

Learning Cardiac Motion Priors for Implicit Neural Representations

Andrew Bell¹[0009–0005–2560–8172], George Webber¹[0009–0003–2573–7690],
Steffen E Petersen²[0000–0003–4622–5160], Andrew P King¹[0000–0002–9965–7015],
Muhummad Sohaib Nazir^{1,3}[0000–0003–2749–1033], and Alistair
Young¹[0000–0001–5702–4220]

¹ School of Biomedical Engineering and Imaging Sciences, King’s College London,
United Kingdom

² William Harvey Research Institute, Queen Mary University of London,
United Kingdom

³ Cardio-Oncology Service, Royal Brompton and Harefield Hospitals, London,
United Kingdom

Abstract. Implicit neural representations (INRs) are well suited to cardiac motion estimation, providing continuous, compact representations of motion fields. However, fitting an INR to each image sequence is time-consuming and sensitive to the optimisation trajectory. Learned priors can help guide optimisation towards plausible motion fields and enable faster adaptation, but learning priors for cardiac motion INRs remains under-explored. In this work, we compare four strategies for learning cardiac motion priors, including a population prior learned by joint optimisation, a consensus prior obtained by weight averaging, auto-decoders, and meta-learning. Using short-axis tagged cardiac magnetic resonance images from the UK Biobank, we evaluate their impact on tracking accuracy, motion behaviour, and adaptation trajectory.

All learned priors substantially improved early adaptation performance compared with random initialisation. While the simple consensus prior was effective, auto-decoders recovered large deformations faster during early adaptation. Meta-learning achieved strong early performance and maintained the best adaptation trajectory over 50 iterations. The code can be found at https://github.com/andrewjackbell/nvf_priors.

Keywords: Implicit Neural Representations · Cardiac Motion Estimation · Meta-learning.

1 Introduction

Implicit neural representations (INRs) are continuous parametrisations of signals using neural networks. INRs have been used to represent a variety of medical data, including images, anatomical shapes, segmentation labels, and physical fields such as motion [12]. INRs are well suited to cardiac motion estimation [4, 2, 6, 9], since physiological constraints allow compact low-dimensional

motion representations, despite motion being defined over high-dimensional image domains [7].

Estimating motion from images is an ill-posed inverse problem [8], requiring additional constraints to recover physiologically plausible solutions [18]. Furthermore, per-case INR optimisation is highly time-consuming and sensitive to the optimisation trajectory. This motivates the use of learned priors [3, 15, 1] which can rapidly recover plausible solutions. However, learning priors for cardiac motion fields remains relatively unexplored, with existing works relying on simple prior formulations [1] or providing only limited analysis of the resulting motion fields [5].

In this work, we investigate learning-based approaches for constructing cardiac motion priors for INRs, including population, consensus, auto-decoder and meta-learning formulations. Using short-axis tagged cardiac magnetic resonance images from the UK Biobank, we compare these formulations to basic initialisation strategies and evaluate their effects on few-step adaptation performance and optimisation trajectory. To our knowledge, this work provides the first systematic evaluation of prior learning strategies for INR-based cardiac motion estimation.

2 Related Work

2.1 Motion Estimation with INRs

An INR can be used to approximate a velocity field $v_\theta : \mathbb{R}^{d+1} \rightarrow \mathbb{R}^d$ mapping d -dimensional spatial coordinate X and time t to a velocity vector $v_\theta(X, t)$. Since ground truth motion is not typically available, previous works optimised network parameters θ according to inter-frame correspondence [18, 9, 10]. This assumes that an image frame I_t should be similar to some reference image I_r after warping according to the motion derived from $v(t)$. Thus, INR parameters θ can be learnt as

$$\theta^* = \arg \min_{\theta} \mathcal{L}_{sim}(I_t, I_r \circ v_\theta(t)), \quad (1)$$

