Single-spoke Motion-compensated Dynamic 3D MRI Reconstruction via Neural Representation

Lixuan Chen, James M. Balter, Liyue Shen✉, and Jeong Joon Park✉

University of Michigan, Ann Arbor MI 48109, USA
{chenlx, jbalter, liyues, jjparkcv}@umich.edu

Abstract. Dynamic 3D Magnetic Resonance Imaging (MRI) is a powerful imaging technique for motion monitoring and tracking, offering both excellent soft-tissue contrast and the ability to capture dynamic changes in tissue. Current reconstruction methods typically assume that multiple spokes share the same motion state. However, this assumption does not align with the complex realities of patient motion and clinical acquisition protocols, often resulting in anatomical discontinuities or blurring artifacts in the reconstructed images. In this work, we propose an unsupervised **S**ingle-s**P**oke motion-compensated **I**mplicit **NE**ural **R**epresentation method (**SPINER**) for dynamic volumetric MRI reconstruction. We address a more challenging yet realistic scenario, single-spoke motion modeling, which assigns a unique motion state for each spoke measurement. To address this highly ill-posed inverse problem, we propose a motion-ignoring static initialization strategy that exploits static anatomical information across all spokes. We find that a good initialization of the canonical volume significantly improves the optimization process and facilitates better dynamic volumetric reconstruction based on implicit neural representation learning. Experiments on abdomen MRI datasets demonstrate that our methods can reconstruct high-quality dynamic volumetric MRI while capturing continuous and accurate motion.

Keywords: Dynamic 3D MRI · Motion Model · Neural Representation.

1 Introduction

Accurately reconstructing volumetric density and tissue motion is a fundamental challenge in magnetic resonance imaging (MRI) [16]. However, the lengthy acquisition time required for dynamic 3D MRI creates an inherent trade-off between temporal and spatial resolution. This limitation makes it difficult to meet clinical requirements for both rapid and precise imaging, particularly in applications such as radiotherapy planning for moving tissues.

To tackle this challenge, many existing methods reconstruct image series by grouping continuously acquired radial spokes into several motion states, either through consecutive spoke grouping [22,6], window-shifting [15,21], or motion signal-based sorting [5,7]. These approaches improve spatial resolution by assuming multiple spokes share a single motion state. However, this assumption

fails to capture precise physiological motion, causing inaccurate motion signal estimation and image blurring artifacts. Our work addresses a more challenging scenario: single-spoke motion model (where "spoke" refers to a stack of spokes from stack-of-stars golden-angle radial sampling) in dynamic 3D MRI reconstruction. By assigning each spoke its own motion state, we better reflect real-world continuous motion. However, this modeling further exacerbates the inherently ill-posedness of the inverse problem of dynamic MRI reconstruction by requiring the recovery of instantaneous images from single-spoke k-space data.

In dynamic MRI reconstruction, motion-compensated methods [9,1,18,10] decompose dynamic images into a canonical (or template) volume and a sequence of canonical-to-observation deformation vector fields (DVFs), which can effectively constrain the solution space of the inverse problem. Recently, implicit neural representation (INR) has demonstrated exceptional ability in representing spatial and temporal scenes [19,17]. Some recent works [12,13] propose to integrate INR with the motion-compensated methods by modeling both canonical volume and DVFs as continuous functions parameterized by neural networks. INR provides memory-efficient representation with inherent regularization effects [14,20], making it well-suited for underconstrained 4D reconstruction problems.

In this work, we proposed SPINER, which is, to the best of our knowledge, the first single-spoke motion-compensated dynamic 3D MRI reconstruction method. Different from existing methods, our formulation assigns a unique motion state for each spoke measurement, leading to more continuous motion reconstruction and reduced artifacts compared with the spoke-sharing strategy from prior works [22,6]. We adopt INR to represent both the canonical volume and DVFs and jointly optimize the two INR models to fit all spoke measurements. However, relying solely on the implicit learning bias of the INR cannot produce satisfactory performance in this extremely ill-posed problem, where only one spoke is available to reconstruct the 3D volume at each time point.

