanically plausible motion estimates and faster convergence for cine CMR registration. Due to the lower resolution of the tagged dataset, we first create a dataset that combines these with cine images to increase the through plane resolution. We use a SIREN as our baseline INR ar-

chitecture, which has demonstrated strong performance in medical image registration tasks (Wolterink et al., 2022). Then, we train these INRs on the tagged and cine data to embed characteristic deformation patterns within their parameters, and aggregate these parameters across subjects to construct a tagged-informed population prior. This prior serves as a physiologically grounded initialization that reduces the optimisation search space, accelerates convergence, and improves the quality of motion estimated from cine CMR. Hence, the main contributions of this work are:

- We introduce the first population-level initialization for INR-based cardiac motion estimation, constructed from tagged and cine CMR. The resulting initialization is lightweight and model-agnostic, with potential to be used in any INR framework for cine motion quantification.
- We show that by using the proposed prior we reduce the INR’s required optimization time by half compared to INRs with a conventional initialization scheme.
- We show that the method generalizes in an independent set from a different center that includes patients with different pathologies unseen by the prior.
- We will release upon publication and approval the prior model weights and accompanying code to facilitate reproducibility and future research.

## 2. Method

We propose a strategy to improve and accelerate INR-based registration of cine CMR by leveraging a population-level prior derived from tagged CMR data (Fig. 1). First, we present the INR method for image registration. Second, we create a tagged-cine dataset and use it to form a population-average prior initialization. Third, we use this prior initialization to warm-start the optimization of cine INRs.

### 2.1. Implicit Neural Representation for Cardiac Motion Quantification

To estimate continuous cardiac motion over the cardiac cycle, we register each cine CMR frame in the cardiac cycle to the final frame of the cine sequence at end-diastole. The registration model is an INR consisting of a multilayer perceptron with sinusoidal activations. (Sitzmann et al., 2020) The multilayer perceptron architecture consist of four input nodes, three fully connected layers of width 256, and three output nodes.

The INR models a continuous four-dimensional displacement field over space and time. Given a spatial location  $x = (x, y, z)$  and a temporal coordinate  $t$ , the network predicts the displacement vector  $u$ :

$$u(x, t) = f_{\theta}(x, t), \quad (1)$$

where  $f$  denotes the INR with trainable parameters  $\theta$ . This formulation provides a smooth and differentiable representation of cardiac motion that enables estimation of displacement at arbitrary spatial locations and time points (López et al., 2023).

Prior to training, each cine volume is cropped to a segmentation-derived bounding box around the heart, and resampled into a canonical orientation to ensure spatial consistency

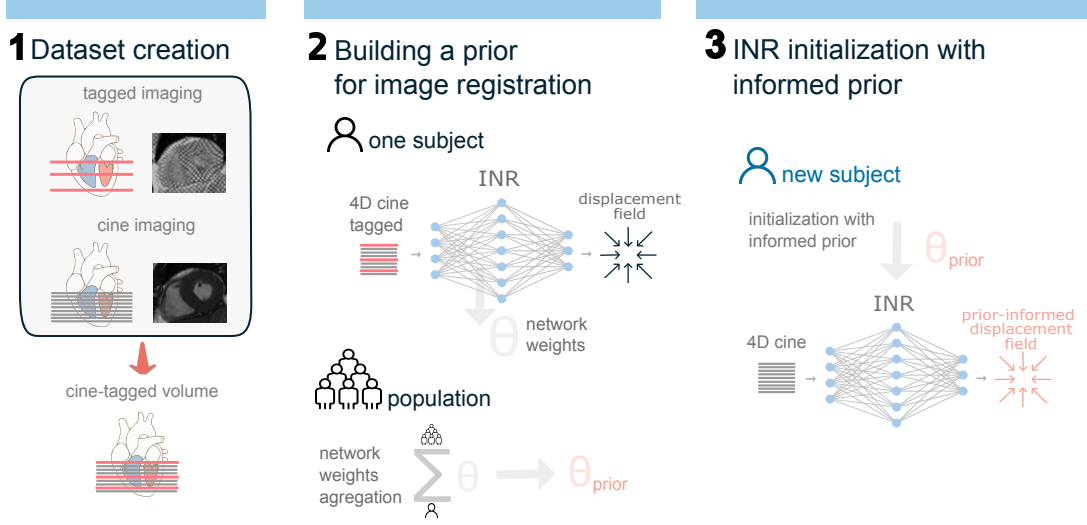


Figure 1: Graphical representation of the proposed tagged-informed registration prior.

