

Physics informed guided diffusion for accelerated multi-parametric MRI reconstruction

Perla Mayo^{1,✉}, Carolin M. Pirkel², Alin Achim¹, Bjoern Menze³, and Mohammad Golbabaee^{1,✉}

¹ University of Bristol, Bristol, United Kingdom
{pm15334, m.golbabaee}@bristol.ac.uk

² GE HealthCare, Munich, Germany

³ University of Zurich, Zurich, Switzerland

Abstract. We introduce MRF-DiPh, a novel physics informed denoising diffusion approach for multiparametric tissue mapping from highly accelerated, transient-state quantitative MRI acquisitions like Magnetic Resonance Fingerprinting (MRF). Our method is derived from a proximal splitting formulation, incorporating a pretrained denoising diffusion model as an effective image prior to regularize the MRF inverse problem. Further, during reconstruction it simultaneously enforces two key physical constraints: (1) k-space measurement consistency and (2) adherence to the Bloch response model. Numerical experiments on in-vivo brain scans data show that MRF-DiPh outperforms deep learning and compressed sensing MRF baselines, providing more accurate parameter maps while better preserving measurement fidelity and physical model consistency—critical for solving reliably inverse problems in medical imaging.

Keywords: quantitative MRI · magnetic resonance fingerprinting · denoising diffusion models · iterative image reconstruction.

1 Introduction

Magnetic Resonance Fingerprinting (MRF) and other transient-state multiparametric mapping techniques [23, 20, 16, 6] have demonstrated significant advantages over traditional steady-state and/or single-parametric quantitative MRI, enabling faster acquisitions for clinical use. However, additional acceleration through the use of compressed sampling or truncated acquisition sequences, increases image reconstruction artifacts and reduces tissue quantification accuracy. Early MRF reconstruction methods approached this problem with iterative dictionary-matching [10, 2, 14], using sparsity and/or low-rank priors [25, 34, 9, 15, 6], but struggled with highly undersampled data. Deep learning, particularly convolutional neural networks (CNNs), has since demonstrated superior performance by learning effective image priors from anatomical datasets [3, 12, 13].

Recently, denoising diffusion models (DDMs) have shown remarkable success in computational imaging [18, 30, 29, 27]. In reconstruction, DDMs can use data-driven conditioning to transform artifact-contaminated images into high-quality

restored outputs [11]. Recent works [19, 8, 35, 22] integrate physical acquisition priors, such as enforcing MRI k-space consistency between diffusion steps [21, 17, 28], to improve generalization and reliability. Despite promising results in single-parametric qMRI [4, 31] and MRF [24], current MRF diffusion-based approaches do not explicitly enforce consistency with the physical acquisition model or scanner measurements, which are critical factors affecting reliability.

Contribution: To address these limitations, we propose MRF-DiPh, a physics-informed diffusion-based MRF reconstruction algorithm. MRF-DiPh uses a pre-trained denoising diffusion model as a spatial image prior to regularize reconstruction, and enforces two essential physical constraints: (1) consistency with k-space measurements and (2) compliance with the Bloch signal-response model. Experiments on retrospectively shortened in vivo brain scans show that MRF-DiPh outperforms existing deep learning and compressed sensing methods, yielding more accurate parameter maps with stronger adherence to physical constraints and scanner measurements.

2 Preliminaries

2.1 The MRF Problem: MRF reconstruction is a *nonlinear* inverse problem:

$$\mathbf{y} \approx \mathcal{A}(\mathbf{x}) \text{ such that } \mathbf{x}_v = \boldsymbol{\rho}_v \text{BLOCH}(\text{T1}_v, \text{T2}_v), \forall v : \text{voxels} \quad (1)$$

Quantitative T1 and T2 tissue property maps (qmaps) have to be estimated for $h \times w$ voxels. $\mathbf{x} \in \mathbb{C}^{s \times h \times w}$ is the time-series of magnetization images (TSMI) that has to be reconstructed. $\text{BLOCH}(\cdot)$ denotes the nonlinear Bloch response model that voxel-wise encodes qmaps into the time-signals (fingerprints) within the TSMI, scaled by the proton density $\boldsymbol{\rho}$. $\mathcal{A}$ is the linear forward acquisition operator that relates TSMI to the undersampled k-space measurements $\mathbf{y} \in \mathbb{C}^{c \times m \times l}$ from m spatial frequency locations across c receiver coil channels, and l time frames. $\mathcal{A}$ encompasses coil sensitivities, a nonuniform FFT, and a linear dimensionality reduction [26, 2, 15] which is commonly used for computationally efficient reconstructions by compressing the time dimension of TSMI into adequately smaller $s \ll l$ time frames. Additionally, the reconstructions employ an MRF *dictionary* as a discretized approximation of the BLOCH model, containing a lookup table, LUT, of d finely sampled T1-T2 values and their precomputed Bloch responses $D \in \mathbb{C}^{d \times s}$ [26]. The Euclidean projections of a TSMI $\mathbf{x}$ onto the Bloch constraints in (1) can then be approximated by dictionary matching [10]:

