

# DD-INR: Dynamics-Driven Implicit Neural Representation for Accelerated Whole-Brain Functional MRI Reconstruction

Qiaoxin Li<sup>1,2</sup>, Caini Pan<sup>1,2,3</sup>, Pierre-Antoine Comby<sup>1,2</sup>, Chaithya Giliyar Radhakrishna<sup>1,2</sup>, and Philippe Ciuciu<sup>1,2\*</sup>

<sup>1</sup> MIND, Inria, Palaiseau, France

<sup>2</sup> Neurospin, CEA Paris Saclay, France

<sup>3</sup> CEA NeuroSpin, Paris-Saclay University, CNRS BAOBAB, Gif-sur-Yvette, France  
philippe.ciuciu@cea.fr

**Abstract.** Accelerated acquisition of fMRI enables enhanced detection of neurovascular (BOLD) activity in the brain, but image reconstruction becomes challenging with high k-space undersampling: Task-evoked BOLD signals are small in magnitude, which traditional anatomical MRI reconstruction methods fail to recover, as they favor spatial accuracy over temporal fidelity. We present DD-INR, a Dynamics-Driver Implicit Neural Representation framework tailored for accelerated fMRI that benefits from incoherent time-varying sampling and a tailored spatio-temporal prior, outperforming traditional methods, demonstrated in simulation and in-vivo acquisition, both in terms of image quality and retrieval of activation patterns.

DD-INR achieves this by splitting the fMRI data into a static background and a temporally varying dynamic component, representing only the dynamics with a dedicated INR, thereby focusing the model's capacity on activation-relevant changes while remaining compact. In general, DD-INR provides a promising framework for accelerated fMRI reconstruction, with the potential to improve the sensitivity and robustness of fMRI studies within practical scan time limits. *The source code is available at <https://github.com/JoosenLi/DD-INR>.*

**Keywords:** functional MRI · Deep Learning · Implicit Neural Representation · Image Reconstruction. · Inverse Problem

## 1 Introduction

Functional Magnetic Resonance Imaging (fMRI) is pivotal for non-invasive brain mapping, yet it faces an inherent trade-off between spatial resolution, temporal resolution, and signal-to-noise ratio (SNR) [11]. High spatiotemporal resolution is critical to capture hemodynamic responses and improve the reliability of downstream tasks [1, 15, 16]. However, acquiring high spatial and temporal resolution

---

\* Corresponding author

fMRI images requires strong k-space undersampling, making image reconstruction an ill-posed inverse problem. In particular, collecting high-speed fMRI data while maintaining full brain coverage requires time-varying k-space sampling patterns such as spiral stack or variations [23, 10].

Solving the image reconstruction problem in fMRI requires introducing prior knowledge into reconstruction via traditional methods such as compressed sensing (CS) with spatial or temporal regularization [2, 7, 23]. More recently, deep learning (DL) based methods have emerged, relying on supervised [17, 27] or self-supervised techniques [20, 13, 14], but their direct use in fMRI remains difficult: high-resolution ground-truth data for training are scarce, volumetric 3D+time data are computationally demanding, and hallucinated anatomical or functional features can be detrimental to downstream fMRI analysis. Solutions such as Plug-and-Play (PnP) have shown promising results in fMRI [9], but remain limited to volume-wise processing.

To overcome data scarcity and generalization issues, Implicit Neural Representation (INR) has gained traction [30, 22, 28, 24, 25]. INRs model images as continuous functions via coordinate-based networks, offering memory efficiency and optimization at reconstruction time from the measured data of a single subject rather than trained offline on a curated cohort. Although INRs have excelled in dynamic cardiac MRI [18, 6, 12, 4, 19], this application focuses mainly on resolving macroscopic anatomical deformations, targeting 2D+time sequences. However, fMRI must capture subtle and possibly fast fluctuations [21] in the intensity of the BOLD signal (transient signal changes of only 2–5% [5]). This mismatch between reconstruction goals and the prohibitive memory requirements of volumetric representations renders the direct application of motion-centric INR approaches poorly suited for functional neuroimaging.


