Non-Linear INR-Based Motion Modeling for 4D Radiotherapy

Johannes B. Gebauer[1,2][0009−0008−0681−268X], Laura E. Büttgen[1,3][0009−0005−0441−0319], Thilo Sentker[1,2][0000−0001−6663−0003]*, and René Werner[1,2][0000−0002−2296−3305]*

[1] Institute for Applied Medical Informatics, University Medical Center Hamburg-Eppendorf, Hamburg, Germany
{j.gebauer, l.buettgen, t.sentker, r.werner}@uke.de
[2] Institute of Computational Neuroscience, University Medical Center Hamburg-Eppendorf, Hamburg, Germany
[3] Department of Radiotherapy and Radiation Oncology, University Medical Center Hamburg-Eppendorf, Hamburg, Germany

Abstract. Accurate motion management via patient-specific motion models is critical for motion compensation and tumor tracking in 4D radiotherapy. Existing approaches typically rely on a two-step pipeline combining deformable image registration (DIR) and linear surrogate-based regression, limiting their ability to capture complex, nonlinear surrogate-to-motion relationships. We propose a novel implicit neural representation (INR) motion model that jointly optimizes DIR and surrogate regression, directly predicting continuous deformation vector fields from arbitrary surrogate states. The framework simultaneously learns forward and backward motion models with enforced cyclic consistency, and adapts the SIREN architecture with anisotropic frequency scaling to account for the differing spatial and temporal dynamics of the input. Evaluated on the DIR-Lab 4DCT dataset, our method achieves a mean TRE of 1.08 mm and DSC of 0.625, outperforming classical and INR-based two-step baselines while yielding strictly zero folding across all cases. Source code is available at https://github.com/IPMI-ICNS-UKE/INR-Motion-Modeling.

Keywords: Motion Modelling · Image Registration · Implicit Neural Representation.

1 Introduction

Respiratory-induced motion makes the treatment of lung and liver tumors clinically challenging [9]. Precise motion management of the patient is therefore crucial for accurately calculating the delivered dose to the tumor and organs at risk (OAR) [20]. Specifically, this requires patient-specific motion modeling: establishing a mathematical relationship between an external surrogate signal,

* These authors contributed equally as last authors.

acquired during treatment, and the internal motion of the region of interest, typically represented as a dense deformation vector field (DVF). Common surrogate signals include optical tracking of a marker block placed on the thorax and real-time imaging [2]. Accurate motion modeling can furthermore benefit the reconstruction of the treatment-day 4D cone-beam CT (4DCBCT) by enabling prior-guided, motion-compensated reconstruction methods [13].

Patient-specific motion models are typically derived from the patient's planning 4DCT, which commonly comprises 10 phase volumes paired with their respective surrogate signals over a median respiratory cycle. The objective of such a motion model is to predict a DVF that maps a reference anatomical state to any target respiratory state guided by the surrogate signal [10]. Constructing such a model first requires pair-wise registration between the reference phase and all other phases. Subsequently, standard modeling approaches assume a linear relationship between the surrogates and the DVFs, and solve for a system matrix via linear regression [17]. Recent works have extended this classic linear model by optimizing the surrogate and allowing it to vary in the superior-inferior (SI) direction [8]. While providing a multi-dimensional surrogate allows linear models to approximate hysteresis, they remain restricted to a simple linear combination of these global inputs. Consequently, they inherently struggle to accurately capture highly complex, non-linear surrogate-motion interactions, such as localized variations in hysteresis or spatially heterogeneous tissue deformations.

A complementary line of work has focused on improving the spatial representation of deformation fields themselves. In this context, implicit neural representations (INRs) have emerged as a flexible and expressive framework for deformable image registration (DIR) [18, 5, 15]. INRs represent dense DVFs as continuous functions rather than relying on discrete 3D voxel grids. In particular, models using sinusoidal representation networks (SIREN) [14] have demonstrated high efficacy for such spatial transformations. Furthermore, INRs have been adapted to model dynamic scenes [12], including applications as internal motion models for 4DCBCT motion-compensated reconstruction [19] and group image reconstruction in cardiac MRI [6]. However, such internal models rely exclusively on the intrinsic ordering of the phase images and do not explicitly incorporate an external surrogate signal to control motion.