where $\mathcal{L}_{sim}$ is an image similarity loss and $\circ$ denotes image warping. This objective does not uniquely identify a solution for θ^* , making the optimisation problem ill-posed. Moreover, in real imaging data, noise, artefacts, and violations of brightness constancy create a mismatch between the image similarity objective and the true underlying motion field. Regularisation terms based on smoothness or mechanics help constrain solutions to plausible motion fields [18], but balancing these terms against image observations remains challenging. Furthermore, fitting a new INR for every image sequence requires an expensive per-case optimisation procedure, which can become prohibitive for high-dimensional cardiac image data. This motivates learning a prior to improve the INR initialisation.

2.2 Learned Priors for Motion INRs

Early works fit INRs to cardiac motion from a random initialisation, using regularisation to encourage plausible solutions [4, 15]. Alvarez-Florez et al. [2] pro-

posed reusing parameters between successive motion fields, which accelerated fitting, but did not leverage population-level motion information. More recently, they learned a ‘consensus’ prior by averaging weights across INRs fit to a training set. This led to faster convergence in new cases, with improved motion plausibility [1]. These results show the benefit of a learned motion prior, but weight averaging is a simple heuristic that does not explicitly model the distribution of cardiac motion.

More principled approaches include auto-decoders [13], which use latent codes to represent inter-subject variation within a distribution, and meta-learning [16], which seeks an optimal initialisation for rapid adaptation to new signals. To our knowledge, meta-learning has not been investigated for INR-based cardiac motion priors, and no systematic comparison of these strategies has been reported.

3 Methods

3.1 Neural Velocity Field

We represent cardiac motion using a neural velocity field (NVF) based on previous work [9]. The NVF is an INR which takes spatial coordinates and time as input, and predicts an instantaneous spatial velocity vector. Formally, we represent d -dimensional cardiac motion as an NVF $v : \mathbb{R}^{d+1} \rightarrow \mathbb{R}^d$, and obtain Lagrangian motion $\varphi_{t_0 \rightarrow t}$ of coordinate X by integrating the ordinary differential equation

$$\frac{d\varphi_{t_0 \rightarrow t}(X)}{dt} = v(\varphi_{t_0 \rightarrow t}(X), t), \quad \varphi_{t_0 \rightarrow t_0}(X) = X,$$

using a second-order Runge-Kutta method with three sub-steps per frame interval. The resulting trajectories define the displacement field between any pair of time points. To encourage periodicity of cardiac motion, we map each time point to a periodic encoding

$$\tau(t) = (\sin(\theta), \cos(\theta)) \quad \theta = \left(\frac{2\pi t}{T} \right),$$

where T is the total number of time frames [9].

3.2 Objective

To optimise the NVF v over an image sequence $I \in \mathbb{R}^{T \times H \times W}$ we assume in-plane appearance-correspondence and use a similarity objective as in Equation 1. The displacement fields required for image warping are obtained by integrating the NVF between adjacent time points. Specifically, we use normalised cross-correlation (NCC) as the similarity loss between the current frame and forward- and backward-warped adjacent frames:

$$\mathcal{L}_{\text{sim}} = -\text{NCC}(I_t, I_{t+1} \circ \varphi_{t+1 \rightarrow t}) - \text{NCC}(I_t, I_{t-1} \circ \varphi_{t-1 \rightarrow t}),$$

We also assume that cardiac motion is spatially and temporally smooth, and define a smoothing loss

$$\mathcal{L}_{\text{smooth}} = \|\nabla_X v\|_F^2 + \left\| \frac{\partial v}{\partial t} \right\|_2^2.$$

