0389

Bilevel Optimized Implicit Neural Representation for Scan-Specific Accelerated MRI Reconstruction

Hongze Yu[1], Jeffrey A. Fessler[1,2,3], and Yun Jiang[2,3]
[1]Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, United States, [2]Department of Radiology, University of Michigan, Ann Arbor, MI, United States, [3]Department of Biomedical Engineering, University of Michigan, Ann Arbor, MI, United States

Synopsis

Keywords: AI/ML Image Reconstruction, AI/ML Image Reconstruction, Implicit Neural Representation

Motivation: Supervised deep learning methods require application-specific training datasets and perform poorly with out-of-distribution data. Scan-specific methods do not require training data, but need careful hyperparameter tuning.

Goal(s): To propose an automatically hyperparameter-optimized, scan-specific deep learning method that reconstructs various highly accelerated MRI acquisitions.

Approach: We developed a self-supervised, bilevel-optimized, implicit neural representation (INR) network. It splits undersampled data into training and validation sets and applies Bayesian optimization for hyperparameter tuning. A multiresolution trainable parametric encoder reconstructs accelerated MRI scans.

Results: Our method achieved performance comparable to oracle-optimized reconstructions, demonstrating the benefits of automatic hyperparameter optimization, and outperformed existing model-based and self-supervised methods.

Impact: By automatically optimizing hyperparameters for scan-specific deep learning, our method reconstructs accelerated MRI scans across diverse protocols with superior image quality. It avoids reliance on training data and complicated task-dependent tuning, enhancing the clinical applicability of deep learning in MRI.

Introduction

Deep learning (DL) based methods[1] can reconstruct highly accelerated MRI scans, but are limited by requiring application-specific large training datasets and by generalizing poorly to out-of-distribution data. Self-supervised deep learning algorithms[2] perform scan-specific reconstruction, but still need complicated hyperparameter tuning based on the acquisition and often provide limited acceleration.
This work developed a Bilevel Optimized Implicit Neural Representation[3] (INR) network for scan-specific MR image reconstruction. The proposed algorithm leverages Gaussian Process regression[4] for optimizing hyperparameters of an INR, enabling tailored reconstruction for various acquisitions. The INR includes a trainable positional encoder for high-dimensional feature embedding and a small multilayer perceptron for decoding. The bilevel optimization takes ≈10 minutes for a 2D Cartesian scan with a matrix size of 384×384. With offline-optimized hyperparameters transferable for the same acquisition, the scan-specific reconstruction takes only 10 seconds. Non-Cartesian reconstruction is approximately 5× slower due to NUFFT computation.

Methods

The proposed algorithm has two levels: (1) lower-level INR-based image reconstruction, and (2) upper-level hyperparameter optimization using dataset splitting[4] and Bayesian Optimization[5] (BayesOpt).
Figure 1(a) shows the method pipeline. The undersampled k-space data is randomly split into an 80% training set and a 20% validation set. A hyperparameter vector β, containing the ℓ_2 regularization strength of network weights λ_{Enc} and λ_{MLP}, learning rate α, and loss weighting controller ϵ, is sampled and used to train the INR, which is then tested on the validation set. The hyperparameter sample with the smallest validation loss is chosen.
The lower-level reconstruction INR, serving as an implicit regularizer, comprises a trainable multiresolution encoder (i.e., Hash Encoder[6]) and a small decoder MLP[7]. For a given β, the lower level is an optimization problem with respect to the network parameters as

$$\mathbf{x}_\theta^* = f_{\theta^*}(\vec{\mathbf{r}}), \qquad \theta^* = \mathrm{argmin}_\theta \|\mathbf{W}_T(y - \mathbf{A}f_\theta(\vec{\mathbf{r}}))\|_2^2, \qquad \text{s.t.} \qquad f_\theta(\vec{\mathbf{r}}) = M_{\theta_{\mathrm{MLP}}}(\gamma_{\theta_{\mathrm{Enc}}}.(\vec{\mathbf{r}})),$$