across subjects. The input coordinates are normalized to  $[-1, 1]$  and uniformly sampled from the cine volume in batches of 10,000 points. The resulting deformation field  $\phi$  represented by the network is expressed through the transformation:

$$\phi(x, t) = x + u(x, t), \quad (2)$$

which maps coordinates from time  $t$  into the reference frame.

During training, the moving cine frames  $I^t$  are warped using the predicted transformation to obtain the corresponding warped frame  $\hat{I}^t$  at time  $t$ :

$$\hat{I}^t(x) = I^t(\phi(x, t)) \quad (3)$$

The warped image  $\hat{I}^t$  is then compared with the reference  $I^{\text{ref}}$  using a normalized cross-correlation (NCC) image similarity loss. To encourage physiologically plausible motion and volume preservation, the loss includes a Jacobian-based regularization term within the myocardium:

$$\mathcal{L}_{\text{Jac}} = \frac{1}{N} \sum_x [(\det J(x) - 1) + (\det J(x)^{-1} - 1)], \quad (4)$$

where  $J = \nabla \phi$  and  $N$  represents the number of sampled coordinates. The total loss is therefore defined as:

$$\mathcal{L}_{\text{registration}} = \mathcal{L}_{\text{NCC}} + \alpha * \mathcal{L}_{\text{Jac}}, \quad (5)$$

where  $\alpha$  controls the strength of the deformation regularization.

After optimization, myocardial strain is derived from the displacement field by computing the Lagrangian strain tensor from the spatial gradients of the deformation field  $\phi$ . Radial and circumferential strain components are then obtained by projecting the strain tensor onto voxel-wise radial directions, defined from the endocardial surface normals, and circumferential directions, defined as the in-plane orthogonal vectors.

## 2.2. Learning a Population Prior from Tagged CMR

To learn a population prior, we first construct a tagged-cine dataset. Tagged CMR provides explicit information about myocardial deformation, making it well suited for learning a biomechanically informed prior. Unlike cine CMR, which typically consists of a full short-axis stack (12 slices, 2D+time), tagged CMR is limited to three 2D+time slices acquired at basal, mid, and apical levels. To capture the deformation information in tagged images without losing through-plane spatial context, we replace each cine slice into the corresponding spatial location of the paired subject’s tagged imaging. The resulting tagged–cine volumes are spatially and temporally aligned, preserving tag-derived deformation while providing the full anatomical context required for training the four-dimensional INRs.

For each subject in the tagged-cine dataset  $s = 1, \dots, N$ , we independently optimize an INR with identical architecture and training scheme as described in Section 2.1. All networks are initialized with the same random seed to ensure that differences in the learned parameters reflect subject-specific deformation patterns rather than initialization variability. After training, the parameter sets  $\theta^{(1)}, \dots, \theta^{(N)}$  are aggregated to form the population prior. For each parameter tensor, we compute the elementwise mean across each subject:

$$\theta_{\text{avg}} = \frac{1}{N} \sum_{s=1}^N \theta^{(s)}. \quad (6)$$

To reduce the influence of outlier parameters caused by noise or optimization instability, we clip each parameter tensor to the 1st and 99th empirical quantiles before averaging. The resulting averaged parameters  $\theta_{\text{avg}}$  capture the typical displacement behaviour across the tagged population. When applied to cine registration,  $\theta_{\text{avg}}$  is used to initialize the INR, after which the network performs image registration as described in Section 2.1. During optimization, all parameters remain fully trainable. Thus, the prior influences only the starting point of the optimization, not the registration formulation itself.

## 2.3. Dataset

For this study, we used two sources of CMR data. The UK Biobank dataset for prior construction and test, and an independent test set from Amsterdam University Medical Center (AUMC) for testing cross-dataset generalization of the prior.

