

PISCO: Self-supervised k-space regularization for improved neural implicit k-space representations of dynamic MRI [★]

Veronika Spieker ^{ID a,b,c,*}, Hannah Eichhorn ^{ID a,b}, Wenqi Huang ^{ID b}, Jonathan K. Stelter ^{ID d},
 Tabita Catalan ^{ID c,e}, Rickmer F. Braren ^{ID d,f}, Daniel Rueckert ^{ID b,d,g,h},
 Francisco Sahli Costabal ^{ID c,e,i}, Kerstin Hammernik ^{ID b}, Dimitrios C. Karampinos ^{ID d,j,k},
 Claudia Prieto ^{ID c,e,l}, Julia A. Schnabel ^{ID a,b,d,h,l}

^a Institute of Machine Learning in Biomedical Imaging, Helmholtz Munich, Munich, Germany

^b School of Computation, Information and Technology, Technical University of Munich (TUM), Munich, Germany

^c Millennium Institute for Intelligent Healthcare Engineering, Santiago, Chile

^d School of Medicine and Health, Klinikum rechts der Isar, Technical University of Munich (TUM), Munich, Germany

^e School of Engineering, Pontificia Universidad Católica de Chile, Santiago, Chile

^f Department of Diagnostic and Interventional Radiology and Department of Nuclear Medicine, University Medical Center Hamburg Eppendorf, Hamburg, Germany

^g Department of Computing, Imperial College London, London, UK

^h Munich Center for Machine Learning (MCML), Munich, Germany

ⁱ Institute for Biological and Medical Engineering, Pontificia Universidad Católica de Chile, Santiago, Chile

^j Laboratory of Magnetic Resonance Imaging Systems and Methods, Ecole Polytechnique Federale de Lausanne (EPFL), Lausanne, Switzerland

^k Center for Biomedical Imaging (CIBM), Lausanne, Switzerland

^l School of Biomedical Imaging and Imaging Sciences, Kings College London, London, UK

ARTICLE INFO

Keywords:

Dynamic MRI reconstruction
 Parallel imaging
 K-space refinement
 Self-supervised learning
 Neural implicit representations
 Non-uniform sampling

ABSTRACT

Neural implicit k-space representations (NIK) have shown promising results for dynamic magnetic resonance imaging (MRI) at high temporal resolutions. Yet, reducing acquisition time, and thereby available training data, results in severe performance drops due to overfitting. To address this, we introduce a novel self-supervised k-space loss function $\mathcal{L}_{\text{PISCO}}$, applicable for regularization of NIK-based reconstructions. The proposed loss function is based on the concept of parallel imaging-inspired self-consistency (PISCO), enforcing a consistent global k-space neighborhood relationship without requiring additional data. Quantitative and qualitative evaluations on static and dynamic MR reconstructions show that integrating PISCO significantly improves NIK representations, making it a competitive dynamic reconstruction method without constraining the temporal resolution. Particularly at high acceleration factors ($R \geq 50$), NIK with PISCO can avoid temporal oversmoothing of state-of-the-art methods and achieves superior spatio-temporal reconstruction quality. Furthermore, an extensive analysis of the loss assumptions and stability shows PISCO's potential as versatile self-supervised k-space loss function for further applications and architectures. Code is available at: <https://github.com/compai-lab/2025-pisco-spieker>

1. Introduction

Magnetic resonance imaging (MRI) suffers from long acquisition times, which can limit its spatial and temporal resolution. This particularly affects dynamic applications, where temporally resolved images are reconstructed by sorting the data into distinct motion states (MS), i.e.

cardiac or respiratory motion states (Zaitsev et al., 2015; Spieker et al., 2023). Yet, the reconstruction of multiple MS reduces the available data per temporal MS, causing undersampling artefacts due to Nyquist theorem violations. Spatial reconstruction quality is commonly recovered by utilizing redundancies through regularization along the temporal dimension (Feng et al., 2016; Terpstra et al., 2023). Nonetheless,

[★] Funding: V.S. and H.E. are partially supported by the Helmholtz Association under the joint research school “Munich School for Data Science - MUDS”. V.S. was supported by the Bayer Foundation Fellowship and the Add-on Fellowship of the Joachim Herz Foundation. This work is supported by the Munich Center of Machine Learning and the DAAD programme Konrad Zuse Schools of Excellence in Artificial Intelligence, both sponsored by the Federal Ministry of Research, Technology and Space.

* Corresponding author.

E-mail address: v.spieker@tum.de (V. Spieker).

the temporal resolution remains constrained by the number of discrete MS. Moreover, in highly accelerated acquisitions, temporal regularization can lead to excessive dynamic smoothing, causing a loss of valuable temporal information. Yet, high temporal resolution and minimized image blurring are essential for clinical applications like assessing cardiac dynamics or detecting malignant lesions in abdominal MRI. Additionally, real-time cine abdominal imaging is essential for applications such as radiation therapy planning, where precise targeting under motion is critical.

Lately, implicit neural representations (INR) have gained attention to learn continuous representations from discrete data (Mildenhall et al., 2020; Sitzmann et al., 2020). Rather than learning a discretized pixel- or voxel-based representation, INRs learn a continuous representation of an object by fitting a multi-layer perceptron (MLP) to map input coordinates to their respective signal values. At inference, any intermediate coordinate value can be queried. This unsupervised training strategy has also been transferred to undersampled MRI reconstruction (Feng et al., 2016; Shen et al., 2022), showing great potential to infer missing spatial information on a subject-specific basis (without the need for training data from other subjects).

Recently, INRs have also been proposed for dynamic MRI, integrating the temporal component as additional input dimension into the MLP. In contrast to recent state-of-the-art dynamic reconstruction approaches (Wang et al., 2022; Chen et al., 2022; Chu et al., 2025), these methods do not require fully sampled ground truth data or large inter-patient datasets for training, making them particularly suitable for data-constrained clinical settings. Methodologically, existing INR-based approaches can be broadly categorized into two classes: On the one hand, some INR-based methods output the object in the image domain, requiring a domain transform within the loss function to compare the prediction with the acquired data in k-space. Huang et al. (2024), Catalán et al. (2025) solve this by using a projection-based loss function to compare the predicted images with the acquired spokes and (Kunz et al., 2024; Shao et al., 2024; Feng et al., 2025) apply the non-uniform Fast Fourier Transformations (NUFFT) on the queried images within each training iteration to compare with the available undersampled k-space. On the other hand, Huang et al. (2023), Spieker et al. (2024b) directly learn a continuous neural implicit representation in k-space (NIK). Compared to the image-based method, this allows for flexible, trajectory-independent training and avoids computationally expensive domain transforms, such as NUFFT.

Like other INR-based methods, NIK is independent of large ground truth datasets but requires retraining for each subject. Thus, minimizing the subject-specific data needed to learn a valid representation is desirable, as it directly reduces acquisition time (i.e., acquiring at higher accelerations). Yet, high accelerations create large k-space gaps, i.e., in radial trajectories. This often leaves the k-space periphery representing high-frequency details undersampled. This challenges the learning of accurate NIK representations and, without proper regularization, may lead to overfitting or unrealistic data filling, resulting in inaccurate, noisy dynamic reconstructions. General learning-based MRI reconstruction methods mitigate overfitting by incorporating regularization, typically applied in the image domain (Ahmad et al., 2020; Huang et al., 2021; Hammernik et al., 2023; Jafari et al., 2023). Image-based INRs also explicitly impose low-rank or total variation constraints on the predicted images (Feng et al., 2025). Yet, enforcing these image domain constraints directly within NIK would undermine the advantages of exclusive training in k-space, and translating these constraints into k-space is not trivial. A regularization method acting purely in k-space would be desirable.

Research into conventional accelerated MRI reconstruction has extensively explored k-space redundancies that arise from the use of multiple coils during MR acquisition, also known as parallel imaging (Hamilton et al., 2017). One widely used method is Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA) (Griswold et al., 2002), which assumes that each k-space point has a linear relation-

ship with its neighboring points across multiple coils. GRAPPA also assumes that this relationship remains consistent throughout the entire k-space. Therefore, in undersampled acquisitions, missing k-space points can be filled in by synthesizing data based on this relationship, which is typically calibrated using a fully sampled autocalibration signal (ACS). This concept has influenced several subsequent works, including methods for constraining conventional iterative reconstructions (Lustig and Pauly, 2010) or advanced trajectories (Seiberlich et al., 2011). Also, learning-based approaches have adopted the neighborhood relationship, e.g., by learning the k-space relationship for k-space interpolation (Akçakaya et al., 2019) or using the calibrated relationship as regularization (Ryu et al., 2021; Spieker et al., 2024b). However, like GRAPPA, these methods require explicit determination of the k-space neighborhood relationship, which is particularly impractical for motion-resolved imaging, as a fully-sampled ACS would be needed for each MS.

To exploit the global k-space relationship in a calibration-free manner, Shin et al. (2014) introduced an additional low-rank constraint on the multi-coil relationships within an iterative reconstruction. Haldar and Zhuo (2016) extended the low-rank concept as a regularization function by computing residuals of SVD-truncated k-space relationships. These approaches demonstrated the potential of calibration-free k-space regularization functions within conventional MR reconstruction. Although the low-rank constraint has also been integrated into convolutional layers for learning-based k-space interpolation (Han et al., 2020; Cui et al., 2023), its application has so far been limited to static undersampled reconstruction and CNN-based frameworks, which lack the continuous representation capability of INRs.

Motivated by the extensive research on k-space neighborhood relationship and with the aim to improve INR-based k-space interpolation using NIK, we have recently introduced a self-supervised calibration-free k-space consistency condition, named parallel imaging-inspired self-consistency (PISCO) (Spieker et al., 2024a). PISCO leverages the intrinsic global k-space relationship without explicitly determining it, and compared to Haldar and Zhuo (2016), does not require expensive SVD decomposition and user-defined truncation settings. We further reformulated PISCO into a self-supervised loss function for k-space-only training using NIK, in the following referred to as $\mathcal{L}_{\text{PISCO-dist}}$. However, this regularizing loss function was only tested during NIK training on free-breathing data for dynamic MRI, without any validation of its assumptions and convergence behaviour.

In this work, we provide a more comprehensive explanation and thorough validation of the basic PISCO condition, investigating various design choices and demonstrating their convergence. This includes a novel residual-based PISCO loss function $\mathcal{L}_{\text{PISCO}}$ with improved optimization properties and enhanced NIK performance compared to the previous version. Compared to Spieker et al. (2024a), we expand PISCO's evaluation by demonstrating its potential for MRI reconstruction in a broader setup, including a different dynamic MRI reconstruction problem (cardiac) as well as solving a distinct undersampled reconstruction problem (k-space optimization independent of NIK). Overall, our contributions are three-fold:

- We present an improved concept of PISCO, extended by a comprehensive analysis of key components such as kernel design, weight solving and consistency quantification.
- We integrate PISCO in a novel self-supervised k-space loss function $\mathcal{L}_{\text{PISCO}}$, validate its convergence behaviour compared to Spieker et al. (2024a) and assess its ability to enhance MRI reconstruction using NIK representations across three distinct in-vivo applications. To the best of our knowledge, $\mathcal{L}_{\text{PISCO}}$ is the first k-space-based self-supervised loss function used in INR-based dynamic MR reconstruction.