where $\mathbf{x}_\theta^* \in \mathbb{C}^N$ denotes the reconstructed image, $y \in \mathbb{C}^{KM}$ is the multi-coil k-space data. The learned INR function $f_\theta : \mathbb{R}^{N\times d} \to \mathbb{R}^{N\times 2}$ maps the coordinates, $\vec{\mathrm{r}}$, to the real and imaginary parts of the target image. The network comprises the pointwise positional encoder, $\gamma_{\theta_{\mathrm{Enc}}}$, and the decoder MLP, $M_{\theta_{\mathrm{MLP}}}$. The self-weighted density compensation matrix, $\mathbf{W} \in \mathbb{R}^{KM\times KM}$, emphasizes the higher frequency components, and updates through iteration as

$$\mathbf{W} = \mathrm{Diag}(|\mathbf{A}x| + \epsilon\mathbf{1})^{-1}\ .$$

Figure 1(b) illustrates the multiresolution Hash Encoder with 16 levels of feature encoding. In each level, the vertices enclosing the target coordinate $\vec{\mathrm{r}}$ are indexed by the spatial hash function to the features stored in the trainable table, H_l. The feature representation of $\vec{\mathrm{r}}$ is bilinearly interpolated and concatenated into a high-dimensional feature vector that feeds into the decoder. The upper-level hyperparameter optimization uses BayesOpt, modeling the hyperparameter sampling and validation loss minimization as Gaussian Process. The overall bilevel optimization problem is

$$\begin{aligned} \beta^* &:= \mathrm{argmin}_\beta L_V(\theta'(\beta)) \quad \text{where} \\ L_V(\theta'(\beta)) &:= \|\mathbf{W}_V(y - \mathbf{A}f_{\theta'(\beta)}(\vec{\mathbf{r}}))\|_2^2 \quad \text{and } \theta'(\beta) := \mathrm{argmin}_\theta L_T(\beta, \theta) \end{aligned}\ .$$

We used 60 upper-level iterations with 4000 lower-level iterations for self-supervised reconstruction. Oracle optimization using fully-sampled data for validation loss was used for performance comparison.
In vivo experiments: Brain imaging was performed using bSSFP and T2-weighted TSE sequences with resolutions of both 1×1 mm^2 and 1.2×1.2 mm^2 on Siemens 3T and 0.55T scanners. Cardiac experiments used a T2-weighted bSSFP sequence with a resolution of 1.4×1.4 mm^2 at 0.55T. Prostate imaging used T2-weighted TSE with a resolution of 1×1 mm^2 at 3T. Retrospective undersampling was used for Cartesian data, and spiral acquisition was simulated using torchkbnufft toolbox[8].

Results

Figure 2 presents the hyperparameter optimization results. Bayesian-optimized INR achieved comparable normalized root-mean-square error (NRMSE), structural similarity index measure (SSIM), and peak signal-to-noise ratio (PSNR) to oracle-optimized INR. Figure 2 (b-d) show that hyperparameters obtained from Bayesian-optimized INR are transferable between different subjects when using the same sequence, but different sequences and field strengths require tailored hyperparameter optimization.
Figure 3 compares our method with model-based and self-supervised INR methods[2]. Our method achieved higher metrics across all anatomies, sampling patterns, and field strengths.
Figure 4 shows an ablation study on the Encoder, Decoder, and loss function weighting. Figures 4(a) and 4(b) demonstrate the major contribution of the Hash Encoder to the reconstruction, and Figure (c) shows that self-weighting is necessary to reconstruct high-resolution details.
Figure 5 presents the impact of example hyperparameters used for optimization. Both λ_{Enc} and ϵ affect blurriness and noise level in reconstruction.

Discussion and Conclusion

The proposed scan-specific bilevel optimized INR effectively reconstructs accelerated MRI scans by automatically optimizing hyperparameters. Results show improved image quality in all metrics compared to existing methods across anatomies, sampling patterns and field strengths, emphasizing the benefits of tailored hyperparameter tuning. With short reconstruction time and good generalizability, this work improves the practical applicability of scan-specific deep learning methods in clinical settings.

