This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Zero-shot Low-field MRI Enhancement via Denoising Diffusion Driven Neural Representation

Xiyue Lin[1*], Chenhe Du[1*], Qing Wu[1], Xuanyu Tian[1], Jingyi Yu[1], Yuyao Zhang[1], and Hongjiang Wei[2] ✉

[1] School of Information Science and Technology, ShanghaiTech University, Shanghai, China
[2] School of Biomedical Engineering, Shanghai Jiao Tong University, Shanghai, China
hongjiang.wei@sjtu.edu.cn

Abstract. Recently, there have been significant advancements in the development of portable low-field (LF) magnetic resonance imaging (MRI) systems. These systems aim to provide low-cost, unshielded, and bedside diagnostic solutions. MRI experiences a diminished signal-to-noise ratio (SNR) at reduced field strengths, which results in severe signal deterioration and poor reconstruction. Therefore, reconstructing a high-field-equivalent image from a low-field MRI is a complex challenge due to the ill-posed nature of the task. In this paper, we introduce diffusion model driven neural representation. We decompose the low-field MRI enhancement problem into a data consistency subproblem and a prior subproblem and solve them in an iterative framework. The diffusion model provides high-quality high-field (HF) MR images prior, while the implicit neural representation ensures data consistency. Experimental results on simulated LF data and clinical LF data indicate that our proposed method is capable of achieving zero-shot LF MRI enhancement, showing some potential for clinical applications.

Keywords: Diffusion models · Low-field MRI · Implicit neural representation.

1 Introduction

Magnetic Resonance Imaging (MRI) is a leading modality in both clinical diagnosis and fundamental research. It is superior to other imaging modalities as the nonionizing, noninvasive, fundamentally 3D, quantitative, and multiparametric [7]. However, the widespread use of high-field (HF) MRI is impeded by the high costs associated with installation, maintenance, and operation. In recent years, there has been intensive development of low-cost MRI hardware at low-field (LF) strengths (<1T) for portable and shielding-free clinical imaging

* Equal contribution.

applications, particularly for brain imaging [18,19,22]. While cost-effective and operationally simple, LF MRI scanners suffer from limitations in spatial resolution, long scanning times, and increased image noise and artifacts due to the low MR signal proportional to the field strength B_0. Addressing these constraints necessitates exploring alternative approaches to enhance LF MRI imaging quality.

Deep learning (DL) is revolutionizing high-field (HF) MR image reconstruction, offering improvements in denoising, artifacts suppression, and reconstruction from undersampled k-space data [20,29]. The growing availability of extensive, high-quality HF MRI data (*e.g.*, from HCP consortium and UK Biobank [16, 26]) is anticipated to drive advancements in LF MRI through DL techniques. Recently, several DL attempts have been made to improve LF MR image quality [9,12,15,17]. For instance, Iglesias *et al.* [9] proposed a domain and resolution agnostic model by mapping any input to high-resolution T1-weighted MR images. Man *et al.* [17] trained a Partial Fourier super-resolution (PF-SR) network using synthetic Partial Fourier sampling (acquiring only a fraction of the k-space data) low-resolution noisy LF data from HF MR images. Nevertheless, these models encounter several challenges: (1) Acquiring paired real LF-HF data is arduous, and even with registration, the precision remains inadequate. (2) Synthetic LF MR images still exhibit domain-shift (out-of-distribution) issues with real LF MR images, leading to a high dependency of supervised learning methods on the data distribution of the training dataset. (3) The scale-specific synthetic degradation paradigm constrains the generalization capability of the model.

Recently, Generative Adversarial Networks (GANs) have gained popularity in low-field MRI enhancement task [1,10]. Despite their ability to generate realistic outputs, GAN-based methods face training challenges, such as mode collapse and instability. As an alternative, denoising diffusion models (DDM) have emerged as state-of-the-art generative models and effective generative priors for unsupervised inverse problem resolution. These models transform data distribution into Gaussian through a forward process of incremental noise perturbation. Trained by aligning the gradient of the log density in the reverse process, they can sample a prior from estimated posterior probabilities when presented with partial and corrupted measurements. The prior facilitates constraining the space towards a feasible solution for ill-posed inverse problems. There have been several methods based on diffusion models for medical imaging reconstruction that achieved notable progress [3, 24]. However, most diffusion-based LF MRI enhancement methods [5, 11] have focused on the conditional settings which are still limited by the difficulty of obtaining the paired LF-HF data and synthetic degradation paradigm. Developing the zero-shot method is crucial to overcome this issue.

