This WACV paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

MRI Reconstruction with Regularized 3D Diffusion Model (R3DM)

Arya Bangun^{*1} Zhuo Cao^{*1} Alessio Quercia^{1,3} Hanno Scharr¹ Elisabeth Pfaehler^{1,2} ¹IAS-8, ²INM-4 Forschungszentrum Jülich, Germany {a.bangun, z.cao, a.quercia, h.scharr, e.pfaehler}@fz-juelich.de

Abstract

Magnetic Resonance Imaging (MRI) is a powerful imaging technique widely used for visualizing structures within the human body and in other fields such as plant sciences. However, there is a demand to develop fast 3D-MRI reconstruction algorithms to show the fine structure of objects from under-sampled acquisition data, i.e., k-space data. This emphasizes the need for efficient solutions that can handle limited input while maintaining high-quality imaging. In contrast to previous methods only using 2D, we propose a 3D MRI reconstruction method that leverages a regularized 3D diffusion model combined with optimization method. By incorporating diffusion-based priors, our method improves image quality, reduces noise, and enhances the overall fidelity of 3D MRI reconstructions. We conduct comprehensive experiments analysis on clinical and plant science MRI datasets. To evaluate the algorithm effectiveness for under-sampled k-space data, we also demonstrate its reconstruction performance with several undersampling patterns, as well as with in- and out-ofdistribution pre-trained data. In experiments, we show that our method improves upon tested competitors.

1. Introduction

In order to speed up the acquisition time, MRI instruments acquire sub-sampled k-space data, a technique where only a fraction of the total k-space data points are sampled during the imaging process. Several attempts have been proposed to develop two-dimensional (2D) and threedimensional (3D) image reconstruction techniques for subsampled k-space, as discussed in [11, 13, 31]. Advancements in 3D MR imaging methods can address the challenges posed by complex anatomical structures of human organs and plant growths. Consequently, the demand for developing 3D MR image reconstruction methods has intensified.

Currently, most works reconstruct a 3D volumetric image by stacking 2D reconstructions because MR images are

* Equal contribution Code: https://jugit.fz-juelich.de/ias-8/r3dm/ acquired slice by slice. This method doesn't consider the inter-dependency between the slices, thus can lead to inconsistencies and artifacts, as discussed in [4, 8, 50]. This particularly affects datasets that have equally distributed information and structures with high continuity on all dimensions, such as roots and vessels [4, 38, 50].

Before the deep learning-based models, which learn the data-driven prior, the model-based iterative reconstruction method proved its effectiveness in the 3D MRI reconstruction problem [15,54]. The problem is formulated as an optimization problem where a data consistency term is applied to ensure fidelity, and a regularisation term, such as the Total Variation (TV) penalty [24] is utilized to provide general prior knowledge of MRI data. Recent developments in generative models, especially the diffusion model, show significant improvement in reconstructing 2D and 3D MR images. This pivots to the idea of combining both data-driven and model-based reconstruction techniques, as presented in [22, 36, 40, 48].

One recent works applying the hybrid approach is DiffusionMBIR [8]. The proposed method can be seen as a two-stage process. First, images are generated from a pretrained 2D diffusion model as in [11]. Then, optimization methods and regularization techniques are applied on the z-axis to produce consistency in the sequence of 3D images. In this approach, the consistency of volumetric images highly depends on regularization. One possibility to directly improve the reconstruction of 3D data is leveraging 3D representations of the data-driven prior. As discussed in [49], 3D representations capture volumetric information more accurately and describe the details of the observed object. This enhancement contributes to a more comprehensive understanding of the underlying data, making the model outperforming its 2D counterpart.

In this article, we propose a reconstruction method that takes into account the 3D nature of MR images. At the same time, it utilizes a hybrid approach that combines the general prior given by the model-based reconstruction techniques with the data-driven prior from the generative model. The workflow can be seen in Figure 1. and our contributions are summarized below:



Figure 1. Workflow of the proposed algorithm generating volumetric data from a random distribution and guiding the reconstruction for specific k-space data to have a unique reconstruction. The reverse diffusion process is performed using a pre-trained diffusion model.

- We propose a Regularized 3D Diffusion Model (R3DM), an optimization algorithm to improve 3D MRI diffusion-based reconstruction.
- We demonstrate the effectiveness and robustness of our method via extensive experiments that compare its reconstruction quality to that of state-of-the-art methods, with in-distribution and out-of-distribution data.
- We study the need of 3D representation and that of the model-based method for 3D MRI reconstruction. We demonstrate in ablation studies that both are important for good performance.

2. Related work

In this section, we discuss the state of the art of MRI reconstruction from the classical to generative model-based approaches.

Model-based methods Conventional methods implemented in current MR scanners for accelerating data acquisition are e.g. parallel imaging and compressed sensing [26, 28, 34]. In parallel imaging, multiple receiver coils are used in the acquisition process, where each coil receives an under-sampled amount of k-space data. By incorporating estimated sensitivity maps for each coil, the information of these coils can then be merged to reconstruct one final image. For this purpose, algorithms such as SENSE and GRAPPA are used [18, 35], where the process to fill the missing information either in image space or k-space is performed. While parallel imaging can accelerate the acquisition time by a factor of 2 or 3, the use of compressed sensing can accelerate the acquisition time by an even higher factor of about 4-8. Compressed sensing uses prior knowledge about the MR signal to iteratively minimize data fidelity and a regularization term and reconstruct the under-sampled kspace data to a high-quality image. With this general prior knowledge, these approaches can be applied to any MR image data. However, to further improve their performance, advanced prior knowledge is needed.

