

Self-Supervised k-Space Regularization for Motion-Resolved Abdominal MRI Using Neural Implicit k-Space Representations

Veronika Spieker^{1,2,3*}, Hannah Eichhorn^{1,2}, Jonathan K. Stelter⁴,
Wenqi Huang⁴, Rickmer F. Braren^{4,5}, Daniel Rueckert^{2,4,6}, Francisco
Sahli Costabal^{3,7}, Kerstin Hammernik², Claudia Prieto^{3,7,8},
Dimitrios C. Karampinos⁴, and Julia A. Schnabel^{1,2,8}

¹ Institute of Machine Learning in Biomedical Imaging, Helmholtz Munich, Germany

² School of Computation, Information & Technology,
Technical University of Munich, Germany

³ Millenium Institute for Intelligent Healthcare Engineering, Chile

⁴ School of Medicine & Health, Klinikum rechts der Isar,
Technical University of Munich, Germany

⁵ German Cancer Consortium (DKTK, Partner Site Munich) &
German Cancer Research Center (DKFZ) Heidelberg, Germany

⁶ Department of Computing, Imperial College London, UK

⁷ School of Engineering, Pontificia Universidad Católica de Chile, Chile

⁸ School of Biomedical Engineering & Imaging Sciences, King's College London, UK

*v.spieker@tum.de

Abstract. Neural implicit k-space representations have shown promising results for dynamic MRI at high temporal resolutions. Yet, their exclusive training in k-space limits the application of common image regularization methods to improve the final reconstruction. In this work, we introduce the concept of parallel imaging-inspired self-consistency (PISCO), which we incorporate as novel self-supervised k-space regularization enforcing a consistent neighborhood relationship. At no additional data cost, the proposed regularization significantly improves neural implicit k-space reconstructions on simulated data. Abdominal in-vivo reconstructions using PISCO result in enhanced spatio-temporal image quality compared to state-of-the-art methods. Code is available at <https://github.com/compai-lab/2024-miccai-spieker>.

Keywords: Dynamic MRI Reconstruction · Parallel Imaging · k-Space Refinement · Self-Supervised Learning · Implicit Neural Representations

1 Introduction

Motion, i.e. induced by respiration in the abdomen, results in non-negligible artefacts in magnetic resonance imaging (MRI). Therefore, motion-sensitive applications such as radiation therapy planning [18] or free-breathing high-resolution diagnostic imaging [22] rely on dynamic MRI reconstruction techniques. These

techniques reconstruct temporally resolved images by binning the data into different respiratory motion states (MS). However, reconstructing multiple MS inherently results in a reduction of data per temporal image, leading to undersampling artifacts due to violation of the Nyquist theorem. Common approaches address this by applying regularization across the temporal dimension to exploit redundancies [6,19]. Yet, the total number of MS is limited, resulting in motion blurring due to low temporal resolution.

Recently, neural implicit k-space representations (NIK) have shown promising results for blurring-free motion-resolved reconstructions [17,9]. Based on the acquired k-space trajectory and a surrogate signal for the current motion state, a multi-layer perceptron (MLP) is trained to predict the k-space signal for a given spatio-temporal input. At inference, any coordinate can be sampled, allowing for flexible spatial sampling patterns at high temporal resolutions. Exclusive training in k-space eliminates the need for costly domain transforms such as non-uniform Fast Fourier Transformations (NUFFT) within each iteration.

Similar to other accelerated reconstruction methods [2,7], NIK learns to fill missing data points within k-space. Yet, without any kind of regularization, the reconstruction model may be prone to overfitting and noisy reconstructions. General learning-based MRI reconstruction methods counteract overfitting by including a regularization, usually enforced on the estimated image [1,8,10]. However, image-based regularization methods are impractical for NIK, as they necessitate sampling of the entire k-space within each training iteration to obtain the image that requires regularization. While k-space-based regularization would be effective, translation of image-based constraints to k-space is not trivial.

Exploring the parallel imaging concept of Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA) [7] reveals a potential spatial neighborhood relationship within k-space itself. This relationship needs to be estimated on a fully-sampled calibration set to be subsequently applied to undersampled k-space regions. Learning-based reconstruction methods already utilize this neighborhood relationship for k-space refinement [12,17]. However, similar to GRAPPA, they first require explicit determination of the k-space relationship. This is impractical for motion-resolved imaging, since calibration needs to be acquired and conducted for every MS.

