Abstract #00027 WORKSHOP ABSTRACT

Joint Implicit Neural Representation for Fast Scan-Specific Magnetic Resonance Fingerprinting

Traditional Poster · 150 min | Session 7: Outside Poster Session · Tuesday, January 13 at 03:30 PM

Hongze Yu [1], **Christopher Keen**[2], **Kaixuan Jin**[1], **Jeffrey A Fessler**[1,2,3], **Yun Jiang**[2,3]

[1]Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, United States of America
[2]Department of Biomedical Engineering, University of Michigan, Ann Arbor, United States of America
[3]Department of Radiology, University of Michigan, Ann Arbor, United States of America

Presenting Author: *Hongze Yu (hongze@umich.edu)*

Introduction

MR Fingerprinting (MRF) enables rapid multiparametric mapping[1,2]. However, the acquisitions are highly undersampled with hundreds of timepoints, making reconstruction ill-posed and slow. Model-based iterative methods[3] and network-based priors[4] can suppress aliasing but have long reconstruction times, limiting clinical applications.

To address these limitations, we propose a scan-specific MRF reconstruction using a gradient-guided joint implicit neural representation (INR[5]). A single INR parameterizes all subspace images derived from the MRF time-series, acting as an implicit spatial prior that suppresses subspace-inconsistent aliasing while preserving fine structure. Our proposed method fits an INR in the image domain, without training data or iterative k-space data consistency needed for other INR-based methods[6-9], reducing reconstruction time. We evaluated our method in quantitative T1 and T2 mapping using the ISMRM/NIST phantom and in vivo brain imaging and compared against maps from low-rank[10] and iterative LLR reconstructions[3].

Methods

Let $\boldsymbol{y} \in \mathbb{C}^{T\times C\times Ns\times Nr}$ be the multi-coil MRF time-series, where T is the number of timepoints, C is the number of coils, and (N_s, N_r) are the number of spiral interleaves and readout samples. We compute a temporal basis $\boldsymbol{U}_K \in \mathbb{C}^{T\times K}$ by low-rank approximation of an MRF dictionary, and form the initial subspace images $\boldsymbol{x}(0)$ by applying adjoint encoding operator $\boldsymbol{A}$, with density compensation $\boldsymbol{W}$:

$$\boldsymbol{x}^{(0)} \triangleq \boldsymbol{A}^{'}\boldsymbol{W}\left(\boldsymbol{U}_K^{'}\boldsymbol{y}\right) \in \mathbb{C}^{K\times N_y\times N_x}.$$

We select K components preserving 99.99% of MRF signal. A single INR ($f_{\boldsymbol{\theta}}: \mathbb{R}^d \to \mathbb{C}^K$) jointly models all K subspace images, mapping each spatial coordinate $\vec{\boldsymbol{r}}$ to the K complex values at that location. This produces a joint feature vector $\boldsymbol{\phi}_{\boldsymbol{\theta}}(\vec{\boldsymbol{r}}) \in \mathbb{R}^F$ at the penultimate layer that encodes anatomy. The kth subspace image is generated from this shared representation by linear combination at a lightweight output head:

$$\boldsymbol{x}_{\boldsymbol{\theta},k}(\vec{\boldsymbol{r}}) = \boldsymbol{w}_k^T\,\boldsymbol{\phi}_{\boldsymbol{\theta}}(\vec{\boldsymbol{r}}) + b_k.$$

This shared architecture (Figure 1a) enforces an implicit prior that is shared among all subspace images.

Fitting is further aided by gradient dominance. The first subspace image has high SNR and minimal aliasing. Early in training, its larger loss gradients force $\boldsymbol{\phi}_{\boldsymbol{\theta}}(\vec{\boldsymbol{r}})$ to encode a clean anatomical representation. Later, more aliased components cannot introduce new structure; they reuse $\boldsymbol{\phi}_{\boldsymbol{\theta}}(\vec{\boldsymbol{r}})$ and adjust contrast through their linear heads. This suppresses subspace-inconsistent aliasing while preserving details. Figure 1b shows $\boldsymbol{\phi}_{\boldsymbol{\theta}}(\vec{\boldsymbol{r}})$ evolving from noise to an aliasing-free shared representation. Separate encoder and MLP weight regularization further limits overfitting.