$$\mathbf{x}_{\text{proj}} \approx \boldsymbol{\rho}^* D(\text{T1}^*, \text{T2}^*) \text{ where, } (\text{T1}^*, \text{T2}^*, \boldsymbol{\rho}^*) \leftarrow \text{DICT-MATCH}(\mathbf{x}) \quad (2)$$

and $\text{DICT-MATCH}(\mathbf{x}) = \arg \min_{\boldsymbol{\rho}, (\text{T1}, \text{T2}) \in \text{LUT}} \|\mathbf{x} - \boldsymbol{\rho} D(\text{T1}, \text{T2})\|_2$, based on voxel-wise dictionary search using exact or approximate algorithms (see [7, 14]).

2.2 Diffusion Models: DDMs generate images from complex distributions $p(\mathbf{x}_0)$ by iteratively refining Gaussian noise samples $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \text{Id})$ over $t =$

$1, \dots, T$ diffusion steps [18, 30, 29, 27]. The forward diffusion process corrupts clean images $\mathbf{x}_0$ into noisy counterparts $\mathbf{x}_t$ through:

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon \quad \text{where, } \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{Id}), \quad t = 1, \dots, T \quad (3)$$

where $\{\bar{\alpha}_t\} \in (0, 1)$ decreases as t increases and defines a steadily lowering signal-to-noise ratio $\text{SNR} := 1/\sigma_t^2 = \bar{\alpha}_t/(1 - \bar{\alpha}_t)$, until $\mathbf{x}_T$ approximates pure noise. Training a DDM involves learning a *noise estimation network* ϵ_θ , minimising the loss $\mathbb{E}_{\mathbf{x}_0, t, \epsilon} [\|\epsilon - \epsilon_\theta(\mathbf{x}_t, t)\|_2^2]$. Once trained, ϵ_θ is used iteratively in a reverse diffusion process to approximate samples from $p(\mathbf{x}_0)$ [29]:

$$\mathbf{x}_{t-1} = \sqrt{\bar{\alpha}_{t-1}} \tilde{\mathbf{x}}_{0,t} + \sqrt{1 - \bar{\alpha}_{t-1} - \eta^2} \epsilon_\theta(\mathbf{x}_t, t) + \eta \epsilon \quad (4)$$

where $\tilde{\mathbf{x}}_{0,t} := \frac{1}{\sqrt{\bar{\alpha}_t}} (\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(\mathbf{x}_t, t))$ estimates $\mathbf{x}_0$ by denoising $\mathbf{x}_t$, $\epsilon_\theta(\mathbf{x}_t, t)$ and $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{Id})$ are the predicted (deterministic) and stochastic noise terms reintroduced to the sample at appropriate scales, and $\eta \in [0, \sqrt{1 - \bar{\alpha}_{t-1}}]$ controls the stochasticity of sampling. For inverse problems, reconstruction guidance can be provided by incorporating *conditional* information into training and sampling e.g, a low-quality reconstruction $\mathbf{x}_c = \mathcal{A}^H \mathbf{y}$ from undersampled measurements $\mathbf{y}$ using the adjoint operator. The network is thus parametrized as $\epsilon_\theta(\mathbf{x}_t, t, \mathbf{x}_c)$ to refine denoising adapted to $\mathbf{x}_c$, approximating samples from distribution $p(\mathbf{x}_0 | \mathbf{x}_c)$.

3 MRF-DiPh Algorithm

We formulate MRF reconstruction as:

$$\arg \min_{\mathbf{x} \in \mathcal{B}} f(\mathbf{x}) + \lambda h(\mathbf{x}) \quad (5)$$

where $f(\mathbf{x}) = \|\mathbf{y} - \mathcal{A}\mathbf{x}\|_2^2$ enforces k-space consistency, $h(\mathbf{x})$ regularizes solutions with an image prior, and the regularization parameter $\lambda > 0$ balances the two terms. The constraint set $\mathcal{B} := \{\mathbf{x} \text{ s.t. } x = \rho D(T1, T2)\}$ ensures Bloch model consistency via an MRF dictionary (section 2.1). Using Half Quadratic Splitting (HQS), we can solve (5) iteratively for $t = T, T-1, \dots, 1$:

$$\tilde{\mathbf{x}}_t = \text{prox}_{\sigma^2 h}(\hat{\mathbf{x}}_t) := \arg \min_{\mathbf{x}} h(\mathbf{x}) + \frac{1}{2\sigma^2} \|\mathbf{x} - \hat{\mathbf{x}}_t\|_2^2, \quad (6)$$

$$\hat{\mathbf{x}}_{t-1} = \arg \min_{\mathbf{x} \in \mathcal{B}} f(\mathbf{x}) + \frac{\mu}{2} \|\mathbf{x} - \tilde{\mathbf{x}}_t\|_2^2, \quad (7)$$

where $\sigma^2 := \lambda/\mu$ and $\mu > 0$ controls convergence. This process decouples image prior enforcement (proximal step (6)) from physical acquisition constraints (7). To incorporate deep learning-based image priors, one could adopt a plug-and-play framework [33, 1] and replace (6) with a deep denoising model pretrained on images to remove additive Gaussian noise of variance $\sigma^2 = 1/\text{SNR}$. Inspired by this idea, we extend HQS with a DDM-based prior:

$$\tilde{\mathbf{x}}_{0,t} = \frac{1}{\sqrt{\bar{\alpha}_t}} \left(\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(\mathbf{x}_t, t) \right), \quad (8)$$

$$\hat{\mathbf{x}}_{0,t} = \arg \min_{\mathbf{x} \in \mathcal{B}} f(\mathbf{x}) + \frac{\mu_t}{2} \|\mathbf{x} - \tilde{\mathbf{x}}_{0,t}\|_2^2, \quad (9)$$

$$\mathbf{x}_{t-1} = \sqrt{\bar{\alpha}_{t-1}} \hat{\mathbf{x}}_{0,t} + \sqrt{1 - \bar{\alpha}_{t-1}} \left(\sqrt{1 - \xi} \hat{\epsilon}_t + \sqrt{\xi} \epsilon \right) \quad (10)$$

where $\hat{\epsilon}_t = \frac{\mathbf{x}_t - \sqrt{\bar{\alpha}_t} \tilde{\mathbf{x}}_{0,t}}{\sqrt{1 - \bar{\alpha}_t}}$ and $\epsilon \sim \mathcal{N}(0, \mathbf{I})$. Eq (8) performs denoising via DDM's pretrained noise estimator ϵ_θ , yielding $\tilde{\mathbf{x}}_{0,t}$, an estimate of the clean image $\mathbf{x}_0$ (section 2.2). Eq. (9) further enforces k-space and Bloch consistency, producing $\hat{\mathbf{x}}_{0,t}$. Finally, (10) reintroduces noise to the current cleaned image estimate to obtain a diffusion sample $\mathbf{x}_{t-1}$ for the next step. The added noise combines a stochastic ϵ and a deterministic term $\hat{\epsilon}_t$, balanced by $\xi \in [0, 1]$, e.g. $\xi = 0$ gives a fully deterministic sampling.⁴ As discussed in section 2.2, ϵ_θ estimates noise at various SNR levels $1/\sigma_t^2$, which increases as sampling progresses, i.e. t decreases. Consequently, $\mu_t := \lambda/\sigma_t^2$ also increases as t decreases.

The constrained subproblem (9) can be solved using the ADMM algorithm [5], which introduces an auxiliary variable $\mathbf{z}_t$, the dual variable (Lagrange multipliers) $\mathbf{v}_t$, and a convergence hyperparameter $\gamma_t > 0$. Early experiments showed a single ADMM iteration is sufficient for updating (9). This single-step ADMM update can be written as follows:

$$\hat{\mathbf{x}}_{0,t} = \arg \min_{\mathbf{x}} f(\mathbf{x}) + \frac{\mu_t}{2} \|\mathbf{x} - \tilde{\mathbf{x}}_{0,t}\|_2^2 + \frac{\gamma_t}{2} \|\mathbf{x} - \mathbf{z}_t + \mathbf{v}_t/\gamma_t\|_2^2, \quad (11)$$

$$\mathbf{z}_{t-1} = \arg \min_{\mathbf{z} \in \mathcal{B}} \|\mathbf{z} - (\hat{\mathbf{x}}_{0,t} + \mathbf{v}_t/\gamma_t)\|_2^2, \quad (12)$$

$$\mathbf{v}_{t-1} = \mathbf{v}_t + \gamma_t (\hat{\mathbf{x}}_{0,t} - \mathbf{z}_{t-1}). \quad (13)$$