**UK Biobank dataset.** We selected 1190 participants from the UK Biobank imaging dataset (Petersen et al., 2016) who underwent both cine and tagged CMR acquisitions during the same visit. We split this data into 890 for the development set, and two test sets. Test set 1 with 200 subjects for cine registration evaluation, and Test set 2 with 100 subjects for determining the optimal prior. Short-axis cine images contain up to 12 slices and 50 temporal frames per sequence. Tagged CMR images consist of three independent 2D+time slices located at basal, mid-ventricular, and apical levels. Segmentations of the left-ventricular (LV) myocardium, LV blood pool (LVBP), and right-ventricular (RV) cavity in cine CMR were generated automatically for this work using a deep learning model (Simonyan and Zisserman, 2014). These segmentations were transferred to the corresponding tagged slices using the affine transformation matrix from the image metadata.

**AUMC dataset.** For independent evaluation, we used a dataset of 855 patients referred for implantable cardioverter-defibrillator (ICD) implantation. The dataset consists of patients suffering from different cardiomyopathies, specifically: hypertrophic cardiomyopathy (HCM), dilated cardiomyopathy (DCM), and ischemic cardiomyopathy (ICM). Each subject underwent a standard short-axis cine CMR examination acquired using balanced steady-state free precession (bSSFP) under breath-hold conditions. A full SAX stack was available for each patient, typically comprising 10–25 slices with 15–25 temporal frames covering one cardiac cycle. Automated segmentations of the LV, LVBP and RV were obtained using an automated deep-learning segmentation pipeline (Sander et al., 2020). Use of this dataset was approved by the institution’s medical ethics committee.

## 2.4. Evaluation

To assess registration performance, we use standard segmentation-based overlap and distance metrics, together with deformation-based measures. All metrics are computed for the LV, RV, and LVBP and reported as the average across structures.

**Segmentation metrics.** We report three widely used measures for the evaluation of image registration through registration: the Dice similarity coefficient, quantifying volumetric overlap; the 95th percentile Hausdorff distance (HD95), capturing boundary alignment; and the average surface distance (ASD), reflecting overall surface agreement. Higher Dice and lower HD95/ASD values indicate better registration performance.

**Deformation metrics.** To characterize the smoothness and biomechanical plausibility of the deformation field, we compute the variability of the Jacobian determinant on the myocardium,  $|\det(J) - 1|$ , which penalizes local expansions or contractions. Lower values indicate more spatially regular, volume-consistent motion. Additionally, local non-invertible deformations are quantified using the folding ratio, defined as the proportion of voxels with negative Jacobian determinant. A lower folding ratio corresponds to more physically plausible deformations.

**Strain curves.** Quantitative geometric metrics describe local deformation properties, but they do not fully capture the temporal coherence or global physiological behaviour of cardiac motion. To complement these metrics, we examine the derived radial and circumferential strain trajectories, offering additional insight into myocardial mechanics.

## 3. Experiments and Results

First, we present the quantitative registration performance results on the UK Biobank and AUMC test sets. Second, we perform a qualitative analysis of the strain. Third, we present an ablation demonstrating the optimal prior settings used for our evaluation.

### 3.1. Results of Registration Performance

The cine registration performance was assessed on Test set 1 from the UK Biobank. Quantitative results are summarized in Table 1. The proposed method was compared with a baseline. The baseline consisted of the INR method initialized with Xavier initialization (Glorot and Bengio, 2010). Across all geometric metrics, the prior-initialized model

Table 1: Registration performance for the UK Biobank Test set 1 (mean  $\pm$  std).

Model	Dice ( $\uparrow$ )	HD95 ( $\downarrow$ )	ASD ( $\downarrow$ )	Jacobian Determinant ( $\downarrow$ )
Baseline	$0.733 \pm 0.097$	$1.933 \pm 1.063$	$0.536 \pm 0.252$	<b><math>0.151 \pm 0.022</math></b>
Prior	<b><math>0.763 \pm 0.073</math></b>	<b><math>1.642 \pm 0.760</math></b>	<b><math>0.461 \pm 0.180</math></b>	$0.170 \pm 0.024$

Table 2: Registration performance for the AUMC test set (mean  $\pm$  std).