Data-driven methods As the aforementioned methods have the limitations of *e.g.* defining the most suitable prior, the use of Deep Learning methods for MR image reconstruction becomes more and more popular [1, 12, 43, 55]. Hereby, traditional Convolutional Neural Networks (CNNs) with different architectures have demonstrated promising results such as e.g. the Automap framework [55] or other deep networks especially designed for MR image reconstruction [14]. Additionally, the effort to accelerate acquisition time by using deep learning has been discussed in [52], where the challenge is to incorporate a machine learningbased approach to reduce measurement time in the MRI system. Several methods such as fastMRI U-Net, and MoDL [1,42,52] outperform conventional methods for the MRI reconstruction for undersampling measurements. While these models with learned prior knowledge perform better than the conventional methods, they only work for in-distribution MR images on which the model is trained. Given the high data requirements of deep learning-based methods, training such models requires an effort to collect high-quality data, which is not always possible.

The development of deep learning-based MR image reconstruction is becoming prevalent due to the development of generative models, such as Generative Adversarial Networks (GAN), Normalizing Flows, Denoising Score Matching, and Diffusion Models [11, 29, 40, 48, 51, 53], where it is demonstrated that those approaches can generate highquality images from under-sampled k-space data. In those approaches, the model learns the data distribution of the specific MR image, e.g., knee, brain. In the inference process, joint reconstructions are proposed by incorporating under-sampled k-space data for a unique reconstruction. Moreover, the Diffusion Model is robust to out-of-distribution data on which it has not been trained [11, 22], which makes it more data efficient than other deep learning-based models.

In this article, we propose a new hybrid method that combines model-based and data-driven methods.

3. Methods and Algorithms

We first briefly review the MRI reconstruction problem and diffusion model in sections 3.1 and 3.2, respectively. In 3.3 we present the proposed method. The notations used throughout this article are summarized below.

Notations Vectors $\mathbf{x} \in \mathbb{C}^{L}$ and matrices are written in bold small-cap and bold big-cap letter $\mathbf{A} \in \mathbb{C}^{K \times L}$, respectively. Volumetric data is written as bold italic letter $\mathbf{A} \in \mathbb{C}^{S \times K \times L}$ or indexing matrix $\mathbf{A}_{s} \in \mathbb{C}^{K \times L}$ for $s \in [S]$. Here the \mathbb{C} and \mathbb{R} are complex and real fields. The set of integers is written as $[N] := \{1, 2, \dots, N\}$. We denote the two-dimensional and one-dimensional Fourier transform by \mathcal{F}_{2D} and \mathcal{F}_{1D} , respectively. For both matrices and vectors, the notation \circ is used to represent element-wise or Hadamard product. For a matrix $\mathbf{X} \in \mathbb{C}^{K \times L}$, the Frobenius norm is denoted by $\|\mathbf{X}\|_{F} := \sqrt{\sum_{k=1}^{K} \sum_{\ell=1}^{L} |x_{k\ell}|^{2}}$.

3.1. 3D MRI Reconstruction

We first present the conventional 2D reconstruction for full sample k-space data and develop the model for undersampled measurement. Afterward, we will discuss the challenge of 3D MRI reconstruction.

Full and Under-sampled Reconstruction. The relation between fully sampled k-space and the image space, i.e., image of interest, in MRI is given by the Fourier transform [25,27,39]. Suppose we acquire S single-coil slices of a fully sampled k-space image $\mathbf{Y}_s \in \mathbb{C}^{N_y \times N_x}$ for $s \in [S]$. Hence, the direct reconstruction in image space can be written as

$$\mathbf{X}_s = \mathcal{F}_{2D}^{-1}\left(\mathbf{Y}_s\right) \quad \text{for} \quad s \in [S]. \tag{1}$$

Accelerating MRI acquisition is achieved by reducing the number of scan lines during k-space acquisition. As a result, under-sampled k-space is recorded with fewer data points. We can write this as

$$\hat{\mathbf{X}}_{s} = \mathcal{F}_{2D}^{-1} \left(\mathbf{M} \circ \mathbf{Y}_{s} \right) \quad \text{for} \quad s \in [S], \tag{2}$$

where the matrix $\hat{\mathbf{Y}}_s = \mathbf{M} \circ \mathbf{Y}_s$ and $\hat{\mathbf{X}}_s$ are the noisy under-sampled k-space and the reconstruction, called zerofilled image for all slices, respectively. Applying a twodimensional inverse Fourier transform directly on undersampled k-space yields low-quality reconstructed images as it results in aliasing and ambiguities.

Challenges in 3D reconstruction. It should be straightforward to construct the 3D image by stacking slice-by-slice reconstruction. However, as discussed in [8], the undersampled k-space could yield inconsistencies in reconstruction for the z- axis direction. Additionally, reconstruction for the data that have sparse structures in all directions, such as roots and angiography, stacking slice-by-slice reconstruction can also yield non-smoothness in all possible directions. One way to address this problem is by adding constraints in the reconstruction algorithm such that the continuity on z - y, x - z, as well as x - y can be pertained. Here, we will use the knowledge of projection of 3D data on the image space and its relation to the k-space by using the relation between 3D k-space and projection of 3D image space, as described in Appendix D, combined with 3D diffusion model. The visualization can be seen in Figure 2

3.2. Diffusion Model

The key idea in diffusion models is to learn the training data distribution, i.e., image, such that we can generate new images corresponding to the learned distribution.

Training. In Denoising Diffusion Probabilistic Models (DDPM) [19, 20], estimating data distribution can be performed by perturbing the training data with a sequence of positive noise scale $0 < \beta_1, \beta_2, \ldots, \beta_T < 1$ as follows

$$\mathbf{x}_t = \sqrt{1 - \beta_t} \mathbf{x}_{t-1} + \sqrt{\beta_t} \mathbf{z}_{t-1}, \text{ for } t \in \{1, 2, \dots, T\}$$

where $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. In this case, the original data $\mathbf{x}_0 \sim p_{\text{data}}(\mathbf{x})$ is perturbed by a random Gaussian vector with increased variance for each time index. Hence, for each time step of perturbation, the neural network model s_{θ} , for instance, the U-Net architecture, is trained to approximate the conditional distribution from perturbed data.