Our method fits an INR by iteratively optimizing its weights using an image-domain $\ell 2$ loss with weight regularization:

$$\hat{\boldsymbol{\theta}}(\boldsymbol{\beta}) = \arg\min_{\boldsymbol{\theta}} \left\|\boldsymbol{x}^{(0)} - f_{\boldsymbol{\theta}}(\vec{\boldsymbol{r}})\right\|_2^2 + \lambda_{\text{Enc}}\left\|\boldsymbol{\theta}_{\text{Enc}}\right\|_2^2 + \lambda_{\text{MLP}}\left\|\boldsymbol{\theta}_{\text{MLP}}\right\|_2^2, \qquad \boldsymbol{x}_{\boldsymbol{\theta}} = f_{\boldsymbol{\theta}}(\vec{\boldsymbol{r}}).$$

The iterative fitting avoids the MRI forward model (e.g., Fourier encoding, gridding), greatly reducing the reconstruction time. Quantitative T1 and T2 maps are obtained by matching the reconstructed subspace images to an MRF dictionary.

Phantom and In vivo experiments: The ISMRM/NIST phantom and three volunteers were scanned using a FISP-based MRF[2] (Variable flip angle, TE=2 ms, TR=9 ms and 1000 timepoints) at 3T (Vida, Siemens). Each timepoint was highly undersampled using a single interleaf (readout duration = 2.21 ms) of a variable-density spiral trajectory that requires 96 interleaves to fully sample the k-space.

We compared the proposed method against low-rank reconstruction[10] and an iterative reconstruction with LLR regularization[3]. All methods were implemented in PyTorch on an NVIDIA A40 GPU.

Results

Figure 2 shows all subspace images during INR fitting. Early iterations collapse to a shared image resembling the first subspace, and later iterations gradually disentangle contrast across subspaces.

Figure 3a shows correlation plots comparing T1/T2 values from the proposed method with 500 timepoints and those from 1000-timepoint low-rank reconstruction. The proposed method (R^2=0.9980) outperforms the 500-timepoint low-rank reconstruction (R^2=0.9089), especially for long T2 (>300 ms). It

also has the lowest coefficient of variations (Fig. 3b): 4.3% T1 and 9.4% T2 vs 4.9% and 19.3% of the 1000-timepoints low-rank reconstruction baseline.

Figure 4 compares subspace images and maps from low-rank reconstruction, LLR, and the proposed method, showing improved image and map quality from the proposed method.

Figure 5 presents representative slices of brain maps. The proposed method suppresses aliasing and avoids the blurring, compared to low-rank reconstruction and LLR. It also produces the lowest T2 standard deviation in white matter. The reconstruction time of the proposed method was ~6.2 s, compared to 158 s for LLR and ~5.6 s for non-iterative low-rank reconstruction.

Discussion and Conclusion

The proposed method learns a shared, scan-specific spatial representation and uses it as an implicit prior to suppress aliasing in highly undersampled MRF. This yields sharp, low-variance T1/T2 maps from as few as 500 timepoints. The approach is fully image-domain, needs no training data, and avoids iterative k-space data consistency, enabling faster reconstruction while improving map quality. These results suggest that reduced acquisition time and near–real-time reconstruction are feasible for MRF without sacrificing quantitative accuracy or precision.

References

1. D. Ma, V. Gulani, N. Seiberlich, K. Liu, J. L. Sunshine, J. L. Duerk, and M. A. Griswold, "Magnetic resonance fingerprinting," Nature, vol. 495, no. 7440, pp. 187–192, 2013. doi.org/10.1038/nature11971.

2. Y. Jiang, D. Ma, N. Seiberlich, V. Gulani, and M. A. Griswold, "MR fingerprinting using fast imaging with steady state precession (FISP) with spiral readout," Magnetic Resonance in Medicine, vol. 74, no. 6, pp. 1621–1631, 2015. doi.org/10.1002/mrm.25559.

3. G. Lima da Cruz, A. Bustin, O. Jaubert, T. Schneider, R. M. Botnar, and C. Prieto, “Sparsity and locally low rank regularization for MR fingerprinting,” Magnetic Resonance in Medicine, vol. 81, no. 6, pp. 3530–3543, 2019. doi.org/10.1002/mrm.27665.

4. J. I. Hamilton, “A self-supervised deep learning reconstruction for shortening the breathhold and acquisition window in cardiac magnetic resonance fingerprinting,” Frontiers in Cardiovascular Medicine, vol. 9, 2022, doi:10.3389/fcvm.2022.928546.

5. V. Sitzmann, J. N. P. Martel, A. W. Bergman, D. B. Lindell, and G. Wetzstein, "Implicit Neural Representations with Periodic Activation Functions," in Advances in Neural Information Processing Systems, vol. 33, pp. 7462–7473, 2020.