Implicit neural representation (INR) is a novel unsupervised zero-shot DL framework that demonstrated a significant potential for the High-Resolution MRI reconstruction problem [27, 28]. Technically, INR trains a multi-layer perceptron (MLP) to represent the reconstructed MR image as a continuous function that maps the spatial coordinates to the image intensities. Due to the contin-

uous representation provided by INR, the single well-trained INR-based model can handle the SR tasks of arbitrary scales and significantly reduce resource consumption. Nonetheless, the continuous prior might hold less significance in scenarios that involve larger data domain-shift gaps without the supplement information of reconstruction images. Thus, existing INR-based methods hardly produce satisfactory performance in the related tasks.

In this work, we propose **DiffDeuR: Diff**usion **D**riven n**eu**ral **R**epresentation, an effective unsupervised framework to reconstruct a high-field-equivalent image from LF MRI, which is a complex challenge due to the ill-posed nature of the task. We decompose the LF MRI enhancement problem into a data consistency subproblem and a prior subproblem and solve them in an iterative framework. To evaluate the performance of our proposed DiffDeuR, we conduct validation experiments using a real clinical LF MR dataset for synthetic and real LF MRI enhancement. The experimental results show that DiffDeuR is capable of recovering implicit information and preserving data fidelity. Furthermore, it surpasses state-of-the-art (SOTA) methodologies in terms of quantitative metrics, qualitative evaluations and downstream task analyses.

2 Method

In this section, we present our modeling of the problem and how we efficiently combined diffusion models and INR and harnessed the strengths of both approaches to achieve LF MRI enhancement. Fig. 1 shows a general overview of our approach. Specifically, our method decouples the original problem into two sub-problems, which are solved using diffusion model and INR, respectively, and optimized iteratively during diffusion sampling. The detailed steps of the method can be found in Algorithm 1. Note that the diffusion model used in this paper follows the form of 'VE-SDE' [25], which is a kind of score-based diffusion model.

2.1 Problem Formulation

The process of degrading HF MR images to LF MR images can be represented by a mapping function, which is formulated as follows:

$$\mathbf{x}_{\mathrm{LF}} = \mathcal{F}(\mathbf{x}_{\mathrm{HF}}), \tag{1}$$

where $\mathbf{x}_{\mathrm{LF}}$ is the LF MR data, $\mathcal{F}$ is the degradation operation that maps HF MR images to low-field and $\mathbf{x}_{\mathrm{HF}}$ is the unknown HF image to be enhanced. Regrettably, the precise expression of $\mathcal{F}$ remains unresolved. Inspired by the insights of previous work [5, 14, 17], we approximate this degradation process as a superposition of downsampling, blurring, and adding noise. The approximate model can be mathematically represented as follows:

$$\mathbf{x}_{\mathrm{LF}} = \mathbf{PSF}(\boldsymbol{L}_k(\mathbf{x}_{\mathrm{HF}})) + \mathbf{n}, \tag{2}$$

where **PSF** is the Gaussian point spread function (PSF) blurring [13], $\boldsymbol{L}_k$ represents the linear downsampling operation with the factor k, $\mathbf{n}$ is the additive