This paper introduces DD-INR, a *Dynamics-Driven Implicit Neural Representation* framework explicitly tailored for accelerated fMRI reconstruction. Unlike generic dynamic INRs, DD-INR splits the 3D+time signal into a static anatomical background, reconstructed from temporally accumulated data, and a time-varying dynamic component, represented by a subject-specific INR. This decomposition allows the INR to focus its capacity on activation-relevant fluctuations rather than redundant anatomical structures that remain constant over time. Furthermore, the differentiability of INRs is leveraged to enforce priors that encode spatial coherence of temporal changes, consistent with the hemodynamic responses observed in fMRI.

## 2 Reconstructing fMRI data with INR

### 2.1 Reconstruction Formulation

Reconstructing a sequence of fMRI volumes from accelerated multi-coil k-space data can be written as:

$$x = \arg \min_{x \in \mathbb{C}^{N_{\text{xyz}} \times T}} \sum_{t=1}^T \|A_t x_t - y_t\|_2^2 + \lambda \mathcal{R}(x), \quad (1)$$



**Fig. 1.** Architecture of proposed DD-INR for accelerated 3D+time (3D+T) fMRI reconstruction with time-varying sampling. 3D+T fMRI data is acquired using a time-varying sampling pattern. For reconstruction, first a static background volume  $\hat{x}_{bg}$  is reconstructed from samples pooled across time into a 3D k-space dataset. Then a SIREN INR is trained on the same subject to reconstruct the dynamic part  $x_{dyn}$  encapsulating the BOLD signal. The training loop is highlighted with red arrows. The dynamic signal is regularized by spatial TV of the temporal derivative of the INR, computed via automatic differentiation of the SIREN block.

where  $A_t$  represents the forward model (e.g. non-uniform Fourier) operator in frame  $t$  (encapsulating the sampling pattern  $\Omega_t$  and the multi-coil encoding), for the associated  $y_t$  k-space data, and  $\lambda\mathcal{R}(x)$  is a regularization term promoting a given prior.

**Dynamics-driven signal factorization.** In task-based fMRI, activation-related BOLD signal variations are small compared to the average intensity of the anatomical background. Directly modeling the full 3D+T signal, therefore, tends to spend most of the model capacity on time-invariant structure rather than on task-relevant dynamics. To better match the reconstruction objective, the signal at spatial coordinate  $r$  and time  $t$  is split into a background and a dynamic part:

$$x(r, t) = \hat{x}_{bg}(r) + x_{dyn}(r, t). \quad (2)$$

The background is estimated from a temporally aggregated 3D k-space dataset. Specifically, the samples and trajectories  $\{y_t, \Omega_t\}_{t=1}^T$  are pooled across time,

yielding  $y_{1:T}$  and a static forward operator  $A_{1:T}$ :

$$\hat{x}_{\text{bg}} = \arg \min_{z \in \mathbb{C}^{N_{\text{xyz}}}} \|A_{1:T}z - y_{1:T}\|_2^2 = A_{1:T}^\dagger y_{1:T}, \quad (3)$$

where  $A_{1:T}^\dagger$  denotes the Moore-Penrose inverse of the aggregated operator.

We now focus on parameterizing  $x_{\text{dyn}}$  using SIREN [26], accounting for spatial and temporal dimensions  $f_\theta(r, t)$ . To improve the sensitivity to local variations in the fMRI signal of  $f_\theta$ , we use Fourier feature mapping embedding [29], separated in space and time:

$$\begin{aligned} x_{\text{dyn}}(r, t) &= f_\theta(r, t) = \text{SIREN}_\theta(\gamma_{\text{xyz}}(r), \gamma_t(t)), \\ \gamma_{\text{xyz}}(r) &= [\sin(2\pi B_{\text{xyz}}r), \cos(2\pi B_{\text{xyz}}r)], \\ \gamma_t(t) &= [\sin(2\pi B_t t), \cos(2\pi B_t t)]. \end{aligned} \quad (4)$$

where  $B_{\text{xyz}}$  and  $B_t$  are random Gaussian matrices. This decoupled design allows us to allocate different bandwidths and embedding dimensions, improving network capacity while reducing the overall parameter count and facilitating the recovery of subtle temporal dynamics.