where (11), equivalent to $\text{prox}_{\frac{1}{\mu_t + \gamma_t} f} \left(\frac{\mu_t \tilde{\mathbf{x}}_{0,t} + \gamma_t \mathbf{z}_t - \mathbf{v}_t}{\mu_t + \gamma_t} \right)$, is a linear least squares term enforcing k-space consistency, which can be efficiently updated by a few conjugate gradient (CG) iterations. Eq (12) is a projection onto the Bloch constraints set, approximable via DICT-MATCH($\hat{\mathbf{x}}_{0,t} + \mathbf{v}_t/\gamma_t$) (see (2)) to give Bloch-consistent TSMI $\mathbf{z}_{t-1}$ and also, the qmaps $\mathbf{q}_{t-1} = (\mathbf{T}1_{t-1}, \mathbf{T}2_{t-1}, \boldsymbol{\rho}_{t-1})$. Finally, (13) updates the ADMM's dual variable $\mathbf{v}_{t-1}$. The ADMM parameter is set as $\gamma_t := \tau \mu_t$, where $\tau > 0$ balances k-space and Bloch consistency constraints. Combining (8)-(13), we obtain MRF-DiPh (Algorithm 1). To accelerate reconstruction, we follow [29] and use a sub-sequence $\{t_k\}_{k=1}^K \subseteq [1, \dots, T]$ of $K \ll T$ time steps, skipping intermediate steps during diffusion sampling.

⁴ The deterministic noise is predicted from the current noisy sample $\mathbf{x}_t$ and the physics-consistent estimate $\hat{\mathbf{x}}_{0,t}$, which differs slightly from the deterministic noise used in (4); replacing $\hat{\mathbf{x}}_{0,t}$ with $\tilde{\mathbf{x}}_{0,t}$ from (8) would recover $\hat{\epsilon}_t = \epsilon_\theta(\mathbf{x}_t, t)$ similar as (4). Instead, we update $\hat{\epsilon}_t$ using physics-consistent $\hat{\mathbf{x}}_{0,t}$ from (9).

Algorithm 1 MRF-DiPh

Require: $f, \epsilon_\theta, \{t_k\}_{k=1}^K, \{\bar{\alpha}_{t_k}\}_{k=1}^K, \lambda, \tau, \xi$.

- 1: Set $\sigma_k^2 := (1 - \bar{\alpha}_{t_k})/\bar{\alpha}_{t_k}$, $\mu_k := \lambda/\sigma_k^2$, $\gamma_k := \tau\mu_k$
- 2: Initialize $\mathbf{x}_K \sim \mathcal{N}(\mathbf{0}, \mathbf{Id})$, $\mathbf{z}_K = \mathbf{v}_K = \mathbf{0}$
- 3: **for** $k = K, \dots, 1$ **do**
- 4: $\tilde{\mathbf{x}}_{0,k} = \frac{1}{\sqrt{\bar{\alpha}_k}} (\mathbf{x}_k - \sqrt{1 - \bar{\alpha}_k} \epsilon_\theta(\mathbf{x}_k, t_k, \mathbf{x}_c))$ //denoised image
- 5: $\hat{\mathbf{x}}_{0,k} = \text{prox}_{\frac{1}{\mu_k + \gamma_k} f} \left(\frac{\mu_k \tilde{\mathbf{x}}_{0,k} + \gamma_k \mathbf{z}_k - \mathbf{v}_k}{\mu_k + \gamma_k} \right)$ //k-space consistent image
- 6: $(\mathbf{z}_{k-1}, \mathbf{q}_{k-1}) \leftarrow \text{DICT-MATCH}(\hat{\mathbf{x}}_{0,k} + \mathbf{v}_k/\gamma_k)$ //Bloch-consistent image & qmaps
- 7: $\mathbf{v}_{t-1} = \mathbf{v}_t + \gamma_t(\hat{\mathbf{x}}_{0,t} - \mathbf{z}_{t-1})$ //dual variable update
- 8: $\hat{\epsilon}_k = \frac{1}{\sqrt{1 - \bar{\alpha}_k}} (\mathbf{x}_k - \sqrt{\bar{\alpha}_k} \mathbf{z}_{k-1})$ //predicted (deterministic) noise
- 9: $\epsilon \sim \mathcal{N}(\mathbf{0}, I)$ //stochastic noise
- 10: $\mathbf{x}_{k-1} = \sqrt{\bar{\alpha}_{k-1}} \mathbf{z}_{k-1} + \sqrt{1 - \bar{\alpha}_{k-1}} (\sqrt{\xi} \epsilon + \sqrt{1 - \xi} \hat{\epsilon}_k)$ //noisy sample for next iter
- 11: **end for**; **return** $\mathbf{x}_{\text{rec}} = \mathbf{z}_0$, $\mathbf{q}_{\text{rec}} = \mathbf{q}_0 := \{\mathbf{T}1_0, \mathbf{T}1_0, \boldsymbol{\rho}_0\}$

3.1 Implementation Details

The source code for our implementation is publicly available at <https://github.com/p-mayo/mrf-diph>. Our approach was developed in PyTorch using the guided diffusion toolbox⁵, following the UNet architecture from [11] at half precision. The UNet consisted of six levels, each with two residual blocks, using channel sizes [128, 128, 256, 256, 512, 512] from highest to lowest feature resolutions, and three attention heads at feature resolutions 28×28 , 14×14 and 7×7 .