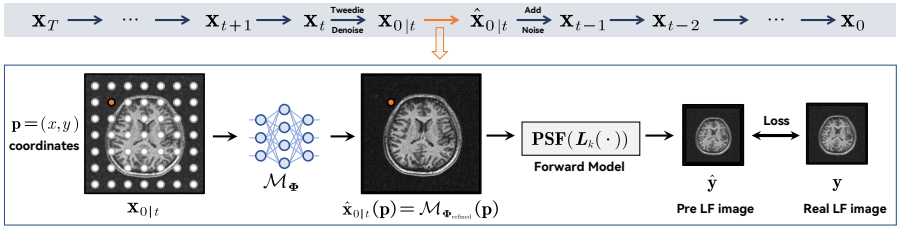

Fig. 1. Overview of the proposed DiffDeuR. The data-prior sub-problem is addressed with reverse diffusion sampling utilizing a pre-trained diffusion model. At the sampling timestep t, the method employs Tweedie denoising to generate a noise-free prior, denoted as $\mathbf{x}_{0|t}$ from $\mathbf{x}_t$. Subsequently, $\hat{\mathbf{x}}_{0|t}$ is represented via INR by solving the data-fidelity sub-problem (indicated by the orange arrow). Next, $\mathbf{x}_{t-1}$ is derived from $\hat{\mathbf{x}}_{0|t}$ by adding noise, enabling the continuation of the reverse diffusion sampling.

noise that follows a Gaussian distribution. The main objective is to recover $\mathbf{x}_{\mathrm{HF}}$ by solving the minimization problem:

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \frac{1}{2\sigma^2} \left\| \mathbf{y} - \mathbf{PSF}(\boldsymbol{L}_k(\mathbf{x})) \right\|^2 + \lambda \cdot \mathcal{R}(\mathbf{x}), \tag{3}$$

where the first data-fidelity term ensures consistency with measurement data $\mathbf{y}$, while the data-prior term $\mathcal{R}(\cdot)$ incorporates prior knowledge to constrain the solution space appropriately. A prevalent method for addressing this optimization problem involves decoupling the data fidelity and prior terms using Half Quadratic Splitting (HQS). This technique introduces an auxiliary variable, $\mathbf{z}$, allowing the problem to be tackled by iteratively solving two distinct sub-problems as follows:

$$\begin{cases} \mathbf{z}_t = \arg\min_{\mathbf{z}_t} \dfrac{1}{2(\sqrt{\lambda/\mu})^2} \left\| \mathbf{z}_t - \mathbf{x}_t \right\|^2 + \mathcal{R}(\mathbf{z}_t) & \text{(4a)} \\ \mathbf{x}_t = \arg\min_{\mathbf{x}_t} \left\| \mathbf{y} - \mathbf{PSF}(\boldsymbol{L}_k(\mathbf{x}_t)) \right\|^2 + \mu\sigma_i^2 \left\| \mathbf{x}_t - \mathbf{z}_{t-1} \right\|^2, & \text{(4b)} \end{cases}$$

where Eq. (4b) defines the data-fidelity sub-problem, concentrating on the identification of a proximal point in relation to $\mathbf{z}_{t-1}$. Meanwhile, as stated in Eq. (4a), from the perspective of Bayesian analysis, the data-prior sub-problem is essentially a denoising problem, which any denoiser can solve.

2.2 Data-Prior Sub-problem Solved by Diffusion Model

As mentioned above, we decompose the LF-MRI enhancement problem into two alternating optimization sub-problems. As the data-prior sub-problem is essentially a Gaussian denoising problem, we propose to use the diffusion model as the denoiser to solve it, and we can rewrite Eq. (4a) as:

$$\mathbf{x}_{0|t} = \mathbf{z}_t \approx \mathbf{x}_t + \sigma_t^2 \mathbf{s}_{\boldsymbol{\theta}^*}(\mathbf{x}_t, t), \tag{5}$$

Algorithm 1 DiffDeuR

Require: LF-MRI $\mathbf{y}$, Pre-trained diffusion model $\mathbf{s}_{\boldsymbol{\theta}^*}$, MLP network $\mathcal{F}_{\boldsymbol{\Phi}}$, Timesteps T, Noise schedule $\{\sigma_t\}_{t=0}^{T}$, Coordinates of images $\mathbf{p}$.