**INR regularization.** In task-based fMRI, the BOLD signal is expected to vary smoothly across neighboring voxels and to evolve coherently within activated regions. We therefore consider two Total-Variation (TV) penalties:

$$\mathcal{R}_{\text{xyz}}(\theta) = \sum_{t=1}^T \text{TV}_{\text{xyz}}(x_{\text{dyn}}(\cdot, t)) \quad \text{and} \quad \mathcal{R}_{\text{xyz},t}(\theta) = \sum_{t=1}^T \frac{1}{t} \text{TV}_{\text{xyz}}(\partial_t x_{\text{dyn}}(\cdot, t)). \quad (5)$$

$\mathcal{R}_{\text{xyz}}$  applies spatial TV directly to the dynamic residual at each time point.  $\mathcal{R}_{\text{xyz},t}$  instead applies spatial TV to the temporal derivative, promoting spatial coherence in the temporal changes while preserving the anatomical structure as well as temporal changes. The intrinsic differentiable property of the INR makes  $\partial_t x_{\text{dyn}}$  directly available by autodifferentiation, and the spatial components required for TV are computed using finite differences to reduce computational cost. Therefore, dynamic reconstruction is obtained by optimizing

$$\mathcal{L}(\theta) = \sum_{t=1}^T \|A_t(\hat{x}_{\text{bg}} + f_\theta(r, t)) - y_t\|_2^2 + \lambda \mathcal{R}(\theta), \quad (6)$$

where  $\mathcal{R}$  denotes  $\mathcal{R}_{\text{xyz}}$  or  $\mathcal{R}_{\text{xyz},t}$  depending on the chosen prior and  $\lambda > 0$  is a regularization hyperparameter.

### 3 Experimental Setup

We evaluated the performance of DD-INR on both simulated and in vivo acquired data. The simulated k-space data were generated using the SNAKE simulator [10], which provides ground-truth hemodynamic activity localized on an

anatomical phantom [3] consisting only of white matter, gray matter, and CSF to remain feasible on consumer hardware. We simulated a 3T-like 3D gradient-echo sequence with parameters ( $TR_{\text{shot}} = 50ms$ ,  $TE = 30ms$ ,  $FA = 12^\circ$ ), a field of view of  $192 \times 192 \times 120 \text{ mm}^3$ , 3mm-isotropic resolution,  $TR_{\text{vol}} = 1s$ , and total scan duration of 240 s, with a block-on/block-off (20s/20s) visual paradigm acquired with a stack-of-spiral trajectory (AF=2 in the  $k_z$  direction). The locations of the spirals in the stack were drawn randomly with a uniform distribution for each frame independently, while enforcing full sampling of the center 20% of  $k_z$  [23].

To evaluate reconstruction performance in a real-world setting, an in vivo experiment was conducted on an healthy volunteer using a 3T scanner (Siemens MAGNETOM Cima.X), following the same acquisition protocol as in simulation. The study was approved by the relevant local and national ethics committees, and written informed consent was obtained prior to participation. Dedicated 44-channel coil sensitivity maps were acquired before the fMRI run. In simulation, SNR was defined in k-space following [10] and tuned to match the multi-coil in vivo acquisition setting.


**Implementation Details** The forward model operators  $A_t$  were implemented using MRI-NUFFT[8]. The background component  $\hat{x}_{\text{bg}}$  is first to be estimated, by aggregating all  $(k, t)$ -space samples and trajectories, following Eq. (3) using 25 iterations of conjugate gradient (CG).

The dynamic component was modeled using a SIREN [26] (3 layers of width 512) with spatial ( $B_{\text{xyz}} \in \mathbb{R}^{256 \times 3}$ ) and temporal encodings ( $B_t \in \mathbb{R}^{64}$ ) and two output channels for the real and imaginary parts.