In this work, we reformulate the concept of k-space relationship independent of calibration data. We exploit the inherent global k-space relationship without the need for explicit determination. Our contributions are three-fold:

1. We introduce PISCO: a novel Parallel Imaging-inspired Self-Consistency for self-supervised k-space refinement that operates independently of any additional training or calibration data.
2. Incorporating the concept of PISCO, we present the first k-space-based regularization loss for neural implicit k-space representations (NIK).
3. Based on a realistic motion simulation as well as in-vivo data, we quantitatively and qualitatively demonstrate the potential of PISCO for motion-resolved abdominal MR reconstruction using NIK.

2 Methods

2.1 k-Space Interpolation using GRAPPA

In MRI, a k-space signal $y = \{y_i \in \mathbb{C}^{N_c} | i = 1, \dots, N_x \cdot N_y\}$ is acquired with N_c coils at spatial k-space coordinates $k = \{k_i \in \mathbb{R}^2 | i = 1, \dots, N_x \cdot N_y\}$ (no time dimension for simplicity). The desired image $x \in \mathbb{C}^{N_x \cdot N_y}$ is then reconstructed by solving the inverse problem $y = Ax$, where the forward operator $A = \mathcal{FS}$ consists of the coil sensitivity maps $\mathbf{S} \in \mathbb{C}^{N_x \cdot N_y}$ and Fourier transform $\mathcal{F}$. For acceleration purposes the acquired k-space is often undersampled, i.e., y is acquired at fewer k_i than required to fulfill the Nyquist theorem. To avoid undersampling artefacts, parallel imaging methods such as GRAPPA [7] leverage the multi-coil nature of MRI data to derive missing values from its neighboring k-space values.

For GRAPPA, a patch of N_n neighboring coordinates is sampled around a missing target $k_i^T \in \mathbb{R}^2$, resulting in a coordinate patch $k_i^P \in \mathbb{R}^{N_n \cdot 2}$. The target signal value $y_i^T \in \mathbb{C}^{N_c}$ is then obtained by linearly combining the neighboring signal values $y_i^P \in \mathbb{C}^{N_n \cdot N_c}$ (Fig. 1A). Stacking N_m pairs of targets and patches, the linear combination can be written as $T = PW$, where $T = [y_1^T, \dots, y_{N_m}^T] \in \mathbb{C}^{[N_m \times N_c]}$, $P = [y_1^P, \dots, y_{N_m}^P] \in \mathbb{C}^{[N_m \times N_n \cdot N_c]}$ and $W \in \mathbb{C}^{[N_n \cdot N_c \times N_c]}$ is the global weight matrix with a total of $N_w = N_n \cdot N_c \cdot N_c$ weights. To determine W , GRAPPA splits the signal space y into a fully sampled auto-calibration signal y_{ACS} and the remaining undersampled k-space y_{US} . Within the fully sampled y_{ACS} , pairs of T_{ACS} and P_{ACS} can be used to solve the least-squares problem:

$$W_{ACS} = \arg \min_W \|P_{ACS}W - T_{ACS}\|_2^2 + \alpha \|W\|_2^2 \quad \text{s.t. } T_{ACS}, P_{ACS} \subseteq y_{ACS} \quad (1)$$

where $\|\cdot\|_2^2$ applies the L2-norm element-wise and α weighs the Tikhonov regularization. Subsequently, undersampled points y_{US}^T are derived using the calibrated weights to compute $y_{US}^T = W_{ACS} \cdot y_{US}^P$, leading to improved reconstructions [7].

2.2 From GRAPPA to PISCO

Acquiring a fully sampled region y_{ACS} to calibrate a weight set W is not always feasible, particularly in dynamic imaging. Hence, we reframe GRAPPA's concept of global spatial k-space relationship to a calibration-free condition based on the following assumption: If a weight set W_{ACS} calibrated on $y_{ACS} \in y$ models the linear relationship for the whole k-space y , then a weight set W_s derived from a random subset $y_s \in y$ should result in the same linear relationship. Consequently, weight sets W_s determined from multiple random subsets $\{y_s \in y | s = 1, \dots, N_s\}$ are expected to converge to the same solution, i.e. $W_i = W_j, \forall i \neq j$. Thus, we propose to refine an interpolated k-space without access to a calibration set by enforcing **Parallel Imaging-inspired Self-Consistency** (PISCO) as follows:

$$W_s = \arg \min_W \|P_s W - T_s\|_2^2 + \alpha \|W\|_2^2 \quad \text{s.t. } T_s, P_s \subseteq y_s \quad (2)$$

$$\text{PISCO: } \sum_{i=1}^{N_s} \sum_{j=1}^{N_s} \text{distance}(W_i, W_j) \xrightarrow{!} 0 \quad (3)$$