Sampling. The pre-trained neural network model δ_{θ^*} is used later to remove the perturbation aspect in order to generate an image similar to the training data distribution. The sampling³ procedure is usually performed, for instance with the Euler-Maruyama scheme, as procedure, as follows:

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{1-\beta_t}} \left(\mathbf{x}_t + \beta_t s_{\boldsymbol{\theta}^*} \left(\mathbf{x}_t, t \right) \right) + \sqrt{\beta_t} \mathbf{z}_t$$

for $t \in \{T, T-1, \dots, 1\},$ (3)

³In this article the term sampling should not be confused with undersampling in MRI definition. Here sampling in the generative model describes generating data from a certain probability distribution.

It should be noted that, the index is reversed because we perform reverse process in the sampling mechanism. For the complete derivation, we refer the interested reader to the literature and references therein [19, 20, 41].

Diffusion Model for MRI Reconstruction Sampling from a pre-trained diffusion model generates MR images randomly. Guidance is needed to reconstruct a high-fidelity image conditioned on the measurement. Previous works [9,10] proposed to solve the reconstruction, or inverse problem in general, using conditional diffusion models. The idea is to alternatively update the sampling and optimization steps so that the transformation of the generated images asymptotically approach to the measurements. We apply the same procedure in our proposed method.

3.3. Proposed Method

In this article, we propose an optimization method for 3D reconstruction by leveraging information about the 2Dprojected image and its relation to the measurement data (k-space). The relation is from the nature of Fourier transform on 3D image space [6,7]. The algorithm allows us to reconstruct 3D MRI data with enhanced continuity on the z-axis using a pre-trained Diffusion Model, i.e. the algorithm is only applied during the sampling process.

3D Representation of data-driven prior As mentioned in Section 1, a 3D representation that captures the volumetric information is needed for an accurate 3D reconstruction. However, training and sampling a 3D Diffusion Model is in general computationally expensive. As highlighted in [46], the 2D U-Net with cross-attention on the 3rd dimension is more efficient than the 3D U-Net version for video data generation. As MR images are taken sequentially, the 3rd dimension of MR images can be regarded as temporal dimension. We called this architecture the 2D+A architecture hereafter. In order to directly reconstruct a 3D volumetric image from a given sub-sampled k-space, the 2D+A diffusion model should be combined with an algorithm that directly processes 3D volumetric data. In Section 5, we compare the performance with and without cross attention to demonstrate its importance.

Optimization problem. To incorporate the measurement process in the image generation in MRI reconstruction, a certain relation between k-space data and the target image space is developed. In this section, we discuss the connection of 3D k-space data and 3D image space.

Suppose we have under-sampled 3D k-space data given as $\hat{\mathbf{Y}}_s \in \mathbb{C}^{N_y \times N_x}$ for all depth slices $s \in [S]$. The zero frequency of k-space at x-axis⁴, i.e., $k_x = 0$, can be written



Figure 2. Relation between the projected image and a slice of kspace from plant roots data

as $\hat{\mathbf{y}}_{s}^{k_{y}} = \hat{\mathbf{Y}}_{s}|_{(k_{x}=0)} \in \mathbb{C}^{N_{y}}$, here we use superscript k_{y} to inform that this vector is obtained for $k_{x} = 0$. Additionally, the projection of the 3D image space on the y- axis is given by summing in x-direction, i.e. $x_{s,i} = \mathcal{P}_{y}(\mathbf{X}_{s})|_{i} =$ $\sum_{j=1}^{N_{x}} x_{s,i,j} \in \mathbb{C}$. The relation between the under-sampled k-space slice and the projection image is given by

$$\hat{\mathbf{y}}_{s}^{k_{y}} = \mathbf{m}^{k_{y}} \circ \mathcal{F}_{1D}(\mathbf{x}_{s}) \in \mathbb{C}^{N_{y}} \text{ for } s \in [S], \quad (4)$$

where \mathcal{F}_{1D} is the one-dimensional Fourier transform and $\mathbf{m}^{k_y} = \mathbf{M}|_{(k_x=0)} \in \mathbb{C}^{N_y}$ is the slice of the under-sampling operator at $k_x = 0$. The same concept can be adopted to obtain the projected image space on the *x*-axis from the projecting function \mathcal{P}_x . The complete derivation is presented in Appendix D. Figure 2 visualizes the relationship of projected volume data and image space and its relation to the Fourier transform. From this information, we can form an optimization problem to reconstruct three-dimensional images from under-sampled k-space, as follows:

$$X^{*} = \underset{X \in \mathbb{C}^{S \times N \times N}}{\operatorname{arg\,min}} \sum_{s=1}^{S} \left\| \hat{\mathbf{Y}}_{s} - \mathbf{M} \circ (\mathcal{F}_{2D} (\mathbf{X}_{s})) \right\|_{F}^{2} + \mathcal{R} (\mathbf{X})$$

subject to
$$\hat{\mathbf{y}}_{s}^{k_{y}} = \mathbf{m}^{k_{y}} \circ \mathcal{F}_{1D} (\mathcal{P}_{y} (\mathbf{X}_{s})) \quad \text{for } s \in [S]$$
$$\underbrace{\hat{\mathbf{y}}_{s}^{k_{x}} = \mathbf{m}^{k_{x}} \circ \mathcal{F}_{1D} (\mathcal{P}_{x} (\mathbf{X}_{s}))}_{\mathcal{G}(\mathbf{X})}$$
(5)

The regularization functions $\mathcal{R}(X)$ used in this article can be written as

$$\alpha \sum_{s=1}^{S} \sum_{i=1}^{N} \sum_{j=1}^{N} |x_{s,i,j}| + \sum_{s=1}^{S} \sum_{i=1}^{N} \sum_{j=1}^{N-1} |x_{s,i,j+1} - x_{s,i,j}|^{2}$$

$$+ \sum_{s=1}^{S} \sum_{i=1}^{N-1} \sum_{j=1}^{N} |x_{s,i+1,j} - x_{s,i,j}|^{2}$$

$$+ \underbrace{\sum_{s=1}^{S} \sum_{i=1}^{N-1} \sum_{j=1}^{N} |x_{s,i+1,j} - x_{s,i,j}|^{2}}_{\text{smoothness}}$$
(6)

and y- axis is given in the middle of row or column space

⁴Since the matrix represents the image per slice, the zero frequency on x-

In general, the optimization problem $\mathcal{G}(X)$ as in (5) describes the objective and the constraint functions to estimate the 3D image space given the relation between its undersampled 3D k-space and the projection image. This ensures the reconstructed image corresponds to the given measurement.