Optimization was conducted with 1000 iterations of ADAM(learning rates  $2 \cdot 10^{-5}$ ). Regularization weights were introduced using a warm-up schedule:  $\lambda = 0$  for the first 100 iterations, then increased linearly for the next 200 iterations (up to  $\lambda = 10^{-2}$  for  $\mathcal{R}_{\text{xyz}}$  and  $\lambda = 2 \cdot 10^{-3}$  for  $\mathcal{R}_{\text{xyz},t}$ )

**Comparison Methods and Metrics** DD-INR was compared against representative baselines spanning direct reconstruction, least squares reconstruction (CG), compressed sensing, model-based learning, and implicit neural representations. For fair comparison, competing methods used the  $\hat{x}_{\text{bg}}$  estimated by DD-INR as initialization or prior input whenever applicable. We compare to the following methods: **NUFFT**: Adjoint NUFFT without regularization, used as a reference baseline. **CG**:  $\hat{x}_{\text{bg}}$ -initialized CG, representing an iterative baseline. **CS**:  $\hat{x}_{\text{bg}}$ -initialized wavelet-regularized CS [2], representing a sparsity-based baseline. **PnP**: Plug-and-play reconstruction, specifically designed for fMRI [9], representing model-based learning approaches that incorporate learned priors. **INR**: Prior-based INR [24, 6], used as a representative INR baseline;  $\hat{x}_{\text{bg}}$  was provided for network initialization.

To analyze the activation detection accuracy, GLM-based  $z$ -score maps were computed and thresholded at  $p < 10^{-3}$ . The resulting binary maps were compared with the simulated ground-truth activation region to compute  $F_2$ -scores,



**Fig. 2.** GLM z-score activation maps from simulated 3D+T data. Columns compare NUFFT, CG, CS, PnP, INR, DD-INR without regularization, DD-INR with spatial regularization, and DD-INR with spatiotemporal regularization. Grayscale images show the magnitude of the first reconstructed frame ( $T_2^*$ ), with three orthogonal views displayed. Cyan contours indicate the simulated ground-truth activation region in SNAKE. Red arrows highlight false-positive or false-negative activations, and blue arrows indicate artifacts in reconstructed  $T_2^*$  images. The  $F_2$ -score is reported for each method.

emphasizing recall and the detection of weak activation signals. Moreover, temporal fidelity was evaluated using the Pearson correlation coefficient (Corr) and normalized root mean square error (NRMSE) between the ROI-averaged reconstructed time series and the expected BOLD response. All metrics were computed over valid slices (defined as slices with more than 10 activated voxels in the ground-truth region) and reported as mean  $\pm$  standard deviation. Corr reflects temporal consistency, while NRMSE quantifies amplitude accuracy.

## 4 Results


Fig. 2 compares the activation maps and the corresponding  $F_2$ -scores in both simulation settings. NUFFT, CG, and CS exhibit noisy reconstructions with limited overlap with the reference activation map. PnP improves statistical sensitivity, but yields spatially diffuse patterns. The naive INR reconstructs anatomically plausible images but shows reduced statistical sensitivity, indicating insufficient recovery of subtle temporal dynamics when modeling the full 3D+T subject without dynamics-centric design.

DD-INR consistently improves alignment with the simulated activation region in the brain and achieves higher  $F_2$ -scores. Spatial regularization improves stability but slightly reduces localization precision, whereas spatiotemporal regularization better preserves activation boundaries.

Fig. 3 presents ROI-averaged time courses within the simulated activation region alongside the expected BOLD response. Most baselines capture coarse stimulus-locked trends but exhibit either pronounced temporal noise (NUFFT, CG, CS) or attenuated response amplitudes (PnP, naive INR). DD-INR provides a substantial improvement in temporal consistency and amplitude preservation.