$\mathbf{x}_T \sim (\mathbf{0}, \sigma_T^2\mathbf{I})$;
for $t = T, \ldots, 1$ **do**
  ▷ *1. Data-prior sub-problem*
  $\mathbf{x}_{t-1} = \mathbf{x}_t + \left(\sigma_t^2 - \sigma_{t-1}^2\right)\mathbf{s}_{\boldsymbol{\theta}^*}\left(\mathbf{x}_t, t\right) + \sqrt{\left(\sigma_t^2 - \sigma_{t-1}^2\right)}\boldsymbol{\epsilon}$, $\boldsymbol{\epsilon} \sim \mathcal{N}\left(\mathbf{0}, \mathbf{I}\right)$
  $\mathbf{x}_{0|t} = \mathbf{x}_t + \sigma_t^2\mathbf{s}_{\boldsymbol{\theta}^*}\left(\mathbf{x}_t, t\right)$ ▷ *Tweedie's denoising*
  ▷ *2. Data-fidelity sub-problem*
  $\boldsymbol{\Phi}_{\text{refined}} = \arg\min_{\boldsymbol{\Phi}} \|\mathcal{M}_{\boldsymbol{\Phi}}(\mathbf{p}) - \mathbf{x}_{0|t}\|^2 + \lambda\frac{\sigma_T^2}{\sigma_t^2}\mathcal{L}_{\text{DF}}(\mathbf{A}(\mathcal{M}_{\boldsymbol{\Phi}}(\mathbf{p})), \mathbf{y})$
  $\hat{\mathbf{x}}_{0|t} = \mathcal{M}_{\boldsymbol{\Phi}_{\text{refined}}}(\mathbf{p})$
  $\mathbf{x}_{t-1} = \hat{\mathbf{x}}_{0|t} + \sigma_{t-1}(\sqrt{\xi}\boldsymbol{\epsilon} + \sqrt{1-\xi}(\mathbf{x}_t - \hat{\mathbf{x}}_{0|t}))$, $\boldsymbol{\epsilon} \sim \mathcal{N}\left(\mathbf{0}, \mathbf{I}\right)$ ▷ *Add noise*
end for
return $\mathbf{x}_0$

which means that the data-prior sub-problem has been transformed into a diffusion denoising process with Tweedie's formula. Then we can solve the data-fidelity sub-problem to get a refined $\hat{\mathbf{x}}_{0|t}$ corrected with measurement $\mathbf{y}$. However, such a solution with only a single iteration of Eq. (4) is inaccurate. Inspired by [23], we use an estimation-correction idea, which is mapping the inaccurate solution to timestep $t-1$ to continue the process of diffusion. Specifically, we follow the setting of [30] and add noise to the inaccurate $\hat{\mathbf{x}}_{0|t}$ with noise level corresponding to timestep $t-1$. The operation mentioned above can be succinctly formulated as follows:

$$\mathbf{x}_{t-1} = \hat{\mathbf{x}}_{0|t} + \sigma_{t-1}(\sqrt{\xi}\boldsymbol{\epsilon} + \sqrt{1-\xi}(\mathbf{x}_t - \hat{\mathbf{x}}_{0|t})) \tag{6}$$

2.3 Data-Fidelity Sub-problem Solved by INR

INR aims to learn a continuous function that maps spatial coordinates to image pixels using MLP. By incorporating a forward model in the MLP optimization process, INR can be used to solve a broad spectrum of imaging inverse problems. Drawing inspiration from [21], we propose a novel model-driven INR framework to solve the data-fidelity sub-problem during the diffusion sampling process. Specifically, considering Eq. (4b), our objective is to ensure that the INR framework satisfies two consistency: 1) prior consistency (PC): utilizing the data prior $\mathbf{x}_{0|t}$ provided by the diffusion model; 2) measurement consistency (MC): maintaining data consistency with the measurements $\mathbf{y}$ (*i.e.*, LF MRI). Technically, we represent the solution $\hat{\mathbf{x}}_{0|t}$ as a continuous function of image coordinate, which can be expressed as below:

$$f : \mathbf{p} = (x, y) \in \mathbb{R}^2 \rightarrow \hat{\mathbf{x}}_{0|t}(\mathbf{p}) \in \mathbb{R}, \tag{7}$$

where $\mathbf{p}$ is any spatial coordinate and $\hat{\mathbf{x}}_{0|t}(\mathbf{p})$ is the corresponding intensity of the solution image at that position. We leverage an MLP network $\mathcal{M}_{\boldsymbol{\Phi}}$ to

approximate the function f by optimizing the following objective function:

$$\boldsymbol{\Phi}_{\text{refined}} = \arg\min_{\boldsymbol{\Phi}} \underbrace{\| \hat{\mathbf{x}}_{0|t} - \mathbf{x}_{0|t} \|^2}_{\text{Prior Consistency}} + \lambda \frac{\sigma_T^2}{\sigma_t^2} \underbrace{\mathcal{L}_{\text{MC}}(\mathbf{A}(\hat{\mathbf{x}}_{0|t}), \mathbf{y})}_{\text{Measurement Consistency}}, \text{ s.t. } \hat{\mathbf{x}}_{0|t}(\mathbf{p}) = \mathcal{M}_{\boldsymbol{\Phi}}(\mathbf{p}) \tag{8}$$

where $\mathbf{A}(\cdot)$ is the forward model $\mathbf{PSF}(\boldsymbol{L}_k(\cdot))$. The PC term minimizes the pixel-wise L2 loss between the prior image $\mathbf{x}_{0|t}$ and the estimated image $\hat{\mathbf{x}}_{0|t}$, ensuring prior consistency. For measurement consistency, we perform the forward model on the estimated image and then minimize the distance between the estimated measurement $\mathbf{A}(\hat{\mathbf{x}}_{0|t})$ and real measurement $\mathbf{y}$. The MC term employed can be expressed as follows:

$$\mathcal{L}_{\text{MC}}(\mathbf{A}\hat{\mathbf{x}}_{0|t}, \mathbf{y}) = \|\mathbf{A}(\hat{\mathbf{x}}_{0|t}) - \mathbf{y}\|^2 + \mu \cdot \text{SSIM}(\mathbf{A}(\hat{\mathbf{x}}_{0|t}), \mathbf{y}) \tag{9}$$

Given the iterative nature of the proposed method, it is imperative to calibrate the contributions of the PC and MC at various timesteps. This calibration is to prevent the model from overfitting to LF measurements and corrupting the final results. The modulation of the contribution weights for PC and MC is achieved through a hyperparameter, denoted as $\lambda \frac{\sigma_T^2}{\sigma_t^2}$. This parameter adjusts the weights according to the noise level corresponding to the timestep in the diffusion sampling process, which has been proved and used in [30].

3 Experiments

3.1 Setup

Dataset. The LF MRI was collected using a 0.2T MR scanner to obtain T1-weighted scans. Two sets of data were collected using different parameters at the time of collection. One with a resolution of 3×3 mm^2 and a thickness of 3 mm, and the other with a resolution of 1×1 mm^2 and a layer thickness of 6 mm. The HF reference MRI was acquired by a 3T uMR 790 MRI scanner of United Imaging, using a T1W, Fast SPoiled GRadient Echo (3D T1 FSP GRE) sequence with resolution $1 \times 1 \times 1$ mm^3 isotropic. All the detailed acquisition parameters can be found in the Supplementary Material. We conducted simulated and real data experiments to evaluate our method comprehensively. For the simulated experiment, we used the degradation model in Eq. (2) on HF reference MRI to generate the simulated LF MRI. For the real experiment, we directly used the acquired LF MRI.

Baseline & Metrics. Four representative image enhancement approaches from different categories are utilized as comparative methods: 1) traditional step-by-step method: denoising with BM3D [6] and super-resolution with Bicubic; 2) GAN-based method: LoHiResGAN [10], which is a GAN-based image-to-image translation method dedicated to LF MRI enhancement; 3) INR-based methods: IREM [28], which is an unsupervised MRI SR method; and 4) diffusion-based

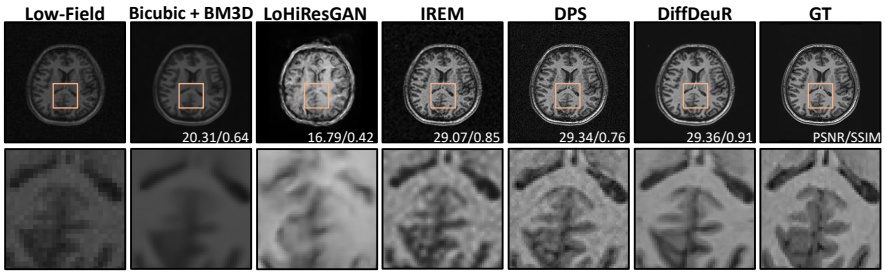

Fig. 2. Qualitative results of methods in comparison on one test sample of synthetic LF MRI data.

method: DPS [2]. To ensure zero-shot scenarios, for LoHiResGAN, we use the officially provided pre-trained model and do not perform any fine-tuning. DPS and ours DiffDeuR using the same score function from [4], which is pre-trained only on 3T MRI.