Building upon these advancements, we propose a SIREN-based patient-specific motion model capable of predicting continuous DVFs inside the lung with respect to a reference state, conditioned on an external surrogate signal. In contrast to classical sequential pipelines, our approach eliminates the need for prior pairwise registrations. Instead, it represents respiratory motion as a continuous implicit function jointly parameterized by spatial location and the surrogate signal.

In summary, our main contributions are threefold: **1** we introduce a novel patient-specific respiratory motion model based on the SIREN framework that predicts continuous, high-resolution DVFs directly from a patient-specific surrogate; **2** we enforce cyclic consistency by simultaneously learning forward and backward transformations and introduce an anisotropic frequency scaling in the first SIREN layer to accommodate the spatial and temporal differences; and **3**

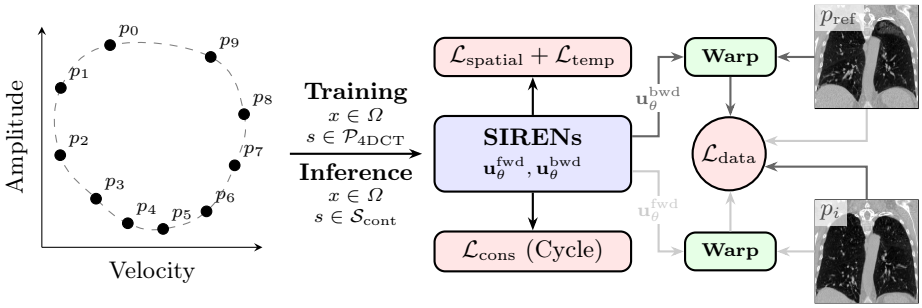

Fig. 1. Overview of the proposed continuous motion modeling framework. Left: The 2D surrogate space, illustrating the discrete 4DCT phases ($\mathcal{P}_{\text{4DCT}}$) used for training and the continuous domain ($\mathcal{S}_{\text{cont}}$) available for inference. Right: Forward ($\mathbf{u}^{\text{fwd}}$) and backward ($\mathbf{u}^{\text{bwd}}$) deformation fields are predicted by individual SIRENs from spatial coordinates $\mathbf{x}$ and surrogate states $\mathbf{s}$. A data loss ($\mathcal{L}_{\text{data}}$) drives the optimization by matching the reference (p_{ref}) and target (p_i) images via forward (light arrows) and backward (dark arrows) warping. Training is further regularized by cycle consistency ($\mathcal{L}_{\text{cons}}$) and spatio-temporal smoothness ($\mathcal{L}_{\text{spatial}}, \mathcal{L}_{\text{temp}}$).

we demonstrate the efficacy of our proposed framework through a comprehensive evaluation on the publicly available DIR-Lab 4DCT dataset [3].

2 Methods

2.1 Problem Formulation and Classical Modeling

A patient-specific respiratory motion model aims to predict the transformation $\mathbf{\Phi}$ from a reference phase to an arbitrary respiratory state parameterized by an external surrogate signal. In this work, we choose the end-exhale (EE) phase (typically corresponding to the 50% phase in a 4DCT) as the reference, as it represents the most reproducible anatomical state. We further represent $\mathbf{\Phi}$ via the displacement field $\mathbf{u}$:

$$\mathbf{\Phi}(\mathbf{x}) = \mathbf{x} + \mathbf{u}(\mathbf{x}), \tag{1}$$

where $\mathbf{x}$ is a coordinate in the image domain Ω. In the following, we use the terms DVF and $\mathbf{u}$ interchangeably.

Motion models can be formulated in two directional variants. In the forward-warping formulation, the DVF maps the reference to the surrogate-driven target state, enabling the tracking of anatomical landmarks. In the backward-warping formulation, the DVF maps the target state back to the reference, allowing images or segmentation masks defined in the reference to be propagated to arbitrary respiratory states, which is essential for dose accumulation.