**Table 1.** Quantitative comparison of ROI time-series recovery in simulation. Correlation coefficient and NRMSE are computed against the expected BOLD response and reported as mean  $\pm$  std over valid slices. Best results are in **bold**, and second-best results are underlined. DD-INR- $R_{xyz,t}$  and DD-INR- $R_{xyz}$  denote DD-INR with spatiotemporal and spatial regularization, respectively.

|                    | DD-INR- $R_{xyz,t}$      | DD-INR- $R_{xyz}$        | INR             | PnP             | CS              | CG              | NUFFT           |
|--------------------|--------------------------|--------------------------|-----------------|-----------------|-----------------|-----------------|-----------------|
| Corr $\uparrow$    | <b>0.909</b> $\pm$ 0.058 | 0.903 $\pm$ 0.054        | 0.40 $\pm$ 0.16 | 0.84 $\pm$ 0.11 | 0.55 $\pm$ 0.13 | 0.39 $\pm$ 0.23 | 0.39 $\pm$ 0.16 |
| NRMSE $\downarrow$ | <b>0.136</b> $\pm$ 0.037 | <u>0.142</u> $\pm$ 0.033 | 0.35 $\pm$ 0.04 | 0.17 $\pm$ 0.05 | 0.30 $\pm$ 0.04 | 0.35 $\pm$ 0.07 | 0.36 $\pm$ 0.04 |



**Fig. 3.** ROI-averaged time series in the simulated activation region. Signals are normalized for visualization.


Tab. 1 further corroborates these observations quantitatively. DD-INR with spatiotemporal regularization achieves the highest correlation and lowest NRMSE.

Fig. 4 presents *in vivo* activation maps under the visual block-design paradigm. CG yields few activations, while CS provides limited improvement with small activation clusters. PnP produces broader activation patterns but exhibits over-smooth  $T_2^*$  images, indicating reduced localization precision. The naive INR baseline reconstructs plausible anatomy, yet completely fails to detect task-locked activations. In contrast, DD-INR shows consistent and spatially coherent activation patterns across views while preserving anatomical structure.

## 5 Discussion and Conclusion

**Simulation vs *in vivo* experiment.** The *in vivo* experiment yielded results consistent with the simulation, despite potentially being further degraded by uncontrollable factors such as physiological noise, scanner imperfections, and off-resonance artifacts. In the simulation, we quantitatively assessed the performance of DD-INR and provided a fair comparison with state-of-the-art fMRI reconstruction methods. Both *in vivo* and simulation studies show that spatial or spatiotemporal regularization improves the detection and strength of brain activation after standard analysis.

**Downstream fMRI inference.** Spatiotemporal regularisation may alter temporal correlations, residual variance, and effective degrees of freedom, which



**Fig. 4.** In vivo visual block-design results.<sup>2</sup> The top-left sub-figure shows the T1-weighted anatomical reference. Subsequent columns compare baselines and DD-INR with spatial / spatiotemporal regularization. Grayscale images correspond to reconstructed  $T_2^*$  magnitude, with overlaid z-score activation maps.

in turn can skew the GLM analysis. To control for this effect, our simulation evaluates GLM-derived activation maps against a known ground truth, in addition to ROI time-series fidelity. Under this setting, DD-INR and its regularizers increase statistical sensitivity while preserving reasonable spatial specificity. Nevertheless, further work is needed to quantify reconstruction-induced activation spread and to understand how such effects impact downstream analysis.

**Spatio-temporal vs Spatial regularization.** Both spatial and spatio-temporal regularization improve BOLD recovery in the INR reconstruction. Spatial-only regularization promotes image-domain smoothness, which can blur activation boundaries, spread task-related dynamics, or partially suppress weak BOLD components. Instead, the spatio-temporal regularization encourages the spatial coherence of temporal variations, allowing neighboring voxels to exhibit consistent dynamic changes while preserving local spatial texture and temporal response amplitude. As a result, it reduces incoherent fluctuations and mitigates the diffuse activation patterns observed with spatial-only regularization.

**Impact of movement.** DD-INR assumes a time-invariant background; therefore, any potential residual head motion has to be accommodated by the INR component, and head motion can violate this assumption. The consistent in vivo activation patterns observed in this work suggest that the proposed factorization is practical under routine fMRI motion levels.