To address the issue, we propose a motion-ignoring static initialization strategy. Specifically, the core idea is to initially disregard motion and exploit all aggregated spoke measurements to reconstruct a motion-ignoring static volume. This reconstructed static volume correctly captures low-frequency and static anatomical information, while motion artifacts remain with blurry structural details. This resulting volume is used to initialize the canonical INR network, which is optimized along with the deformation fields. This initialization not only provides low-frequency anatomical information to the canonical network but also allows the DVF network to focus on extracting temporal changes. As the DVF network captures the changes, blurring artifacts are reduced and the canonical network gradually learns intricate details and high-frequency structures.

To evaluate the performance of SPINER, we conduct extensive experiments on our in-house abdominal MRI datasets. The results demonstrate that our method can reconstruct 3D volumes corresponding to each spoke measurement, without suffering from motion blurs induced by spoke-sharing. Additionally, we conduct ablation studies to evaluate the effectiveness of motion-ignoring static initialization and single-spoke motion modeling.

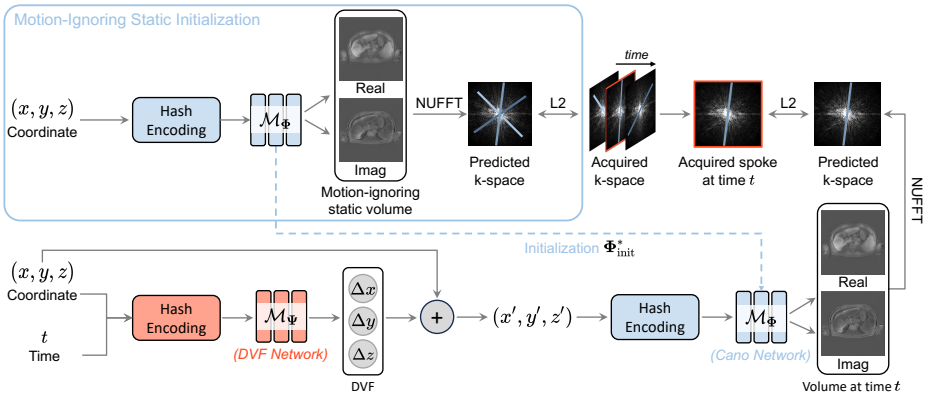

Fig. 1. Overview of the SPINER. (a) **Motion-ignoring static initialization.** The network $\mathcal{M}_\Phi$ maps spatial coordinates (x, y, z) to the complex-value MRI image and minimize the difference between predicted and acquired k-space, yielding Φ^*_{init} as a powerful initialization for the canonical network. (b) **Joint optimization canonical and DVF network.** The DVF network $\mathcal{M}_\Psi$ takes (x, y, z, t) and outputs the displacement $(\Delta x, \Delta y, \Delta z)$. These displacements transform the canonical coordinates, which are then passed to the canonical network to obtain the volume at time t. We jointly optimize both networks by minimizing the loss between predicted and acquired k-space data at time t.

2 Method

2.1 Problem Formulation

The measurement process of dynamic volumetric MRI can be formulated as:

$$\mathbf{y}_t = \mathbf{P}_t \mathcal{F} \mathbf{x}_t + \mathbf{n}_t, \;\; t = 0, \ldots, T, \tag{1}$$

where $\mathbf{X} = [\mathbf{x}_0, \ldots, \mathbf{x}_T]$ denotes the dynamic 3D image, $\mathcal{F}$ denotes the Fourier transform, $\mathbf{P}_t$ is the sampling pattern at t, and $\mathbf{y}_t$ is the acquired k-space data at time t. Recovering the dynamic 3D MRI $\mathbf{X}$ from the acquired k-space signal $\{\mathbf{y}_t\}_{t=0}^T$ is a highly ill-posed inverse problem. To address this, motion-compensated methods [9,1,18,10] propose to model the 4D image $\mathbf{X}$ as a 3D canonical volume $\mathbf{m}$, supplemented by a sequential DVFs $\mathbf{U} = [\mathbf{u}_0, \ldots \mathbf{u}_T]$. As such, the measurement process of dynamic 3D MRI can be modeled as:

$$\mathbf{y}_t = \mathbf{P}_t \mathcal{F} (\mathbf{u}_t \circ \mathbf{m}) + \mathbf{n}_t, \tag{2}$$

where $\mathbf{x}_t = \mathbf{u}_t \circ \mathbf{m}$ represents the warped volume at time t. The canonical volume $\mathbf{m}$ and the DVFs $\mathbf{U}$ can be recovered by solving the above inverse problem in Eq. 2. Subsequently, the desired dynamic 3D images $\mathbf{X}$ are obtained by applying the DVFs to the canonical volume (*i.e.*, $\mathbf{U} \circ \mathbf{m}$).