Our base model employed a conditional DDM trained on paired TSMIs $\{\mathbf{x}_c = \mathcal{A}^H \mathbf{y}, \mathbf{x}_0 = \mathbf{x}_{\text{ref}}\}$. Low-aliasing reference TSMIs $\mathbf{x}_{\text{ref}}$ were generated from k-space measurements $\mathbf{y}_{\text{ref}}$ acquired through extended MRF scans [20]. To simulate faster acquisitions, $\mathbf{y}_{\text{ref}}$ was subsampled by truncation along the time dimension, producing k-space data $\mathbf{y}$, from which highly artifact-contaminated condition images $\mathbf{x}_c$ were obtained. Both $\mathbf{x}_c$ and $\mathbf{x}_{\text{ref}}$ were time-compressed with $s = 5$ [26]. Complex-valued data were processed by concatenating the real and imaginary parts along the channel dimension, and images were range-normalized to $[-1, 1]$. Our framework also supports an unconditional DDM, trained solely on $\{\mathbf{x}_0 = \mathbf{x}_{\text{ref}}\}$, without data-driven conditioning/guidance. We explored this approach in Section 4. Training used $T = 1000$ diffusion steps, $\bar{\alpha}_t = \prod_{s=0}^{t-1} (1 - \beta_s)$ with linearly spaced β_t ($\beta_0 = 0.0001$, $\beta_T = 0.02$), for 100k iterations, using ADAM optimizer with learning rate 10^{-4} , batch size 32, and data augmentation with random vertical/horizontal spatial flips. For inference, our base model used experimentally tuned parameters $K = 30$, $\lambda = 10^{-4}$, $\tau = 0.01$, $\xi = 1$, CG maximum iterations 5 and initialized by the previous update of line 5 in Algorithm 1.

4 Numerical Experiments

Setup: The dataset consists of anonymized 2D brain MRF scans from healthy adult volunteers, obtained with informed consent in compliance with the Ger-

⁵ <https://github.com/openai/guided-diffusion>

man Act on Medical Devices, and acquired using the Steady State Precession (FISP) sequence and flip angles from [20] for $l = 1000$ time frames, repetition/echo/inversion times of 10/1.908/18 ms, non-Cartesian k-space sampling with variable density spiral readouts, image size 230×230 , 1mm in-plane resolution and 5mm slice thickness. Scans were performed on the same 3T GE MR750w scanner with 8-channel receive-only RF head coils. Models were trained and tested for the reconstruction of $R = 5$ fold accelerated scans by retrospective truncation of the FISP sequence length to $l = 200$. The dataset includes 8 subjects with 15 axial slices each, split 75%-25% over mutually exclusive subjects –6 for training, 2 held-out for testing– ensuring no data leakage. Reference qmaps were reconstructed using LRTV [15] from full-length ($l = 1000$) scans, from which the reference TSMIs were estimated using BLOCH model in (1). An MRF dictionary containing $d = 95k$ atoms, was simulated from EPG [32] and used in methods that required quantitative mapping. Except the SCQ baseline, other methods used a time-domain linear dimensionality reduction with $s = 5$ from [26]. Reconstruction performance was evaluated using Mean Average Percentage Error (MAPE) for skull-stripped qmaps, channel-averaged Normalized Root Mean Squared Error (NRMSE) for TSMI, and NRMSE between k-space measurements ($\mathbf{y}$) and predictions ($\mathbf{y}_{\text{rec}} := \mathcal{A}\mathbf{x}_{\text{rec}}$), where applicable. Experiments ran on an NVIDIA GeForce RTX 4090 GPU.