- The potential of PISCO is quantitatively and qualitatively demonstrated on static as well as multiple dynamic MRI applications using NIK, highlighting its ability to notably improve highly accelerated reconstructions using k-space-based INRs.

2. Background

This section introduces the parallel imaging-based k-space interpolation method GRAPPA (Griswold et al., 2002) and the k-space based dynamic reconstruction method NIK (Huang et al., 2023). GRAPPA lays the foundation for our proposed self-supervised k-space condition PISCO, applicable as regularization for NIK.

2.1. GRAPPA

An MR image $x \in \mathbb{C}^{N_x \cdot N_y}$ is reconstructed by solving the inverse problem $y = Ax + n$, where $y = \{y_i \in \mathbb{C}^{N_c} | i = 1, \dots, N_x \cdot N_y\}$ is the k-space signal acquired with N_c coils at spatial k-space coordinates $k = \{k_i \in \mathbb{R}^2 | i = 1, \dots, N_x \cdot N_y\}$ (for simplicity, the time dimension is only introduced in the next subsection), $A = FS$ is the forward operator including coil sensitivity maps $S \in \mathbb{C}^{N_x \cdot N_y \cdot N_c}$ and Fourier transform F and n is noise. In practice, only a part of the k-space y is acquired (i.e. less k_i sampled) to reduce acquisition time, which results in an undersampled dataset that violates the Nyquist theorem. To minimize subsequent undersampling artefacts, GRAPPA (Griswold et al., 2002) utilizes the multi-coil setup of MRI to estimate absent k-space values based on surrounding data points (see Fig. 1 left). In detail, for a missing target location $k_i^T \in \mathbb{R}^2$ a coordinate patch $k_i^P \in \mathbb{R}^{N_n \cdot 2}$ of N_n neighboring coordinates is sampled. The missing target signal value $y_i^T \in \mathbb{C}^{N_c}$ can then be estimated using a linear combination of the neighboring signal values $y_i^P \in \mathbb{C}^{N_n \cdot N_c}$. N_m target/patch pairs can be stacked to create the linear equation system $T = PW$, where $T = [y_1^T, \dots, y_{N_m}^T] \in \mathbb{C}^{[N_m \times N_c]}$ are the stacked target points, $P = [y_1^P, \dots, y_{N_m}^P] \in \mathbb{C}^{[N_m \times N_n \cdot N_c]}$ are the stacked neighboring patches, 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. Then, the relationship W_{ACS} can be determined within a fully sampled auto-calibration signal y_{ACS} by solving the following regularized least squares problem (Griswold et al., 2002):

$$W \triangleq W_{ACS} = \arg \min_W \|P_{ACS} W - T_{ACS}\|_2^2 + \alpha \|W\|_2^2. \quad (1)$$

Here, $T_{ACS}, P_{ACS} \subseteq y_{ACS}$, $\|\cdot\|_2^2$ applies the L2-norm element-wise and α weighs the Tikhonov regularization. Since GRAPPA assumes one global neighborhood relation for the whole k-space, W_{ACS} can also be considered equivalent (\triangleq) to the global neighborhood relationship W , applicable to the remaining undersampled k-space y_{US} . Thus, the determined weights W_{ACS} can be applied to the acquired neighboring k-space samples $P_{US} \subseteq y_{US}$ to derive the missing target samples $T_{US} \subseteq y_{US}$, i.e., by computing $T_{US} = P_{US} \cdot W_{ACS}$, as shown in Fig. 1.

2.2. NIK

Reformulating the MRI reconstruction problem from Section 2.1 to the dynamic case with the temporal dimension t results in k-space coordinates $k = \{k_i = [k_x, k_y, t] \in \mathbb{R}^3 | i = 1, \dots, N_x \cdot N_y \cdot N_t\}$ and multiple temporal images $x = \{x_t | t = 1, 2, \dots, N_t\}$. This problem can be solved using neural implicit k-space representations (NIK), which enable binning-free dynamic MR reconstructions at high temporal resolutions (Huang et al., 2023). NIK consists of a model G_θ (usually a multi-layer perceptron) that maps input coordinates k_i to multi-coil k-space signal values y_i (Spieker et al., 2024b). During training, the model is fitted on a patient-specific basis on pairs of acquired k-space coordinates k_i^{acq} and its corresponding signal values y_i^{acq} :

$$\theta^* = \arg \min_\theta \mathcal{L}_{DC}(G_\theta(k_i^{acq}), y_i^{acq}) \quad (2)$$

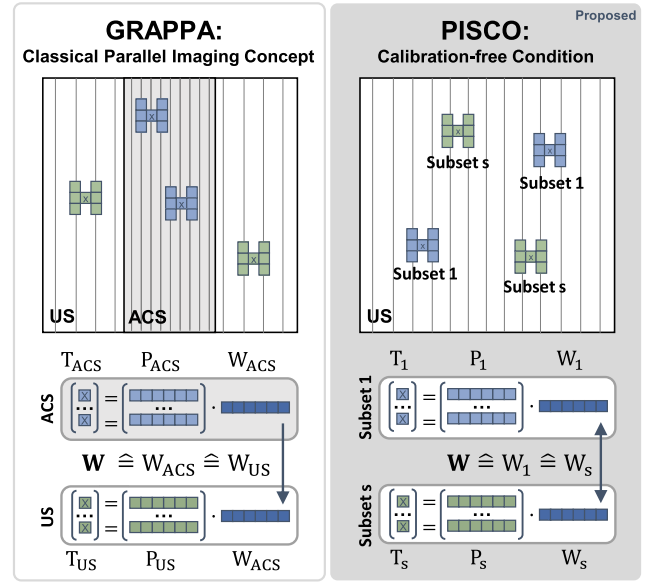


Fig. 1. From GRAPPA (Griswold et al., 2002) to PISCO (proposed): GRAPPA (left) calibrates a weight vector W_{ACS} on a fully sampled autocalibration k-space (ACS, gray area, Eq. (1)), where gray lines indicate sampled k-space points. This vector determines the global neighborhood relationship W and is applied to derive target points in the undersampled k-space (US, white area). For PISCO (right), multiple random subsets of targets and patches (T_s and P_s) are sampled to solve for W_s (Eq. (3)). The parallel-imaging inspired self-consistency (PISCO) condition states that for an ideal k-space, all these weight vectors should equally converge to one single global neighborhood relationship W (Eq. (4)).

where \mathcal{L}_{DC} is a data consistency loss and θ^* the optimized network parameters. At inference, any coordinate k_i can be inputted into the fitted model G_{θ^*} . The final dynamic reconstructions x_t are obtained by sampling Cartesian k-space coordinates \hat{k} at any time point t and transforming them to image space, i.e., $x_t = S^H F^{-1} G_\theta(\hat{k})$.

3. Methods

3.1. From GRAPPA to PISCO

GRAPPA (Griswold et al., 2002) requires a fully sampled ACS for calibration of the global weight matrix W . Yet, this is not always available, particularly in dynamic imaging. Hence, we reformulate the global spatial k-space relationship concept to a calibration-free condition based on a similar assumption as (Spieker et al., 2024a): If a weight set W_{ACS} calibrated on $y_{ACS} \in y$ models the global linear relationship W for the whole k-space y , then a weight set W_s derived from a random subset $y_s \in y$ should equally result in a global linear relationship within an ideal k-space:

$$W \triangleq W_s = \arg \min_W \|P_s W - T_s\|_2^2 + \alpha \|W\|_2^2 \quad (3)$$

where $T_s, P_s \subseteq y_s$. Considering one global linear k-space relationship, multiple weight sets W_1, \dots, W_{N_s} derived from various random subsets $\{y_s \in y | s = 1, \dots, N_s\}$ (visualized in green and blue in Fig. 1) are expected to converge to the same solution. Thus, without access to fully-sampled calibration data, Parallel Imaging-inspired Self-Consistency (PISCO) of an undersampled k-space can be enforced as follows:

$$W_1 \triangleq \dots \triangleq W_{N_s} (\triangleq W). \quad (4)$$

3.2. PISCO as self-supervised k-space regularization

Next to checking the consistency of a given k-space, PISCO can also serve as a self-supervised regularization objective function by quantifying the PISCO condition within a loss function \mathcal{L}_{PISCO} . Therefore, a

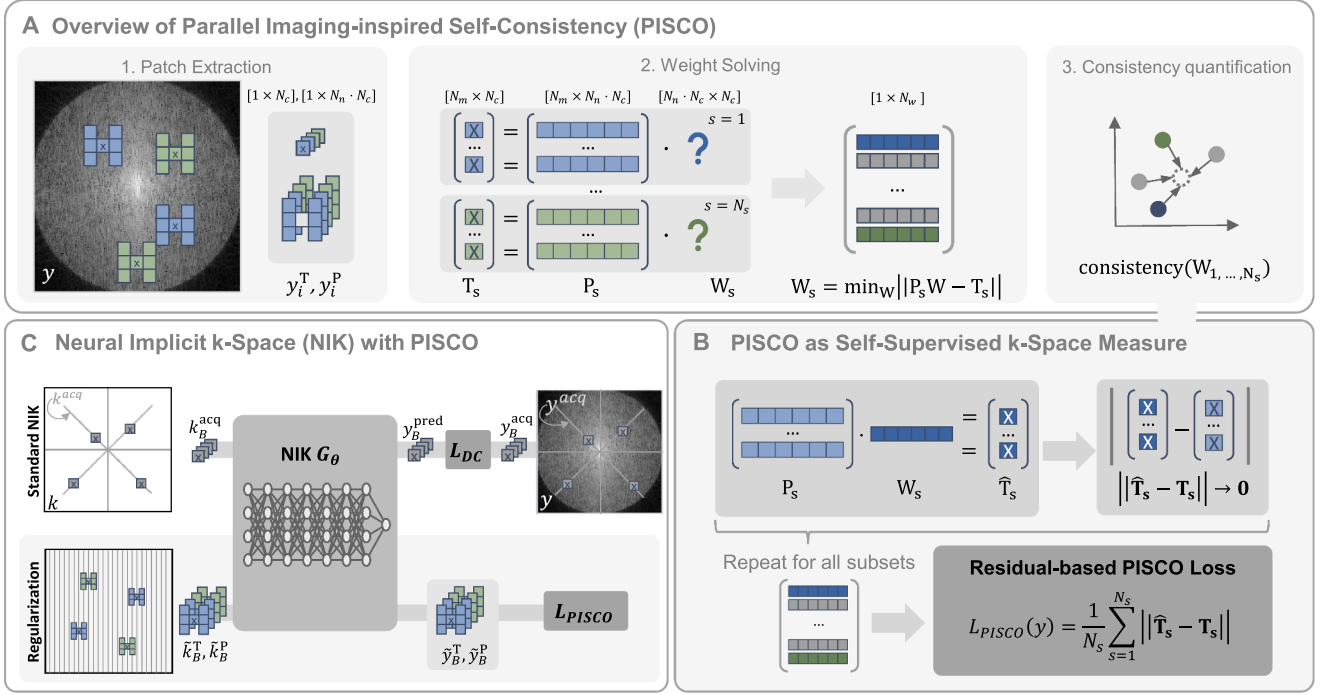


Fig. 2. Method overview (A/B/C corresponding to Section 3A/B/C). For simplicity the coil dimension is not visualized, but indicated by the matrix sizes. (A) PISCO: 1. From a k-space y , multiple subsets (one color = one subset) including pairs of target y_i^T and patches y_i^P are sampled. The patches consist of neighboring values, whereas any sampling shape can be assumed (see Section 4.2.1). 2. Each subset is reshaped to a linear equation system and used to solve for the neighborhood relationship vector W_s using Eq. (3) (Section 4.2.2). 3. The PISCO self-consistency of y can be quantified using all subset weight vectors, e.g. either computing the distance (Spieker et al., 2024a) or the residual (proposed, visualized in B). (B) Residual-based PISCO as self-supervised k-space measure: For each subset s , W_s is used to estimate the targets \hat{T}_s exclusively from their neighbors P_s and the residual to the actual targets T_s of y is computed. The PISCO loss is determined as the weighted sum of all these residuals. (C) Integration of PISCO into NIK training: The "supervised" top represents the classical NIK-training (Huang et al., 2023), where actually acquired coordinates are sampled, predicted and compared to the acquisition signal using L_{DC} . With PISCO, the perceptual field of NIK can be extended during training, because any patches (independent of acquisition trajectory) can be sampled and used for the self-supervised L_{PISCO} computation. In the present study, patches are sampled from a Cartesian grid, with alternating undersampling in the x- and y-dimension (only in y-dimension visualized).