Existing methods [9,22] typically assume that continuously acquired spoke lines share the same motion stage, which may not adequately capture the real-world motion. To address this issue, we introduce a single-spoke motion model, where a canonical-to-observation deformation is applied to every spoke.

2.2 Single-spoke Motion-compensated Dynamic 3D Reconstruction

Fig. 1 shows the pipeline of the proposed SPINER. First, we represent the canonical volume $\mathbf{m}$ and DVFs $\mathbf{U}$ as neural implicit fields. However, reconstructing 3D volume from single-spoke k-space data is extremely ill-posed. Thus, we propose a motion-ignoring static initialization to capture static anatomical information for constraining the solution space. After initializing the canonical network, we jointly optimize the canonical volume and DVFs.

Neural Representations of Canonical Volume and DVFs. We represent the 3D canonical volume as a function f of the spatial coordinates $\mathbf{p}$ to the complex-valued MRI image:

$$f : \mathbf{p} = (x, y, z) \in \mathbb{R}^3 \rightarrow \mathbf{m}(\mathbf{p}) \in \mathbb{C}, \tag{3}$$

where $\mathbf{m}(\mathbf{p})$ is the intensity of the 3D canonical volume. The function f can be learned with an MLP network $\mathcal{M}_{\mathbf{\Phi}}$ (referred to as the *canonical network*). To facilitate the representation of high-frequency signals, we use hash encoding [11] to map the coordinates into a high-dimensional feature vector $\gamma(\mathbf{p})$. We model the 4D DVF $\mathbf{U} = [\mathbf{u}_0, \dots \mathbf{u}_T]$ as a function g of spatiotemporal coordinates $(\mathbf{p}, t)$ as follows:

$$g : (\mathbf{p}, t) \in \mathbb{R}^4 \rightarrow \mathbf{u}_t(\mathbf{p}) = (\Delta x, \Delta y, \Delta z) \in \mathbb{R}^3, \tag{4}$$

where $(\Delta x, \Delta y, \Delta z)$ denotes the displacement on the spatial domain, corresponding to the time t. We use an MLP network $\mathcal{M}_{\mathbf{\Psi}}$ (referred to as the *DVF network*) to fit g after mapping $(\mathbf{p}, t)$ into $\gamma(\mathbf{p}, t)$ via hash encoding [11].

Motion-ignoring Static Initialization. Although the single-spoke motion model better approximates the realistic motion continuity, it further exacerbates the ill-conditioning of this reconstruction problem. To constrain the search space and stabilize optimization, we propose a motion-ignoring static initialization strategy, as illustrated in the blue box of Fig. 1. Since all spokes redundantly sample the static regions, we disregard motion and exploit all aggregated spoke measurements to initialize the canonical network $\mathcal{M}_{\mathbf{\Phi}}$. Incorporating the forward model of dynamic 3D MRI, we can learn the initialization by calculating the relative ℓ_2 loss between the predicted and acquired measurement data:

$$\mathbf{\Phi}^*_{\text{init}} = \arg\min_{\mathbf{\Phi}} \mathcal{L}\left(\mathbf{P}\mathcal{F}\mathcal{M}_{\mathbf{\Phi}}, \mathbf{Y}\right) + \lambda \cdot \text{TV}(\mathcal{M}_{\mathbf{\Phi}}), \tag{5}$$

where $\mathbf{Y} = [\mathbf{y}_0, \dots, \mathbf{y}_T]$ denotes the acquired k-space at at all time and $\text{TV}(\cdot)$ is the spatial total variation (TV) regularization. Due to the inherent learning bias

of neural networks towards low-frequency signal patterns [14,20], the network fully exploits the redundancy in static regions to represent a 3D volume with high quality in those areas, even though motion artifacts may still be present in dynamic regions. The learned parameters $\mathbf{\Phi}^*_{\text{init}}$ serve as a powerful initialization for the canonical network.