Baseline methods: Our method was evaluated against SVDMMRF [26], MRF-ADMM [2], LRTV [15], SCQ [12] and MRF-IDDPM [24]. **SVDMMRF** uses $\mathbf{x}_{\text{rec}} := \mathcal{A}^H \mathbf{y}$ to reconstruct the TSMIs (equivalent to the conditions $\mathbf{x}_c$ in our approach) which are then passed to DICT-MATCH for quantitative mapping. **MRF-ADMM** uses the ADMM algorithm to enforce k-space consistency and Bloch constraints (DICT-MATCH) during reconstruction. Here we did not employ additional regularisation. **LRTV** employs k-space consistency and a Total Variation image regularization using convex optimisation. **SCQ** is a CNN-based deep learning method that maps aliasing-contaminated TSMIs from undersampled acquisitions to restored T1/T2 maps. It uses a fully-connected network to reduce the TSMI dimension, then employs two UNets to estimate T1 and T2 maps separately. We implemented SCQ following [12] for architecture, and trained the models for 1k epochs using data augmentation (patching, vertical/horizontal flips), ADAM optimizer, MSE loss, and a learning rate of 0.005–parameters that performed best for our data. **MRF-IDDPM** is a conditional DDM model that does not employ physics-driven guidance during sampling. We assess its sampling performance using the same network ϵ_θ trained for MRF-DiPh.

Ablations: The parameters λ and τ were tuned to reach a satisfactory balance between image prior and physics guidance. We found $\lambda = 10^{-4}$, and $\gamma \in [0.01, 0.1]$ provided competitive performance. To assess the effect of ξ , we tested two MRF-DiPh sampling configurations, one used an even mix of deterministic and stochastic noise ($\xi = 0.5$, Mode A), whereas a second used purely deterministic noise ($\xi = 0$, Mode B), and found no significant impact (Table 1).

Table 1. Reconstruction metrics (averaged over the test dataset) for TSMI, T1 and T2 maps, as well as the k-space fitting errors for different methods.

Method	MAPE (%) ↓		NRMSE (%) ↓	
	T1	T2	TSMI	K-Space
SVDMRF	20.01	144.27	57.10	99.93
MRF-ADMM	20.30	68.51	30.90	18.69
LRTV	19.84	39.01	37.95	11.72
SCQ	8.76	22.61	—	—
MRF-IDDPM	8.45	22.54	27.26	36.06
MRF-DiPh (base)	6.75	18.40	18.65	22.82
MRF-DiPh (A)	6.80	18.41	18.70	22.52
MRF-DiPh (B)	7.15	18.63	18.64	22.40
MRF-DiPh (C)	7.17	18.82	18.79	22.12
MRF-DiPh (D)	11.32	29.78	25.36	19.17

Two additional modes were evaluated: Mode C, which combined DDM sampling with k-space consistency but without dictionary matching, and Mode D, which employed an unconditional DDM denoiser trained separately without conditioning images $\mathbf{x}_c$, while keeping other training parameters the same as our base model. To improve sampling speed, we inspected the maximum number of CG iterations in $prox_f$, and the number sampling steps K . Results of these investigations are in Tables 1 and 2.

5 Results

MRF-DiPh outperforms the tested baselines across all reconstruction metrics (Table 1), improving T1 MAPE by $\sim 2\%$, T2 MAPE by $\sim 4.2\%$, and TSMI NRMSE by $\sim 8.6\%$. Figure 1 compares the reconstructed T1 and T2 maps. While most baselines perform well on T1, T2 estimation is notably harder in short MRF sequences. At 5x acceleration, differences are more evident in T2 maps (Figure 1, Table 1). Zoom-ins show MRF-IDDPM underestimating T2 in white matter, SCQ lacking clear/sharp tissue boundaries on T2 for deep brain tissue regions, and other baselines failing on T2—consistent with trends in Table 1. Table 1 shows that MRF-DiPh outperforms the purely data-driven MRF-IDDPM not only in quantitative mapping, but also in k-space fidelity by $\sim 13\%$ less errors, leveraging physics-informed guidance during reconstruction. Enforcing only k-space consistency along with DDM (ablated case C) improved performance over MRF-IDDPM, while adding Bloch consistency (MRF-DiPh base) further enhanced qmap reconstructions. A trade-off exists with k-space consistency—weak or absent image priors can lead to overfitting, introducing artifacts from noisy undersampled data (see LRTV, ADMM-MRF).

In our experiments, the unconditional model (MRF-DiPh D) improved upon SVDMRF, ADMM-MRF and LRTV methods. However, it underestimated T2 maps, particularly around grey matter and CSF areas (Fig. 1), falling short of conditional/supervised learning baselines. The embedded conditioning in our




Fig. 1. Reconstructed T1-T2 maps and their zoomed-in absolute percentage error maps (rows), using different methods (columns) for a representative brain slice in test set (electronic zoom recommended).

Table 2. MRF-DiPh reconstruction time vs. accuracy for the options: number of sampling steps (K), and maximum CG iterations. Results for the test image in Fig 1.