Commonly, classical motion models are constructed in a sequential two-step process. First, ground-truth DVFs, denoted as $U \in \mathbb{R}^{m \times n_{\text{ph}}}$, are calculated between the reference phase and all other n_{ph} phases via DIR, where $m = 3N_{\text{vox}}$

represents the number of image voxels multiplied by the spatial dimensions. Each column of U corresponds to a flattened DVF for one respiratory phase. Subsequently, a linear relationship between the DVFs and the corresponding surrogate signal $S \in \mathbb{R}^{n_{\text{ind}} \times n_{\text{ph}}}$ (with n_{ind} being the dimensionality of the surrogate) is assumed:

$$U = BS, \tag{2}$$

where $B \in \mathbb{R}^{m \times n_{\text{ind}}}$ is the desired system matrix. Equation 2 can be solved in a least-squares sense via multivariate linear regression, as described by Wilms et al. [17].

2.2 Implicit Motion Model

We model respiratory motion as a continuous displacement field

$$\mathbf{u}(\mathbf{x}, \mathbf{s}) \in \mathbb{R}^3, \tag{3}$$

where $\mathbf{x} \in [-1, 1]^3$ denotes the normalized spatial position and $\mathbf{s} \in [-1, 1]^{n_{\text{ind}}}$ represents the respiratory surrogate signal. Specifically, we implement $\mathbf{u}$ as a SIREN [14]

$$\mathbf{u}_{\text{non-linear}}(\mathbf{x}, \mathbf{s}) = f_\theta([\mathbf{x}, \mathbf{s}]), \tag{4}$$

with $f_\theta : \mathbb{R}^{3+n_{\text{ind}}} \to \mathbb{R}^3$, and θ being the model parameters optimized during training. The full model framework is illustrated in Fig. 1.

An important hyperparameter of SIRENs is their frequency scaling parameter ω_0 in their activations $\sin(\omega_0 \cdot \mathbf{W}\mathbf{x} + \mathbf{b})$ [14]. It is commonly set to a large value (e.g., $\omega_0 = 30$), to enable the network to represent high-frequency components across the input domain $[-1, 1]$. Acknowledging the anisotropy of our input, i.e. high-frequency components in space and smooth variations in the surrogate, we set ω_0 individually per input channel. Specifically, we use $\omega_{0,\mathbf{x}} = 30$ for spatial and $\omega_{0,\mathbf{s}} = \pi/2$ for surrogate inputs. The latter ensures that the surrogate inputs only span half a period over the input domain.

Ablation Models We train the following ablation models to disentangle the effect of joint spatio-temporal optimization, the nature of the surrogate relationship and the impact of INR-based DIR. First, we train a (temporal)-linear SIREN:

$$\mathbf{u}_{\text{linear}}(\mathbf{x}, \mathbf{s}) = b_\theta(\mathbf{x})\mathbf{s}, \tag{5}$$

where $b_\theta : \mathbb{R}^3 \to \mathbb{R}^{3 \times n_{\text{ind}}}$ represents a spatially continuous analogue of the system matrix B in Equation 2, enforcing a strictly linear relationship with respect to the surrogate signal.

We then build two variants of the two-step procedure as described in Equation 2. The required pair-wise DVFs are, for the first model, acquired via classical

variational registration (VarReg) [4], and for the second with DIR-SIRENs as proposed in [18]. The latter is then sampled on a regular grid before the second step.

2.3 Training Objective

The network parameters are optimized by minimizing a joint training objective, defined as a weighted sum of four loss components:

$$\mathcal{L} = \mathcal{L}_{\text{data}} + \lambda_{\text{spatial}}\mathcal{L}_{\text{spatial}} + \lambda_{\text{temp}}\mathcal{L}_{\text{temp}} + \lambda_{\text{cons}}\mathcal{L}_{\text{cons}}, \tag{6}$$

where λ_{spatial}, λ_{temp}, and λ_{cons} are hyperparameters controlling the relative influence of each regularization term. The individual terms are defined as follows:
(1) Data term ($\mathcal{L}_{\text{data}}$): We utilize the masked normalized cross-correlation (NCC) to evaluate the similarity between the deformed reference and the target phase images strictly within the lung regions.
(2) Spatial regularization ($\mathcal{L}_{\text{spatial}}$): To ensure biologically plausible deformations, we penalize local folding and enforce spatial smoothness using a Jacobian determinant penalty and a Laplacian penalty, respectively:

$$\mathcal{L}_{\text{spatial}} = \|\det(I + \nabla\mathbf{u}) - 1\|_2^2 + \lambda_{\text{Lap}} \|\Delta\mathbf{u}\|_2^2 . \tag{7}$$