Jointly Optimization Canonical Volume and DVFs. After initializing the canonical network, our goal is to jointly recover the canonical 3D volume $\mathbf{m}$ and the 4D DVF $\mathbf{U} = [\mathbf{u}_0, \dots \mathbf{u}_T]$ The content outside the box of Fig. 1 demonstrates the optimization process. By passing the spatiotemporal coordinates $(\mathbf{p}, t)$ into the DVF network $\mathcal{M}_{\mathbf{\Psi}}$, we obtain the displacement vector $\mathbf{u}_t(\mathbf{p}) = (\Delta x, \Delta y, \Delta z)$. Applying these deformation fields to the canonical coordinates, we obtain the transformed coordinates $\mathbf{p}' = \mathbf{u}_t(\mathbf{p}) \circ \mathbf{p} = (x', y', z')$:

$$x' = x + \Delta x, \quad y' = y + \Delta y, \quad z' = z + \Delta z. \tag{6}$$

Benefiting from the inherent consistency bias of INR, we obtain the transformed volume by passing $\mathbf{p}' = (x', y', z')$ to the canonical network $\mathcal{M}_{\mathbf{\Phi}}$. Finally, leveraging the dynamic 3D MRI forward model with the sample trajectory $\mathbf{P}_t$ at time t, we jointly optimize the canonical and DVF network by calculating the relative ℓ_2 loss between the predicted and acquired k-space data at time t:

$$\mathbf{\Phi}^*, \mathbf{\Psi}^* = \underset{\mathbf{\Phi}, \mathbf{\Psi}}{\operatorname{argmin}} \mathcal{L}(\mathbf{P}_t \mathcal{F} \mathcal{M}_{\mathbf{\Phi}}(\mathcal{M}_{\mathbf{\Psi}} \circ \mathbf{p}), \mathbf{y}_t) + \alpha \cdot \text{TV}(\mathcal{M}_{\mathbf{\Psi}}), \tag{7}$$

where $\mathcal{M}_{\mathbf{\Psi}} \circ \mathbf{p} = \mathbf{p}'$ denotes the transformed coordinates and $\text{TV}(\cdot)$ represents the TV regularization applied to the estimated DVFs.

The joint optimization of $\mathbf{\Phi}$ and $\mathbf{\Psi}$ introduces indeterminacy that makes the problem more challenging. To overcome this issue, we use the optimized parameters $\mathbf{\Phi}^*_{\text{init}}$ of motion-ignoring static initialization learning as the good initialization of the parameters $\mathbf{\Phi}$ of the canonical network $\mathcal{M}_{\mathbf{\Phi}}$. As a result, the canonical network leverages low-frequency spatial information to guide the learning of the DVF network. The gradually learned DVFs further enables the canonical network to reduce the motion artifacts and capture fine details. Finally, the DVF network $\mathcal{M}_{\mathbf{\Psi}}$ accurately estimates the motion transformation, while the canonical network $\mathcal{M}_{\mathbf{\Phi}}$ represents a high-quality canonical volume.

3 Experiments

3.1 Experiments Setup

Datasets and Pre-processing. Under an institutional review board-approved protocol, two 10-min DCE-MRI scans were performed using a work-in-progress golden-angle stack-of-stars spoiled gradient echo sequence [2,3] with fat suppression, yielding a total of 3500 radial stacks of spokes. We applied the hierarchical motion modeling method [23] to determine the continuous respiratory position