Acknowledgements

This study was supported by NIH grants R37CA263583, R01CA284172, and Siemens Healthineers.

References

1. F. Knoll, K. Hammernik, C. Zhang, S. Moeller, T. Pock, D. K. Sodickson, and M. Akcakaya, “Deep-Learning Methods for Parallel Magnetic Resonance Imaging Reconstruction: A Survey of the Current Approaches, Trends, and Issues,” IEEE Signal Processing Magazine, vol. 37, no. 1, pp. 128–140, 2020.
2. 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.
3. 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.
4. B. Yaman, S. A. H. Hosseini, and M. Akcakaya, “Zero-Shot Self-Supervised Learning for MRI Reconstruction,” in Proceedings of the International Conference on Learning Representations, 2022. https://openreview.net/forum?id=085y6YPaYjP
5. P. I. Frazier, “A Tutorial on Bayesian Optimization,” arXiv preprint arXiv:1807.02811, July 2018.
6. T. Müller, A. Evans, C. Schied, and A. Keller, “Instant Neural Graphics Primitives with a Multiresolution Hash Encoding,” ACM Transactions on Graphics, vol. 41, no. 4, pp. 102:1–102:15, July 2022.
7. F. Murtagh, “Multilayer perceptrons for classification and regression,” in Neurocomputing, vol. 2, no. 5, pp. 183–197, 1991.
8. M. J. Muckley, R. Stern, T. Murrell, and F. Knoll, “TorchKbNufft: A High-Level, Hardware-Agnostic Non-Uniform Fast Fourier Transform,” in Proceedings of the ISMRM Workshop on Data Sampling & Image Reconstruction, 2020.
9. B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis,” in Proceedings of the European Conference on Computer Vision (ECCV), vol. 12346, pp. 405–421, 2020.

Figures

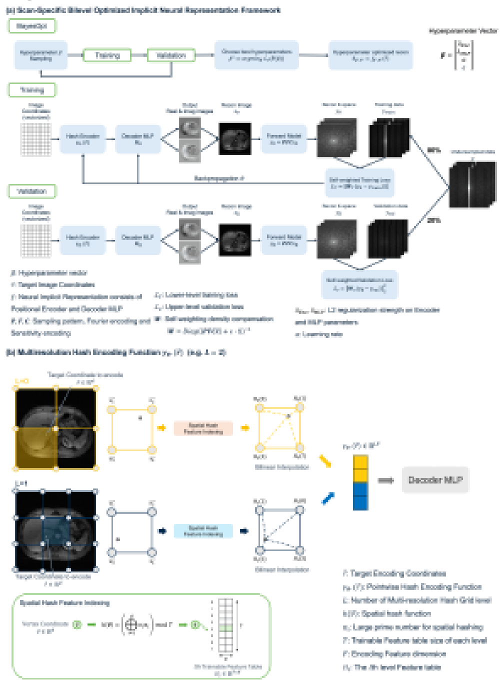

Figure 1.(a) The proposed scan-specific Bilevel Optimized Implicit Neural Representation framework with upper-level hyperparameter optimization stage (BayesOpt and Validation) and lower-level image reconstruction stage (Training). (b) Example illustration of a 2-leveled Trainable Multiresolution Hash Encoder; in practice 16 levels are used.