Fig. 1. Overview. A: Multiple pairs of targets and surrounding neighbors y_i^T, y_i^P are sampled and randomly sorted into subsets and solved for the linear relationship W_s (Eq. 2). PISCO aims to minimize the distance between all W_s . For simplicity, the coil dimension N_c is not visualized, but included in matrix dimensions. B: PISCO for regularization of NIK. Any points can be sampled, but L_{DC} can only be compared to measured k-space y_{meas} (gray lines). L_{PISCO} refines independent of y_{meas} .

2.3 Dynamic MRI Reconstruction using NIK with PISCO

In dynamic MRI reconstruction, the image reconstruction problem from Sec. 2.1 is expanded by a temporal dimension t , i.e. k-space coordinates k consist of $k_i = [k_x, k_y, t]$ and result in multiple temporal images $x = \{x_t | t = 1, 2, \dots, N_t\}$.

Within each training iteration of NIK, a coordinate batch $k_B = \{k_i\}_{i=1}^{N_b}$ is sampled from acquired coordinates k_{meas} to predict $y_B^{pred} = G_\theta(k_B)$ (Fig. 1B). Then, the model is optimized with a data consistency loss $\mathcal{L}_{DC}$ comparing the predicted y_B^{pred} and measured signal y_B^{meas} . This limits NIK’s training strategy to actually acquired points. By adding the concept of PISCO during training we enable further self-supervision to coordinates that have not been acquired, which is especially beneficial for addressing large gaps in the outer regions of k-space in radial trajectories. As in Sec. 2.1, we assign k_B as targets k_B^T and sample spatial neighbors k_B^P . To account for the dynamic nature, we sort the coordinates along t to allow for separation into N_s temporally adjacent subsets [3] with N_m coordinates each. The resulting T_s and P_s are used to solve for $[W_s]_{s=1}^{N_s}$ (Eq. 2). To fulfill Eq. 3, we introduce the PISCO loss:

$$\mathcal{L}_{PISCO} = \frac{1}{N_s^2} \sum_{i=1}^{n_s} \sum_{j=1}^{n_s} \mathcal{L}_{dist}(W_i, W_j) \quad (4)$$

where $\mathcal{L}_{dist}$ is chosen to be the complex difference $\mathcal{L}_C^1 = ||Re(W_i - W_j)||_1 + ||Im(W_i - W_j)||_1$. After pre-training NIK for E_{pre} epochs, the self-supervised consistency is enforced by optimizing the model’s weights θ with $\mathcal{L}_{DC}$ and $\mathcal{L}_{PISCO}$

Algorithm 1 Neural Implicit k-Space Representation with PISCO

Require: Acquired k-space coordinates k^{meas} and signal values y^{meas} , total epochs E , pretraining epochs E_{pre} , NIK architecture G_θ , initial model parameters θ_0

- 1: **for** $e = 0$ to E **do**
- 2: **for** batch (k_B, y_B) in (k^{meas}, y^{meas}) **do** ▷ Sample batch of coordinates
- 3: $y_B^{pred} \leftarrow G_{\theta_e}(k_B)$ ▷ Predict k-space with NIK
- 4: $\theta_{e+1} \leftarrow L_{DC}(y_B^{pred}, y_B)$ ▷ Update model with data consistency loss
- 5: **if** $e > E_{pre}$ **then** ▷ Apply PISCO regularization
- 6: $k_B^T, k_B^P \leftarrow k_B$, patches around k_B ▷ Sample coordinate patches
- 7: $y_B^T, y_B^P \leftarrow G_{\theta_e}(k_B^T), G_{\theta_e}(k_B^P)$
- 8: $[P_s]_{s=1}^{N_s}, [T_s]_{s=1}^{N_s} \leftarrow$ divide y_B^T, y_B^P into multiple subsets ▷ Sec. 2.2
- 9: $[W_s]_{s=1}^{N_s} \leftarrow$ solve each subset for weights with regularization α ▷ Eq. 2
- 10: $\theta_{e+1} \leftarrow \lambda \cdot L_{PISCO}([W_s]_{s=1}^{N_s})$ ▷ Eq. 4
- 11: **return** Learned NIK model G_θ