proper computation of the consistency of all the solved subset weight vectors W_s is required (Fig. 2A.3). In Spieker et al. (2024a), the PISCO of a k-space y is quantified by summing the complex distance $\mathcal{L}_C^1(\Delta) = \|Re(\Delta)\|_1 + \|Im(\Delta)\|_1$ between all estimated weight vectors:

$$\mathcal{L}_{PISCO\text{-}dist} = \frac{1}{N_s^2} \sum_{i=1}^{N_s} \sum_{j=1, j \neq i}^{N_s} \mathcal{L}_C^1(W_i - W_j). \quad (5)$$

To avoid the risk of convergence to an unfeasible global solution and improve the condition's stability, we alternatively propose to enforce PISCO by measuring the fit of a determined weight set as follows (Fig. 2B): Consider a target and patch matrix $T_s, P_s \subseteq y_s$ with more target/patch pairs N_m than unknown weights N_w , i.e., an overdetermined linear equation system (LES) when solving for W_s with Eq. (3). For an ideal k-space, the determined W_s should consistently model the global neighborhood relationship over the whole k-space. Thus, the neighborhood-derived k-space $\hat{T}_s = P_s W_s$ should approximate (\rightarrow) the original targets T_s :

$$\|\hat{T}_s - T_s\|_p \rightarrow 0. \quad (6)$$

The computation of the residual with any p -norm enables the quantification of PISCO-res for a given k-space, e.g. the larger the residual the lower the self-consistency and vice versa. This condition can be rephrased as regularization loss $\mathcal{L}_{PISCO\text{-}res}$ for a given k-space y (Eq. (7)):

$$\mathcal{L}_{PISCO\text{-}res}(y) = \frac{1}{N_s} \sum_{s=1}^{N_s} (\underbrace{\|\mathcal{K}_P(y_s)W_s\|_2}_{\hat{T}_s} - \underbrace{\|\mathcal{K}_T(y_s)\|_2}_{T_s}) \quad (7)$$

where $y_s \in y$, \mathcal{K}_P and \mathcal{K}_T represent kernels extracting the patches and targets from y_s according to Section 2.1 and W_s is the estimated neigh-

borhood relationship (Eq. (3)). Both PISCO quantification methods, $\mathcal{L}_{PISCO\text{-}res}$ and $\mathcal{L}_{PISCO\text{-}dist}$, can then be integrated as \mathcal{L}_{PISCO} into the reconstruction objective \mathcal{L}_{RECON} (Eq. (8)):

$$\mathcal{L}_{RECON}(y) = \mathcal{L}_{DC}(y) + \lambda \cdot \mathcal{L}_{PISCO}(y) \quad (8)$$

where \mathcal{L}_{DC} can be any loss enforcing data consistency and λ is a weighting factor. In the following, PISCO-res will refer to the proposed residual-based version and PISCO-dist to the distance-based version from Spieker et al. (2024a).

3.3. NIK with PISCO

NIK's are fitted to the acquired data y^{acq} (Section 2.2), which limits the training strategy to coordinates $k^{acq} \subseteq k$ from the acquisition's trajectory within the desired k-space k . By incorporating PISCO, additional coordinates $\tilde{k} \subseteq k$ can be included in the training process, where \tilde{k} may extend beyond k^{acq} . This self-supervised inclusion of coordinates expands the perceptual field during the training procedure, making it particularly beneficial for large gaps within k-space, e.g. in radial trajectories.

The complete training pipeline of NIK with PISCO is summarized in Algorithm 1. When training NIK with PISCO, a batch of acquired coordinates k_B^{acq} as well as further coordinate batches of target \tilde{k}_B^T and respective patch neighbors \tilde{k}_B^P are inputted to NIK (line 2 and 5, respectively). Due to NIK's continuous sampling nature, any target and kernel for neighbor extraction can be chosen from k (line 4), e.g. as shown in Fig. 2 in a Cartesian manner. As in standard NIK (Huang et al., 2023), the predicted k-space y_B^{pred} is compared with the respective acquired data y_B^{acq} for \mathcal{L}_{DC} computation (line 3). Simultaneously, \tilde{y}_B^T

Table 1
Summary of acquisition parameters and evaluation settings.

		Upper leg	Cardiac cine ^a	Abdominal
Acquisition	Pulse Sequence	spoiled GRE	bSSFP cine	spoiled GRE
	Trajectory	pseudo golden angle stack-of-stars	radial (binned)	pseudo golden angle stack-of-stars
	Resolution	$1.5 \times 1.5 \times 3 \text{ mm}^3$	$1.8 \times 1.8 \text{ mm}^2$	$1.5 \times 1.5 \times 3 \text{ mm}^3$
	Number of coils N_c	27	15–18	26
	Number of k-space points [$N_{spokes} \times N_{FE}$] per slice	704×448	4900×414 (total) 196×414 (per MS)	1341×536
	Acceleration Factor per MS	R1	R1 (per MS)	$\sim R0.6/R2.4/R30$ (1/4/50MS)
Evaluation	Reconstructed MS	1	25	4/50 ^b
	Reconstruction matrix	268×268	$208 \times 208^\circ$, cropped to 104×104	268×268
	Reference reconstruction	XD-GRASP1 for R1 ($\lambda_{TV} = 0.1$)	XD-GRASP25 for R1 ($\lambda_{TV} = 0.01$)	-
	Tested accelerations ^d	R5, R10, R20 ($\lambda_{TV} = 0.1$)	R15, R26, R52 ($\lambda_{TV} = 0.3$)	$\sim R30$ ($\lambda_{TV} = 0.1$)

^a Data from [El-Rewaidy \(2020\)](#),

^b 4MS as in [Feng et al. \(2016\)](#), 50MS for high temporal resolution

^c after zero-filling removal ([Shimron et al., 2022](#))

^d with λ_{TV} for accelerated XD-GRASP reconstructions

Algorithm 1 Pseudocode for training NIK with PISCO.

Required: Acquired k-space coordinates k^{acq} and signals y^{acq} , NIK architecture G_θ , initial model parameters θ_0 , total epochs E , pretraining epochs E_{pre} , total subsets N_s , weight solving regularization α , PISCO weighting factor λ

- 1: **for** $e = 0$ to E **do** sample batch (k_B^{acq}, y_B^{acq}) in (k^{acq}, y^{acq})
- 2: $y_B^{pred} \leftarrow G_{\theta_e}(k_B^{acq})$ \triangleright Predict k-space with NIK
- 3: $\mathcal{L}_{DC} \leftarrow \mathcal{L}_{DC}(y_B^{pred}, y_B)$ \triangleright **Compute DC loss, Eq. (2)**
- 4: $\tilde{k}_B^T, \tilde{k}_B^P \leftarrow \tilde{k}$, patches around \tilde{k} \triangleright Sample any targets in k-space and surrounding patch coordinates
- 5: $\tilde{y}_B^T, \tilde{y}_B^P \leftarrow G_{\theta_e}(\tilde{k}_B^T), G_{\theta_e}(\tilde{k}_B^P)$ \triangleright Query target and patch k-space values
- 6: $[P_s]_{s=1}^{N_s}, [T_s]_{s=1}^{N_s} \leftarrow \tilde{y}_B^T, \tilde{y}_B^P$ \triangleright Sort target and patches into N_s subsets with N_m pairs each
- 7: $[W_s]_{s=1}^{N_s} \leftarrow \arg \min_W \|P_s W - T_s\|_2^2 + \alpha \|W\|_2^2$ \triangleright Solve each subset for weights, [Eq. \(3\)](#)
- 8: $\mathcal{L}_{PISCO} \leftarrow \frac{1}{N_s} \sum_{s=1}^{N_s} \|\hat{T}_s - T_s\|_2$ \triangleright **Compute PISCO loss, Eq. (7)**
- 9: $\lambda \leftarrow 0$ if $e < E_{pre}$, else λ \triangleright Apply solely DC loss if preconditioning
- 10: $\theta_{e+1} \leftarrow \theta_e - \nabla_{\theta}(\mathcal{L}_{DC} + \lambda \mathcal{L}_{PISCO})$ \triangleright Update weights with DC and PISCO loss
- 11: **end for**
- 12: **return** Learned PISCO-NIK model G_θ

and \tilde{y}_B^P are sorted into N_s subsets with $N_m = f_{od} \cdot N_w$ patch pairs each (line 6), where $f_{od} > 1$ is a factor to ensure an overdetermined LES. The resulting subsets of patch pairs are then processed to compute the respective weight vectors (line 7), the resulting \mathcal{L}_{PISCO} (line 8, [Eq. \(7\)](#)) and the overall objective ([Eq. \(8\)](#)) to update the model weights (line 10). Due to the random initialization of MLPs, NIK's early predictions resemble pure noise without structural detail, also making PISCO ineffective. Thus, pre-conditioning the NIK model with only the data consistency loss for a predefined number of epochs E_{pre} is recommended to capture rough structural information before applying PISCO as refinement. This is achieved by setting $\lambda = 0$ until E_{pre} (line 9).

4. Experimental setup

4.1. Data

Experiments are conducted on three different MR datasets of varying complexity, with detailed acquisition parameters for all datasets listed in [Table 1](#). In-house data was acquired at 3T (Ingenia Elition X, Philips Healthcare, Best, The Netherlands) after local ethics committee

approval. Coil sensitivities are estimated with ESPIRiT ([Uecker et al., 2014](#)). The datasets evaluated are:

Upper leg (quasi-static/no motion) A radial pseudo golden angle stack-of-stars in-house acquisition of the upper leg, an anatomy where we assume little to no motion. This quasi-static acquisition is used to validate PISCO's potential independent of the time dimension. We use the 3D stack-of-stars acquisition from a single subject to extract multiple 2D slices, as our method is currently developed and evaluated on a slice-by-slice basis as a proof-of-concept. Slices are extracted by applying an inverse Fourier transform (IFFT) along the kz direction (feet-head axis) to reduce the computational cost of the reconstruction problem [Feng \(2022\)](#). No additional pre-processing steps were performed except for scanner-specific corrections to read the raw data and a hamming filter along the Cartesian axis to reduce Gibbs ringing. Retrospective undersampling is conducted by discarding spokes from the acquisition.