Algorithm	1 Regularized 3D Diffusion Model (R3DM)
·		/

1: Initialization:

- Timesteps T and optimization iterations m
- Volume sampling from prior distribution $X_T \in$ $\mathbb{R}^{S \times N \times N} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
- 2: for each reverse iteration t = T to 1 do
- Generate $\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if t > 1 else $\mathbf{Z} = \mathbf{0}$ 3:
- 4: Update sample (3)

$$\boldsymbol{X}^{(0)} = \frac{1}{\sqrt{1 - \beta_t}} \left(\boldsymbol{X}_t + \beta_t \boldsymbol{\vartheta}_{\boldsymbol{\theta}^*} \left(\boldsymbol{X}_t, t \right) \right) + \sqrt{\beta_t} \boldsymbol{Z}_t$$

- 5: for each iteration i = 1 to m do
 - Solving the optimization problem in (5)

$$\mathbf{X}^{(i)} = \mathcal{G}\left(\mathbf{X}^{(i-1)}\right)$$

- end for 7: $\boldsymbol{X}_{t-1} = \boldsymbol{X}^{(m)}$
- 8:

9: end for

6:

10: return $X_0 \in \mathbb{C}^{S \times N \times N}$

Regularization In addition, we use regularization for each element of volumetric data in terms of ℓ_1 -norm that enforces sparsity structure, left side constraint in (6), and the approximated total variation norm that enforces smoothness, right side constraint in (6). These terms are used for imposing the result with certain structures, i.e., sparsity or smoothness [32, 37]. The importance of the regularization terms is discussed in Section 5. In practice, the implementation can be performed in terms of alternating minimization combined with proximal methods for the regularization function. Note that the ℓ_1 -norm constraint is non-smooth, the proximal method is used to avoid calculating its derivative. The analytic proximal of ℓ_1 -norm is the well known soft thresholding function [5,16,32], which is also presented in theAppendix E. Hence, the algorithm can be seen as applying a proximal function on the gradient update of the loss function. This is done by applying all constraints in the optimization formula (5) as a single loss equation. This process is also called proximal gradient method [5,32]. The implementation detail will be presented in Appendix E.

Regularized 3D Diffusion Model (R3DM) The above optimization steps are inserted in the sampling process of a pre-trained Diffusion Model. Therefore, the proposed Regularized 3D Diffusion Model (R3DM), consists of two parts. The Diffusion Model part provides the data-related prior knowledge. The optimization method incorporates the k-space measurement to ensure the fidelity of the generated data. In addition, the regularisation term in (6) gives general prior knowledge of the MRI data, i.e. sparsity and smoothness. The overall algorithm is described in Algorithm 1. The algorithm starts with generating a 3D MR image via the sampling process of DDPM, which is guided by the optimization method in (5) to incorporate the under-sampled k-space. I.e. the optimization method is guided such that a specific 3D MR image is reconstructed from the k-space measurement. It should be noted that the superscript (i)and subscript index t refers to the iteration index between both optimization update and the DDPM sampling. Hence, the algorithm performs DDPM sampling and optimization updates alternately.

4. Experiments

In this section, we discuss the evaluation of proposed algorithms in terms of several metric, such as structural similarity index (SSIM) and peak signal-to-noise ratio (PSNR).

4.1. Experiment setup

Datasets In this paper, we use one dataset consisting of MR images of the knee from the fastMRI dataset [52] for training, and three datasets for testing our approach, namely MR images of the knee from the fastMRI dataset [52], the BRATS brain [2, 3, 30], and plant roots [33, 45]. For the knee data, we use the central slices to avoid noise-only images. Therefore the input data dimension is $N \times N \times S =$ $320\times320\times16$ with voxel size of 0.5 mm \times 0.5 mm \times 3 mm. All 973 training volumetric knee data from fastMRI are used for training. The brain dataset BRATS has a dimension of $240 \times 240 \times 155$. The voxel size of each image is 1 mm³. For plant roots data the full k-space data are acquired with a single coil MR scanner. The k-space data as well as the reconstructed images yield the image dimension $192 \times 192 \times 100$ and voxel size 0.5 mm \times 0.5 mm \times 0.99 mm. To evaluate in-distribution reconstruction, we use 30 randomly chosen volumetric data provided in the validation set of fastMRI knee data. For out-of-distribution, we evaluate the algorithm with 30 and 25 randomly chosen volumetric data for BRATS and plant roots, respectively. During the evaluation of BRATS data, the noise-only slices are removed.

Diffusion Model Training Procedure The experiments conducted below utilize the model that was trained on the fastMRI knee dataset [52]. We trained the model for 70,000 iterations using ℓ_1 loss and a learning rate of 10^{-4} . Due to



Figure 3. Single slice from the volume reconstruction of file1000758 (top) and file1001862 (bottom) from fastMRI knee data. The numbers on the top right corner represent the PSNR/SSIM of the slice. The subplots on the lower left corner represent the difference map between the reconstruction and ground truth. The color range is between -0.02 (bluish) and 0.02 (reddish). Note that the volumetric ground truth data has been normalized. The subplots on the lower right are a zoomed-in view.

Table 1. Mean of SSIM and PSNR from fastMRI knee data. The **bold** and <u>underline</u> represent the best and second best result.