in an alternating manner, with the latter’s impact weighted by λ . Additional robustness against outliers in T and P is enhanced with Thikonov regularization weight α when solving for W . Moreover, the number of samples per equation system is increased to $N_m = f_{od} \cdot N_w$, with $f_{od} > 1$ determining the ratio of unknowns N_w to available equations N_m . The integration of $\mathcal{L}_{PISCO}$ in NIK’s training procedure is summarized as pseudo-code in Alg. 1.

3 Experimental Setup

3.1 Data

Due to the lack of ground-truth data for motion-resolved imaging [16], we create a dynamic free-breathing *simulation* using the XCAT phantom [13]. We generate water/fat/susceptibility maps [5,11] for 100 time points t_{MS} within one breathing cycle and simulate complex images $\mathbf{x}_t$ based on a water-fat model [21] (echo time $T_e=1.4\text{ms}$). After applying six simulated coil sensitivity maps, radial motion-free multi-coil k-space data Y_{mf} is obtained using NUFFT for acceleration factors $R=1,2,3$. Then, Y_{mf} is merged into one motion-affected k-space Y_{ma} based on a respiratory navigator t extracted from the lung-liver edge.

For in-vivo validation, datasets were acquired using a pseudo golden angle stack-of-star trajectory (voxel size= $1.5 \times 1.5 \times 3 \text{mm}^3$, flip angle= 10°) at 3T (Ingenia Elition X, Philips Healthcare) after local ethics committee approval (*Anonymous). We obtained a *dynamic* free-breathing dataset of the abdomen, as well as a respiratory-gated acquisition from the same subject as reference. To validate the regularization independent of motion, a *static* dataset of the thigh region was acquired and retrospectively accelerated ($R=1,2,3$). Each k-space consists of 26 coils (N_C), 536 frequency encoding steps (N_{FE}), 1341/537/537 spokes (N_{PE}) for dynamic/gated/static. Coil sensitivities are estimated with ESPIRiT [20].

3.2 Training and Inference

We adapt NIK’s [9] architecture using 4 layers, 512 hidden features, high-dynamic range loss as $\mathcal{L}_{DC}$, SIREN activations [15], batch size of 10k and use STIFF feature encoding [4]. We correct the respiratory signal t for linear drifts and rescale it to $[0,0.5]$. For PISCO-NIK, L_{PISCO} is included after $E_{pre} = 200$ and weighted to match the magnitude of L_{DC} , i.e., $\lambda=0.01/0.1$ for simulation/in-vivo data. For the neighbors P, a kernel of size $[3,3]$ with $\delta x/\delta y = N_{FE}^{-1}$ is sampled around the target T resulting in $N_n = 8$, as shown in Fig. 1. Hence, N_w for the phantom is 288 ($8 \cdot 6 \cdot 6$) and for the subject would be 5408 ($8 \cdot 26 \cdot 26$). To increase the number of possible sets for the subject ($N_s \propto N_w^{-1}$), we reduce the number of output coils solved for in each iteration, thereby lowering N_w to 624 ($8 \cdot 26 \cdot 3$). Weight solving for L_{PISCO} is regularized with $\alpha=1e-4$ and $f_{OD}=1.1$. Both losses are optimized with separate Adam optimizers (lr=3e-5). All models are run for 1000 epochs total (NVIDIA RTX A6000, using Python 3.10.1/PyTorch 1.13.1).

3.3 Evaluation

Dynamic reconstructions using L_{PISCO} (PISCO-NIK) are compared to vanilla NIK (same hyperparameters, only L_{DC}), inverse NUFFT for 1 or 4MS (INUFFT/INUFFT4) and state-of-the-art motion-resolved XD-GRASP [6] using the standard 4 MS as well as 50 MS to match the temporal resolution of NIK (XD-GRASP4 and XD-GRASP50). Quantitative metrics are computed in reference to the motion-resolved k_{mf} (simulation) and INUFFT-R1 (static). Calculated metrics include peak signal-to-noise ratio (PSNR) and spatial feature similarity (FSIM) of xy -reconstructions at 50 time points. To evaluate temporal performance, FSIM-t is calculated for the simulation on xt/yt with y/x fixed.