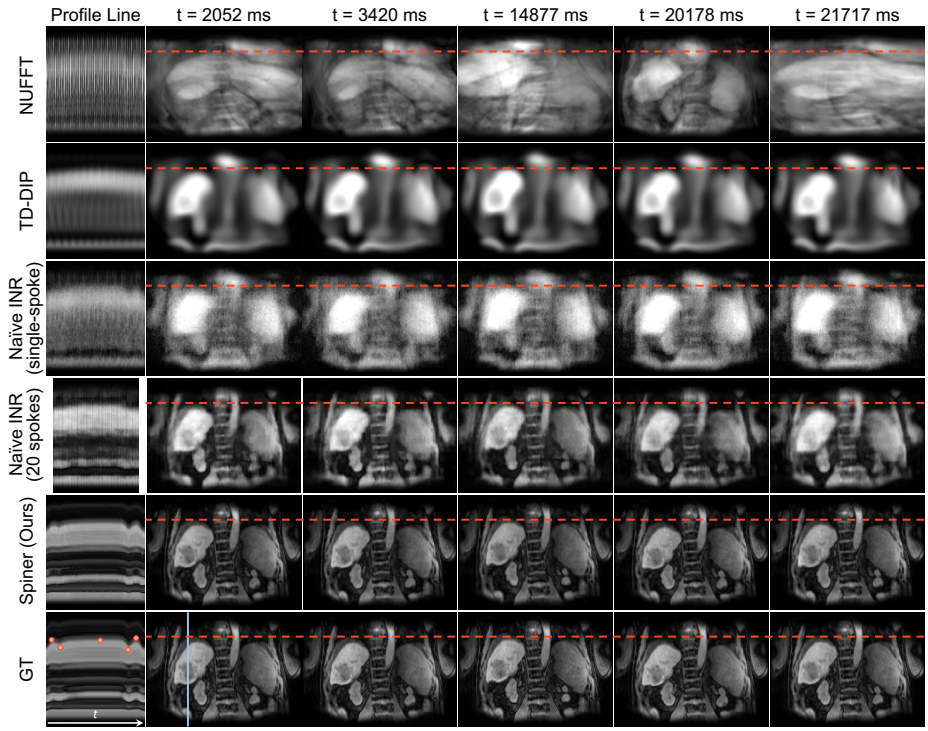

Fig. 2. Qualitative results of dynamic 3D MRI images reconstructed by the compared methods. The first column shows the blue line profile over time. The other columns show the reconstructed images at the time points corresponding to the orange points on the ground truth profile. Orange dashed lines help visualize the motion changing.

for each stack of spokes and to jointly reconstruct image volumes via view sharing across 21 breathing states. These volumes were then aligned to the exhale position volume to yield DVFs. Each spoke's DVF was derived by linearly interpolating its respiratory position between the two bounding states in this 21-state vector field. By applying these DVFs to the static volume, we generated the ground truth 4D dynamic MRI. From these two ground truth 4D dynamic MRI datasets, we extracted five subsets, each comprising 140 radial stacks of spokes, for further simulation. We employed a golden-angle stack-of-stars sampling pattern, acquiring only one stack of spoke per time-specific volume, with spoke length $= 448$ and $k_z = 96$.

Compared Methods and Metrics. We compare our method with NUFFT [8], TD-DIP [22], and Naïve INR [4]. TD-DIP [22] is applied in a single-spoke setting, while Naïve INR [4] is evaluated under both single-spoke and window-shift conditions, using a window size of 20 (i.e., 20 spokes per time point). For the

Table 1. Quantitative comparison of dynamic 3D images by all methods. The best performance is highlighted in **bold**, and the second best is underlined.

	NUFFT [8]	TD-DIP [22]	Naïve INR [4] (single-spoke)	Naïve INR [4] (20 spokes)	SPINER (w/o Init.)	SPINER (w/ Init.)
PSNR	10.64	23.83	19.16	27.13	32.59	**38.99**
SSIM	0.3309	0.6539	0.2113	0.7339	0.7564	**0.9664**