6. R. Feng, Q. Wu, J. Feng, H. She, C. Liu, Y. Zhang, and H. Wei, “IMJENSE: Scan-Specific Implicit Representation for Joint Coil Sensitivity and Image Estimation in Parallel MRI,” IEEE Transactions on Medical Imaging, vol. 43, no. 4, pp. 1539–1553, 2024. doi: 10.1109/TMI.2023.3342156

7. J. Feng, R. Feng, Q. Wu, X. Shen, L. Chen, X. Li, L. Feng, J. Chen, Z. Zhang, C. Liu, Y. Zhang, and H. Wei, "Spatiotemporal Implicit Neural Representation for Unsupervised Dynamic MRI Reconstruction," IEEE Transactions on Medical Imaging, vol. 44, no. 5, pp. 2143–2156, 2025. doi: 10.1109/TMI.2025.3526452

8. C. Li, J. Li, J. Zhang, E. Solomon, A. V. Dimov, P. Spincemaille, T. D. Nguyen, M. R. Prince, and Y. Wang, "Navigator motion-resolved MR fingerprinting using implicit neural representation: Feasibility for free-breathing three-dimensional whole-liver multiparametric mapping," Magnetic Resonance in Medicine, 2025, advance online publication. doi: 10.1002/mrm.70063

9. H. Yu, J. A. Fessler, and Y. Jiang, "Bilevel Optimized Implicit Neural Representation for Scan-Specific Accelerated MRI Reconstruction," arXiv preprint arXiv:2502.21292, January 2025.

10. D. F. McGivney, E. Pierre, D. Ma, Y. Jiang, H. Saybasili, V. Gulani, and M. A. Griswold, "SVD Compression for Magnetic Resonance Fingerprinting in the Time Domain," IEEE Transactions on Medical Imaging, vol. 33, no. 12, pp. 2311–2322, 2014. doi.org/10.1109/TMI.2014.2337321.

Figures and Tables

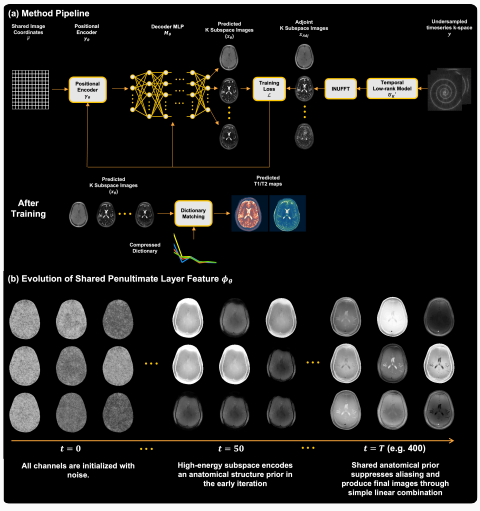

Scan for high-resolution version

Figure 1: *(a) Proposed reconstruction pipeline. (b). Evolution of the shared penultimate-layer feature representation $\bm{\phi}_{\bm{\theta}}(\vec{\bm{r}})$ (channels 1-9). Features are first shaped by the high-energy subspace (k with largest energy, e.g., k=1) and then reused as an anatomical structure prior to constrain lower-energy subspaces (k≥3), suppressing aliasing. All subspace images are produced by linear combiantion of $\bm{\phi}_{\bm{\theta}}(\vec{\bm{r}})$ at output head.*

Animated Figure 2

This figure contains animation/video content. To view the animation, scan the QR code with your mobile device or visit the online version of this abstract.

Figure 2: *Method Illustration (2D Brain MRF reconstruction). The animation shows the evolution of subspace images and T1/T2 maps. Early high-energy subspace (e.g., k=1) defines a shared anatomical prior; low-energy subspace images (e.g., k≥3) gain contrast while preserving anatomical structure and suppressing aliasing artifacts.*

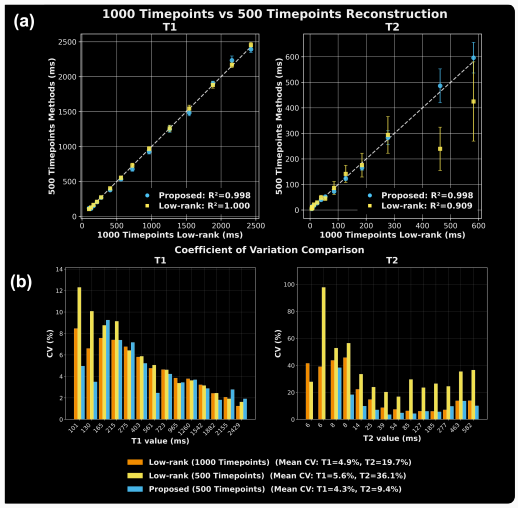

Scan for high-resolution version

Figure 3: *ISMRM/NIST phantom results (variable-density spiral, 2.21 ms readout, 48–96× undersampled). (a) Correlation of T1/T2 between the proposed/low-rank reconstruction (500 timepoints) and 1000-timepoints low-rank reconstruction reference, (b) coefficient of variation comparison. The proposed method has the lowest mean coefficient of variation (T1: 4.3%, T2 9.4%), even lower than the 1000-timepoints reference (T1: 4.9%, T2 19.3%).*