Cardiac cine (cardiac motion) A public cardiac cine dataset ([El-Rewaidy, 2020](#)) with binned fully-sampled data as reference (retrospective ECG-triggering). For 30 subjects in total, zero-padding is removed before processing to avoid implicit data crimes ([Shimron et al., 2022](#)) and retrospective undersampling is conducted to match a uniform distribution of the k-space spokes per MS ([El-Rewaidy et al., 2021](#)).

Abdominal (respiratory motion) A free-breathing radial pseudo golden angle stack-of-stars in-house acquisition of the abdomen and, for comparison, a gated acquisition of the same subject. The acquisition is performed on a single subject and undergoes the same preprocessing steps as the upper leg dataset. Due to the self-navigating nature of the acquisition protocol, the respiratory signal - representing the patient's current breathing position - can be extracted from the data itself. Specifically, the temporal signal is derived by applying principal component analysis (PCA) to projections of the central k-space values, following the method described in [Feng et al. \(2016\)](#). The acquisition itself is prospectively undersampled at a factor of R2.4/R30 for 4MS/50MS, respectively.

4.2. PISCO design choices

Within PISCO, multiple design choices can be made to ensure and improve PISCO convergence behaviour. In the following, we explain major design choices for each module in the PISCO computation pipeline: (1) patch extraction ([Fig. 2A.1](#)), (2) weight solving ([Fig. 2A.2](#)) and (3) PISCO consistency quantification ([Fig. 2A.3](#)). Therefore, we experimentally validate the effect of different design choices independently of NIK on a sample frame for a **cardiac cine** subject by simulating ideal k-space data, i.e., generating signal values for all required k-space coordinates using *torchkbnufft* ([Muckley et al., 2020](#)).

4.2.1. Kernel design

The combination of PISCO and NIK allows for an arbitrary selection of target points T as well as any kernel design \mathcal{K}_P to extract the neighboring patches P , since it is not necessary that the kernel points are actually sampled within the training dataset. The target points T are selected to be on a Cartesian grid to focus the model's attention to sparsely sampled regions. The neighboring points P could be chosen arbitrarily, yet the assumed global consistency needs to be ensured. Therefore, we investigated three spatial kernel geometries extracting P around T : (a) Cartesian kernels, as originally used in GRAPPA (Griswold et al., 2002), (b) radial kernels, as proposed in Seiberlich et al. (2011) for radial trajectories and (c) equidistant radial kernels, which mitigate unequal kernel spacing at different radii of the radial sampling. We test kernels of shape $[3 \times 2]$, as originally proposed for GRAPPA (Griswold et al., 2002), opting for the smaller configuration to reduce the number of unknown weights. The neighbor distance is set to $\delta = 2 \cdot N_{FE}^{-1}$ in Cartesian or radial coordinates for the equidistant and radial kernel, respectively. To avoid temporal blurring due to merging of multiple time points (Breuer et al., 2005), all patches within one subset are sampled from one time point. The global consistency assumption is then validated by sampling multiple random subsets y_s from the artificially generated k-space with the different kernel geometries. Then, the resulting subset weight vectors W_s can be compared and - since we assume an ideal k-space - all weight vectors should resemble each other.

4.2.2. Weight solving

The computation of the PISCO residual requires the LES to be over-determined by a factor $f_{od} = N_m/N_w$ larger than 1. To ensure robustness of the solution albeit the overdetermination, i.e. to outliers, regularization of the weight vector magnitude is included and weighted by a factor α , as stated in Eq. (3). Empirical evaluation yielded feasible values of $\alpha = 1e-4$ and $f_{od} = 1.1$ for all datasets.

Another challenge in weight solving arises from the high dynamic range of k-space magnitudes (varying from center to the periphery k-space). Mixing patches with these different magnitudes results in a poorly scaled LES, posing an instability risk. Thus, the randomly sampled patch pairs are first sorted according to their spatial distance to the k-space center (i.e., $\sqrt{k_x^2 + k_y^2}$) and only then, divided into subsets. By grouping patch pairs with similar k-space center distanced, the magnitude variance within the LES for each individual subset is minimized. Additionally, k-space points within a small radius r around $k_x/k_y = 0$ are removed to avoid inclusion of individual high magnitudes (in our case $r = 10 \cdot N_{FE}^{-1}$, requires adjustment if kernel size and distance are increased). Using the k-space obtained with *torchkbnfft*, we validate the effect of patch sorting on the weight estimates.

4.2.3. PISCO consistency measure

As presented in Section 3.2, the self-consistency of a k-space can be measured in several ways. Yet, for ideal convergence behaviour during training the loss must be monotonically decreasing when approximating an ideal k-space. To test the convergence behaviour, we create "non-ideal" k-space data in two separate ways: (1) introducing noise directly in k-space to test NIK's robustness to noisy MLP outputs, as it predicts k-space values directly; and (2) adding noise in the image domain before applying the Fourier transform to ensure the regularizer remains sensitive to image-domain noise. We add zero-centered complex Gaussian noise at different standard variations σ , thereby determining the noise/corruption level. Then, both PISCO measures, distance-based (Spieker et al., 2024a) and residual-based (proposed), are evaluated dependent on the noise level.

4.3. Validation of PISCO convergence

To ensure convergence of PISCO independent of NIK, we provide a proof-of-concept by solving a simplified version of the PISCO reconstruction problem in Eq. (8). Instead of \mathcal{L}_{DC} of NIK, we replace the data

Table 2

Summary of feature processing and loss parameters.

		Upper leg	Cardiac	Abdominal
Features	Drift correction	–	✗	✓
	Navigator scaling	–	[0,1]	[0,0.5]
	σ	6	6	1
PISCO	λ	0.05	0.01-0.15	0.01
	N_s	20	30	15
	f_{OD} / α	1.1 / 1e-4		

consistency component to fit a fully-sampled k-space y with an under-sampling mask M to the actual acquired k-space \bar{y} and solve the reconstruction problem:

$$\hat{y} = \arg \min_y (\|My - \bar{y}\|_1 + \lambda \cdot \mathcal{L}_{PISCO}(y)). \quad (9)$$

A Cartesian k-space is simulated from a cardiac cine slice (random retrospective undersampling for $R=2$, 4 % of center lines kept). Optimization is conducted for a total of 500 epochs for two kernel shapes, $[3 \times 2]$ and $[5 \times 4]$, with \mathcal{L}_{PISCO} weighted by $\lambda = 5e-4$. The model is preconditioned for the first 100 epochs, remaining parameters are defined in Section 4.2.

4.4. Training NIK with PISCO

We adapt NIK's (Huang et al., 2023) architecture using 4 layers, 512 hidden features, high-dynamic range loss as \mathcal{L}_{DC} , SIREN activations with $\omega=20$ (Sitzmann et al., 2020), batch size of 10k and use STIFF feature encoding (Catalán et al., n.d.) with σ as initialization of the feature distribution. Depending on the motion pattern, we rescale the navigator t , and correct linear drifts for abdominal imaging (see Table 2).

For NIK with PISCO, $\mathcal{L}_{PISCO-dists}$ or $\mathcal{L}_{PISCO-res}$ is applied after $E_{pre} = 1000$, as rough structural details were observed for all datasets at this stage (Section 3.3) and the loss curve gradient began to decline. For each dataset and acceleration factor, the PISCO regularization weight, λ , was determined empirically. To this end, multiple candidate values were evaluated, and the one offering a suitable trade-off between under- and over-smoothing was selected (see Supplementary). For the cardiac cine dataset, this selection was performed on a representative subject for each acceleration factor and subsequently applied to the remaining cases. The final values are documented in Table 2.

The patch sampling is conducted as follows: Targets for PISCO computation originate from a Cartesian grid ($\delta_{k_x}/\delta_{k_y} = N_{FE}^{-1}$) and neighbors are sampled according to kernel design. To account for undersampling in both, x-/y-dimension, the Cartesian kernel design is applied alternating in both directions. Since a large number of coils N_c within a dataset results in a large number of weight parameters N_w , more patches are needed to reliably solve for these parameters within a single subset. Rather than keeping the number of subsets N_s fixed - which would increase computational overhead and lead to denser patch sampling for datasets with large N_c - we adapt N_s accordingly. Specifically, N_s is empirically selected to fit GPU memory consumption and use approximately all sampled patches within one epoch. Weight solving parameters are defined as in Section 4.2.2.

All models are jointly optimized for a total of 5000 epochs (NVIDIA RTX A6000, Python 3.10.9/PyTorch 1.13.1) with Adam (lr = 1e-5), with amsgrad enabled to encounter convergence issues due to the high dynamic range of k-space (Reddi et al., 2019). As both NIK and its PISCO extensions are patient-specific self-supervised methods, optimization is performed independently for each subject and slice. The models are optimized directly on the undersampled acquired data. For the cardiac subject, training durations were approximately 17 min for standard NIK, 42 min for PISCO-dist, and 22 min for PISCO-res (see Supplementary). After training, reconstructions are generated by querying the individually trained model on a Cartesian grid at the desired time points, with inference taking only milliseconds per frame.