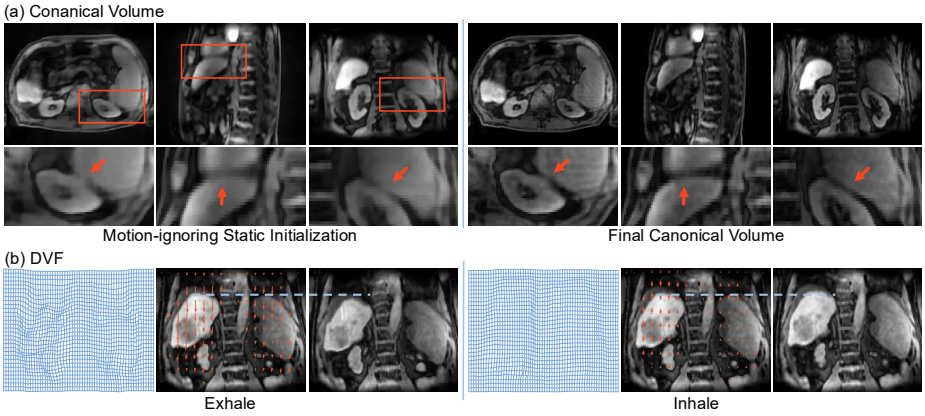

Fig. 3. Qualitative results of canonical volumes and DVFs. (a) The canonical volume of motion-ignoring initialization (left) and the canonical volume estimation after optimization (right). The volume is shown in three orthogonal views. (b) Illustration of DVF (first column), canonical volume with quiver vectors indicating the DVF direction (second column) and warped image (third column).

reconstructed MRI images, we employ peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) as quantitative evaluation metrics.

Implementation Details. The canonical MLP network $\mathcal{M}_{\mathbf{\Phi}}$ consists of five FC layers with a width of 128, while the DVF network $\mathcal{M}_{\mathbf{\Psi}}$ contains two FC layers with a width of 32. The hyperparameter of TV regularization is set as $\lambda = 0.1$ and $\alpha = 0.4$.

3.2 Results

Comparisons of Dynamic 3D Image Reconstruction. Fig. 2 shows the qualitative results of reconstruction. NUFFT appears as a projection at a specific angle due to the single-spoke motion model. Both TD-DIP and Naïve INR capture the overall organ structure but lack finer details, with Naïve INR exhibiting more noise. In contrast, our method reconstructs dynamic 3D MRI with high spatiotemporal resolution. Regarding motion capture, the ground truth profile line clearly shows a small-amplitude breathing cycle in the early time and a

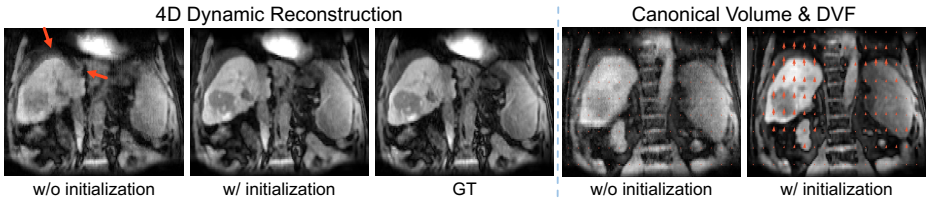

Fig. 4. Comparisons of SPINER with or without motion-ignoring static initialization.

larger-amplitude cycle later on. The profile lines of NUFFT and TD-DIP show that they fail to accurately represent motion, while the profile line of Naïve INR shows that it captures only the larger-amplitude motion. In contrast, our approach precisely captures both types of motion, demonstrating superior performance in dynamic imaging. Table 1 shows the quantitative results aligning with the above qualitative comparisons.

Results of Learned Canonical Volume and DVFs. Fig. 3.(a) shows that during the motion-ignoring static initialization, the INR network learns a volume with anatomical structures but also exhibits motion blurring, visible as unclear boundaries. The subsequent optimization process progressively corrects these motion artifacts, enabling the canonical network to ultimately represent a high-quality static image free from motion-induced blurring. Fig. 3.(b) shows the learned DVFs. For both the inhale and exhale phases, the resulting DVFs are consistent with the expected organ movements, demonstrating our method accurately capture respiratory motion.

3.3 Ablation Studies

Effectiveness of Motion-ignoring Static Initialization. Fig. 4 shows the qualitative results with and without initialization. Without initialization, the DVF network fails to capture the motion, causing the canonical network to ultimately learn a volume with motion artifacts. This comparison demonstrates that the motion-ignoring static initialization helps the DVF network focus on capturing dynamics, and the learned DVF further enables the canonical network to reduce motion artifacts.