We utilize two standard metrics to quantitatively assess the performance of the compared methods: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM) and Learned Perceptual Image Patch Similarity (LPIPS).

3.2 Results

Table 1. Performance of simulated LF data with all compared method results on the test set. The best and second performances are highlighted in red and blue, respectively.

Method	PSNR	SSIM	LPIPS
Bicubic+BM3D [6]	20.708±0.606	0.6747±0.008	0.2665±0.024
LoHiResGAN [10]	18.015±0.707	0.3909±0.013	0.2718±0.015
IREM [28]	27.019±0.255	0.8236±0.004	0.1839±0.010
DPS [2]	27.717±0.359	0.8198±0.007	0.1632±0.012
DiffDeuR (Ours)	28.130±0.264	0.9146±0.004	0.0905±0.007

Validation Study on the Simulated Low Field Experiment Dataset. In Table 1, we show the quantitative results of all methods on the test set of the simulation low-field data. To sum up, our DiffDeuR achieves the best performance across all experimental conditions. Fig. 2 demonstrates the qualitative evaluation results of all the compared methods. Our DiffDeuR method outperforms others in preserving image details and maintaining consistency with the original textures. In contrast, LoHiResGAN produces unsatisfactory results with

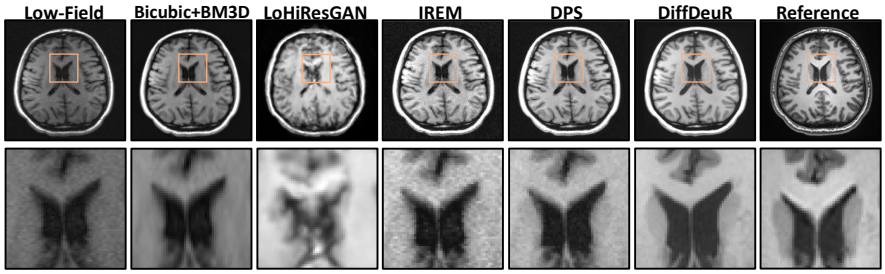

Fig. 3. Qualitative results of methods in comparison on one test sample of the real clinical LF MRI data with a resolution of 1×1 mm^2.

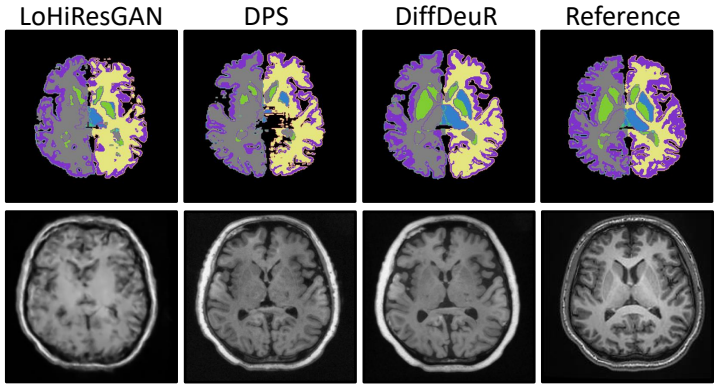

Fig. 4. Qualitative results of the whole brain segmentation of the compared method on one test sample of the real clinical LF MRI data with a resolution of 3×3 mm^2.

distorted images. While IREM and DPS methods show some consistency with the original images, artifacts remain. Bicubic with BM3D only provides interpolation without addressing the low-field to high-field domain shift.

Performance on the Real Experiment Dataset. Fig. 3 illustrates the qualitative evaluation results of all the compared methods on the representative test sample with a resolution of 1×1 mm^2. Evidently, our DiffDeuR method exhibits a remarkable capacity to recover the discernible content of brain structures. Upon the zoomed-in regions, DiffDeuR and DPS can recover particular implicit information by leveraging the advantages of the diffusion prior. In contrast, LoHiResGAN is significantly constrained by data distribution, leading to distorted reconstructions. Also, Fig. 4 demonstrates the whole brain segmentation results of compared methods on a sample with a resolution of 3×3 mm^2 using FastSurfer [8], which is an efficient resolution-independent segmentation

method. Our DiffDeuR reconstruction outperforms DPS and approaches the HF reference segmentation results.