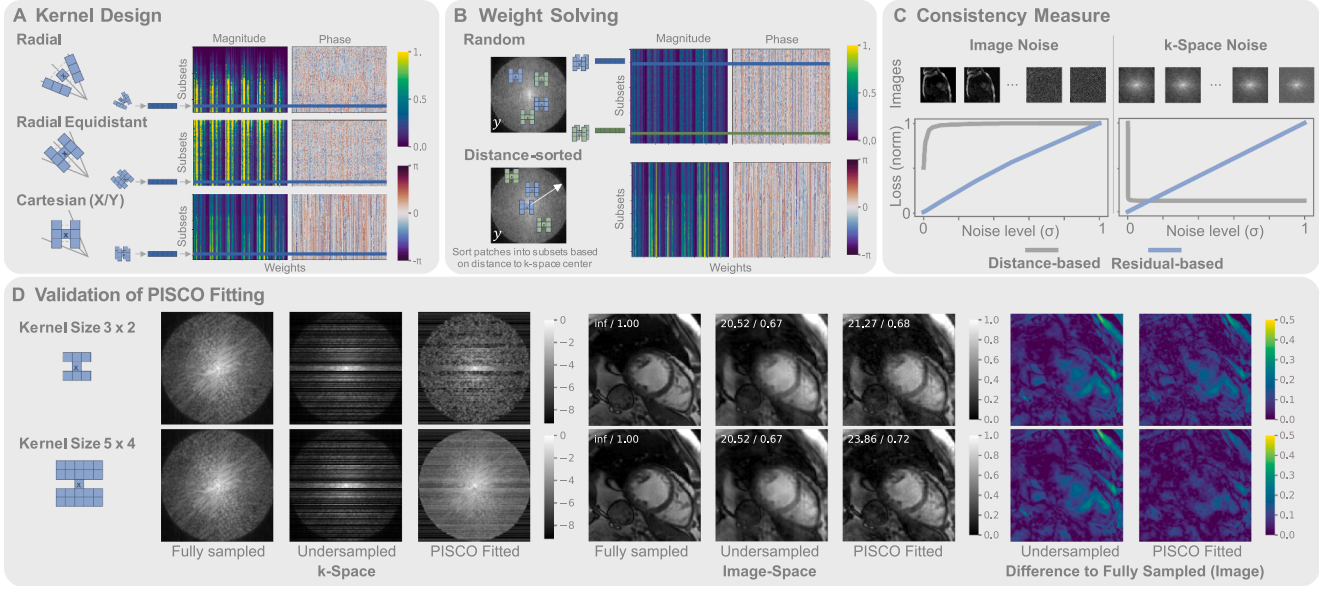


Fig. 3. PISCO validation results. (A) Kernel Design: For each kernel geometry, multiple subsets of patches (one exemplary shown in blue) are sampled and solved for the weights (dark blue vector). The magnitude and phase of all weight vectors are stacked in the plot. To validate the PISCO condition, all weight vectors should result in the same solution, e.g. a vertical pattern is expected. Only the Cartesian kernel - shown in y-direction, but equal result in x-direction (90° turned) - fulfills this condition. (B) Weight Solving: Solution of subset weights (magnitude and phase) are stacked vertically (green and blue vectors are two examples). Top: Subsets consist of randomly sampled patches. Bottom: The sampled patches are sorted from the center to the outside of the k-space (white arrow) and then into subsets. This results in subsets with patches of similar k-space magnitude, leading to less noisy weight vector solutions overall. (C) Consistency Measure: Non-ideal k-spaces are simulated by adding increasing noise in image-space (left) and k-space (right). On the corresponding k-spaces, the distance-based (Spieker et al., 2024a) and residual-based (proposed) PISCO loss is computed and normalized by the maximum loss value. For image noise, both losses increase monotonically, while the residual-based loss is consistently sensitive. For k-space noise, only the residual-based loss monotonically increases, making it feasible for optimization towards the ideal k-space (where $\sigma = 0$). (D) Validation of PISCO Fitting: K-Space (left) and image (middle) results for PISCO fitted undersampled k-spaces. Within k-space, inclusion of PISCO allows for derivation of missing k-space lines without any additional knowledge. Independent of the kernel size, reconstruction results lead to reduced undersampling using PISCO compared to the original undersampled image, visible qualitatively in the images as well as the quantitative PSNR/FSIM results (middle) and in difference images (right). A larger kernel (5x4, bottom) increases the perceptible field and fills more missing k-space lines, but also requires more computational power due to more unknown kernel weights N_w .

4.5. Baseline comparisons and evaluation metrics

To evaluate the performance of our proposed PISCO regularization, we compare the following INR-based approaches: standard NIK (Huang et al., 2023), ICoNIK (Spieker et al., 2024b), PISCO-dist (Spieker et al., 2024a) and PISCO-res (NIK with $\mathcal{L}_{\text{PISCO-dist}}$ and $\mathcal{L}_{\text{PISCO-res}}$, respectively). Further, we compute the inverse NUFFT (INUFFT) as well as the state-of-the-art motion-resolved reconstruction method XD-GRASP (Feng et al., 2016). For the latter two, the number of motion states (MS) is defined depending on the dataset, i.e., 1MS for static upper leg, 25MS for cardiac cine, and 4/50MS for abdominal (in the following referred to as INUFFTMS or XD-GRASPMS). We conduct the XD-GRASP reconstruction using a conjugate gradient algorithm with line search and, for each dataset, perform a grid search on a representative subject to determine the TV regularization weight λ_{TV} (see Table 1).

As a further baseline, we include TD-DIP (Yoo et al., 2021), a recent unsupervised dynamic reconstruction method based on deep image priors (DIP). Since the CNN-based method is not resolution-agnostic and the original version only outputs 128x128 (TD-DIP128), we further evaluate an adapted version matching our dataset resolution without the need for additional interpolation (TD-DIP512). Remaining training parameters are defined as in the original work.

Additionally, quantitative evaluation is performed using peak signal-to-noise ratio (PSNR) (Wang et al., 2004) and feature similarity (FSIM) (Zhang et al., 2011), which have been shown to correlate well with radiological evaluation (Marchetto et al., 2025). To evaluate the spatial reconstruction performance, FSIM is computed on all spatial images at all time points (FSIM-spat). Further, to evaluate the dynamic performance, FSIM is computed on all temporal profiles, i.e. all $x/y/z$ slices

(FSIM-temp). Before metric computation, images are clipped to their 99th percentile and normalized to [0,1], corresponding to the recommended preprocessing in Marchetto et al. (2025) for image evaluation. Since each subject was reconstructed using all methods, the study follows a repeated-measures design. Therefore, we use the Friedman test to assess overall differences between methods, followed by pairwise comparisons using the Wilcoxon signed-rank test with False Discovery Rate (FDR) correction for multiple comparisons.

5. Results

5.1. PISCO design choices

Kernel design Fig. 3A shows the visualization of the weight vector solutions W_s (individual weights on the horizontal axes) for multiple randomly sampled subsets $s = 1 \dots N_s$ (stacked on vertical axes) for all three kernel designs. Since the weights were computed on an ideal k-space, all subset weight vectors should be the same according to PISCO, i.e., a vertical pattern should be visible with as little variation in the vertical direction as possible. No consistent vertical pattern is visible for the radial as well as radial equidistant kernel. Yet, only the Cartesian kernel results in consistent magnitudes and phases for all subsets, confirming the applicability of PISCO as self-supervised consistency measure with this kernel design.

Weight solving The subset weight vector solutions with and without distance sorting from the center to periphery k-space are shown in Fig. 3B. Both weight vector solutions indicate that including Tikhonov regularization and the overdetermination factor result in stable and consistent solutions. Yet, more noise in both, magnitude and phase of the

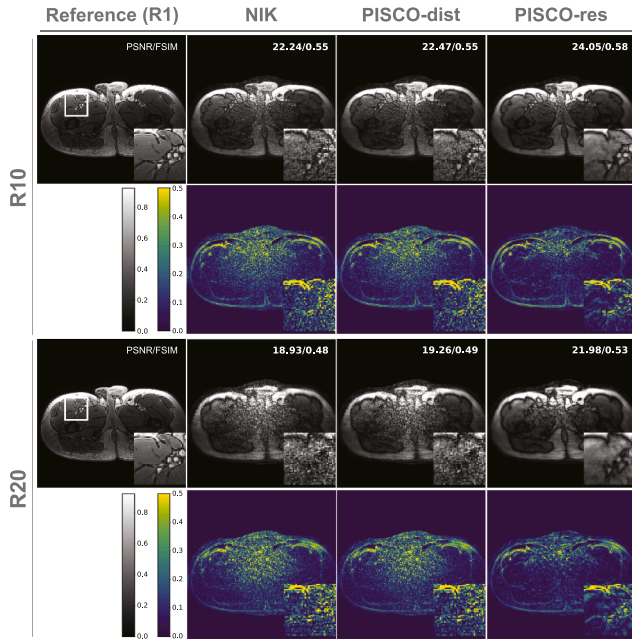


Fig. 4. Static upper leg: Qualitative and quantitative results for two exemplary acceleration factors (R10/20). NIK results in increasing noise, particularly towards the center of the reconstruction. PISCO-dist marginally improves reconstructions, but still remains noisy. With the proposed PISCO-res a better implicit representation is learned, that results in reduced noise and sharper structures (PSNR and FSIM \uparrow).

weight vector solution, is recognizable when no frequency sorting is applied. Avoiding poorly scaled LES by sorting the patch pairs according to their k-space distance results in less variance of the W_s , and a more stable PISCO condition.

Consistency measure As shown in Fig. 3C, both PISCO measures (distance and residual-based) increase monotonically with rising noise level in image-space. Yet, the linear increase of residual PISCO may offer additional stability when optimizing with the objective of a noise-free solution compared to the distance-based metric, where the gradient is sensitive to the amount of noise added. For k-space added noise, the distance-based measure decreases with increasing k-space noise, which makes this measure infeasible for k-space noise reduction optimization. In contrast, a proportional relationship of residual PISCO to k-space based noise can be observed, making it a suitable metric for further optimization. To further validate this finding, we include reconstructions using the distance-based loss (PISCO-dist) in the in-vivo results.

5.2. Validation of PISCO convergence

The undersampled k-space fitting results using PISCO with two different kernel sizes are shown in Fig. 3D. Without any additional information, PISCO is capable of filling the undersampled gaps within k-space (Fig. 3D left). The larger kernel enables a larger receptive field and consequently, less remaining gaps within k-space. The corresponding reconstructions show sharper results and higher quantitative values when including PISCO (Fig. 3D middle) as well as reduced undersampling artefacts, as visible in the difference images (Fig. 3D right). Note that this reconstruction does not leverage any temporal redundancy yet, but the improvement is solely based on including the neighborhood relationship with PISCO. Also, some gaps remain unfilled since they are never updated and defaulted to zero within the simplified optimization problem. This problem does not occur with NIK, since gaps are filled with noisy predictions, which PISCO is capable to correct.

5.3. PISCO for NIK regularization

Upper leg Quantitative and qualitative results are visualized in Fig. 4 for two acceleration factors. NIK's reconstruction performance noticeably degrades with increasing acceleration, i.e. reduced amount of training data. PISCO-dist marginally improves reconstructions and metrics, but remains noisy. In contrast, the proposed PISCO-res results in less noisy reconstructions than both, NIK and PISCO-dist. Also, it recovers vessel details more reliably, represented in higher PSNR and FSIM, respectively. Again, no temporal redundancy is leveraged yet.

Cardiac cine Quantitative results of the cardiac cine dataset (Fig. 5) show that PISCO consistently outperforms NIK and PISCO-dist at all acceleration factors in both, spatial and temporal metrics. Particularly at high acceleration factors (R52/R104 or 4/2 spokes per frame), NIK's and PISCO-dist's spatial and temporal performance drastically decay. In contrast, PISCO enables spatial reconstruction quality similar or better to the reference method XD-GRASP25 (PSNR/FSIM-spat) and additionally, models the temporal dynamics better at these high accelerations (FSIM-temp). Among the INR-based methods, PISCO-res achieves reconstruction performance most comparable to TD-DIP512, while being more memory-efficient and resolution-agnostic.

Similar observations can be made in the qualitative reconstructions of one exemplary subject (Fig. 6). All motion-resolved reconstruction methods encounter the strong undersampling artefacts visible in INUFFT. XD-GRASP results in spatial smoothness by regularizing over the temporal dimension, which introduces blurring, particularly with reduced data (R52/R104). The DIP-based methods encounter this over-smoothing, yet the higher resolution version (TD-DIP512) is required to avoid noisy reconstructions. NIK, ICoNIK and PISCO-dist result in noisy spatiotemporal reconstructions, particularly with increasing acceleration factors. With PISCO, an improved neural k-space representation could be learned, that is spatially smooth, and recovers temporal dynamics even at 2 spokes/frame (R104). At high acceleration rates (R52/R104), PISCO surpasses the state-of-the-art XD-GRASP25 in capturing temporal detail (FSIM-temp) while achieving comparable spatial smoothness (PSNR/FSIM-spat). Further, PISCO-res is the only method that results in comparable performance as TD-DIP512. Note that in this case, only 25 time points were analyzed due to the use of binned reference data. Yet, PISCO's temporal resolution can further be increased by sampling more temporal points. Nonetheless, at accelerations like R104, the resulting images do not yet reach diagnostic quality but may be valuable for intermediate applications, such as motion estimation.