Effectiveness of Single-spoke Motion Model. The third row of Fig. 2 (Naïve INR with single-spoke) shows that Naïve INR can capture large motion in the later phase, although the spatial resolution is low. In contrast, the fourth row of Fig. 2 (Naïve INR with 20 spokes) shows that increasing the number of spokes improves image quality but also causes blurring artifacts and makes motion less discernible. These results demonstrate that the proposed single-spoke motion model can effectively represent the real-world motion.

4 Conclusion and Discussion

This work proposes SPINER, a novel single-spoke motion-compensated dynamic 3D reconstruction method. Unlike existing methods, SPINER introduces a single-spoke motion model and addresses this highly ill-posed problem by proposing the motion-ignoring static initialization strategy that exploits the redundancy of static anatomical information across all spokes. Experimental results demonstrate that SPINER reconstruct images with high spatiotemporal resolution while accurately capturing motion. Despite the success of the current framework, which only considers respiratory motion, it can be further extended to reconstruct dynamic 3D MRI with more complex, composite motion.

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

References

1. Baik, D., Yoo, J.: Dynamic-aware spatio-temporal representation learning for dynamic mri reconstruction. arXiv preprint arXiv:2501.09049 (2025)
2. Block, K.T., Chandarana, H., Milla, S., Bruno, M., Mulholland, T., Fatterpekar, G., Hagiwara, M., Grimm, R., Geppert, C., Kiefer, B., et al.: Towards routine clinical use of radial stack-of-stars 3d gradient-echo sequences for reducing motion sensitivity. Investigative Magnetic Resonance Imaging **18**(2), 87–106 (2014)
3. Chandarana, H., Block, K.T., Winfeld, M.J., Lala, S.V., Mazori, D., Giuffrida, E., Babb, J.S., Milla, S.S.: Free-breathing contrast-enhanced t1-weighted gradient-echo imaging with radial k-space sampling for paediatric abdominopelvic mri. European radiology **24**, 320–326 (2014)
4. Feng, J., Feng, R., Wu, Q., Shen, X., Chen, L., Li, X., Feng, L., Chen, J., Zhang, Z., Liu, C., Zhang, Y., Wei, H.: Spatiotemporal implicit neural representation for unsupervised dynamic mri reconstruction. IEEE Transactions on Medical Imaging pp. 1–1 (2025). https://doi.org/10.1109/TMI.2025.3526452
5. Feng, L., Axel, L., Chandarana, H., Block, K.T., Sodickson, D.K., Otazo, R.: Xd-grasp: Golden-angle radial mri with reconstruction of extra motion-state dimensions using compressed sensing. Magnetic Resonance in Medicine **75**(2), 775–788 (2016). https://doi.org/https://doi.org/10.1002/mrm.25665, https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.25665
6. Feng, L., Grimm, R., Block, K.T., Chandarana, H., Kim, S., Xu, J., Axel, L., Sodickson, D.K., Otazo, R.: Golden-angle radial sparse parallel mri: Combination of compressed sensing, parallel imaging, and golden-angle radial sampling for fast and flexible dynamic volumetric mri. Magnetic Resonance in Medicine **72**(3), 707–717 (2014). https://doi.org/https://doi.org/10.1002/mrm.24980, https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.24980
7. Feng, L., Tyagi, N., Otazo, R.: Mrsigma: Magnetic resonance signature matching for real-time volumetric imaging. Magnetic Resonance in Medicine **84**(3), 1280–1292 (2020). https://doi.org/https://doi.org/10.1002/mrm.28200, https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.28200
8. Fessler, J.A., Sutton, B.P.: Nonuniform fast fourier transforms using min-max interpolation. IEEE transactions on signal processing **51**(2), 560–574 (2003)