4 Conclusion

This paper proposes DiffDeuR, an unsupervised DL model for low-field MRI enhancement tasks. The proposed DiffDeuR follows the HQS framework to decompose the low-field MRI enhancement problem into a data fidelity sub-problem and a distribution prior sub-problem. Then, the two sub-problems are respectively addressed by INR and a pre-trained diffusion model. The comprehensive evaluation, spanning qualitative, quantitative, and downstream task analyses on synthetic and real datasets, validates the superior performance of our DiffDeuR model compared with SOTA methods in low-field MRI enhancement. This reinforces the effectiveness of our unsupervised DL approach in this domain.

Acknowledgments. This study is supported by the National Natural Science Foundation of China (No. 62071299).

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

References

1. Chen, Y., Shi, F., Christodoulou, A.G., Xie, Y., Zhou, Z., Li, D.: Efficient and accurate mri super-resolution using a generative adversarial network and 3d multi-level densely connected network. In: International conference on medical image computing and computer-assisted intervention. pp. 91–99. Springer (2018)
2. Chung, H., Kim, J., Mccann, M.T., Klasky, M.L., Ye, J.C.: Diffusion posterior sampling for general noisy inverse problems. In: The Eleventh International Conference on Learning Representations (2022)
3. Chung, H., Ryu, D., McCann, M.T., Klasky, M.L., Ye, J.C.: Solving 3d inverse problems using pre-trained 2d diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22542–22551 (2023)
4. Chung, H., Sim, B., Ryu, D., Ye, J.C.: Improving diffusion models for inverse problems using manifold constraints. Advances in Neural Information Processing Systems **35**, 25683–25696 (2022)
5. Cui, Z.X., Liu, C., Cao, C., Liu, Y., Cheng, J., Zhu, Q., Zhu, Y., Wang, H., Liang, D.: Meta-learning enabled score-based generative model for 1.5 t-like image reconstruction from 0.5 t mri. arXiv preprint arXiv:2305.02509 (2023)
6. Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K.: Image denoising by sparse 3-d transform-domain collaborative filtering. IEEE Transactions on image processing **16**(8), 2080–2095 (2007)
7. Edelman, R.R., Warach, S.: Magnetic resonance imaging (2). N Engl J Med **328**(11), 785–791 (Mar 1993)
8. Henschel, L., Conjeti, S., Estrada, S., Diers, K., Fischl, B., Reuter, M.: Fastsurfer-a fast and accurate deep learning based neuroimaging pipeline. NeuroImage **219**, 117012 (2020)