(3) Temporal regularization ($\mathcal{L}_{\text{temp}}$): To encourage smooth motion trajectories over time, we apply a Laplacian smoothness penalty with respect to the surrogate inputs for the non-linear model. For the linear baseline model, we penalize the L_2-norm of the first derivatives with respect to the surrogates.
(4) Consistency term ($\mathcal{L}_{\text{cons}}$): As demonstrated by van Harten et al. [15], training of the DIR-SIRENs can be stabilized by enforcing forward-backward cycle consistency. We adopt these findings and simultaneously train both directions and formulate the following L_2 penalty on the residual displacement:

$$\mathcal{L}_{\text{cons}} = \left\|\mathbf{u}^{\text{fwd}}(\mathbf{x}, \mathbf{s}) + \mathbf{u}^{\text{bwd}}\big(\mathbf{x} + \mathbf{u}^{\text{fwd}}(\mathbf{x}, \mathbf{s}), \mathbf{s}\big)\right\|_2^2 . \tag{8}$$

2.4 Implementation Details

Our framework is implemented in PyTorch [11], utilizing its autograd functionality to analytically compute all spatial and temporal gradients required for the training objective. During each training step, we randomly sample 1,000 continuous spatial coordinates per phase from the union of the phase-specific lung masks. We apply intra-voxel jitter to these coordinates to fully leverage the continuous spatial domain of the INRs. For the data loss computation, image intensities at these continuous coordinates are obtained via trilinear interpolation from the discrete voxel grids. We set individual weights in Equation 6 to $\lambda_{\text{spatial}} = 0.01$, $\lambda_{\text{temp}} = 0.1$ and $\lambda_{\text{cons}} = 0.1$.

The network parameters are optimized using the AdamW optimizer with an initial learning rate of 10^{-4}. The learning rate is modulated by a cosine annealing

schedule following a 20-step linear warmup. Each model is trained for a total of 10,000 iterations, requiring approximately 15 minutes of training. However, since we observed earlier convergence in many cases, this duration could be reduced by introducing early stopping criteria. All SIRENs are configured with 3 hidden layers, each having a width of 256. We set $\omega_0 = 30$, excluding the anisotropic scaling described in Sec. 2.2.

2.5 Dataset and Surrogate Extraction

We evaluate our proposed method on the publicly available DIR-Lab 4DCT dataset [3], which consists of 10 patient cases. Each case comprises 10 distinct 3D respiratory phase images covering a full breathing cycle. In this dataset, phase 0 represents the maximum/end-inhale (EI) state, and phase 5 represents the maximum/end-exhale (EE) state. Consequently, phases 1 through 4 capture the exhalation, while phases 6 through 9 capture the inhalation.

We aim to evaluate our model on a 2D surrogate $\mathbf{s} \in \mathbb{R}^2$, often given by the displacement amplitude and its temporal derivative of an external marker block, allowing for the differentiation between inhalation and exhalation. Because the original DIR-Lab surrogate traces are unavailable, we derive an internal signal as a replacement. Specifically, we extract the phase-wise lung volumes from automatically predicted lung masks and use these as the signal amplitudes. We then assume an equal phase-spacing of the images, allowing us to extract the derivative.

2.6 Evaluation Metrics

Registration accuracy is commonly assessed via the target registration error (TRE) computed on paired ground-truth landmarks. The DIR-Lab dataset provides 300 expert-annotated landmarks per case, available only at the extreme respiratory states. For intermediate phases, only automatically propagated landmarks are available and exclusively for the exhalation, precluding a balanced evaluation across the full respiratory cycle. We therefore adopt lung vessel maps as a dense anatomical proxy, with vessel probability maps extracted automatically using TotalSegmentator [16].

To avoid interpolation artifacts, we apply the predicted DVFs directly to continuous vessel probability maps via trilinear interpolation before binarizing at a 0.5 threshold. We report DSC and Average Symmetric Surface Distance (ASSD), the latter being more robust than DSC for thin tubular structures, as well as TRE evaluated on the 300 extreme-phase landmarks. To assess spatial regularity, we compute the Jacobian determinant $|J|$ of each transformation, reporting the fraction of voxels with $|J| \leq 0$ as an indicator of non-physiological folding and $\mathrm{std}(|J|)$ as a measure of transformation smoothness.