MRF-DiPh	Base	K				CG		
		5	10	20	50	1	10	20
Runtime (s)	44.17	8.83	15.88	30.03	72.29	30.02	61.03	91.79
TSMI NRMSE	15.76	16.74	14.99	15.55	16.06	15.73	15.87	15.70
(T1+T2)/2 MAPE	10.42	14.20	11.93	10.62	10.22	10.37	10.45	10.33

Table 3. Average reconstruction errors (MAPE T1 + MAPE T2)/2 over the test set throughout training iterations for MRF-DiPh and baseline MRF-IDDPM.

Checkpoint $\times 10^3$	10	20	30	40	50	60	70	80	90	100
MRF-DiPh	145.81	23.70	14.41	13.30	12.87	12.62	12.48	12.44	12.57	12.58
MRF-IDDPM	340.51	347.01	405.90	404.77	225.16	128.07	60.53	25.22	19.06	15.50

base model was able to resolve this challenging task. Therefore, exploring the application of unconditional MRF-DiPh in other acceleration schemes – such as further k-space subsampling, rather than truncating the sequence length – requires further investigation.

The runtimes MRF-DiPh base model reconstructions took ~ 44 seconds, slightly longer but comparable to the iterative baselines LRTV and ADMM-MRF. The three major steps, DDM denoising, k-space consistency ($prox_f$) and DICT-MATCH took approximately 7%, 53% and 40% of the total runtime. Reconstruction time can be decreased by reducing the number of CG iterations in $prox_f$ (e.g. to 1), or lower sampling steps K (e.g. to 20), allowing a controlled trade-off with accuracy (see Table 2). Exploring faster/approximate DICT-MATCH alternatives [7, 14] could further enhance efficiency. Our model took 55.5 hours to complete 100k training iterations. However, analysis of the performance

checkpoints reveals that MRF-DiPh, utilizing physics-based guidance, achieves competitive test accuracy early in training—significantly faster than the purely data-driven MRF-IDDP model. This advantage could have significant implications for 3D/high-dimensional imaging tasks, where training becomes highly resource-intensive.

6 Conclusions

This work introduced MRF-DiPh, a diffusion-based model for reconstructing MRF data with more accurate tissue parameter estimations and improved measurement fidelity. The proposed method intertwines the sampling steps of the diffusion model with measurement and Bloch consistency regularizations. Our experiments demonstrate MRI-DiPh’s potential with robust reconstructions that incorporate physics-driven guidance. Extending this work to 3D MRF imaging, pathological/disease cases, or anatomical regions beyond the brain remains an important direction for future research.

Acknowledgments. This work was supported by the EPSRC grant EP/X001091/1.

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

References

1. Ahmad, R., et al.: Plug-and-play methods for magnetic resonance imaging: Using denoisers for image recovery. *IEEE Signal Processing Magazine* **37**(1), 105–116 (2020)
2. Assländer, J., et al.: Low rank alternating direction method of multipliers reconstruction for MR fingerprinting. *Magnetic resonance in medicine* **79**(1), 83–96 (2018)
3. Balsiger, F., et al.: Magnetic resonance fingerprinting reconstruction via spatiotemporal convolutional neural networks. In: *Machine Learning for Medical Image Reconstruction*. pp. 39–46. Springer (2018)
4. Bian, W., Jang, A., Zhang, L., Yang, X., Stewart, Z., Liu, F.: Diffusion Modeling with Domain-conditioned Prior Guidance for Accelerated MRI and qMRI Reconstruction. *IEEE Transactions on Medical Imaging* pp. 1–1 (2024)
5. Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J.: vol. 3. Now Publishers, Inc. (2011)
6. Cao, X., et al.: Optimized multi-axis spiral projection MR fingerprinting with subspace reconstruction for rapid whole-brain high-isotropic-resolution quantitative imaging. *Magnetic Resonance in Medicine* **88**(1), 133–150 (2022)
7. Cauley, S.F., et al.: Fast group matching for mr fingerprinting reconstruction. *Magnetic resonance in medicine* **74**(2), 523–528 (2015)
8. 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* (2023)