		SSIM (↑)				PSNR (†)			
Mask Types	Methods	3D	Axial	Sagittal	Coronal	3D	Axial	Sagittal	Coronal
	GRAPPA [18]	0.868	0.733	0.783	0.805	31.33	29.15	26.27	27.53
	U-Net [23]	0.839	0.687	0.743	0.770	30.36	28.01	25.24	26.46
Uniform	MoDL [1]	0.867	0.721	0.792	0.807	31.12	28.92	26.12	27.28
(2x, 0.15)	DiffusionMBIR [8]	0.850	0.700	0.764	0.783	30.53	28.23	25.48	26.67
	Proposed	0.909	0.818	0.859	0.870	32.89	30.63	27.80	29.02
	GRAPPA	0.794	0.572	0.672	0.698	28.81	26.68	23.88	25.05
	U-Net	0.789	0.561	0.661	0.692	<u>29.25</u>	<u>26.93</u>	24.21	<u>25.37</u>
Gaussian	MoDL	0.795	0.566	0.678	0.701	29.01	26.81	24.07	25.24
(8x, 0.08)	DiffusionMBIR	0.761	0.514	0.631	0.656	28.08	25.78	23.13	24.27
	Proposed	0.823	0.626	0.722	0.743	29.83	27.61	24.84	26.00

the large size of the 3D data, the batch size is set to 1 while a gradient accumulation technique is applied to ensure training stability. In general, we follow the parameter setting given in [19] for training a U-Net to estimate the data distribution.

Description of the comparative methods We perform a comparison with several existing algorithms for MRI reconstruction, namely GRAPPA [18], fastMRI U-Net [23], MoDL [1], and DiffusionMBIR [8]. For each method, we perform reconstruction with the same under-sampling operator, i.e., masking. In contrast to all other methods, GRAPPA is a classical reconstruction algorithm that is applied directly for image reconstruction of under-sampled k-space data. We report method details in the Appendix A.

4.2. Results

Volumetric reconstruction is performed with undersampled k-space data from the specified under-sampling operator. For each algorithm, we calculate the image quality metrics SSIM and PSNR [17, 21, 47]. We report the average of those metrics for the testing datasets. For our optimization method $\mathcal{G}(\mathbf{X})$ in (5), we use the learning rate $\eta = 0.01$, number of iteration m = 10, scaling for ℓ_1 -norm $\alpha = 0.02$ and approximated total variation tv = 1, as a result of a hyper-parameter grid search, reported in Appendix **B**. In what follows, we will discuss the reconstruction performance for in-distribution and out-of-distribution data.

In Distribution In-distribution reconstruction measures the performance of the network model on data from the same distribution as the training data. We perform an evaluation of in-distribution reconstruction by using pre-trained fastMRI knee data for fastMRI U-Net, MoDL, DiffusionM-BIR, and our method. In contrast to this, GRAPPA works directly without pre-training. Table 1 shows the mean SSIM and PSNR values for all methods applied on several under-

			SSIM (\uparrow)			PSNR (†)				
Dataset	Mask Types	Methods	3D	Axial	Sagittal	Coronal	3D	Axial	Sagittal	Coronal
		GRAPPA [18]	0.966	0.878	0.795	0.797	44.40	37.39	32.21	32.16
		U-Net [23]	0.853	0.731	0.630	0.629	37.71	30.81	26.01	25.93
	Uniform	MoDL [1]	0.963	0.872	0.778	0.789	43.47	36.52	31.54	31.63
	(2x, 0.15)	DiffusionMBIR [8]	0.908	0.666	0.504	0.494	40.18	32.93	27.79	27.55
		Proposed	0.975	0.902	0.835	0.830	46.92	39.68	34.54	34.30
Roots		GRAPPA	<u>0.945</u>	0.802	<u>0.666</u>	<u>0.660</u>	<u>41.83</u>	<u>34.88</u>	29.87	29.78
		U-Net	0.885	0.725	0.580	0.573	38.32	31.59	26.94	26.96
	Gaussian	MoDL	0.941	0.797	0.656	0.653	41.09	34.24	29.41	29.52
	(8x, 0.08)	DiffusionMBIR	0.876	0.588	0.418	0.409	38.58	31.34	26.21	25.97
		Proposed	0.954	0.823	0.698	0.687	44.11	36.93	31.82	31.59
		GRAPPA	<u>0.950</u>	0.927	0.891	0.892	37.70	35.60	<u>31.48</u>	<u>33.58</u>
		U-Net	0.806	0.675	0.691	0.653	36.69	33.23	30.93	31.74
	Uniform	MoDL	0.950	0.927	0.887	0.883	37.84	35.21	31.81	32.70
	(2x, 0.15)	DiffusionMBIR	0.947	0.867	0.859	0.858	41.12	<u>37.31</u>	36.17	36.36
		Proposed	0.971	0.958	0.932	0.938	39.43	38.18	33.29	36.50
BRATS		GRAPPA	0.894	0.848	<u>0.834</u>	<u>0.779</u>	33.97	31.40	27.80	29.48
		U-Net	0.742	0.610	0.654	0.545	<u>34.69</u>	31.30	<u>28.98</u>	<u>29.68</u>
	Gaussian	MoDL	0.872	0.821	0.817	0.735	33.93	31.01	27.95	28.88
	(8x, 0.08)	DiffusionMBIR	<u>0.909</u>	0.796	0.780	0.779	37.21	33.54	31.79	32.27
		Proposed	0.932	0.908	0.856	0.863	33.72	<u>31.43</u>	27.61	29.85

Table 2. Mean of SSIM and PSNR from BRATS and plant roots data. The **bold** and <u>underline</u> represent the best and second-best results.

sampling measurements. Apart from directly evaluating the SSIM and PSNR per volume, we also evaluate those metrics for all 2D slices of x, y, and z axes.