Abdominal Fig. 7 shows reconstructions of exemplary slices of the abdominal data. The reference unpaired prospectively GATED acquisition, common in a clinical setting, shows a spatial smooth image, but still appears slightly blurry at the organ edges and lacks temporal information overall. Motion-resolved reconstructions with 4MS retain undersampling artefacts (INUFFT4) or lose temporal information (XD-GRASP4). Increasing the temporal resolution by binning to 50 MS (XD-GRASP50) improves the dynamic depiction, but suffers from noise and undersampling. TD-DIP results in smoothing of the temporal domain, but loses any vessel detail. NIK, ICoNIK, PISCO-dist and PISCO-res achieve high spatiotemporal resolution, with PISCO-res further smoothing results both spatially and temporally.

General training details Examining the training details (see Supplementary S.2), PISCO-res not only enhances dynamic reconstruction performance but also significantly reduces training time compared to the previous PISCO-dist variant (22 vs. 42 min). While DIP-based methods offer faster training overall, their parameter count and memory usage scale with output resolution—leading to substantial increases at higher resolutions. In contrast, INR-based methods provide resolution-agnostic reconstructions with a single model and maintain a minimal memory footprint, requiring only a small fraction of the storage used by TD-DIP models (e.g., 4.3MB vs. 1443MB).




Fig. 5. Cardiac cine: Quantitative reconstruction results of 30 subjects for acceleration factors R15/R26/R52/R104 (with $\lambda=0.01/0.05/0.1/0.15$ for PISCO-dist/PISCO-res reconstruction, respectively). For R15, XD-GRASP25 offers least noisy reconstruction (PSNR) with high spatial and temporal resolution (FSIM-spat/FSIM-temp). For R26, NIK, PISCO-dist and PISCO-res lead to similar temporal results. At even larger acceleration factors, XD-GRASP25 performance decreases, as does NIK rapidly. PISCO-dist outperforms NIK, but does not reach XD-GRASP25. Yet, inclusion of the proposed PISCO-res significantly improves temporal structure compared to all other methods ($R>26$). At the same time, spatial reconstruction quality is maintained (R52) or even significantly improved (R104) compared to XD-GRASP25. All comparisons, except those marked with "N", are statistically significant (Friedman test followed by Wilcoxon signed-rank test with False Discovery Rate (FDR) correction as a post-hoc test, with significance set at $p<0.05$).

6. Discussion

Based on the concept of parallel imaging-inspired self-consistency (PISCO), we have proposed a novel k-space consistency measure which can be determined in a self-supervised manner. With multiple ablation/simulation studies, we have validated the convergence behaviour of PISCO, and hence, its applicability as objective function within dynamic MR reconstruction, i.e. using NIK. Without the need for any additional data, we have verified PISCO's potential to learn improved neural implicit representation, resulting in enhanced spatial and temporal image quality. Furthermore, the proposed PISCO-res loss outperforms our previous version, PISCO-dist, in computation times, decreasing the required time by almost 50 %.

PISCO for improved MR reconstruction. While NIK was originally developed for dynamic reconstruction, PISCO is not limited to this application. We have shown for both, simple k-space fitting (Section 3D) as well as learning a **static upper leg** NIK (Fig. 4), PISCO enables reduction of undersampling artefacts. Still, performance improvement will always be limited since no additional information except the neighborhood constrain is available. Exploiting additional forms of redundancy, i.e. given by the temporal dimension or multiple echoes (Spieker et al., 2024c), is expected to further enhance PISCO's regularization capabilities.

In our **cardiac cine** example, we show that incorporating the temporal dimension enhances representation learning for higher acceleration factors (e.g., achieving $R\geq 26$ in cardiac cine imaging versus $R20$ in upper leg imaging). Yet, NIK also faces challenges when very little information is available per frame ($R\geq 52$, or ≤ 4 spokes per frame), likely due to overfitting to noise in unsampled k-space regions. A similar behavior is observed with ICoNIK, which—although inspired by the convolutional GRAPPA concept—explicitly estimates kernel weights. This approach becomes unstable at high acceleration factors, failing to properly regularize the reconstruction. In contrast, the proposed PISCO-res addresses the overfitting noise by (1) including the unknown points in the training procedure and (2) enforcing consistency throughout all the k-space points without relying on explicit neighborhood definitions. This leads to visible improvements in both k-space and image-space. For lower acceleration factors, XD-GRASP remains a viable comparison but tends to oversmooth at $R\geq 52$ (Fig. 6), sacrificing temporal information. In contrast, PISCO-res maintains temporal resolution without such trade-offs. With PISCO-res, INR-based reconstruction approaches the performance of TD-DIP, while retaining its key advantage of supporting arbitrary spatial and temporal resolutions—unlike DIP-based methods, which require retraining for each new setting. However, in this cardiac cine scenario, PISCO's performance is constrained by the binned and discontinuous

nature of the dataset, which introduces residual motion blur during training. The advantage of PISCO over TD-DIP methods becomes more evident in more challenging, highly undersampled settings, as demonstrated in the abdominal data (see Fig. 7). Future studies on real time cardiac data are anticipated to further demonstrate PISCO-NIK's improved efficiency in this particular case.

In the **abdominal** case, although continuous data is available, the irregularity of respiratory motion compared to cardiac motion presents additional challenges. Acquired spokes are unequally distributed between end-exhale and end-inhale phases. Also, higher resolution and expanded field of view (FOV) requirements increase k-space gaps in radial acquisitions, complicating reconstruction. This presents visible challenges for XD-GRASP4 (Fig. 7), which prioritizes spatial resolution in abdominal imaging but sacrifices dynamic information. Compared to all other methods - both DIP- and INR- based - PISCO-res enables improved depiction of fine anatomical structures. Overall, the enhanced visualization of fine anatomical structures - such as small vessels and ductal networks - highlights PISCO's potential to improve the detection of malignant lesions, which in abdominal imaging can be as small as 1-2mm. Although we have not yet directly demonstrated sub-2mm lesion conspicuity, the enhanced visualization of these microvascular and ductal surrogates strongly suggests its potential to reveal lesions at or below the resolution limits of conventional imaging. Systematic, prospective validation of PISCO's impact on lesion detection, as well as its potential to enhance temporal resolution for motion-sensitive applications such as radiation therapy, will be subject to future work.

General PISCO design. Within our PISCO validation studies (Section 5.1), we have shown a feasible kernel, weight solving and consistency measure design to leverage the global neighborhood relationship for k-space optimization.

The kernel design plays a critical role in the design of PISCO, as only the Cartesian kernel exhibited the desired global consistency properties (Fig. 3A). In contrast, radial kernels demonstrate higher variability between subset weight vectors, consistent with the observations of Seiberlich et al. (2011), who noted that radial kernels are ideally calibrated in local segments rather across the entire k-space. To eliminate the effect of varying distance of the radial kernels points (e.g., tighter spacing near the center than at the periphery), an equidistant kernel was explored. Yet, similar inconsistencies were observed, emphasizing the importance of maintaining a consistent kernel orientation for self-consistency.

Kernel flexibility remains regarding kernel size and distance. In the k-space fitting example (Fig. 3D), a larger kernel resulted in even better reconstruction results due to the enlarged perceptive field, incorporating more information into the optimization problem. However, applying an




Fig. 6. Cardiac cine: Qualitative reconstruction results of one subject for four acceleration factors. The reference reconstruction (R1) is shown on the top right (xy), with the white dotted line marking the slice where the temporal profile (yt) is extracted. Videos of the complete dynamic reconstruction (xyt) can be found in the Supplementary. At lower acceleration factors (R15/26), XD-GRASP appears to have sufficient data to recover spatial information, yet at higher accelerations (R52/R104) it starts to oversmooth in the temporal domain (white arrow). The DIP methods retain temporal information, yet the higher output resolution (TD-DIP512) is required for blurring-free spatial images. While the INR-based NIK starts to be noisy at R26, PISCO-res can correct the representation to approximate TD-DIP's performance. When acceleration increases, NIK, ICoNIK and PISCO-dist suffer from increasing noise in the spatially and temporally (see difference images and blue arrows, respectively). Particularly for high accelerations (R52 +), PISCO-res results in similar spatial and improved temporal quality as XD-GRASP25 (white arrows).

increased kernel size also results in a larger computational overhead, since the number of weights N_w proportionally increases, and hence, the required number of patches to solve one subset. A more efficient handling of the coil dimension, since $N_w \propto N_c^2$, as potential solution remains subject to further research. A computationally efficient solution for weight solving would also open possibilities for advanced physical modeling with PISCO in the temporal dimension, such as motion, time-dependent MR field imperfections (Wang et al., 2019; Abraham et al., n.d.) or phase modeling (Halder and Zhuo, 2016).

Regarding the consistency measure, we have validated the improved convergence behaviour of the residual-based PISCO measure vs. the previously proposed distance-based loss (PISCO-dist in reconstruction

results). The reduced sensitivity to image noise (Fig. 3C) is also evident in the in-vivo reconstructions (Fig. 4/6), where PISCO-dist offers only marginal denoising improvements, likely due to suboptimal optimization. An additional benefit of the proposed PISCO-res measure is the reduced computational effort, since expensive distance calculations between all weight vectors (with thousands of complex numbers) are avoided.