**Relation to previous INR reconstructions.** To our knowledge, DD-INR is the first INR framework explicitly tailored for accelerated fMRI. Most existing INR-based dynamic MRI reconstructions model the full spatiotemporal signal (or a latent subspace that captures both anatomy and dynamics) and are primarily developed for motion-dominated applications [18, 12, 6, 4], often with periodic dynamics (e.g., cardiac CINE). Chen et al. [6] leverage temporally aggregated measurements to stabilize training, but the INR is still required to model both static anatomy and temporal variations. In contrast, in fMRI, the goal is to recover weak BOLD fluctuations superimposed on a strong static background, rather than to estimate the fine anatomical structure or motion. DD-INR

is therefore explicitly problem-driven: It reconstructs a static background from temporally aggregated data and uses the INR capacity to represent only the dynamic residual, which better matches the objective of activation recovery.

**Conclusions.** This work introduces DD-INR, a dynamic-driven implicit neural representation framework for accelerated fMRI that adapts signal representation, network architecture, and regularization to the dynamics of fMRI, concentrating model capacity on activation-related fluctuations. Experiments on simulated and in vivo data demonstrate improved activation detection and temporal fidelity under acceleration, highlighting the value of dynamics-centric modeling for more sensitive and robust accelerated fMRI.

**Acknowledgments.** This work is funded by WP2 (Optimizing Workflows Efficiency) of The European Joint Virtual Lab on Artificial Intelligence, Data Analytics and Scalable Simulation (AIDAS), an initiative of Forschungszentrum Jülich (FZJ) and the French Alternative Energies and Atomic Energy Commission (CEA).

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

## References

1. Aggarwal, H., Al-Shikhley, L., Thirion, B.: Across-subject ensemble-learning alleviates the need for large samples for fMRI decoding. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 35–45. Springer (2024)
2. Amor, Z., Comby, P.A., Ster, C.L., Vignaud, A., Ciuciu, P.: Non-cartesian non-fourier fMRI imaging for high-resolution retinotopic mapping at 7 tesla. In: 2023 IEEE 9th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP). p. 201–205 (2023)
3. Aubert-Broche, B., Griffin, M., Pike, G., Evans, A., Collins, D.: Twenty New Digital Brain Phantoms for Creation of Validation Image Data Bases **25**(11), 1410–1416. <https://doi.org/10.1109/TMI.2006.883453>
4. Baik, D., Yoo, J.: Dynamic-aware spatio-temporal representation learning for dynamic MRI reconstruction. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 174–184. Springer (2025)
5. Chang, C., Thomason, M.E., Glover, G.H.: Mapping and correction of vascular hemodynamic latency in the bold signal. *Neuroimage* **43**(1), 90–102 (2008)
6. Chen, L., Balter, J.M., Shen, L., Park, J.J.: Single-spoke motion-compensated dynamic 3d MRI reconstruction via neural representation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 513–522. Springer (2025)
7. Chiew, M., Graedel, N.N., Miller, K.L.: Recovering task fMRI signals from highly under-sampled data with low-rank and temporal subspace constraints. *NeuroImage* **174**, 97–110 (2018)
8. Comby, P.A., Daval-Fr erot, G., Pan, C., Tanabene, A., Oudjman, L., Cencini, M., Ciuciu, P.: MRI-nufft: Doing non-cartesian MRI has never been easier. *Journal of Open Source Software* **10**(108), 7743 (2025)