Cohort	Model	Dice ( $\uparrow$ )	HD95 ( $\downarrow$ )	ASD ( $\downarrow$ )	Jacobian Determinant ( $\downarrow$ )
HCM	Baseline	$0.752 \pm 0.116$	$2.329 \pm 2.759$	$0.619 \pm 1.226$	<b><math>0.124 \pm 0.027</math></b>
	Prior	<b><math>0.777 \pm 0.099</math></b>	<b><math>2.034 \pm 2.639</math></b>	<b><math>0.548 \pm 1.208</math></b>	$0.135 \pm 0.033$
DCM	Baseline	$0.781 \pm 0.091$	$2.208 \pm 4.708$	$0.617 \pm 2.633$	<b><math>0.113 \pm 0.024</math></b>
	Prior	<b><math>0.798 \pm 0.080</math></b>	<b><math>2.042 \pm 4.667</math></b>	<b><math>0.573 \pm 2.592</math></b>	$0.120 \pm 0.028$
ICM	Baseline	$0.762 \pm 0.119$	$2.011 \pm 1.890$	$0.518 \pm 0.517$	<b><math>0.121 \pm 0.027</math></b>
	Prior	<b><math>0.772 \pm 0.133</math></b>	<b><math>1.850 \pm 1.932</math></b>	<b><math>0.494 \pm 0.665</math></b>	$0.128 \pm 0.030$

consistently outperformed the baseline. Dice similarity increased by approximately 4%, indicating improved volumetric alignment. Boundary-based metrics also showed substantial gains: HD95 decreased by roughly 15% and ASD by about 14%, reflecting more accurate contour alignment and smoother spatial correspondence. The folding ratio remained at zero for both models, confirming the absence of topological violations across all test cases. The Jacobian determinant increased (around 12%) when using the prior, reflecting the influence of the strong initialization.

In addition, to assess generalization across centers, scanners, and pathologies, we evaluated the model on the AUMC test set. Quantitative results for each cardiomyopathy are reported in Table 2. Across all cardiomyopathies, the prior-initialized INR outperformed the baseline on the segmentation metrics, mirroring the trends observed in the UK Biobank Test set 1. Improvements were most pronounced for the HCM group, where Dice increased by roughly 3% and both HD95 and ASD decreased noticeably. On the other hand, the Jacobian Determinant presented lower values for the baseline method. For the DCM and ICM groups, improvements were more modest but remained consistent across metrics.

### 3.2. Qualitative strain analysis

We conducted a qualitative evaluation to study the biomechanical plausibility, beyond deformation based metrics, of the derived strain trajectories across models. Figure 2 illustrates the qualitative strain patterns for representative DCM and HCM cases. The prior-initialized INR produces visibly smoother and more physiologically consistent radial and circumferential strain trajectories for both DCM and HCM cases. In particular, the prior reduces noise in the trajectory and preserves the expected systolic peak and diastolic relaxation pattern. Additionally, differences in trajectories between endocardial and epicardial strains are more consistent and plausible.