4 Results

Quantitative reconstruction results of 20 slices of the simulation for R=1,2,3 are shown in Fig. 2. Videos of dynamic reconstructions are included in Suppl.




Fig. 2. Quantitative results for 20 motion-affected simulation slices accelerated by R=1,2,3. All comparisons, except those marked with "N", are statistically significant (Wilcoxon signed rank test with False Discovery Rate correction at $p < 0.05$).




Fig. 3. *Static* in-vivo reconstructions (thighs) for accelerations $R=1,2,3$. Exclusion of the temporal component enables comparison to INUFFT-R1, showing increased PSNR using PISCO regularization. PISCO-NIK sharpens reconstructions (blue arrows and increased FSIM) and reduces ringing artefacts of NIK-R2 (green arrows).

Fig. 1. Both neural implicit representations, NIK and PISCO-NIK, clearly outperform XD-GRASP50 (i.e. with same temporal resolution). XD-GRASP4 results in higher PSNR due to lower temporal resolution, but significantly lower spatial FSIM due to residual blurring. Particularly with increasing acceleration ($R=2/R=3$), PISCO-NIK results in a significant increase of PSNR (up to 1.1dB), FSIM and FSIM-t (up to 0.01/0.02).




Fig. 4. *Dynamic* in-vivo reconstructions for two abdominal slices, with spatial xy -image (fixed t) and temporal yt -image (fixed x at white dotted line). Complete reconstruction videos in Suppl. Fig. 2. Temporally resolved reconstructions are not available for the gated reference and limited for INUFFT4/XD-GRASP4 (blue arrows). XD-GRASP50 results in streaking, but more precise vessel structure, likely due to less motion blurring (white arrow). NIK enables improved temporal resolution, whereas PISCO additionally temporally smoothens (blue arrow) while maintaining spatial sharpness (white arrow).

Static reconstruction results are shown in Fig. 3. Similar to the simulation, including PISCO results in higher PSNR and leads to sharper results, specifically with increasing acceleration factors. Note that for the static case all data are considered as one MS, therefore naturally resulting in lower undersampling for INUFFT compared to the dynamic case with multiple MS. In Fig. 4, motion-resolved reconstructions of the dynamic in-vivo data demonstrate the benefit of neural implicit representations regarding temporal resolution (*yt* images of NIK and PISCO-NIK, also visible in videos in Suppl. Fig. 2). Since no ground truth reconstruction is available for dynamic imaging, an unpaired gated scan is included as comparison. It results in sharp spatial images, but no temporal resolution. XD-GRASP4 smoothens results compared to INUFFT4, but residual motion blurring and discrete jumps in the temporal dimension persist (blue arrow). XD-GRASP50 encounters these visible jumps, but results in streaking. NIK allows for high temporally resolved but noisy reconstructions and the proposed PISCO regularization enforces smoother results with sharper vessel structures.

5 Discussion and Conclusion

In this work, we have presented the novel concept of parallel imaging-inspired self-consistency (PISCO), which leverages global k-space relationship without the use of calibration data to explicitly determine it. We have shown the seamless integration of PISCO as self-supervised k-space regularization in NIK-based reconstructions, eliminating the need for additional data or training steps. Both simulated and in-vivo dynamic reconstruction results illustrate that PISCO enhances spatial and temporal denoising while preserving temporal resolution and sharp vessel structures. Additionally, we have verified PISCO’s spatial denoising potential on a static example, highlighting its capacity to learn improved neural representations resulting in enhanced image quality.

Generally, NIK’s flexibility in sampling and PISCO’s calibration-free nature allow for a flexible kernel design. This is particularly beneficial for radial acquisitions, where GRAPPA calibration with a radial kernel would demand multiple ACS acquisitions due to the limitation of neighborhood relationships to small radial sections [14]. In contrast, our work illustrates how the proposed PISCO-NIK approach effectively utilizes Cartesian relationships to refine radial acquisitions.