A further assumption is that the myocardium is approximately incompressible, and that the underlying motion field preserves local topology. Computing deformation gradients from integrated trajectories at each step is computationally expensive. We therefore employ a first-order approximation based on the instantaneous velocity Jacobian $\tilde{F} = \mathbb{I} + \Delta t \nabla_X v$, where $\mathbb{I}$ is the identity matrix and Δt is the temporal interval between adjacent frames. We then encourage volume and topology preservation with respective penalties:

$$\mathcal{L}_{\text{volume}} = (\det(\tilde{F}) - 1)^2, \quad \mathcal{L}_{\text{fold}} = \text{ReLU}(-\det(\tilde{F}))^2,$$

The final objective $\mathcal{L}$ is given by the weighted sum of the loss terms according to hyperparameters $\lambda = \{\lambda_{\text{sim}}, \lambda_{\text{smooth}}, \lambda_{\text{volume}}, \lambda_{\text{fold}}\}$.

3.3 Learned Prior Formulations

In this section, we abstract the NVF objective as a function and describe the formulation of each prior in this notation. Let $\mathcal{F}$ denote a distribution of functions and let $T = \{f_i^*\}_{i=1}^N$ be a training set sampled from $\mathcal{F}$. Each function $f_i^* : X \rightarrow Y$ is observed at coordinates $\{x_k\}_{k=1}^K \subset X$, and reconstruction quality is measured using an objective function $\mathcal{L} : Y \times Y \rightarrow \mathbb{R}$. The goal is to learn prior parameters that enable rapid adaptation to unseen functions sampled from $\mathcal{F}$.

Population Prior As a simple baseline, we define a population prior θ^P as the parameters obtained by jointly optimising a single INR across all training instances:

$$\theta^P = \arg \min_{\theta} \mathbb{E}_{f^* \sim \mathcal{F}} \left[\sum_{k=1}^K \mathcal{L}(f_{\theta}(x_k), f^*(x_k)) \right].$$

Consensus Prior We define a consensus prior θ^C as the mean of independently optimised parameters for each training function:

$$\theta^C = \frac{1}{N} \sum_{i=1}^N \arg \min_{\theta} \sum_{k=1}^K \mathcal{L}(f_{\theta}(x_k), f_i^*(x_k)).$$

Although parameter alignment is not guaranteed across independently trained INRs, consensus priors can nevertheless be meaningful when training starts from a common initialisation [1].

Auto-decoder Prior Each training instance is assigned a latent code $z_i \in \mathbb{R}^Z$. For each hidden layer l , a linear modulation network m_l maps the latent code to layer-wise scale and shift parameters, which modulate the periodic MLP

activations [11]. Network parameters θ , modulation parameters κ , and subject-latent matrix $z \in \mathbb{R}^{N \times \mathcal{Z}}$ are jointly optimised:

$$(\theta^*, z^*, \kappa^*) = \arg \min_{\theta, z, \kappa} \sum_{f_i^* \in T} \sum_{k=1}^K \mathcal{L}(f_{\theta, z_i, \kappa}(x_k), f_i^*(x_k)) + \lambda_z \|z_i\|_2^2,$$

where λ_z is a latent-space regularisation weight. During adaptation to an unseen case, the decoder and modulation parameters are fixed and a new latent code is optimised.

Meta-learned Prior We use second-order Meta-SGD to jointly learn an initialisation θ_0 and layer-wise adaptation step sizes $\gamma = \{\gamma_l\}_{l=1}^L$. Let θ^ϵ denote the parameters obtained after adapting θ^* to a function $f^* \sim \mathcal{F}$ for ϵ inner-loop adaptation steps. The objective is to minimise the post-adaptation loss:

$$(\theta^*, \gamma) = \arg \min_{(\theta^*, \gamma)} \mathbb{E}_{f^* \sim \mathcal{F}} \left[\sum_{k=1}^K \mathcal{L}(f_{\theta^\epsilon}(x_k), f^*(x_k)) \right].$$

During adaptation, parameters are updated according to

$$\theta_l^{(s+1)} = \theta_l^{(s)} - \gamma_l \nabla_{\theta_l^{(s)}} \mathcal{L}(f_{\theta^{(s)}}, f^*), \quad l = 1, \dots, L,$$

where s is the adaptation step and γ_l is the learnable step size associated with layer l .