9. Comby, P.A., Terris, M., Vignaud, A., Ciuciu, P.: Plug-and-play reconstruction for 3d non-cartesian fMRI data. In: 2025 33rd European Signal Processing Conference (EUSIPCO). pp. 1015–1019. IEEE (2025)
10. Comby, P.A., Vignaud, A., Ciuciu, P.: Snake: A modular realistic fMRI data simulator from the space-time domain to k-space and back. *Imaging Neuroscience* **3**, IMAG-a (2025)
11. Faro, S.H., Mohamed, F.B., Law, M., Ulmer, J.T.: *Functional neuroradiology: principles and clinical applications*. Springer Science & Business Media (2011)
12. Feng, J., Feng, R., Wu, Q., Shen, X., Chen, L., Li, X., Feng, L., Chen, J., Zhang, Z., Liu, C., et al.: Spatiotemporal implicit neural representation for unsupervised dynamic MRI reconstruction. *IEEE Transactions on Medical Imaging* **44**(5), 2143–2156 (2025)
13. Gu, H., Zhang, C., Yu, Z., Rettenmeier, C., Stenger, V.A., Akçakaya, M.: Non-cartesian self-supervised physics-driven deep learning reconstruction for highly-accelerated multi-echo spiral fMRI. In: 2024 IEEE International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2024)
14. Hu, C., Li, C., Wang, H., Liu, Q., Zheng, H., Wang, S.: Self-supervised learning for MRI reconstruction with a parallel network training framework. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 382–391. Springer (2021)
15. Hu, D., Han, K., Cheng, J., Li, G.: Consecutive-contrastive spherical u-net: Enhancing reliability of individualized functional brain parcellation for short-duration fMRI scans. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 88–98. Springer (2024)
16. Huang, J., Chen, N., Qiu, A.: Topological cycle graph attention network for brain functional connectivity. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 723–732. Springer (2024)
17. Huang, W., Ke, Z., Cui, Z.X., Cheng, J., Qiu, Z., Jia, S., Ying, L., Zhu, Y., Liang, D.: Deep low-rank plus sparse network for dynamic mr imaging. *Medical Image Analysis* **73**, 102190 (2021)
18. 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: *International Conference on Information Processing in Medical Imaging*. pp. 548–560. Springer (2023)
19. Huang, W., Wang, N., Liao, C., Lin, Y., Gao, M., Rueckert, D., Setsompop, K.: Physics-informed implicit neural representations for joint b0 estimation and echo planar imaging. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 481–490. Springer (2025)
20. Korkmaz, Y., Cukur, T., Patel, V.M.: Self-supervised MRI reconstruction with unrolled diffusion models. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 491–501. Springer (2023)
21. Lewis, L.D., Setsompop, K., Rosen, B.R., Polimeni, J.R.: Fast fMRI can detect oscillatory neural activity in humans. *Proceedings of the national academy of sciences* **113**(43), E6679–E6685 (2016)
22. Li, Q., Liang, D., Li, Y.: Acceleration: Sequentially-scanning dual-energy computed tomography imaging using high temporal resolution image reconstruction and temporal extrapolation. In: 2024 IEEE International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2024)
23. Petrov, A.Y., Herbst, M., Andrew Stenger, V.: Improving temporal resolution in fMRI using a 3d spiral acquisition and low rank plus sparse (l+s) reconstruction **157**, 660–674 (2017)

24. Shen, L., Pauly, J., Xing, L.: Nerp: implicit neural representation learning with prior embedding for sparsely sampled image reconstruction. *IEEE transactions on neural networks and learning systems* **35**(1), 770–782 (2022)
25. Shen, X., Feng, J., Feng, R., Zhang, Y., Wei, H.: Imj-plus: Implicit representation for dynamic MRI and coil sensitivity joint reconstruction using low-rank plus sparse regularization. In: *2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI)*. pp. 1–5. IEEE (2025)
26. 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)
27. Sun, J., Li, H., Xu, Z., et al.: Deep admm-net for compressive sensing MRI. *Advances in neural information processing systems* **29** (2016)
28. Xu, S., Früh, M., Hammernik, K., Lingg, A., Kübler, J., Krumm, P., Rueckert, D., Gatidis, S., Küstner, T.: Self-supervised feature learning for cardiac cine mr image reconstruction. *IEEE Transactions on Medical Imaging* (2025)
29. Zheng, J., Ramasinghe, S., Lucey, S.: Rethinking positional encoding. *arXiv preprint arXiv:2107.02561* (2021)
30. Zhu, Y., Liu, Y., Zhang, Y., Liang, D.: Implicit neural representation for medical image reconstruction. *Physics in Medicine & Biology* **70**(12), 12TR01 (2025)