While some parameters, such as those for weight solving, were generally applicable across datasets, others require adaptation to specific applications. One example is the weighting factor for the PISCO loss, λ , which has been adapted by observing PISCO's loss magnitude and behaviour during optimization. Extending the proposed residual-based

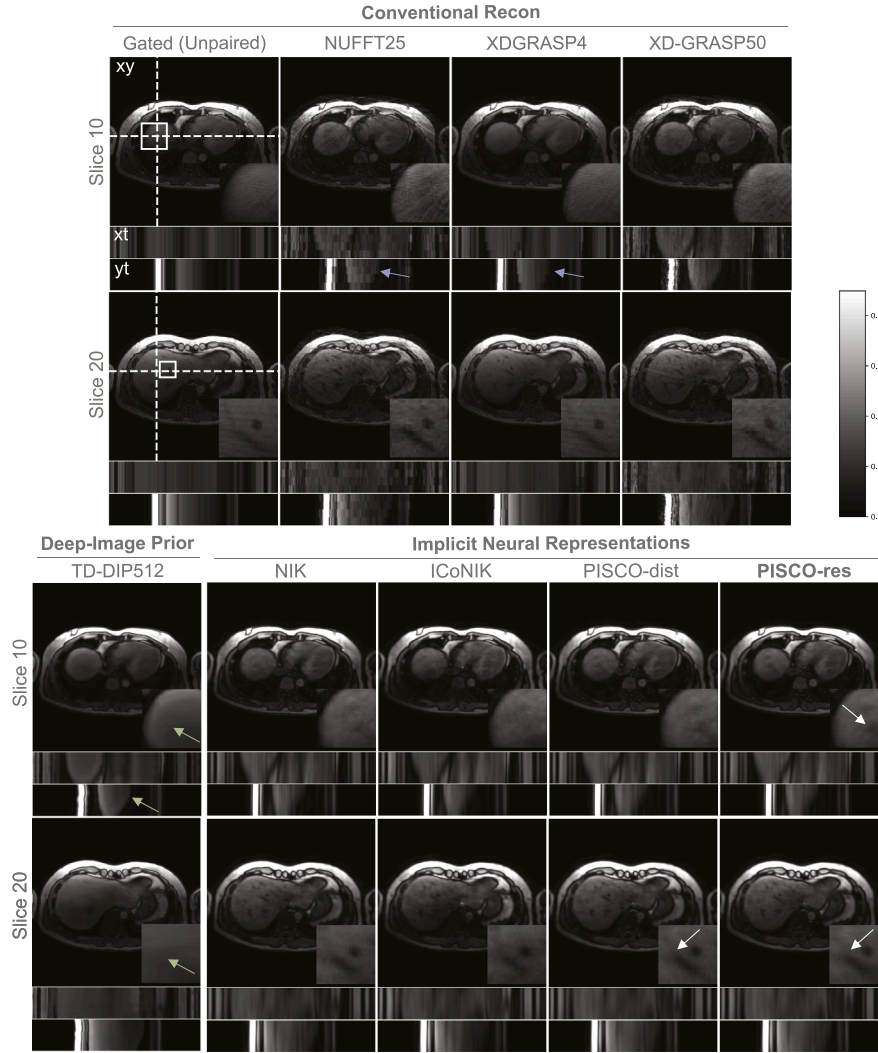


Fig. 7. Abdomen - respiratory resolved: Qualitative reconstruction results for two slices of a subject are shown: the xy-image depicts spatial reconstruction, while temporal profiles (marked by a white dotted line) indicate temporal reconstructions (xt/yt) below. Another examples as well as videos of the complete dynamic reconstruction (xyt) can be found in the Supplementary. The gated reconstruction (unpaired to the rest) does not offer any temporal information and remains blurry (zoom-in). Although XD-GRASP50 improves spatial resolution compared to NUFFT4, its temporal resolution remains limited (blue arrows). Increasing the number of MS (XD-GRASP50) results in better temporal resolution, but noisier spatiotemporal reconstructions. TD-DIP512 smoothens in the temporal domain, but loses any detail in the spatial dimension (green arrow). NIK and ICoNIK also reduce the noise in the temporal dimension, but remain noisy spatially. Both PISCO versions further smoothen spatiotemporal reconstructions, while the proposed PISCO-res maintains clearer vessel structures (white arrows)..

loss to automatically adapt to the dataset - considering factors such as kernel design, N_c , and acceleration - could further enhance PISCO's out-of-the-box applicability.

The observed NIK reconstruction improvements at high acceleration factors suggest PISCO's potential to shorten clinical protocol times. While these highly accelerated acquisitions may not yet achieve diagnostic-quality reconstructions, PISCO has shown significant enhancements over standard NIK-based dynamic reconstructions without requiring additional data. This can serve as a foundation for use cases where high temporal resolution is critical, i.e. it may open new avenues for rapid motion estimation, which could be beneficial for subsequent analysis and downstream tasks. Moreover, its efficient design also allows flexible integration with other regularization methods, offering opportunities for further improvements in reconstruction quality. This could include, but is not limited to, extensions that incorporate low-rank priors or constraints [Haldar and Zhuo \(2016\)](#). However, identifying the optimal combinations and integration strategies would require further investigation, which we consider a valuable direction for future work.

Limitations. From a **method design** perspective, we have demonstrated that the PISCO framework can effectively enhance reconstruction quality and explored several design variants. Yet, some design aspects remain open for further investigation, e.g., larger kernel sizes may offer potential benefits but are currently constrained by computational resources and implementation limitations. The PISCO weighting factor was selected empirically and the feasible number of subsets only approximated. As a k-space quality metric, PISCO's performance is inherently influenced by these design choices. While parameters were manually fine-tuned in this study, a more automated and adaptive configuration would be desirable for clinical deployment and will be investigated in future works. Importantly, PISCO provides a mechanism to evaluate the quality of k-space estimates but does not inherently guide convergence toward an optimal reconstruction. Therefore, it must be integrated with data consistency constraints. Our current implementation evaluates this integration in two reconstruction settings, but future work could explore its application within a broader range of reconstruction frameworks.

The proposed integration of PISCO with NIK leads to further limitations inherent to INR-based approaches. Although INRs do not depend

on external training data, they require per-subject optimization, leading to longer runtimes compared to convolutional methods like TD-DIP. Optimization of training speed remains crucial to fully leverage the advantages of the resolution-agnostic and lightweight MLP architecture of INRs. Beyond reducing runtime, such improvements are key to overcome the current limitation of 2D slice-wise training, as efficient 3D reconstruction is essential for clinical applicability.

Further limitations should be acknowledged looking at the **datasets**. A key advantage of INRs is their ability to model high temporal resolution; however, this requires access to a reliable and highly resolved temporal signal to fully leverage their potential. In our experiments, the cardiac dataset was retrospectively discretized into 25 cardiac states through binning, which inherently limits temporal continuity. The abdominal dataset retained continuous temporal information and demonstrated better performance, highlighting the potential of INRs in such settings. However, the use of self-navigation introduces uncertainties in the temporal timestamps, making accurate modeling critical. Furthermore, respiratory motion is generally more variable and less periodic than cardiac motion, which may necessitate more advanced temporal encoding strategies, such as accounting for hysteresis effects and intra-patient variability.

In terms of **evaluation**, our experiments covered multiple anatomies but were largely restricted to radial acquisitions with navigators and a single retrospective Cartesian example. A broader evaluation across acquisition strategies, anatomical regions, and contrast settings is needed to better understand PISCO's generalizability. Furthermore, although the current work demonstrates promising technical performance, future studies should assess clinical utility through dedicated evaluation using clinical parameters and expert reader assessments.

7. Conclusion

With PISCO, we have demonstrated how a conventional parallel imaging concept can be adapted into a self-supervised consistency measure that enhances learning-based MR reconstruction. Its calibration-free and flexible design allows for seamless integration into the training process, making it a promising method for application in other anatomies or k-space based reconstruction techniques.

CRedit authorship contribution statement

Veronika Spieker: Writing – original draft, Validation, Data curation, Writing – review & editing, Visualization, Software, Methodology, Conceptualization, Project administration, Formal analysis; **Hannah Eichhorn:** Methodology, Conceptualization, Writing – review & editing, Formal analysis; **Wenqi Huang:** Conceptualization, Writing – review & editing, Formal analysis; **Jonathan K. Stelter:** Validation, Conceptualization, Data curation, Writing – review & editing; **Tabita Catalan:** Writing – review & editing, Software, Validation; **Rickmer F. Braren:** Writing – review & editing, Resources, Supervision, Formal analysis; **Daniel Rueckert:** Supervision, Writing – review & editing, Conceptualization; **Francisco Sahli Costabal:** Conceptualization, Supervision, Writing – review & editing; **Kerstin Hammernik:** Writing – review & editing, Supervision, Conceptualization, Methodology; **Dimitrios C. Karampinos:** Supervision, Writing – review & editing, Resources, Conceptualization, Project administration; **Claudia Prieto:** Supervision, Conceptualization, Writing – review & editing, Project administration; **Julia A. Schnabel:** Conceptualization, Supervision, Writing – review & editing, Resources, Project administration.

Data availability

The code will be publicly available upon publication. The authors used public data (accessible as noted in manuscript) as well as in-house data for evaluation purposes.

Declaration of competing interest

Editor of MedIA - J.A.S. Editor of MedIA - D.R. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at [10.1016/j.media.2025.103890](https://doi.org/10.1016/j.media.2025.103890)