4 Experiments

4.1 Data

The UK Biobank (UKB) is a population-based dataset with health, genomic, and imaging data. A subset of participants underwent cardiac magnetic resonance (CMR) imaging, details of which are given in [14]. In this work, we use short-axis tagged CMR, including apical, mid-ventricular and basal slices, each having 20 time frames. Reference point tracks and segmentations of the left ventricle (LV) myocardium for tagged CMR slices were obtained using semi-automated analysis (CIM Tag2D, University of Auckland) as described previously [6, 19]. Images and annotations were cropped to the left ventricle with 20 pixels of padding and rescaled to consistent dimensions. In this work, we sample slices randomly from existing data splits [6] which contain no subject overlap. From the respective sets we sample 500 training slices (469 subjects), 100 validation slices (98 subjects), and 100 test slices (93 subjects).

4.2 Implementation details

All priors initialised a SIREN network with three hidden layers of 256 nodes [17]. During training and adaptation, 512 random spatial coordinates were sampled per frame, yielding $K = 10,240$ spatio-temporal coordinates per iteration used

to fit the NVF. For meta-learning, an additional 2048 query coordinates per frame were sampled to evaluate post-adaptation performance. The NVF regularisation weights were set to $\lambda_{\text{smooth}} = 0.01$, $\lambda_{\text{volume}} = 10^{-3}$, and $\lambda_{\text{fold}} = 10^{-3}$; the latent regularisation weight was $\lambda_z = 10^{-4}$ for the auto-decoder. Adaptation learning rates were tuned on the validation set to minimise the objective after five adaptation steps, resulting in learning rates of 10^{-4} (population), 10^{-5} (consensus), and 3×10^{-2} (auto-decoder). Meta-learning was trained using $\epsilon = 5$ inner-loop adaptation steps to match the evaluation protocol, and learned its own layer-wise learning rates. To investigate the effect of latent dimensionality, the auto-decoder was trained with various latent dimensions $\mathcal{Z} \in \{32, 64, 128, 256, 512, 1024\}$.

4.3 Evaluation

Priors were evaluated by fitting an NVF to each test case, initialised from the corresponding prior. We compared performance after 5 iterations and the longer-term adaptation trajectory over 50 iterations. Performance was assessed using tracking accuracy, segmentation overlap, and mechanical plausibility. Specifically, the fitted NVFs were used to warp a set of reference points across the cardiac cycle, which were compared to manually tracked points using mean Euclidean distance (displacement error). Furthermore, the reference segmentation mask was warped according to the fitted NVFs and compared to manually segmented time-varying masks with Dice score. Mechanical plausibility was assessed using the mean of deformation Jacobian determinants and the mean Green–Lagrange strains. These are averaged over all points in the left ventricular myocardium, and strains are rotated to circumferential and radial directions for clinical interpretation.

5 Results and Discussion

Implicit neural representations are often limited by the lengthy iterative process required to fit data. Learning a prior from training data can overcome this limitation by initialising the neural network with parameters that allow fast adaptation to new data. Furthermore, in an ill-posed problem such as motion reconstruction, the optimisation may converge to different solutions depending on the chosen initialisation. For this reason, a prior can recover motion with higher accuracy than a random initialisation, even over long adaptation horizons.

Table 1 shows that all learned priors substantially outperformed random initialisation after five adaptation steps, reducing the mean displacement error from 3.45 mm to 1.66–1.85 mm. Meta-learning achieved the lowest overall displacement error (1.66 mm), whereas the auto-decoder prior had the lowest end-systolic displacement error (2.11 mm). All methods produced deformation fields with mean Jacobian determinants close to 1 in the LV and no cases of folding. While metrics between prior formulations were relatively similar after five

Table 1. Motion estimation results after **five** adaptation steps. Best values are shaded.