3 Results

For clarity, we restrict the evaluation to the backward models, as we observed similar trends for both directions. Quantitative results are shown in Table 1. Our proposed full implicit SIREN outperforms all ablation models across all registration accuracy metrics. Regarding topological plausibility, it is the only model to achieve strictly zero folding, while all other models exhibit near-zero but non-zero rates (order 10^{-6}). DVF smoothness as measured by std($|J|$) is consistent across all SIREN-based methods.

Fig. 2 provides a per-case, per-phase breakdown of DSC. As expected, all models perform best for small respiratory motion, evident in higher DSC values for phases adjacent to the reference (p_5) and for cases with limited overall motion (e.g., cases 1 and 2). The most performance gains occur at two transitions: first, from classical to INR-based DIR, and second, moving from the temporal-linear assumption to our fully joint SIREN approach.

Finally, Fig. 3 visualizes the behavior of the SIREN motion models in the surrogate phase space. The linear SIREN produces a globally uniform displacement field, paired with a steady increase of spatial irregularity. Standard (isotropic) SIREN (Full (Iso.)) exhibits temporal overfitting, leading to deviations outside the physiological breathing cycle. Introducing our proposed anisotropic frequency scaling (Full (Aniso.)) suppresses these artifacts, yielding a smooth temporal motion model.

Table 1. Quantitative comparison of motion models evaluated on lung vessel maps and landmarks. DSC↑, ASSD↓ [mm], TRE↓ [mm], % of $|J|_{\leq 0}$ ↓ and std($|J|$) ↓.

Model		DSC ↑	ASSD ↓	TRE ↓	% of $\lvert J\rvert_{\leq 0}$ ↓	std($\lvert J\rvert$) ↓
2-Step Seq.	Classical	0.562	0.816	1.521	6.0e−7	0.127
	INR	0.571	0.786	1.341	3.5e−6	**0.072**
Joint SIREN	Linear	0.573	0.785	1.404	2.4e−6	**0.072**
	Full	**0.625**	**0.682**	**1.084**	**0.0**	0.080

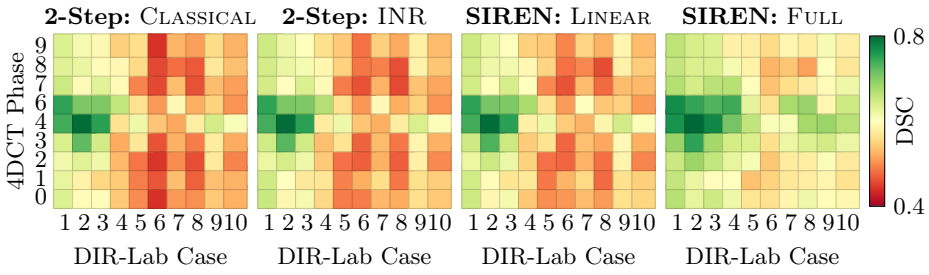

Fig. 2. Phase-wise DSC for all cases and models. Phase 5 serves as the fixed reference and is thus omitted. Our proposed approach (Full) mitigates the performance drops observed in challenging cases across all baseline models.