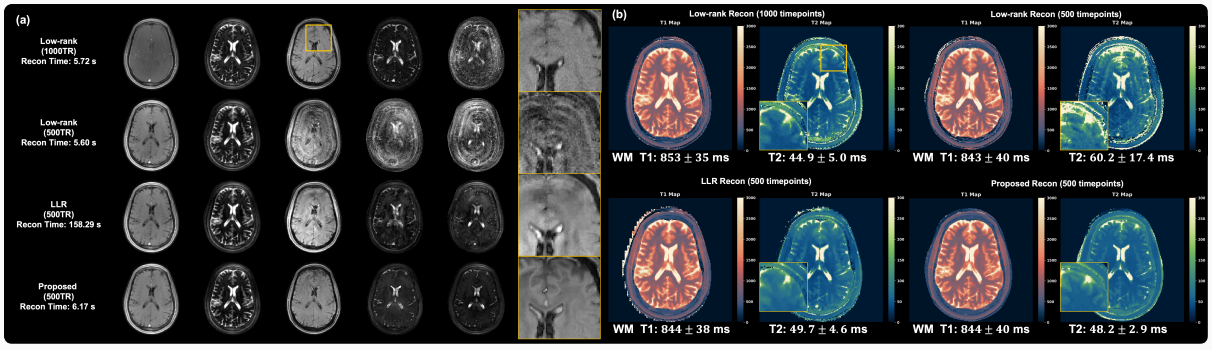

Scan for high-resolution version

Figure 4: *Single-slice 2D brain MRF reconstruction comparing low-rank adjoint (500/1000 timepoints), LLR (500 timepoints), and the proposed method (500 timepoints). (a) Subspace images with per-method runtime. (b) T1/T2 maps with white-matter means/SDs. The proposed method suppresses subspace-inconsistent aliasing, preserves anatomy, and yields the lowest WM-T2 standard deviation. Regularization strength of iterative LLR reconstruction was tuned by grid search.*

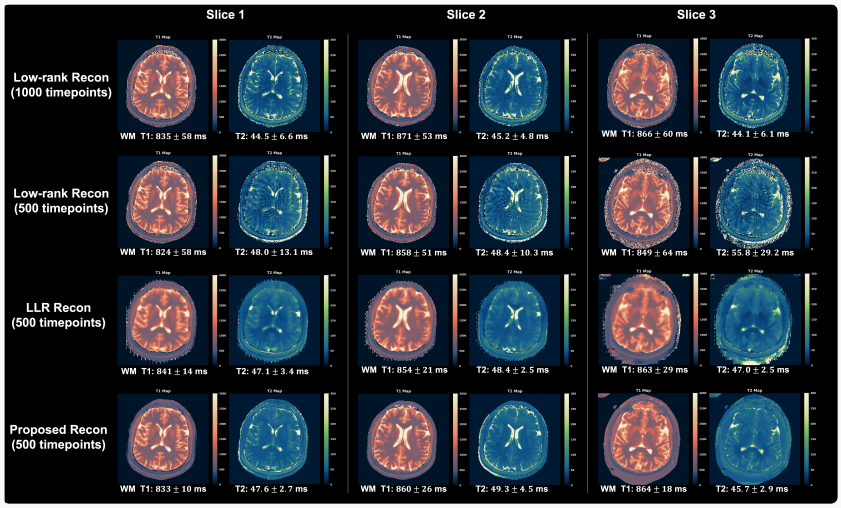

Scan for high-resolution version

Figure 5: *Multi-slice 2D brain MRF (three representative slices, 2.21 ms readout). The proposed method suppresses aliasing and blur versus low-rank and LLR, preserves sharper anatomical structure, and yields the lowest white-matter T2 standard deviation. Runtimes: proposed ~6.2 s; LLR 158 s; low-rank adjoint (non-iterative) ~5.6 s. Regularization strength of iterative LLR reconstruction was tuned by grid search.*