Prior	Error ↓ (mm)	ES Error ↓ (mm)	Dice ↑ (LVM)	det(F)
Random	3.45 ± 0.51	4.62 ± 0.77	0.850	0.99
Population	1.85 ± 0.49	2.69 ± 0.97	0.911	0.96
Consensus	1.82 ± 0.49	2.62 ± 0.89	0.916	0.97
Auto-decoder	1.80 ± 0.48	2.11 ± 0.74	0.917	0.94
Meta-learning	1.66 ± 0.42	2.32 ± 0.75	0.921	0.97

Table 2. Motion estimation results after **50** adaptation steps. Best values are shaded.

Prior	Error ↓ (mm)	ES Error ↓ (mm)	Dice ↑ (LVM)	det(F)
Random	1.61 ± 0.38	2.11 ± 0.74	0.922	0.98
Population	1.63 ± 0.41	1.94 ± 0.72	0.919	0.96
Consensus	1.47 ± 0.37	1.92 ± 0.69	0.927	0.98
Auto-decoder	1.72 ± 0.41	2.21 ± 0.68	0.918	0.95
Meta-learning	1.43 ± 0.35	1.86 ± 0.65	0.928	0.97



**Fig. 1.** Adaptation trajectory of each prior (left) and various latent dimensions for auto-decoders (right). Measured is displacement error (top) and ES (end-systolic) displacement error (bottom).

Table 3. Training time of each prior on 500 slices using an RTX 4090 GPU

Prior	Training time (min)
Random Init	0
Population	3.3
Consensus	19.3
Auto-decoder	53.4
Meta-learning	116.4

adaptation steps, the differences became more apparent over longer adaptation trajectories.

Although the meta-learned prior was optimised for five-step adaptation, performance continued to improve well beyond this horizon (Fig. 1). In contrast, the auto-decoder prior converged rapidly but plateaued after approximately ten iterations. Changing the latent dimensionality had little effect on the adaptation trajectory (Fig. 1, right panels), suggesting that the plateau arises from the modulation mechanism rather than the latent-space capacity. While both simple priors had good early performance, the population prior started to deteriorate after 30–40 iterations while the consensus prior continued to improve. Since the population prior is optimised to represent an average motion field, it may represent a compromise solution from which further adaptation is more difficult. Table 2 shows the full results after 50 iterations, revealing that meta-learning had the best performance across all metrics for this longer budget. Unlike the other priors, meta-learning explicitly optimises parameters for efficient adaptation to new subjects, which likely explains its superior performance. However, the chosen formulation used second-order gradients, which incurred a greater computational cost and training time relative to the other priors (see Table 3).

**Fig. 2.** Circumferential (left) and radial (right) mean LV strain curves of each prior after five adaptation steps.




Fig. 3. End-systolic displacement magnitude of each prior, adapting to a mid-ventricular test case for 0, 5 and 50 steps. The numbers underneath each image represent the error against manual ground truth (GT) in mm. *Interpolated from sparse points. Images reproduced by kind permission of the UK Biobank ©

Fig. 2 shows that all learned priors recovered smooth physiological strain patterns after five adaptation steps. However, the auto-decoder produced larger strain magnitudes than the other methods, particularly near end-systole, consistent with its lower end-systolic displacement error. Fig. 3 further shows that the auto-decoder recovers large deformations more rapidly than the other priors but exhibits little improvement with prolonged adaptation. In contrast, the population prior produces noticeably noisier displacement fields after 50 adaptation steps compared with a random initialisation. Although the population prior provides a strong starting point, these results indicate that it becomes less favourable for longer-term optimisation.