9. Kettelkamp, J., Romanin, L., Piccini, D., Priya, S., Jacob, M.: Motion compensated unsupervised deep learning for 5d mri. In: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2023. pp. 419–427. Springer Nature Switzerland, Cham (2023)
10. Lingala, S.G., DiBella, E., Jacob, M.: Deformation corrected compressed sensing (dc-cs): A novel framework for accelerated dynamic mri. IEEE Transactions on Medical Imaging **34**(1), 72–85 (2015). https://doi.org/10.1109/TMI.2014.2343953
11. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. ACM Trans. Graph. **41**(4), 102:1–102:15 (Jul 2022). https://doi.org/10.1145/3528223.3530127, https://doi.org/10.1145/3528223.3530127
12. Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. ICCV (2021)
13. Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-NeRF: Neural Radiance Fields for Dynamic Scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2020)
14. Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F., Bengio, Y., Courville, A.: On the spectral bias of neural networks. In: Chaudhuri, K., Salakhutdinov, R. (eds.) Proceedings of the 36th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 97, pp. 5301–5310. PMLR (09–15 Jun 2019), https://proceedings.mlr.press/v97/rahaman19a.html
15. Riederer, S.J., Tasciyan, T., Farzaneh, F., Lee, J.N., Wright, R.C., Herfkens, R.J.: Mr fluoroscopy: Technical feasibility. Magnetic Resonance in Medicine **8**(1), 1–15 (1988). https://doi.org/https://doi.org/10.1002/mrm.1910080102, https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.1910080102
16. Stemkens, B., Paulson, E.S., Tijssen, R.H.N.: Nuts and bolts of 4d-mri for radiotherapy. Physics in Medicine & Biology **63**(21), 21TR01 (oct 2018). https://doi.org/10.1088/1361-6560/aae56d, https://dx.doi.org/10.1088/1361-6560/aae56d
17. Tian, X., Chen, L., Wu, Q., Du, C., Shi, J., Wei, H., Zhang, Y.: Unsupervised self-prior embedding neural representation for iterative sparse-view ct reconstruction. Proceedings of the AAAI Conference on Artificial Intelligence **39**(7), 7383–7391 (Apr 2025). https://doi.org/10.1609/aaai.v39i7.32794, https://ojs.aaai.org/index.php/AAAI/article/view/32794
18. Usman, M., Atkinson, D., Odille, F., Kolbitsch, C., Vaillant, G., Schaeffter, T., Batchelor, P.G., Prieto, C.: Motion corrected compressed sensing for free-breathing dynamic cardiac mri. Magnetic Resonance in Medicine **70**(2), 504–516 (2013). https://doi.org/https://doi.org/10.1002/mrm.24463, https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.24463
19. Wu, Q., Du, C., Tian, X., Yu, J., Zhang, Y., Wei, H.: Moner: Motion correction in undersampled radial MRI with unsupervised neural representation. In: The Thirteenth International Conference on Learning Representations (2025), https://openreview.net/forum?id=0dnqG1fYpo
20. Xu, Z.Q.J., Zhang, Y., Luo, T., Xiao, Y., Ma, Z.: Frequency principle: Fourier analysis sheds light on deep neural networks. arXiv preprint arXiv:1901.06523 (2019)
21. Yang, H., Hong, K., Baraboo, J.J., Fan, L., Larsen, A., Markl, M., Robinson, J.D., Rigsby, C.K., Kim, D.: Grasp reconstruction amplified with view-

sharing and kwic filtering reduces underestimation of peak velocity in highly-accelerated real-time phase-contrast mri: A preliminary evaluation in pediatric patients with congenital heart disease. Magnetic Resonance in Medicine **91**(5), 1965–1977 (2024). https://doi.org/https://doi.org/10.1002/mrm.29974, https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.29974

22. Yoo, J., Jin, K.H., Gupta, H., Yerly, J., Stuber, M., Unser, M.: Time-dependent deep image prior for dynamic mri. IEEE Transactions on Medical Imaging **40**(12), 3337–3348 (2021). https://doi.org/10.1109/TMI.2021.3084288
23. Zhang, Y., Kashani, R., Cao, Y., Lawrence, T.S., Johansson, A., Balter, J.M.: A hierarchical model of abdominal configuration changes extracted from golden angle radial magnetic resonance imaging. Physics in Medicine & Biology **66**(4), 045018 (feb 2021). https://doi.org/10.1088/1361-6560/abd66e, https://dx.doi.org/10.1088/1361-6560/abd66e