We observe that in Uniform masking, the proposed algorithm outperforms other methods in terms of SSIM and PSNR. Given the $8 \times$ acceleration factor applied in Gaussian masking, there is a decrease in both overall SSIM and PSNR when compared to Uniform masking with a $2 \times$ acceleration factor due to the reduced amount of information available in the k-space data for the reconstruction process. For SSIM evaluation the proposed algorithm still performs better than other algorithms. A single slice taken from the 3D reconstruction of fastMRI knee data for different each mask type is presented in Figure 3. Overall, all methods yield visual similar results, however with clearly noticeable differences on close inspection. For $2 \times$ as well as for 8× acceleration we observe that DiffusionMBIR reconstructs somewhat noisy-looking images, while fastMRI U-Net yields crisp-looking results with the highest contrast. However, as the metric scores support, the high contrast surpasses the contrast of the ground truth and this may indicate mild hallucination of structures. In addition, it smooths low-contrast regions generating a mildly cartoonish impression. GRAPPA produces blurrier images, especially in the $8 \times$ case. Quantitatively, the proposed method is closest to the ground truth.

Out of Distribution An out-of-distribution (OOD) evaluation shows the robustness of algorithms for the reconstruction task. As discussed in Section 1, general prior knowledge can be applied to any MRI data, however, with limited representativeness. While deep learning-based methods usually only work for in-distribution data, the hybrid Diffusion Model has been shown to be robust to OOD data where it has not been trained [11, 22]. The proposed algorithm utilizes both conventional optimization methods and diffusion models and is therefore expected to achieve good performance in OOD evaluation.

We here measure the performance of all methods applied to OOD data, i.e., to reconstruct images that are different from the training images. Table 2 presents the evaluation of average Structural Similarity Index Measure (SSIM) and Peak Signal-to-Noise Ratio (PSNR) for both plant roots and BRATS data. In BRATS case, the fastMRI U-Net performs the worst compared to all other algorithms. Meanwhile, the Diffusion MBIR struggles to perform for plant roots data. The proposed method effectively enhances the reconstruction, specifically improving the structure and feature properties of the images. Since each slice in the roots data only represents the sparse distribution of a root, we perform the maximum intensity projection on each axis to visualize the fine roots structure. When assessing PSNR for BRATS data, Diffusion MBIR mostly outperforms other methods. This observation suggests that the Diffusion MBIR primarily contributes to reducing the absolute error

		SSIM (\uparrow)				PSNR (†)				
Mask Types	Methods	3D	Axial	Sagittal	Coronal	3D	Axial	Sagittal	Coronal	
Uniform	2D + A [20]	0.815	0.648	0.742	0.756	28.41	26.61	23.44	24.70	
(2x, 0.15)	2D [19]	0.769	0.579	0.684	0.698	26.88	24.57	21.93	23.29	
Gaussian	2D + A	0.673	0.404	0.567	0.584	25.46	23.44	20.54	21.65	
(8x, 0.08)	2D	0.558	0.304	0.454	0.464	22.79	20.38	17.82	19.02	

Table 3. Mean of SSIM and PSNR for 2D + A and stacking 2D Architecture to generate 3D knee images. The **bold** represents the best result. The regularization terms are not activated.

between pixels and the ground truth than features and structures of images. We report qualitative results for the BRATS and plant roots data Appendix F.



Figure 4. Ablation study to investigate the effect of regularization for reconstruction in terms of SSIM (\uparrow) metric for uniform mask. The optimization iteration m = 10 and the learning rate 0.01.

5. Ablation studies

We perform ablation studies to highlight the effects of regularization on the optimization method and to compare the impact of the 2D + A architecture versus the stacking 2D architecture in generating 3D images

5.1. Importance of regularization

In Figure 4, we demonstrate that both regularization terms are essential for improving reconstruction based on the SSIM metric across all datasets. In the roots dataset, dominated by sparse regions, the ℓ_1 -norm significantly influences performance more than the approximated total variation (tv). Conversely, in the knee dataset, both the approximated total variation and ℓ_1 -norm exhibit similar effects on performance. The result for the Gaussian mask is presented in Section C.

5.2. Importance of 3D representation

While the 2D + A architecture directly generates 3D volumetric images, the 2D architecture constructs 3D volumetric images by stacking the 2D image slice by slice for both Uniform and Gaussian masks. It can be seen that, as discussed in Section 3.1, stacking 2D images to generate a 3D image leads to inconsistency as shown in Table 3, yielding low SSIM metric for all combination axes.

6. Conclusion

In this study, we present a novel 3D reconstruction approach by integrating 2D + A diffusion model and proposed anoptimization method as a guiding constraint for refining the diffusion model's sampling process. Our proposed method enables a pre-trained 3D diffusion model to faithfully reconstruct MR images from both in- and out-of-distribution input data. This capability has been validated both quantitatively and qualitatively through comprehensive comparisons with existing methods. For fastMRI Knee and plant roots datasets, the proposed method outperforms all other tested methods.

Moving forward, we plan to investigate various diffusion model architectures to further enhance the performance and capabilities of the proposed method for a broader range of applications. Moreover, we aspire to augment the quantity and diversity of data utilized by the model, thereby enriching its generalization capabilities. Additionally, we are keen on enabling parallelization techniques to facilitate real-time inference, thus advancing the practical applicability of the models in real-world scenarios.

Acknowledgment

This work was supported by the President's Initiative and Networking Funds of the Helmholtz Association of German Research Centres [Grant HighLine ZT-I-PF-4-042]. The authors gratefully acknowledge computing time on the supercomputer JURECA [44] at Forschungszentrum Jülich under grant no. delia-mp.