This study has limitations that should be considered when interpreting the results. Firstly, the motion estimation task was limited to 2D tagged CMR images, which have modality-specific image properties and artefacts. Secondly, due to computational limitations, the models were adapted for fixed iteration counts to assess early and longer-term adaptation performance. Therefore, fully converged performance was not evaluated in this study. Finally, the UK Biobank consists primarily of healthy individuals, limiting the diversity of motion represented in the learned priors. Priors learned from more diverse disease cohorts may exhibit different adaptation behaviour.

6 Conclusion

Learned priors substantially improved the speed and quality of INR-based cardiac motion estimation. A consensus prior, created by averaging INR weights, was a simple and effective initialisation. Auto-decoders were able to recover large deformations with just five adaptation steps, while meta-learning had the strongest optimisation trajectory over a 50-step adaptation horizon. These results show that the choice of learned prior dictates both adaptation speed and optimisation trajectory.

Acknowledgments. This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) Doctoral Training Partnership [EP/W524475/1] and EPSRC project grant [Z5336762]. This study was conducted using the UK Biobank resource under access application 2964.

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

Bibliography

- [1] Alvarez-Florez, L., Haddou, S.B., Tjong, F.V.Y., Igum, I.: Tagged-Informed Prior for Motion Quantification in Cine CMR Using Implicit Neural Representations. In: Proc. 9th Int. Conf. Med. Imaging Deep Learn. pp. 2003–2018. PMLR (May 2026), <https://proceedings.mlr.press/v315/alvarez-florez26a.html>
- [2] Alvarez-Florez, L., Sander, J., Bourfiss, M., Tjong, F.V.Y., Velthuis, B.K., Išgum, I.: Deep Learning for Automatic Strain Quantification in Arrhythmogenic Right Ventricular Cardiomyopathy. In: STACOM. pp. 25–34. Springer Nature Switzerland, Cham (2024). https://doi.org/10.1007/978-3-031-52448-6_3
- [3] Amiranashvili, T., Lüdke, D., Li, H.B., Zachow, S., Menze, B.H.: Learning continuous shape priors from sparse data with neural implicit functions. *Medical Image Analysis* **94**, 103099 (May 2024). <https://doi.org/10.1016/j.media.2024.103099>
- [4] Arratia López, P., Mella, H., Uribe, S., Hurtado, D.E., Sahli Costabal, F.: WarpPINN: Cine-MR image registration with physics-informed neural networks. *Medical Image Analysis* **89**, 102925 (Oct 2023). <https://doi.org/10.1016/j.media.2023.102925>
- [5] Banus, J., Delaloye, A., M. Gordaliza, P., Georgantas, C., van Heeswijk, R.B., Richiardi, J.: NIMOSEF: Neural Implicit Motion and Segmentation Functions. In: Med. Image Comput. Comput. Assist. Interv. – MICCAI 2025. pp. 444–454. Springer Nature Switzerland, Cham (2026). https://doi.org/10.1007/978-3-032-05141-7_43
- [6] Bell, A., Choi, Y.K., Petersen, S.E., King, A., Nazir, M.S., Young, A.A.: Implicit neural representations of intramyocardial motion and strain. In: STACOM. pp. 150–159. Springer Nature Switzerland, Cham (2026). https://doi.org/10.1007/978-3-032-17734-6_15
- [7] Chandrashekara, R., Rao, A., Sanchez-Ortiz, G.I., Mohiaddin, R.H., Rueckert, D.: Construction of a statistical model for cardiac motion analysis using nonrigid image registration. In: Inf. Process. Med. Imaging. pp. 599–610. Springer Berlin Heidelberg, Berlin, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45087-0_50
- [8] Fischer, B., Modersitzki, J.: Ill-posed medicine - An introduction to image registration. *Inverse Probl.* **24** (Jun 2008). <https://doi.org/10.1088/0266-5611/24/3/034008>
- [9] Garzia, S., Rygiel, P., Dummer, S., Cademartiri, F., Celi, S., Wolterink, J.M.: Neural Fields for Continuous Periodic Motion Estimation in 4D Cardiovascular Imaging. In: STACOM. pp. 378–389. Springer Nature Switzerland, Cham (2025). https://doi.org/10.1007/978-3-031-87756-8_37
- [10] Lowes, M.M., Pedersen, J.J., Hansen, B.S., Kofoed, K.F., Sermesant, M., Paulsen, R.R.: Implicit Neural Representations for Registration of Left