Adaptation of kernel design to the undersampling pattern may further enhance the effectiveness of PISCO. Notably, improved performance can be observed for accelerated reconstructions, particularly for $R=2$. This improvement can be attributed to the kernel design used, which samples one adjacent neighbor in each direction to the available acquisition points. Thereby, it effectively covers points not present with $R=2$. For higher acceleration factors, larger k-space gaps persist, potentially requiring a larger kernel size to focus PISCO’s attention during model training at these gaps. Kernel design as well as hyperparameter tuning for PISCO are subject to further investigation.

By introducing a regularization term for neural implicit k-space representations for the first time, we have improved its potential for motion-resolved

abdominal imaging. Yet, training times and the need for re-training of neural implicit k-space representations remain a challenge for MR reconstruction [16]. Moreover, the reliability of the learned representation depends on a reliable surrogate motion signal, introducing uncertainties. Future research addresses these uncertainties as well as methods to accelerate the training procedure.

Concluding, we have demonstrated how a conventional parallel imaging concept can be leveraged as self-supervised regularization for learning-based reconstruction. Due to its calibration-free and flexible design, our proposed method PISCO can be seamlessly integrated into the clinical workflow, making it an attractive regularization method for application in further anatomies or k-space-based reconstruction methods.

Acknowledgments. V.S. and H.E. are partially supported by the Helmholtz Association under the joint research school "Munich School for Data Science - MUDS".

Disclosure of Interests. V.S. received the Bayer Foundation Fellowship during part of this work.

References

1. Ahmad, R., Bouman, C.A., Buzzard, G.T., Chan, S., Liu, S., Reehorst, E.T., Schniter, P.: Plug-and-Play Methods for Magnetic Resonance Imaging: Using Denoisers for Image Recovery. *IEEE Signal Processing Magazine* **37**(1), 105–116 (2020). <https://doi.org/10.1109/msp.2019.2949470>
2. Akçakaya, M., Moeller, S., Weingärtner, S., Uğurbil, K.: Scan-specific robust artificial-neural-networks for k-space interpolation (RAKI) reconstruction: Database-free deep learning for fast imaging. *Magnetic Resonance in Medicine* **81**(1), 439–453 (2019). <https://doi.org/10.1002/mrm.27420>
3. Breuer, F.A., Kellman, P., Griswold, M.A., Jakob, P.M.: Dynamic autocalibrated parallel imaging using temporal GRAPPA (TGRAPPA). *Magnetic Resonance in Medicine* **53**(4), 981–985 (2005). <https://doi.org/10.1002/mrm.20430>
4. Catalán, T., Courdurier, M., Osses, A., Botnar, R., Costabal, F.S., Prieto, C.: Unsupervised reconstruction of accelerated cardiac cine MRI using Neural Fields (25072023), <http://arxiv.org/pdf/2307.14363v1>
5. Collins, C.M., Yang, B., Yang, Q.X., Smith, M.B.: Numerical calculations of the static magnetic field in three-dimensional multi-tissue models of the human head. *Magnetic resonance imaging* **20**(5), 413–424 (2002). [https://doi.org/10.1016/s0730-725x\(02\)00507-6](https://doi.org/10.1016/s0730-725x(02)00507-6)
6. 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/10.1002/mrm.25665>
7. Griswold, M.A., Jakob, P.M., Heidemann, R.M., Nittka, M., Jellus, V., Wang, J., Kiefer, B., Haase, A.: Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magnetic Resonance in Medicine* **47**(6), 1202–1210 (2002). <https://doi.org/10.1002/mrm.10171>
8. Hammernik, K., Kustner, T., Yaman, B., Huang, Z., Rueckert, D., Knoll, F., Akçakaya, M.: Physics-Driven Deep Learning for Computational Magnetic Resonance Imaging: Combining physics and machine learning for improved medical imaging. *IEEE Signal Processing Magazine* **40**(1), 98–114 (2023). <https://doi.org/10.1109/msp.2022.3215288>
9. Huang, W., Li, H.B., Pan, J., Cruz, G., Rueckert, D., Hammernik, K.: Neural Implicit k-Space for Binning-Free Non-Cartesian Cardiac MR Imaging. In: Frangi, A., de Bruijne, M., Wassermann, D., Navab, N. (eds.) *Information Processing in Medical Imaging*. pp. 548–560. Springer Nature Switzerland, Cham (2023)
10. Jafari, R., Do, R.K.G., LaGratta, M.D., Fung, M., Bayram, E., Cashen, T., Otazo, R.: GRASPNET: Fast spatiotemporal deep learning reconstruction of golden-angle radial data for free-breathing dynamic contrast-enhanced magnetic resonance imaging. *NMR in biomedicine* **36**(3), e4861 (2023). <https://doi.org/10.1002/nbm.4861>
11. Maril, N., Collins, C.M., Greenman, R.L., Lenkinski, R.E.: Strategies for shimming the breast. *Magnetic Resonance in Medicine* **54**(5), 1139–1145 (2005). <https://doi.org/10.1002/mrm.20679>
12. Ryu, K., Alkan, C., Choi, C., Jang, I., Vasanawala, S.: K-space refinement in deep learning MR reconstruction via regularizing scan specific SPIRiT-based self consistency. In: *2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*. IEEE (2021). <https://doi.org/10.1109/iccvw54120.2021.00446>