References

- Hemant K Aggarwal, Merry P Mani, and Mathews Jacob. Modl: Model-based deep learning architecture for inverse problems. *IEEE transactions on medical imaging*, 38(2):394–405, 2018. 2, 6, 7
- [2] Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin S Kirby, John B Freymann, Keyvan Farahani, and Christos Davatzikos. Advancing the cancer genome atlas glioma mri collections with expert segmentation labels and radiomic features. *Scientific data*, 4(1):1–13, 2017. 5
- [3] Spyridon Bakas, Mauricio Reyes, Andras Jakab, Stefan Bauer, Markus Rempfler, Alessandro Crimi, Russell Takeshi Shinohara, Christoph Berger, Sung Min Ha, Martin Rozycki, et al. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. *arXiv preprint arXiv:1811.02629*, 2018. 5
- [4] Abhirup Banerjee, Julià Camps, Ernesto Zacur, Christopher M Andrews, Yoram Rudy, Robin P Choudhury, Blanca Rodriguez, and Vicente Grau. A completely automated pipeline for 3d reconstruction of human heart from 2d cine magnetic resonance slices. *Philosophical Transactions of the Royal Society A*, 379(2212):20200257, 2021. 1
- [5] Jérôme Bolte, Shoham Sabach, and Marc Teboulle. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming*, 146(1-2):459–494, 2014. 5, 13
- [6] Ronald N Bracewell. Strip integration in radio astronomy. Australian Journal of Physics, 9(2):198–217, 1956. 4, 12
- [7] Ronald N Bracewell. Numerical transforms. *Science*, 248(4956):697–704, 1990. 4, 12
- [8] Hyungjin Chung, Dohoon Ryu, Michael T McCann, Marc L Klasky, and Jong Chul Ye. Solving 3d inverse problems using pre-trained 2d diffusion models. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 22542–22551, 2023. 1, 3, 6, 7, 12
- [9] Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. *Advances in Neural Information Processing Systems*, 35:25683–25696, 2022. 4
- [10] Hyungjin Chung, Byeongsu Sim, and Jong Chul Ye. Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 12413–12422, 2022. 4
- [11] Hyungjin Chung and Jong Chul Ye. Score-based diffusion models for accelerated mri. *Medical image analysis*, 80:102479, 2022. 1, 2, 3, 7
- [12] Salman Ul Hassan Dar, Muzaffer Özbey, Ahmet Burak Çatlı, and Tolga Çukur. A transfer-learning approach for accelerated mri using deep neural networks. *Magnetic resonance in medicine*, 84(2):663–685, 2020. 2
- [13] Nicholas Dwork, Corey A Baron, Ethan MI Johnson, Daniel O'Connor, John M Pauly, and Peder EZ Larson. Fast variable density poisson-disc sample generation with directional

variation for compressed sensing in mri. *Magnetic resonance imaging*, 77:186–193, 2021. 1, 12

- [14] Taejoon Eo, Yohan Jun, Taeseong Kim, Jinseong Jang, Ho-Joon Lee, and Dosik Hwang. Kiki-net: cross-domain convolutional neural networks for reconstructing undersampled magnetic resonance images. *Magnetic resonance in medicine*, 80(5):2188–2201, 2018. 2
- [15] Jeffrey A Fessler. Model-based image reconstruction for mri. IEEE signal processing magazine, 27(4):81–89, 2010. 1
- [16] Simon Foucart and Holger Rauhut. A Mathematical Introduction to Compressive Sensing. Springer Science & Business Media, 2013. 5, 13
- [17] Daniel Gourdeau, Simon Duchesne, and Louis Archambault. On the proper use of structural similarity for the robust evaluation of medical image synthesis models. *Medical Physics*, 49(4):2462–2474, 2022. 6
- [18] Mark A Griswold, Peter M Jakob, Robin M Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, and Axel Haase. Generalized autocalibrating partially parallel acquisitions (grappa). *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 47(6):1202–1210, 2002. 2, 6, 7
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 3, 4, 6, 8
- [20] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J. Fleet. Video Diffusion Models. *ICLR 2022 Workshop Deep Generative Models for Highly Structured Datas*, 2022. 3, 4, 8
- [21] Alain Hore and Djemel Ziou. Image quality metrics: Psnr vs. ssim. In 2010 20th international conference on pattern recognition, pages 2366–2369. IEEE, 2010. 6
- [22] Ajil Jalal, Marius Arvinte, Giannis Daras, Eric Price, Alexandros G Dimakis, and Jon Tamir. Robust compressed sensing mri with deep generative priors. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 14938–14954. Curran Associates, Inc., 2021. 1, 3, 7
- [23] Kyong Hwan Jin, Michael T. McCann, Emmanuel Froustey, and Michael Unser. Deep convolutional neural network for inverse problems in imaging. *IEEE Transactions on Image Processing*, 26(9):4509–4522, 2017. 6, 7, 12
- [24] Shantanu H Joshi, Antonio Marquina, Stanley J Osher, Ivo Dinov, John D Van Horn, and Arthur W Toga. Mri resolution enhancement using total variation regularization. In 2009 IEEE international symposium on biomedical imaging: from nano to macro, pages 161–164. IEEE, 2009. 1
- [25] Florian Knoll, Kerstin Hammernik, Chi Zhang, Steen Moeller, Thomas Pock, Daniel K Sodickson, and Mehmet Akcakaya. Deep-learning methods for parallel magnetic resonance imaging reconstruction: A survey of the current approaches, trends, and issues. *IEEE signal processing magazine*, 37(1):128–140, 2020. 3
- [26] David J Larkman and Rita G Nunes. Parallel magnetic resonance imaging. *Physics in Medicine & Biology*, 52(7):R15, 2007. 2