9. Lima da Cruz, G., Bustin, A., Jaubert, O., Schneider, T., Botnar, R.M., Prieto, C.: Sparsity and locally low rank regularization for MR fingerprinting. *Magnetic resonance in medicine* **81**(6), 3530–3543 (2019)
10. Davies, M., Puy, G., Vandergheynst, P., Wiaux, Y.: A compressed sensing framework for magnetic resonance fingerprinting. *Siam journal on imaging sciences* **7**(4), 2623–2656 (2014)
11. Dhariwal, P., Nichol, A.: Diffusion models beat GANs on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)
12. Fang, Z., et al.: Deep learning for fast and spatially constrained tissue quantification from highly accelerated data in magnetic resonance fingerprinting. *IEEE transactions on medical imaging* **38**(10), 2364–2374 (2019)
13. Fatania, K., Pirkel, C.M., Menzel, M.I., Hall, P., Golbabae, M.: A plug-and-play approach to multiparametric quantitative MRI: image reconstruction using pre-trained deep denoisers. In: *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*. pp. 1–4. IEEE (2022)
14. Golbabae, M., Chen, Z., Wiaux, Y., Davies, M.: CoverBLIP: accelerated and scalable iterative matched-filtering for magnetic resonance fingerprint reconstruction. *Inverse Problems* **36**(1), 015003 (2019)
15. Golbabae, M., et al.: Compressive mri quantification using convex spatiotemporal priors and deep encoder-decoder networks. *Medical image analysis* **69**, 101945 (2021)
16. Gómez, P.A., et al.: Rapid three-dimensional multiparametric mri with quantitative transient-state imaging. *Scientific reports* **10**(1), 13769 (2020)
17. Güngör, A., et al.: Adaptive diffusion priors for accelerated MRI reconstruction. *Medical image analysis* **88**, 102872 (2023)
18. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
19. Jalal, A., Arvinte, M., Daras, G., Price, E., Dimakis, A.G., Tamir, J.: Robust compressed sensing mri with deep generative priors. In: *Advances in Neural Information Processing Systems*. vol. 34, pp. 14938–14954. Curran Associates, Inc. (2021)
20. Jiang, Y., Ma, D., Seiberlich, N., Gulani, V., Griswold, M.A.: MR fingerprinting using fast imaging with steady state precession (FISP) with spiral readout. *Magnetic resonance in medicine* **74**(6), 1621–1631 (2015)
21. Korkmaz, Y., Cukur, T., Patel, V.M.: Self-supervised mri reconstruction with unrolled diffusion models. In: *Medical Image Computing and Computer Assisted Intervention*. pp. 491–501. Springer Nature Switzerland, Cham (2023)
22. Liu, J., et al.: DOLCE: A model-based probabilistic diffusion framework for limited-angle ct reconstruction. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 10498–10508 (2023)
23. Ma, D., et al.: Magnetic resonance fingerprinting. *Nature* **495**(7440), 187–192 (2013)
24. Mayo, P., Pirkel, C.M., Achim, A., Menze, B.H., Golbabae, M.: Denoising diffusion probabilistic models for magnetic resonance fingerprinting. *arXiv preprint arXiv:2410.23318* (2024)
25. Mazor, G., Weizman, L., Tal, A., Eldar, Y.C.: Low-rank magnetic resonance fingerprinting. *Medical physics* **45**(9), 4066–4084 (2018)
26. McGivney, D., et al.: SVD compression for magnetic resonance fingerprinting in the time domain. *IEEE transactions on medical imaging* **33**(12), 2311–2322 (2014)
27. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: *International conference on machine learning*. pp. 8162–8171. PMLR (2021)

28. Peng, C., Guo, P., Zhou, S.K., Patel, V.M., Chellappa, R.: Towards performant and reliable undersampled MR reconstruction via diffusion model sampling. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 623–633. Springer (2022)
29. Song, J., Meng, C., Ermon, S.: Denoising Diffusion Implicit Models. In: International Conference on Learning Representations (2021)
30. 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 (2021)
31. Wang, S., Ma, H., Hernandez-Tamames, J.A., Klein, S., Poot, D.H.J.: qMRI Diffuser: Quantitative T1 Mapping of the Brain Using a Denoising Diffusion Probabilistic Model. In: Mukhopadhyay, A., Oksuz, I., Engelhardt, S., Mehro, D., Yuan, Y. (eds.) Deep Generative Models. pp. 129–138. Springer Nature Switzerland, Cham (2025)
32. Weigel, M.: Extended phase graphs: dephasing, RF pulses, and echoes-pure and simple. *Journal of Magnetic Resonance Imaging* **41**(2), 266–295 (2015)
33. Zhang, K., Li, Y., Zuo, W., Zhang, L., Van Gool, L., Timofte, R.: Plug-and-play image restoration with deep denoiser prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(10), 6360–6376 (2022)
34. Zhao, B., et al.: Improved magnetic resonance fingerprinting reconstruction with low-rank and subspace modeling. *Magnetic resonance in medicine* **79**(2), 933–942 (2018)
35. Zhu, Y., et al.: Denoising diffusion models for plug-and-play image restoration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 1219–1229 (June 2023)