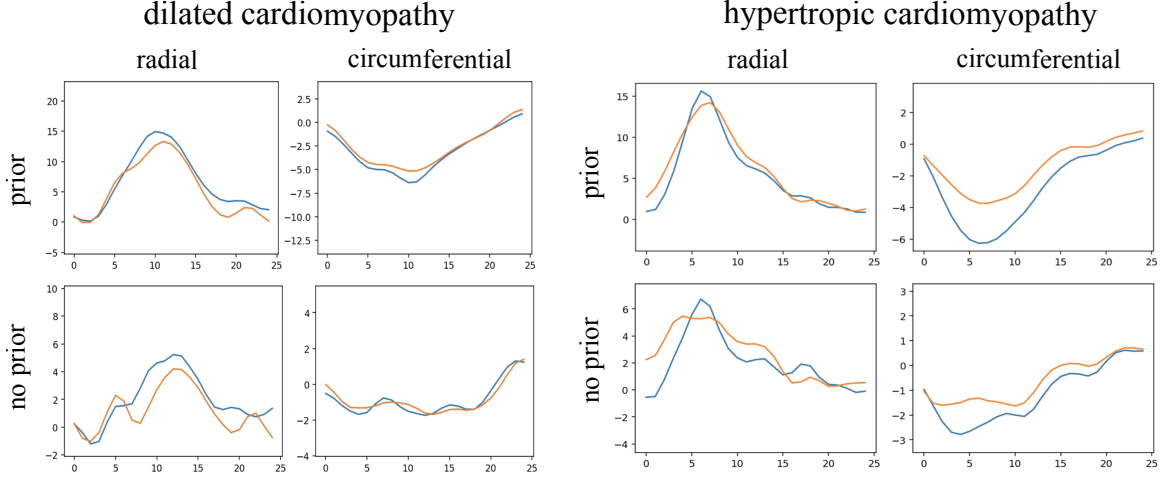


Figure 2: Radial and circumferential strain curves for a patient with dilated cardiomyopathy (DCM, left) and hypertrophic cardiomyopathy (HCM, right). Top row: registration using the tagged-informed INR prior. Bottom row: baseline INR without prior. The prior produces smoother and more physiologically plausible endocardial and epicardial strain trajectories. Endocardium is depicted in color blue and epicardium in orange.

Table 3: Results on the UK Biobank Test set 2 showing the registration performance (mean  $\pm$  std) across baseline method and priors derived from different population sizes. N refers to the number of subjects used to create the prior initialization. Arrows indicate whether higher ( $\uparrow$ ) or lower ( $\downarrow$ ) is better.

Model	Dice ( $\uparrow$ )	HD95 ( $\downarrow$ )	ASD ( $\downarrow$ )	Jacobian Determinant ( $\downarrow$ )
Baseline	$0.732 \pm 0.097$	$1.969 \pm 1.154$	$0.544 \pm 0.259$	<b><math>0.148 \pm 0.021</math></b>
Prior N=50	<b><math>0.763 \pm 0.074</math></b>	$1.678 \pm 0.882$	$0.466 \pm 0.190$	$0.169 \pm 0.023$
Prior N=200	$0.762 \pm 0.074$	$1.684 \pm 0.882$	$0.467 \pm 0.189$	$0.169 \pm 0.023$
Prior N=450	<b><math>0.763 \pm 0.074</math></b>	<b><math>1.677 \pm 0.875</math></b>	<b><math>0.465 \pm 0.189</math></b>	$0.169 \pm 0.023$
Prior N=890	$0.762 \pm 0.074$	$1.686 \pm 0.891$	$0.466 \pm 0.191$	$0.169 \pm 0.024$

### 3.3. Influence of Initialization Strategy

The ablation study evaluates how the size of the population used to construct the initialization prior, and the number of training epochs, affect registration performance. We compared four priors derived from populations of size  $N = 50, 200, 450, 890$  against the baseline model presented in the previous sections, and trained each configuration for 20, 50, 250, 500, and 1,000 epochs. This analysis was performed on Test set 2 from the UK biobank.

Figure 3 shows the mean Dice across subjects as a function of training epochs and prior size. All population priors lead to substantially faster performance gains than the baseline, with convergence curves that are nearly identical across prior sizes. This indicates that even small cohorts (e.g.,  $N = 50$ ) provide an initialization strong enough to guide the INR effectively, and increasing the population size does not meaningfully influence performance.






Figure 3: Effect of prior size and training epochs on registration performance, measured by mean Dice on UK Biobank Test set 2. Priors accelerate early performance gains, with only marginal improvements beyond approximately 500 epochs.

The training time analysis shows that after roughly 500 epochs, prior-initialized models reach performance comparable to or exceeding the baseline trained for 1,000 epochs.

To further investigate the effect of the size of the prior, we report registration accuracy across multiple metrics. Table 3 shows the results for different INR models trained for 500 epochs. All prior initializations provide a clear improvement over the baseline, with increases of approximately 3–4 percentage points in Dice, and consistent reductions in HD95 and ASD. The Jacobian determinant metric exhibits slightly higher variability for prior-based models, likely reflecting the stronger initial constraints imposed by the population initialization. The folding ratio, not presented on the table, remains near zero for all methods, indicating that none of the priors introduces folding artefacts. Differences between prior sizes remain insignificant and consistent across metrics.