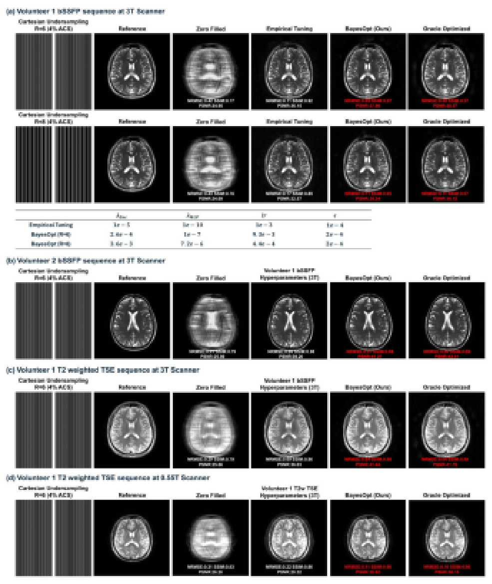

Figure 2. Demonstration of hyperparameter optimization using Bayesian Optimization. (a) Comparisons of bSSFP acquisition of volunteer 1 at 3T. (b) Comparisons of bSSFP acquisition of volunteer 2 at 3T, showing hyperparameters are transferable to a different subject for the same imaging sequence. (c) Comparison of T2w TSE acquisition of volunteer 1 at 3T, showing that different sequences need tailored hyperparameter optimization. (d) Comparisons of T2w TSE acquisition of volunteer 1 at 0.55T, showing different field strengths need tailored hyperparameter optimization.

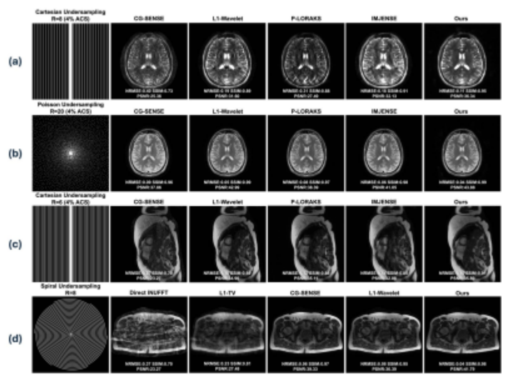

Figure 3. Method comparisons (a) Brain volunteer at 3T using bSSFP of 1x1 mm^2 resolution. (b) Brain volunteer at 3T using T2w TSE of 1.2x1.2 mm^2 resolution. (c) Cardiac volunteer at 0.55T using T2w bSSFP of 1.4x1.4 mm^2 resolution. (d) Prostate volunteer at 3T using T2w TSE of 1x1 mm^2 resolution. PSNR is in dB.

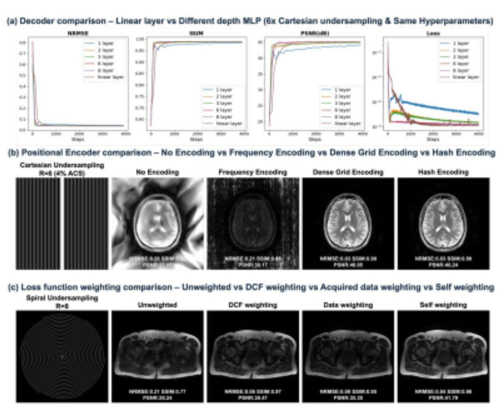

Figure 4. Ablation study. (a) Decoder comparison with fixed Hash Encoder (i.e., Linear layer vs 1/2/3/6/8 layer MLP), (b) Encoder comparison with fixed Decoder MLP (i.e., No Encoding vs Frequency Encoding[9] vs Dense Grid vs Hash Encoder), (c) Loss function weighting comparison (i.e., Unweighted vs DCF weighting vs Acquired data weighting vs Self weighting).

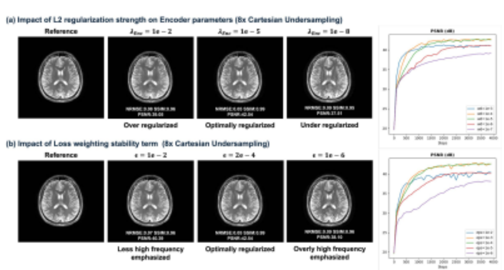

Figure 5. Illustrating choosing hyperparameters for Bayesian Optimization; (a) ℓ_2 regularization strength for Hash Encoder parameters, (b) Self-weighting loss stability value that controls the emphasis on higher frequency k-space components.

Proc. Intl. Soc. Mag. Reson. Med. 33 (2025)
0389