- Ventricle Myocardium During a Cardiac Cycle. In: STACOM. pp. 172–182. Springer-Verlag, Marrakesh, Morocco (2025). https://doi.org/10.1007/978-3-031-87756-8_17
- [11] Mehta, I., Gharbi, M., Barnes, C., Shechtman, E., Ramamoorthi, R., Chandraker, M.: Modulated Periodic Activations for Generalizable Local Functional Representations. In: ICCV. pp. 14194–14203 (Oct 2021). <https://doi.org/10.1109/ICCV48922.2021.01395>
 - [12] Molaei, A., Aminimehr, A., Tavakoli, A., Kazerouni, A., Azad, B., Azad, R., Merhof, D.: Implicit Neural Representation in Medical Imaging: A Comparative Survey. In: 2023 IEEE CVF Int. Conf. Comput. Vis. Workshop ICCVW. pp. 2373–2383. IEEE, Paris, France (Oct 2023). <https://doi.org/10.1109/ICCVW60793.2023.00252>
 - [13] Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In: 2019 IEEE CVF Conf. Comput. Vis. Pattern Recognit. CVPR. pp. 165–174 (Jun 2019). <https://doi.org/10.1109/CVPR.2019.00025>
 - [14] Petersen, S.E., Matthews, P.M., Francis, J.M., Robson, M.D., Zemrak, F., Boubertakh, R., Young, A.A., Hudson, S., Weale, P., Garratt, S., Collins, R., Piechnik, S., Neubauer, S.: UK Biobank’s cardiovascular magnetic resonance protocol. *J Cardiovasc Magn Reson* **18**, 8 (Feb 2016). <https://doi.org/10.1186/s12968-016-0227-4>
 - [15] Shen, C., Zhu, H., Zhou, Y., Liu, Y., Yi, S., Dong, L., Zhao, W., Brady, D.J., Cao, X., Ma, Z., Lin, Y.: Continuous 3D Myocardial Motion Tracking via Echocardiography. *IEEE Trans. Med. Imaging* **43**(12), 4236–4252 (Dec 2024). <https://doi.org/10.1109/TMI.2024.3419780>
 - [16] Sitzmann, V., Chan, E., Tucker, R., Snively, N., Wetzstein, G.: MetaSDF: Meta-Learning Signed Distance Functions. In: Adv. Neural Inf. Process. Syst. vol. 33, pp. 10136–10147. Curran Associates, Inc. (2020), https://papers.neurips.cc/paper_files/paper/2020/file/731c83db8d2ff01bdc000083fd3c3740-Paper.pdf
 - [17] Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit Neural Representations with Periodic Activation Functions. In: Adv. Neural Inf. Process. Syst. vol. 33, pp. 7462–7473. Curran Associates, Inc. (2020), <https://proceedings.neurips.cc/paper/2020/hash/53c04118df112c13a8c34b38343b9c10-Abstract.html>
 - [18] Wolterink, J.M., Zwienenberg, J.C., Brune, C.: Implicit Neural Representations for Deformable Image Registration. In: Proc. 5th Int. Conf. Med. Imaging Deep Learn. pp. 1349–1359. PMLR (Dec 2022), <https://proceedings.mlr.press/v172/wolterink22a.html>
 - [19] Young, A.A., Kraitchman, D.L., Dougherty, L., Axel, L.: Tracking and finite element analysis of stripe deformation in magnetic resonance tagging. *IEEE Trans Med Imaging* **14**(3), 413–421 (1995). <https://doi.org/10.1109/42.414605>