

A probabilistic denoising diffusion-based framework for even higher accelerated quantitative MRI

Perla Mayo¹, Carolin M. Pirkel², Alin Achim¹, Bjoern H. Menze³ and Mohammad Golbabaee¹

¹University of Bristol, Bristol, United Kingdom

²GE HealthCare, Munich, Germany

³University of Zurich, Zurich, Switzerland

Synopsis

Motivation: Fast quantitative MRI using highly accelerated acquisitions like in MRF comes at the cost of severe aliasing artifacts that needs to be resolved.

Goal: Addressing undersampling artifacts and quantifying uncertainties in quantitative maps to pave the way to even shorter acquisitions e.g. in MRF.

Approach: Introducing the first probabilistic diffusion-based framework for the example of MRF reconstruction, advancing state-of-the-art-deep learning techniques for more accurate quantitative mapping with tools to assess uncertainties.

Results: Quantitative and qualitative evaluations show that our diffusion-based approach outperforms state-of-the-art in producing more accurate tissue parameters. Uncertainty maps exhibit correlations between areas of large variance with areas of large errors.

Impact

Our proposed approach enables the efficient use of Improved Denoising Diffusion Probabilistic Models for reconstructing highly accelerated quantitative MRI acquisitions, such as Magnetic Resonance Fingerprinting, leading to more accurate tissue parameter estimations.

Introduction

Compressed sampling strategies significantly reduce MRI scan times, and in quantitative MRI, such as Magnetic Resonance Fingerprinting (MRF)¹, enable fast multiparametric mapping. However, this comes at the cost of strong aliasing artifacts that need to be addressed through image reconstruction techniques. State-of-the-art algorithms are deep learning-based^{2,3,4,5,6}, focusing either on the reconstruction of MRF image timeseries data⁶ or directly the quantitative maps². Recently, denoising diffusion probabilistic models (DDPM)^{7,8} have been introduced to medical imaging applications^{9,10,11}, however, their potential for MRF problem remains unexplored. The high dimensionality of MRF data, combined with limited dataset sizes, presents significant challenges, as DDPMs are both computationally intensive and data-hungry. Our work aims to address them. We introduce a novel efficient patch-based DDPM approach for quantitative MRI and demonstrate it for the example of MRF reconstruction. We show that our method, MRF-IDDPM, outperforms current baselines and has great potential for further shortening of MRF acquisitions without loss of parameter encoding information. Findings are validated on in-vivo brain MRF scans.

Methods

To estimate the quantitative maps $\mathbf{q} = \{\text{T1}, \text{T2}, \text{Proton Density (PD)}\}$, we reconstruct \mathbf{x} the timeseries of magnetization images (TSMI) from the undersampled MRF k-space data \mathbf{y} :

$$\mathbf{y} \approx A(\mathbf{x}), \quad \text{s.t. } \mathbf{x}_p = \text{PD}_p \cdot B(\text{T1}_p, \text{T2}_p), \quad \forall p: \text{pixels} \quad (\text{eq1})$$

Here A is the linear acquisition operator, comprising nonuniform-FFT, coil sensitives and SVD reduction⁹. B is the pixel-wise nonlinear Bloch response. To reconstruct \mathbf{x} , we propose a conditional DDPM pipeline. Given a dataset of clean (reference) and degraded (condition) TSMI images $(\mathbf{x}_0, \mathbf{x}_c)$ a network $\epsilon_\theta(\mathbf{x}_t, t, \mathbf{x}_c)$ is trained to capture coarse-to-fine image details in \mathbf{x}_0 by estimating Gaussian noise added in latent variables \mathbf{x}_t during the *forward diffusion process*:

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}; \quad \alpha_t = 1 - \beta_t; \quad \bar{\alpha}_t = \prod_{j=0}^t \alpha_j; \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{Id})$$

Where β_t defines a noise schedule over $t=1, \dots, T$ diffusion steps. Guided by the gridding reconstructions $\mathbf{x}_c = A^H \mathbf{y}$ containing severe aliasing artifacts, the network learns to separate noise from image content, optimizing the loss $\mathbb{E}_{\mathbf{x}_0, t, \boldsymbol{\epsilon}} \left[\|\boldsymbol{\epsilon} - \epsilon_\theta(\mathbf{x}_t, t, \mathbf{x}_c)\|_2^2 \right]$.