- [27] Dana J Lin, Patricia M Johnson, Florian Knoll, and Yvonne W Lui. Artificial intelligence for mr image reconstruction: an overview for clinicians. *Journal of Magnetic Resonance Imaging*, 53(4):1015–1028, 2021. 3
- [28] Michael Lustig, David L Donoho, Juan M Santos, and John M Pauly. Compressed sensing mri. *IEEE signal processing magazine*, 25(2):72–82, 2008. 2
- [29] Morteza Mardani, Enhao Gong, Joseph Y Cheng, Shreyas S Vasanawala, Greg Zaharchuk, Lei Xing, and John M Pauly. Deep generative adversarial neural networks for compressive sensing mri. *IEEE transactions on medical imaging*, 38(1):167–179, 2018. 2
- [30] Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34(10):1993–2024, 2014. 5
- [31] Matthew J Muckley, Bruno Riemenschneider, Alireza Radmanesh, Sunwoo Kim, Geunu Jeong, Jingyu Ko, Yohan Jun, Hyungseob Shin, Dosik Hwang, Mahmoud Mostapha, et al. Results of the 2020 fastmri challenge for machine learning mr image reconstruction. *IEEE transactions on medical imaging*, 40(9):2306–2317, 2021. 1
- [32] Neal Parikh, Stephen Boyd, et al. Proximal algorithms. Foundations and trends[®] in Optimization, 1(3):127–239, 2014. 5, 13
- [33] Daniel Pflugfelder, Ralf Metzner, Dagmar van Dusschoten, Rüdiger Reichel, Siegfried Jahnke, and Robert Koller. Noninvasive imaging of plant roots in different soils using magnetic resonance imaging (mri). *Plant Methods*, 13(1):1–9, 2017. 5
- [34] Klaas P Pruessmann. Encoding and reconstruction in parallel mri. NMR in Biomedicine: An International Journal Devoted to the Development and Application of Magnetic Resonance In vivo, 19(3):288–299, 2006. 2
- [35] Klaas P Pruessmann, Markus Weiger, Markus B Scheidegger, and Peter Boesiger. Sense: sensitivity encoding for fast mri. Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine, 42(5):952–962, 1999. 2
- [36] Tran Minh Quan, Thanh Nguyen-Duc, and Won-Ki Jeong. Compressed sensing mri reconstruction using a generative adversarial network with a cyclic loss. *IEEE transactions on medical imaging*, 37(6):1488–1497, 2018. 1
- [37] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena*, 60(1-4):259–268, 1992. 5
- [38] Hannes Schulz, Johannes A Postma, Dagmar van Dusschoten, Hanno Scharr, and Sven Behnke. Plant root system analysis from mri images. In Computer Vision, Imaging and Computer Graphics. Theory and Application: 7th International Joint Conference, VISIGRAPP 2012, Rome, Italy, February 24-26, 2012, Revised Selected Papers, pages 411– 425. Springer, 2013. 1
- [39] Charles P Slichter. Principles of magnetic resonance, volume 1. Springer Science & Business Media, 2013. 3

- [40] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. In *International Conference on Learning Representations*, 2021. 1, 2
- [41] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2020. 4
- [42] Anuroop Sriram, Jure Zbontar, Tullie Murrell, Aaron Defazio, C Lawrence Zitnick, Nafissa Yakubova, Florian Knoll, and Patricia Johnson. End-to-end variational networks for accelerated mri reconstruction. In Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part II 23, pages 64–73. Springer, 2020. 2
- [43] Maarten L Terpstra, Matteo Maspero, Federico d'Agata, Bjorn Stemkens, Martijn PW Intven, Jan JW Lagendijk, Cornelis AT Van den Berg, and Rob HN Tijssen. Deep learningbased image reconstruction and motion estimation from undersampled radial k-space for real-time mri-guided radiotherapy. *Physics in Medicine & Biology*, 65(15):155015, 2020. 2
- [44] Philipp Thörnig. Jureca: Data centric and booster modules implementing the modular supercomputing architecture at jülich supercomputing centre. *Journal of large-scale research facilities JLSRF*, 7:A182–A182, 2021. 8
- [45] Dagmar van Dusschoten, Ralf Metzner, Johannes Kochs, Johannes A Postma, Daniel Pflugfelder, Jonas Bühler, Ulrich Schurr, and Siegfried Jahnke. Quantitative 3d analysis of plant roots growing in soil using magnetic resonance imaging. *Plant physiology*, 170(3):1176–1188, 2016. 5
- [46] Vikram Voleti, Alexia Jolicoeur-Martineau, and Chris Pal. Mcvd-masked conditional video diffusion for prediction, generation, and interpolation. Advances in neural information processing systems, 35:23371–23385, 2022. 4
- [47] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6
- [48] Jeffrey Wen, Rizwan Ahmad, and Philip Schniter. A conditional normalizing flow for accelerated multi-coil mr imaging. 2023. 1, 2
- [49] Liang Wu, Shunbo Hu, and Changchun Liu. Denoising of 3d brain mr images with parallel residual learning of convolutional neural network using global and local feature extraction. *Computational Intelligence and Neuroscience*, 2021:1– 18, 2021. 1
- [50] Wei Wu, Saurabhi Samant, Gijs de Zwart, Shijia Zhao, Behram Khan, Mansoor Ahmad, Marco Bologna, Yusuke Watanabe, Yoshinobu Murasato, Francesco Burzotta, et al. 3d reconstruction of coronary artery bifurcations from coronary angiography and optical coherence tomography: feasibility, validation, and reproducibility. *Scientific reports*, 10(1):18049, 2020. 1
- [51] Hanlu Yang. POCS Augmented CycleGan for MR Image Reconstruction. Temple University, 2020. 2

- [52] Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J Muckley, Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, et al. fastmri: An open dataset and benchmarks for accelerated mri. *arXiv* preprint arXiv:1811.08839, 2018. 2, 5, 12
- [53] Pengyue Zhang, Fusheng Wang, Wei Xu, and Yu Li. Multichannel generative adversarial network for parallel magnetic resonance image reconstruction in k-space. In *Medical Image Computing and Computer Assisted Intervention– MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part I, pages 180–188. Springer, 2018. 2*
- [54] Yihang Zhou, Haifeng Wang, Congcong Liu, Binyu Liao, Ye Li, Yanjie Zhu, Zhangqi Hu, Jianxiang Liao, and Dong Liang. Recent advances in highly accelerated 3d mri. *Physics* in Medicine & Biology, 68(14):14TR01, 2023. 1
- [55] Bo Zhu, Jeremiah Z Liu, Stephen F Cauley, Bruce R Rosen, and Matthew S Rosen. Image reconstruction by domaintransform manifold learning. *Nature*, 555(7697):487–492, 2018. 2