#### 4. Discussion

In this work, we proposed a tagged-informed population prior for INR-based cine CMR registration and demonstrated that it substantially accelerates convergence while improving registration accuracy. By averaging the parameters of subject-specific INRs trained on tagged CMR, we obtained a population template that captures typical myocardial motion patterns and provides a strong initialization for subsequent cine registration. To our knowledge, this is the first study to exploit tagged CMR in this way, bridging its informative myocardial deformation to more widely available cine imaging.

The evaluation showed that the proposed prior consistently outperformed the baseline INR initialized with standard weights on the segmentation metrics. Convergence curves revealed that prior-based models reached near-optimal performance after about 500 epochs,

whereas the baseline required roughly twice as many epochs to reach comparable accuracy. These results suggest that the prior effectively guides optimization into a favourable region of the parameter space, reducing the burden of per-subject INR training. Future work should explore the practical implications of this improvement, including whether such reductions in computation time can help bridge the gap toward real-time cardiac motion estimation in clinical workflows.

The results generalized well to the test dataset, composed of cine CMR images from ICD patients, a different population than the one used to construct the prior, acquired at a different institution. Improvements were particularly consistent for HCM, suggesting that the tagged-informed prior can transfer beyond the population from which it was learned. In contrast, gains were more modest for DCM and ICM, whose deformation patterns are substantially altered by global or regional hypokinesia and scar tissue. Since the prior was derived from a general population, its inductive bias is naturally better aligned with preserved or hyperdynamic contraction, such as in HCM, and less representative of the severely abnormal motion commonly seen in DCM and ICM. This suggests that an important limitation of the current approach is that a population prior learned may not fully capture the whole spectrum of pathological motion. Future work should explore whether mixed-population priors better reflect the variability present across different cardiac conditions.

While the quantitative metrics evaluated in this study capture important aspects of registration performance, our qualitative analysis of myocardial strain revealed additional benefits of the proposed prior. The prior-initialized INRs produced smoother, more physiologically plausible strain trajectories, with clearer systolic peaks and more consistent separation between endocardial and epicardial layers. These findings suggest that the advantages of the proposed initialization extend beyond geometric accuracy, and influence the biomechanical fidelity of the estimated motion fields. Future work should investigate these effects in greater depth, including a characterization of strain-derived biomarkers and their potential clinical utility in detecting subtle abnormalities in myocardial mechanics.

Interestingly, the ablation study revealed that performance improvements saturated quickly with the size of the population used to build the cine tagged-informed prior. Even a prior constructed from only 50 subjects yielded substantial gains over the baseline. This robustness is encouraging from a practical standpoint, as centers with limited images available may be able to construct an effective prior. Future work could explore whether the number of subjects used to build the population prior can be further reduced and how such changes may affect the quality of the registration.

## 5. Conclusion

We introduced a cine tagged population prior for INR-based cine CMR registration. We showed that the proposed prior provides a strong and effective initialization for cine registration, accelerating training, and improving registration performance and strain estimation. Moreover, we demonstrated that indicating that the learned prior generalizes across institutions, offering a practical path towards making INR-based cardiac motion analysis more efficient and scalable.

## Acknowledgments

This research has been conducted using the UK Biobank Resource under application number 24711.