During inference, the *reverse process* approximates a sample from the conditional distribution $q(\mathbf{x}_0 | \mathbf{x}_c)$ by iteratively denoising:

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(\mathbf{x}_t, t, \mathbf{x}_c) \right) + \sigma_t \mathbf{z}, \quad t = T, T-1, \dots, 1; \quad \mathbf{x}_T, \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{Id})$$

We used an IDDPM⁸ strategy for training and inference, enabling $\hat{\mathbf{x}}_0 \approx \mathbf{x}_0$ in far fewer $K \ll T$ steps by also learning variances $\sigma_t = \sigma_\theta(\mathbf{x}_t, t, \mathbf{x}_c)$. Our network is trained on randomly-cropped TSMI patches $\{\mathbf{x}_c^{(i)} = \text{Crop}(\mathbf{x}_c), \mathbf{x}_0^{(i)} = \text{Crop}(\mathbf{x}_0)\}$ of size 64x64. $\mathbf{x}_t^{(i)}$ are also treated as patches. This allows for faster training without compromising inference speed nor accuracy while also providing more variability in the training dataset. Our architecture is a UNet adopted from IDDPM⁸, trained with $T=1000$, linear noise schedule ($\beta_0 = 0.0001, \beta_T = 0.02$). Inference (reconstruction) used $K=50$ interpolated steps⁸.

We evaluated our approach on 2D healthy brain FISP-MRF data, acquired using the variable flip angle schedule¹² and TR/TE/TI of 10/1.908/18ms, $l=1000$ repetitions, variable density spiral sampling, matrix size of 230x230 with 1mm in-plane resolution and 5mm slice thickness. We compare our approach against SVD-MRF¹³, LRTV¹⁴, a DRUNet¹⁵ network trained for TSMI reconstruction, and SCQ² trained for direct parameter map estimation (SCQ does not produce TSMIs). To reconstruct \mathbf{q} , dictionary matching¹ is used on the output of all methods except SCQ. Methods are evaluated on retrospectively truncated scans with acceleration factor 5 ($l=200$). Our dataset consists of 8 subjects, 15 axial slices each. We use 6 subjects for training and 2 for evaluation. The reference parameter maps were reconstructed using LRTV¹⁴ from full-length FISP acquisitions. TSMI references were estimated from the reference parameter maps following the Bloch response model (eq1). Reconstruction metrics are reported in Table1. Figure1 shows the estimated T1/T2 parameter maps, Figure2 the reconstructed TSMIs, and Figure3 the uncertainty maps for T1/T2 parameter maps.

Results

Table1 shows that MRF-IDDPM outperforms baselines across all metrics, with an improvement of ~1.7% and ~3.5% in MAPE of T1 and T2, respectively. Qualitative assessment of Figure1 and Figure2 further support the improved performance on reconstructed T1/T2 maps and the TSMIs. The power of MRF-IDDPM becomes more evident in Figure2, where our approach is capable of reconstructing the last (i.e. weakest) two TSMI subspace components (Table1 also shows ~15% less TSMI NRMSE using MRF-IDDPM). We hypothesize such performance can be attributed to the attention modules in our architecture and DDPM's iterative refinement, which allows for progressive restoration of details and more precise recovery compared to CNNs' single-pass. Most importantly, it is demonstrated that despite the shortened TSMI and the thereby reduced encoding of the relaxometry processes, T1 and T2 parameter information are better restored. Figure3 presents the reconstruction uncertainty maps generated by IDDPM-MRF, showing a clear correlation with the reconstruction error maps.