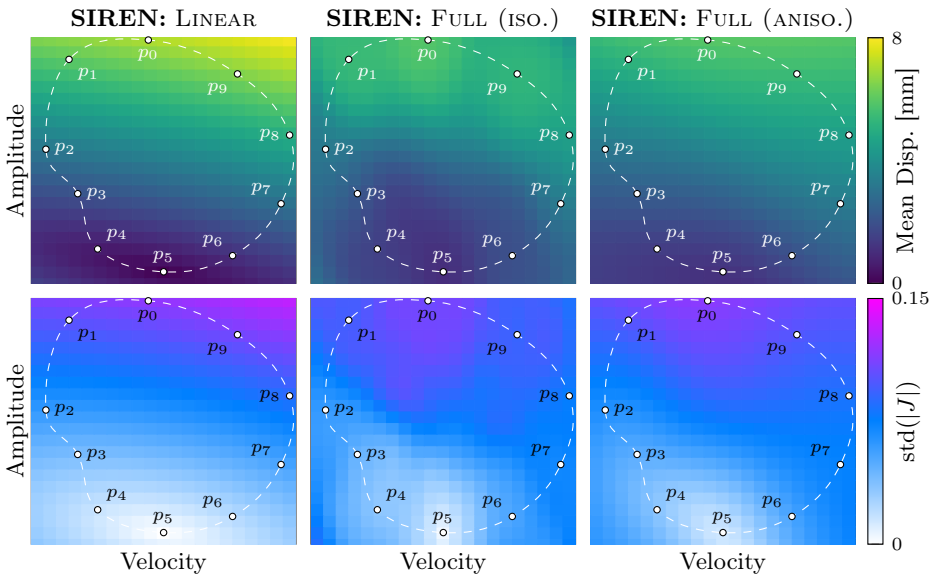

Fig. 3. Continuous motion model behavior in the surrogate phase space. Top row: Mean displacement [mm]. Bottom row: DVF regularity (std($|J|$)). The dashed line represents the respiratory trajectory with discrete 4DCT phases (p_0–p_9) for case 9 and reference phase 5. High-frequency artifacts both inside and outside the observed breathing cycle are effectively suppressed by our proposed anisotropic frequency scaling.

4 Discussion

In this work, we proposed an INR-based motion model for the lung, predicting dense deformation vector fields conditioned on a 2D surrogate signal. Case-specific models were trained and evaluated on the 10 open-access DIR-Lab 4DCTs. INR-based DIR yields a modest improvement over classical registration, while joint spatio-temporal optimization alone does not surpass the two-step pipeline. The largest performance gain arises when joint optimization is combined with a nonlinear surrogate-to-motion relationship, as in the full implicit SIREN. The strictly zero folding and generally smooth deformations of our model further highlight the efficacy of SIREN-based registration with continuous regularization, consistent with prior work [18, 5, 15]. Our proposed anisotropic frequency scaling further ensures smoothness in the surrogate space.

Several aspects remain to be addressed before clinical deployment. First, this work relied on an internal lung-volume signal as a surrogate substitute, since the original DIR-Lab traces are unavailable. Translating the method to a true external signal is therefore an important next step. We expect the model to generalize well here, as its nonlinear formulation is not constrained to the linear manifold of a potentially sub-optimal external signal. Second, the stability of the model when transferred from the planning 4DCT to the treatment day remains an

open question, given known instability issues for classical baselines [1]. Here, fine-tuning on a pre-treatment cone-beam CT is a possible mitigation strategy. Finally, extension to higher-dimensional surrogates, such as real-time imaging during treatment [7], represents a promising direction toward fully online adaptive radiotherapy workflows.

Disclosure of Interests. R.W. receives a research grant from Siemens Healthineers.

References

1. Büttgen, L.E., Werner, R., Gauer, T.: Stability analysis of patient-specific 4DCT- and 4DCBCT-based correspondence models. Medical Physics **51**(9), 5890–5900 (2024). https://doi.org/10.1002/mp.17304
2. Caillet, V., Booth, J.T., Keall, P.: IGRT and motion management during lung SBRT delivery. Physica Medica **44**, 113–122 (Dec 2017). https://doi.org/10.1016/j.ejmp.2017.06.006
3. Castillo, R., Castillo, E., Guerra, R., Johnson, V.E., McPhail, T., Garg, A.K., Guerrero, T.: A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets. Physics in Medicine & Biology **54**(7), 1849 (Mar 2009). https://doi.org/10.1088/0031-9155/54/7/001
4. Ehrhardt, J., Schmidt-Richberg, A., Werner, R., Handels, H.: Variational Registration. In: Handels, H., Deserno, T.M., Meinzer, H.P., Tolxdorff, T. (eds.) Bildverarbeitung Für Die Medizin 2015. pp. 209–214. Springer, Berlin, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46224-9_37
5. Fischer, B., Modersitzki, J.: Curvature Based Image Registration. Journal of Mathematical Imaging and Vision **18**(1), 81–85 (Jan 2003). https://doi.org/10.1023/A:1021897212261
6. 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: Camara, O., Puyol-Antón, E., Sermesant, M., Suinesiaputra, A., Zhao, J., Wang, C., Tao, Q., Young, A. (eds.) Statistical Atlases and Computational Models of the Heart. Workshop, CMRxRecon and MBAS Challenge Papers. pp. 378–389. Springer Nature Switzerland, Cham (2025). https://doi.org/10.1007/978-3-031-87756-8_37
7. Hindley, N., Keall, P.J.: An open-source deep learning framework for respiratory motion monitoring and volumetric imaging during radiation therapy. Medical Physics **52**(7), e18015 (2025). https://doi.org/10.1002/mp.18015
8. Huang, Y., Eiben, B., Thielemans, K., McClelland, J.R.: Resolving Variable Respiratory Motion From Unsorted 4D Computed Tomography (Jun 2024). https://doi.org/10.48550/arXiv.2407.00665
9. Keall, P.J., Mageras, G.S., Balter, J.M., Emery, R.S., Forster, K.M., Jiang, S.B., Kapatoes, J.M., Low, D.A., Murphy, M.J., Murray, B.R., Ramsey, C.R., Van Herk, M.B., Vedam, S.S., Wong, J.W., Yorke, E.: The management of respiratory motion in radiation oncology report of AAPM Task Group 76a). Medical Physics **33**(10), 3874–3900 (Oct 2006). https://doi.org/10.1118/1.2349696
10. McClelland, J.R., Hawkes, D.J., Schaeffter, T., King, A.P.: Respiratory motion models: A review. Medical Image Analysis **17**(1), 19–42 (Jan 2013). https://doi.org/10.1016/j.media.2012.09.005

11. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: PyTorch: An Imperative Style, High-Performance Deep Learning Library, pp. 8026–8037. No. 721, Curran Associates Inc., Red Hook, NY, USA (Dec 2019)
12. Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-NeRF: Neural Radiance Fields for Dynamic Scenes (Nov 2020). https://doi.org/10.48550/arXiv.2011.13961
13. Rit, S., Wolthaus, J.W.H., van Herk, M., Sonke, J.J.: On-the-fly motion-compensated cone-beam CT using an a priori model of the respiratory motion. Medical Physics **36**(6Part1), 2283–2296 (2009). https://doi.org/10.1118/1.3115691
14. Sitzmann, V., Martel, J.N.P., Bergman, A.W., Lindell, D.B., Wetzstein, G.: Implicit Neural Representations with Periodic Activation Functions (Jun 2020). https://doi.org/10.48550/arXiv.2006.09661
15. van Harten, L.D., Stoker, J., Išgum, I.: Robust Deformable Image Registration Using Cycle-Consistent Implicit Representations. IEEE Transactions on Medical Imaging **43**(2), 784–793 (Feb 2024). https://doi.org/10.1109/TMI.2023.3321425
16. Wasserthal, J., Breit, H.C., Meyer, M.T., Pradella, M., Hinck, D., Sauter, A.W., Heye, T., Boll, D.T., Cyriac, J., Yang, S., et al.: Totalsegmentator: robust segmentation of 104 anatomic structures in ct images. Radiology: Artificial Intelligence **5**(5), e230024 (2023)
17. Wilms, M., Werner, R., Ehrhardt, J., Schmidt-Richberg, A., Schlemmer, H.P., Handels, H.: Multivariate regression approaches for surrogate-based diffeomorphic estimation of respiratory motion in radiation therapy. Physics in Medicine & Biology **59**(5), 1147 (Feb 2014). https://doi.org/10.1088/0031-9155/59/5/1147
18. Wolterink, J.M., Zwienenberg, J.C., Brune, C.: Implicit Neural Representations for Deformable Image Registration. In: Proceedings of The 5th International Conference on Medical Imaging with Deep Learning. pp. 1349–1359. PMLR (Dec 2022)
19. Zha, R., Zhang, Y., Li, H.: NAF: Neural Attenuation Fields for Sparse-View CBCT Reconstruction. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2022. pp. 442–452. Springer Nature Switzerland, Cham (2022). https://doi.org/10.1007/978-3-031-16446-0_42
20. Zhang, Q., Pevsner, A., Hertanto, A., Hu, Y.C., Rosenzweig, K.E., Ling, C.C., Mageras, G.S.: A patient-specific respiratory model of anatomical motion for radiation treatment planning. Medical Physics **34**(12), 4772–4781 (2007). https://doi.org/10.1118/1.2804576