13. Segars, W.P., Sturgeon, G., Mendonca, S., Grimes, J., Tsui, B.M.W.: 4D XCAT phantom for multimodality imaging research. *Medical Physics* **37**(9), 4902–4915 (2010). <https://doi.org/10.1118/1.3480985>
14. Seiberlich, N., Ehses, P., Duerk, J., Gilkeson, R., Griswold, M.: Improved radial GRAPPA calibration for real-time free-breathing cardiac imaging. *Magnetic Resonance in Medicine* **65**(2), 492–505 (2011). <https://doi.org/10.1002/mrm.22618>
15. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit Neural Representations with Periodic Activation Functions. *Advances in Neural Information Processing Systems* **33**, 7462–7473 (2020)
16. Spieker, V., Eichhorn, H., Hammernik, K., Rueckert, D., Preibisch, C., Karampinos, D.C., Schnabel, J.A.: Deep Learning for Retrospective Motion Correction in MRI: A Comprehensive Review. *IEEE Transactions on Medical Imaging* (2023). <https://doi.org/10.1109/TMI.2023.3323215>
17. Spieker, V., Huang, W., Eichhorn, H., Stelter, J., Weiss, K., Zimmer, V.A., Braren, R.F., Karampinos, D.C., Hammernik, K., Schnabel, J.A.: ICoNIK: Generating Respiratory-Resolved Abdominal MR Reconstructions Using Neural Implicit Representations in k-Space. In: Mukhopadhyay, A., Oksuz, I., Engelhardt, S., Zhu, D., Yuan, Y. (eds) *Deep Generative Models. MICCAI 2023. Lecture Notes in Computer Science*, vol 14533. Springer, Cham, pp. 183–192 (2023). https://doi.org/10.1007/978-3-031-53767-7_{_}18
18. Stemkens, B., Paulson, E.S., Tijssen, R.H.N.: Nuts and bolts of 4D-MRI for radiotherapy. *Physics in medicine and biology* **63**(21), 21TR01 (2018). <https://doi.org/10.1088/1361-6560/aae56d>
19. Terpstra, M.L., Maspero, M., Verhoeff, J.J.C., van den Berg, C.A.T.: Accelerated respiratory-resolved 4D-MRI with separable spatio-temporal neural networks. *Medical Physics* **50**(9), 5331–5342 (2023). <https://doi.org/10.1002/mp.16643>
20. Uecker, M., Lai, P., Murphy, M.J., Virtue, P., Elad, M., Pauly, J.M., Vasanawala, S.S., Lustig, M.: ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA. *Magnetic Resonance in Medicine* **71**(3), 990–1001 (2014). <https://doi.org/10.1002/mrm.24751>
21. Yu, H., Shimakawa, A., McKenzie, C.A., Brodsky, E., Brittain, J.H., Reeder, S.B.: Multiecho water-fat separation and simultaneous R2* estimation with multifrequency fat spectrum modeling. *Magnetic Resonance in Medicine* **60**(5), 1122–1134 (2008). <https://doi.org/10.1002/mrm.21737>
22. Zaitsev, M., Maclaren, J., Herbst, M.: Motion artifacts in MRI: A complex problem with many partial solutions. *Journal of Magnetic Resonance Imaging* **42**(4), 887–901 (2015). <https://doi.org/10.1002/jmri.24850>