## References

- Laura Alvarez-Florez, Jörg Sander, Mimount Bourfiss, Fleur V. Y. Tjong, Birgitta K. Velthuis, and Ivana Išgum. Deep learning for automatic strain quantification in arrhythmogenic right ventricular cardiomyopathy. In Oscar Camara, Esther Puyol-Antón, Maxime Sermesant, Avan Suinesiaputra, Qian Tao, Chengyan Wang, and Alistair Young, editors, *Statistical Atlases and Computational Models of the Heart. Regular and CMR $\times$ Recon Challenge Papers*, pages 25–34, Cham, 2024. Springer Nature Switzerland. ISBN 978-3-031-52448-6.
- Bob D De Vos, Floris F Berendsen, Max A Viergever, Hessam Sokooti, Marius Staring, and Ivana Išgum. A deep learning framework for unsupervised affine and deformable image registration. *Medical image analysis*, 52:128–143, 2019.
- Marat Fudim, Frederik Dalgaard, Mouhammad Fathallah, Ami E Iskandrian, and Salvator Borges-Neto. Mechanical dyssynchrony: How do we measure it, what it means, and what we can do about it. *Journal of Nuclear Cardiology*, 28(5):2174–2184, 2021.
- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, 2010.
- El-Sayed H Ibrahim. Myocardial tagging by cardiovascular magnetic resonance: evolution of techniques—pulse sequences, analysis algorithms, and applications. *Journal of Cardiovascular Magnetic Resonance*, 13(1):36, 2011.
- Chamin Hewa Koneputugodage, Yizhak Ben-Shabat, Sameera Ramasinghe, and Stephen Gould. Vi<sup>3</sup>nr: Variance informed initialization for implicit neural representations. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 13477–13486, 2025.
- Torben Lange and Andreas Schuster. Quantification of myocardial deformation applying cmr-feature-tracking—all about the left ventricle? *Current Heart Failure Reports*, 18(4):225–239, 2021.
- Pablo Arratia López, Hernán Mella, Sergio Uribe, Daniel E Hurtado, and Francisco Sahli Costabal. WarpPINN: Cine-MR image registration with physics-informed neural networks. *Medical Image Analysis*, page 102925, 2023.
- Manuel A Morales, Maaïke Van den Boomen, Christopher Nguyen, Jayashree Kalpathy-Cramer, Bruce R Rosen, Collin M Stultz, David Izquierdo-Garcia, and Ciprian Catana. DeepStrain: a deep learning workflow for the automated characterization of cardiac mechanics. *Frontiers in Cardiovascular Medicine*, 8:730316, 2021.

- Steffen E Petersen, Paul M Matthews, Jane M Francis, Matthew D Robson, Filip Zemrak, Redha Boubertakh, Alistair A Young, Sarah Hudson, Peter Weale, Steve Garratt, et al. Uk biobank’s cardiovascular magnetic resonance protocol. *Journal of cardiovascular magnetic resonance*, 18(1):8, 2016.
- Chen Qin, Shuo Wang, Chen Chen, Huaqi Qiu, Wenjia Bai, and Daniel Rueckert. Biomechanics-informed neural networks for myocardial motion tracking in mri. In *International conference on medical image computing and computer-assisted intervention*, pages 296–306. Springer, 2020.
- Jörg Sander, Bob D de Vos, and Ivana Išgum. Automatic segmentation with detection of local segmentation failures in cardiac MRI. *Scientific Reports*, 10(1):21769, 2020.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020.
- Roshan Reddy Upendra, Brian Jamison Wentz, Suzanne M Shontz, and Cristian A Linte. A convolutional neural network-based deformable image registration method for cardiac motion estimation from cine cardiac mr images. In *2020 Computing in Cardiology*, pages 1–4. IEEE, 2020.
- Jelmer M Wolterink, Jesse C Zwienenberg, and Christoph Brune. Implicit neural representations for deformable image registration. In Ender Konukoglu, Bjoern Menze, Archana Venkataraman, Christian Baumgartner, Qi Dou, and Shadi Albarqouni, editors, *Proceedings of The 5th International Conference on Medical Imaging with Deep Learning*, volume 172 of *Proceedings of Machine Learning Research*, pages 1349–1359. PMLR, 06–08 Jul 2022.
- Qiang Zhang, Anastasia Fotaki, Sona Ghadimi, Yu Wang, Mariya Doneva, Jens Wetzl, Jana G Delfino, Declan P O’Regan, Claudia Prieto, and Frederick H Epstein. Improving the efficiency and accuracy of cardiovascular magnetic resonance with artificial intelligence—review of evidence and proposition of a roadmap to clinical translation. *Journal of Cardiovascular Magnetic Resonance*, 26(2):101051, 2024.