9. Iglesias, J.E., Billot, B., Balbastre, Y., Magdamo, C., Arnold, S.E., Das, S., Edlow, B.L., Alexander, D.C., Golland, P., Fischl, B.: Synthsr: A public ai tool to turn heterogeneous clinical brain scans into high-resolution t1-weighted images for 3d morphometry. Science advances **9**(5), eadd3607 (2023)
10. Islam, K.T., Zhong, S., Zakavi, P., Chen, Z., Kavnoudias, H., Farquharson, S., Durbridge, G., Barth, M., McMahon, K.L., Parizel, P.M., et al.: Improving portable low-field mri image quality through image-to-image translation using paired low- and high-field images. Scientific Reports **13**(1), 21183 (2023)
11. Kim, S., Tregidgo, H.F., Eldaly, A.K., Figini, M., Alexander, D.C.: A 3d conditional diffusion model for image quality transfer–an application to low-field mri. arXiv preprint arXiv:2311.06631 (2023)
12. Koonjoo, N., Zhu, B., Bagnall, G.C., Bhutto, D., Rosen, M.: Boosting the signal-to-noise of low-field mri with deep learning image reconstruction. Scientific reports **11**(1), 8248 (2021)
13. Liang, Z.P., Lauterbur, P.C.: Principles of magnetic resonance imaging. SPIE Optical Engineering Press Belllingham, WA (2000)
14. Lin, H., Figini, M., D'Arco, F., Ogbole, G., Tanno, R., Blumberg, S.B., Ronan, L., Brown, B.J., Carmichael, D.W., Lagunju, I., et al.: Low-field magnetic resonance image enhancement via stochastic image quality transfer. Medical Image Analysis **87**, 102807 (2023)
15. Lin, H., Figini, M., Tanno, R., Blumberg, S.B., Kaden, E., Ogbole, G., Brown, B.J., D'Arco, F., Carmichael, D.W., Lagunju, I., et al.: Deep learning for low-field to high-field mr: image quality transfer with probabilistic decimation simulator. In: Machine Learning for Medical Image Reconstruction: Second International Workshop, MLMIR 2019, Held in Conjunction with MICCAI 2019, Shenzhen, China, October 17, 2019, Proceedings 2. pp. 58–70. Springer (2019)
16. Littlejohns, T.J., Holliday, J., Gibson, L.M., Garratt, S., Oesingmann, N., Alfaro-Almagro, F., Bell, J.D., Boultwood, C., Collins, R., Conroy, M.C., et al.: The uk biobank imaging enhancement of 100,000 participants: rationale, data collection, management and future directions. Nature communications **11**(1), 2624 (2020)
17. Man, C., Lau, V., Su, S., Zhao, Y., Xiao, L., Ding, Y., Leung, G.K., Leong, A.T., Wu, E.X.: Deep learning enabled fast 3d brain mri at 0.055 tesla. Science Advances **9**(38), eadi9327 (2023)
18. Mazurek, M.H., Cahn, B.A., Yuen, M.M., Prabhat, A.M., Chavva, I.R., Shah, J.T., Crawford, A.L., Welch, E.B., Rothberg, J., Sacolick, L., et al.: Portable, bedside, low-field magnetic resonance imaging for evaluation of intracerebral hemorrhage. Nature communications **12**(1), 5119 (2021)
19. O'Reilly, T., Teeuwisse, W.M., de Gans, D., Koolstra, K., Webb, A.G.: In vivo 3d brain and extremity mri at 50 mt using a permanent magnet halbach array. Magnetic resonance in medicine **85**(1), 495–505 (2021)
20. Qu, L., Zhang, Y., Wang, S., Yap, P.T., Shen, D.: Synthesized 7t mri from 3t mri via deep learning in spatial and wavelet domains. Medical image analysis **62**, 101663 (2020)
21. 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 (2022)
22. Sheth, K.N., Mazurek, M.H., Yuen, M.M., Cahn, B.A., Shah, J.T., Ward, A., Kim, J.A., Gilmore, E.J., Falcone, G.J., Petersen, N., et al.: Assessment of brain injury using portable, low-field magnetic resonance imaging at the bedside of critically ill patients. JAMA neurology **78**(1), 41–47 (2021)

23. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: International Conference on Learning Representations (2020)
24. Song, Y., Shen, L., Xing, L., Ermon, S.: Solving inverse problems in medical imaging with score-based generative models. In: International Conference on Learning Representations (2021)
25. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: International Conference on Learning Representations (2020)
26. Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., Consortium, W.M.H., et al.: The wu-minn human connectome project: an overview. Neuroimage **80**, 62–79 (2013)
27. Wu, Q., Li, Y., Sun, Y., Zhou, Y., Wei, H., Yu, J., Zhang, Y.: An arbitrary scale super-resolution approach for 3d mr images via implicit neural representation. IEEE Journal of Biomedical and Health Informatics **27**(2), 1004–1015 (2022)
28. Wu, Q., Li, Y., Xu, L., Feng, R., Wei, H., Yang, Q., Yu, B., Liu, X., Yu, J., Zhang, Y.: Irem: High-resolution magnetic resonance image reconstruction via implicit neural representation. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VI 24. pp. 65–74. Springer (2021)
29. Zhang, H., Zhang, Y., Wu, Q., Wu, J., Zhen, Z., Shi, F., Yuan, J., Wei, H., Liu, C., Zhang, Y.: Self-supervised arbitrary scale super-resolution framework for anisotropic mri. arXiv preprint arXiv:2305.01360 (2023)
30. Zhu, Y., Zhang, K., Liang, J., Cao, J., Wen, B., Timofte, R., Van Gool, L.: Denoising diffusion models for plug-and-play image restoration. arXiv preprint arXiv:2305.08995 (2023)