References

- Abraham, D., Nishimura, M., Cao, X., Liao, C., Setsompop, K., 2023. Implicit representation of GRAPPA kernels for fast MRI reconstruction. <https://arxiv.org/pdf/2310.10823>.
- Ahmad, R., Bouman, C.A., Buzzard, G.T., Chan, S., Liu, S., Reehorst, E.T., Schniter, P., 2020. Plug-and-play methods for magnetic resonance imaging: using denoisers for image recovery. *IEEE Signal Process. Mag.* 37 (1), 105–116. <https://doi.org/10.1109/msp.2019.2949470>
- Akçakaya, M., Moeller, S., Weingärtner, S., Uğurbil, K., 2019. Scan-specific robust artificial-neural-networks for k-space interpolation (RAKI) reconstruction: database-free deep learning for fast imaging. *Magn. Reson. Med.* 81 (1), 439–453. <https://doi.org/10.1002/mrm.27420>
- Breuer, F.A., Kellman, P., Griswold, M.A., Jakob, P.M., 2005. Dynamic autocalibrated parallel imaging using temporal GRAPPA (TGRAPPA). *Magn. Reson. Med.* 53 (4), 981–985. <https://doi.org/10.1002/mrm.20430>
- Catalán, T., Courdurier, M., Osses, A., Botnar, R., Costabal, F.S., Prieto, C., 2023. Unsupervised reconstruction of accelerated cardiac cine MRI using Neural Fields. <https://arxiv.org/pdf/2307.14363v1>.
- Catalán, T., Courdurier, M., Osses, A., Fotaki, A., Botnar, R., Sahli-Costabal, F., Prieto, C., 2025. Unsupervised reconstruction of accelerated cardiac cine MRI using neural fields. *Comput. Biol. Med.* 185, 109467. <https://doi.org/10.1016/j.combiomed.2024.109467>
- Chen, Z., Chen, Y., Xie, Y., Li, D., Christodoulou, A.G., 2022. Data-consistent non-cartesian deep subspace learning for efficient dynamic MR image reconstruction. In: 2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI). IEEE, pp. 1–5. <https://doi.org/10.1109/isbi52829.2022.9761497>
- Chu, J., Du, C., Lin, X., Zhang, X., Wang, L., Zhang, Y., Wei, H., 2025. Highly accelerated MRI via implicit neural representation guided posterior sampling of diffusion models. *Med. Image Anal.* 100, 103398. <https://doi.org/10.1016/j.media.2024.103398>
- Cui, Z.-X., Jia, S., Cao, C., Zhu, Q., Liu, C., Qiu, Z., Liu, Y., Cheng, J., Wang, H., Zhu, Y., Liang, D., 2023. K-UNN: K-space interpolation with untrained neural network. *Med. Image Anal.* 88, 102877. <https://doi.org/10.1016/j.media.2023.102877>
- El-Rewaady, H., 2020. Replication Data for: Multi-Domain Convolutional Neural Network (MD-CNN) For Radial Reconstruction of Dynamic Cardiac MRI. Harvard Dataverse. <https://doi.org/10.7910/DVN/C13WB6>
- El-Rewaady, H., Fahmy, A.S., Pashakhanloo, F., Cai, X., Kucukseymen, S., Csics, I., Neisius, U., Haji-Valizadeh, H., Menze, B., Nezafat, R., 2021. Multi-domain convolutional neural network (MD-CNN) for radial reconstruction of dynamic cardiac MRI. *Magn. Reson. Med.* 85 (3), 1195–1208. <https://doi.org/10.1002/mrm.28485>
- Feng, J., Feng, R., Wu, Q., Shen, X., Chen, L., Li, X., Feng, L., Chen, J., Zhang, Z., Liu, C., Zhang, Y., Wei, H., 2025. Spatiotemporal implicit neural representation for unsupervised dynamic MRI reconstruction. *IEEE Trans. Med. Imaging*, 1. <https://doi.org/10.1109/tmi.2025.3526452>
- Feng, L., 2022. Golden-Angle radial MRI: basics, advances, and applications. *J. Magn. Reson. Imaging* 56 (1), 45–62. <https://doi.org/10.1002/jmri.28187>
- Feng, L., Axel, L., Chandarana, H., Block, K.T., Sodickson, D.K., Otazo, R., 2016. XD-GRASP: Golden-angle radial MRI with reconstruction of extra motion-state dimensions using compressed sensing. *Magn. Reson. Med.* 75 (2), 775–788. <https://doi.org/10.1002/mrm.25665>
- Griswold, M.A., Jakob, P.M., Heidemann, R.M., Nittka, M., Jellus, V., Wang, J., Kiefer, B., Haase, A., 2002. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn. Reson. Med.* 47 (6), 1202–1210. <https://doi.org/10.1002/mrm.10171>
- Haldar, J.P., Zhuo, J., 2016. P-LORAKS: low-rank modeling of local k-space neighborhoods with parallel imaging data. *Magn. Reson. Med.* 75 (4), 1499–1514. <https://doi.org/10.1002/mrm.25717>
- Hamilton, J., Franson, D., Seiberlich, N., 2017. Recent advances in parallel imaging for MRI. *Prog. Nucl. Magn. Reson. Spectrosc.* 101, 71–95. <https://doi.org/10.1016/j.pnmrs.2017.04.002>
- Hammernik, K., Kustner, T., Yaman, B., Huang, Z., Rueckert, D., Knoll, F., Akçakaya, M., 2023. Physics-driven deep learning for computational magnetic resonance imaging: combining physics and machine learning for improved medical imaging. *IEEE Signal Process. Mag.* 40 (1), 98–114. <https://doi.org/10.1109/msp.2022.3215288>
- Han, Y., Sunwoo, L., Ye, J.C., 2020. K-space deep learning for accelerated MRI. *IEEE Trans. Med. Imaging* 39 (2), 377–386. <https://doi.org/10.1109/TMI.2019.2927101>
- Huang, W., Ke, Z., Cui, Z.-X., Cheng, J., Qiu, Z., Jia, S., Ying, L., Zhu, Y., Liang, D., 2021. Deep low-Rank plus sparse network for dynamic MR imaging. *Med. Image Anal.* 73, 102190. <https://doi.org/10.1016/j.media.2021.102190>
- Huang, W., Li, H.B., Pan, J., Cruz, G., Rueckert, D., Hammernik, K., 2023. 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*. Springer Nature Switzerland, Cham, pp. 548–560.
- Huang, W., Spieker, V., Xu, S., Cruz, G., Prieto, C., Schnabel, J., Hammernik, K., Kuestner, T., Rueckert, D., 2024. Subspace implicit neural representations for real-time cardiac cine MR imaging. <http://arxiv.org/pdf/2412.12742v1>.
- Jafari, R., Do, R. K.G., LaGratta, M.D., Fung, M., Bayram, E., Cashen, T., Otazo, R., 2023. GRASPNET: fast spatiotemporal deep learning reconstruction of golden-angle radial data for free-breathing dynamic contrast-enhanced magnetic resonance imaging. *NMR Biomed.* 36 (3), e4861. <https://doi.org/10.1002/nbm.4861>
- Kunz, J.F., Ruschke, S., Heckel, R., 2024. Implicit neural networks with fourier-feature inputs for free-breathing cardiac MRI reconstruction. *IEEE Trans. Comput. Imaging* 10, 1280–1289. <https://doi.org/10.1109/tci.2024.3452008>
- Lustig, M., Pauly, J.M., 2010. SPIRiT: iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magn. Reson. Med.* 64 (2), 457–471. <https://doi.org/10.1002/mrm.22428>
- Marchetto, E., Eichhorn, H., Gallichan, D., Schnabel, J.A., Ganz, M., 2025. Agreement of image quality metrics with radiological evaluation in the presence of motion artifacts. *Magn. Reson. Mater. Phys., Biol. Med.*, 1–12 <https://link.springer.com/article/10.1007/s10334-025-01266-y>. <https://doi.org/10.1007/s10334-025-01266-y>
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R., 2020. NeRF: representing scenes as neural radiance fields for view synthesis. *Springer, Cham*, pp. 405–421. https://link.springer.com/chapter/10.1007/978-3-030-58452-8_24.
- Muckley, M.J., Stern, R., Murrell, T., Knoll, F., 2020. TorchKbNufft: a high-level, hardware-agnostic non-uniform fast fourier transform. In: *ISMRM Workshop on Data Sampling & Image Reconstruction*. Vol. 22. <https://mmuckley.github.io/assets/publications/2020muckleytorchkbnufft.pdf>.
- Reddi, S.J., Kale, S., Kumar, S., 2019. On the convergence of adam and beyond. <http://arxiv.org/abs/1904.09237>
- Ryu, K., Alkan, C., Choi, C., Jang, I., Vasanawala, S., 2021. 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. <https://doi.org/10.1109/iccvw54120.2021.00446>
- Seiberlich, N., Ehse, P., Duerk, J., Gilkeson, R., Griswold, M., 2011. Improved radial GRAPPA calibration for real-time free-breathing cardiac imaging. *Magn. Reson. Med.* 65 (2), 492–505. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3012142/>. <https://doi.org/10.1002/mrm.22618>
- Shao, H.-C., Mengke, T., Deng, J., Zhang, Y., 2024. 3D cine-magnetic resonance imaging using spatial and temporal implicit neural representation learning (STINR-MR). *Phys. Med. Biol.* 69 (9). <https://doi.org/10.1088/1361-6560/ad33b7>
- Shen, L., Pauly, J., Xing, L., 2022. NeRP: implicit neural representation learning with prior embedding for sparsely sampled image reconstruction. *IEEE Trans. Neural Netw. Learn. Syst.* PP. <https://doi.org/10.1109/TNNLS.2022.3177134>
- Shimron, E., Tamir, J.I., Wang, K., Lustig, M., 2022. Implicit data crimes: machine learning bias arising from misuse of public data. *Proc. Natl. Acad. Sci. U.S.A.* 119 (13), e2117203119. <https://doi.org/10.1073/pnas.2117203119>
- Shin, P.J., Larson, P. E.Z., Ohliger, M.A., Elad, M., Pauly, J.M., Vigneron, D.B., Lustig, M., 2014. Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion. *Magn. Reson. Med.* 72 (4), 959–970. <https://doi.org/10.1002/mrm.24997>
- Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G., 2020. Implicit neural representations with periodic activation functions. *Adv. Neural Inf. Process. Syst.* 33, 7462–7473.
- Spieker, V., Eichhorn, H., Hammernik, K., Rueckert, D., Preibisch, C., Karampinos, D.C., Schnabel, J.A., 2023. Deep learning for retrospective motion correction in MRI: a comprehensive review. *IEEE Trans. Med. Imaging* PP. <https://doi.org/10.1109/TMI.2023.3323215>
- Spieker, V., Eichhorn, H., Stelter, J.K., Huang, W., Braren, R.F., Rueckert, D., Sahli Costabal, F., Hammernik, K., Prieto, C., Karampinos, D.C., Schnabel, J.A., 2024a. Self-supervised k-space regularization for motion-resolved abdominal MRI using neural implicit k-space representations. In: Dou, Q., Feragen, A., Giannarou, S., Glocker, B., Lekadir, K., Linguraru, M.G., Schnabel, J.A. (Eds.), *Medical Image Computing and Computer Assisted Intervention - MICCAI 2024*. Springer International Publishing AG, Cham. Vol. 15007 of *Lecture Notes in Computer Science*, pp. 614–624. https://doi.org/10.1007/978-3-031-72104-5_59
- Spieker, V., Huang, W., Eichhorn, H., Stelter, J., Weiss, K., Zimmer, V.A., Braren, R.F., Karampinos, D.C., Hammernik, K., Schnabel, J.A., 2024b. 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*. Springer Nature Switzerland, Cham. Vol. 14533 of *Lecture Notes in Computer Science*, pp. 183–192. https://doi.org/10.1007/978-3-031-53767-7_18
- Spieker, V., Stelter, J., Huang, W., Eichhorn, H., Weiss, K., Braren, R., Zimmer, V., Hammernik, K., Prieto, C., Karampinos, D., Schnabel, J.A., 2024c. DE-NIK: leveraging dual-echo data for respiratory-resolved abdominal MR reconstructions using neural implicit k-space representations. In: *2024 ISMRM & ISMRT Annual Meeting & Exhibition*. Vol. 32. <https://doi.org/10.58530/2024/0014>
- Terpstra, M.L., Maspero, M., Verhoeff, J. J.C., van den Berg, C. A.T., 2023. Accelerated respiratory-resolved 4D-MRI with separable spatio-temporal neural networks. *Med. Phys.* 50 (9), 5331–5342. <https://doi.org/10.1002/mp.16643>
- Uecker, M., Lai, P., Murphy, M.J., Virtue, P., Elad, M., Pauly, J.M., Vasanawala, S.S., Lustig, M., 2014. ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA. *Magn. Reson. Med.* 71 (3), 990–1001. <https://doi.org/10.1002/mrm.24751>
- Wang, F., Dong, Z., Reese, T.G., Bilgic, B., Katherine Manhard, M., Chen, J., Polimeni, J.R., Wald, L.L., Setsompop, K., 2019. Echo planar time-resolved imaging (EPTI). *Magn. Reson. Med.* 81 (6), 3599–3615. <https://doi.org/10.1002/mrm.27673>
- Wang, S., Ke, Z., Cheng, H., Jia, S., Ying, L., Zheng, H., Liang, D., 2022. DIMENSION: dynamic MR imaging with both k-space and spatial prior knowledge obtained via multi-supervised network training. *NMR Biomed.* 35 (4), e4131. <https://doi.org/10.1002/nbm.4131>
- Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P., 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* 13 (4), 600–612. <https://doi.org/10.1109/tip.2003.819861>
- Yoo, J., Jin, K.H., Gupta, H., Yerly, J., Stuber, M., Unser, M., 2021. Time-dependent deep image prior for dynamic MRI. *IEEE Trans. Med. Imaging* 40 (12), 3337–3348. <https://doi.org/10.1109/tmi.2021.3084288>
- Zaitsev, M., Maclaren, J., Herbst, M., 2015. Motion artifacts in MRI: a complex problem with many partial solutions. *J. Magn. Reson. Imaging* 42 (4), 887–901. <https://doi.org/10.1002/jmri.24850>
- Zhang, L., Zhang, L., Mou, X., Zhang, D., 2011. FSIM: a feature similarity index for image quality assessment. *IEEE Trans. Image Process.* 20 (8), 2378–2386. <https://doi.org/10.1109/tip.2011.2109730>