Conclusions

We introduced a novel IDDPM-based model for reconstruction of parameter maps from further accelerated MRF data. Qualitative and quantitative results demonstrated IDDPM-MRF outperforms established algorithms for MRF reconstruction with respect to parameter accuracy and image quality. The proposed patch-wise approach enables more efficient training and inference pipeline.

Acknowledgements

This work has been carried out under the EPSRC grant EP/X001091/1

References

1. Dan Ma, Vikas Gulani, Nicole Seiberlich, Kecheng Liu, Jeffrey L Sunshine, Jeffrey L Duerk, and Mark A Griswold. Magnetic resonance fingerprinting. *Nature*, 495(7440):187–192, 2013.
2. Zhenghan Fang, Yong Chen, Mingxia Liu, Lei Xiang, Qian Zhang, Qian Wang, Weili Lin, and Dinggang Shen. 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.
3. Fabian Balsiger, Amaresha Shridhar Konar, Shivaprasad Chikop, Vimal Chandran, Olivier Scheidegger, Sairam Geethanath, and Mauricio Reyes. "Magnetic resonance fingerprinting reconstruction via spatiotemporal convolutional neural networks." In *Machine Learning for Medical Image Reconstruction: First International Workshop, MLMIR 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Proceedings 1*, pp. 39-46. Springer International Publishing, 2018.
4. Mohammad Golbabaee, Dongdong Chen, Pedro A. Gómez, Marion I. Menzel, and Mike E. Davies. "Geometry of deep learning for magnetic resonance fingerprinting." In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7825-7829. IEEE, 2019.

5. Refik Soyak, Ebru Navruz, Eda Ozgu Ersoy, Gastao Cruz, Claudia Prieto, Andrew P. King, Devrim Unay, and Ilkay Oksuz. "Channel attention networks for robust MR fingerprint matching." *IEEE Transactions on Biomedical Engineering* 69, no. 4 (2021): 1398-1405.
6. Ketan Fatania, Carolin M. Pirkl, Marion I. Menzel, Peter Hall, and Mohammad Golbabaee. "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.
7. Jonathan Ho, Ajay Jain, and Pieter Abbeel. "Denoising diffusion probabilistic models." *Advances in neural information processing systems* 33 (2020): 6840-6851.
8. Alexander Quinn Nichol, and Prafulla Dhariwal. "Improved denoising diffusion probabilistic models." In *International conference on machine learning*, pp. 8162-8171. PMLR, 2021.
9. Junde Wu, Rao Fu, Huihui Fang, Yu Zhang, Yehui Yang, Haoyi Xiong, Huiying Liu, and Yanwu Xu. "Medsegdiff: Medical image segmentation with diffusion probabilistic model." In *Medical Imaging with Deep Learning*, pp. 1623-1639. PMLR, 2024.
10. Jiaming Liu, Rushil Anirudh, Jayaraman J. Thiagarajan, Stewart He, K. Aditya Mohan, Ulugbek S. Kamilov, and Hyojin Kim. "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.
11. Mengting Huang, Thanh Nguyen-Duc, Martin Soellradl, Daniel F. Schmidt, and Roland Bammer. "PADMr: Patch-Based Denoising Diffusion Probabilistic Model for Magnetic Resonance Imaging Reconstruction." In *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 1-5. IEEE, 2024.
12. Yun Jiang, Dan Ma, Nicole Seiberlich, Vikas Gulani, and Mark A. Griswold. "MR fingerprinting using fast imaging with steady state precession (FISP) with spiral readout." *Magnetic resonance in medicine* 74, no. 6 (2015): 1621-1631.
13. Debra F. McGivney, Eric Pierre, Dan Ma, Yun Jiang, Haris Saybasili, Vikas Gulani, and Mark A. Griswold. "SVD compression for magnetic resonance fingerprinting in the time domain." *IEEE transactions on medical imaging* 33, no. 12 (2014): 2311-2322.
14. Mohammad Golbabaee, Guido Buonincontri, Carolin M. Pirkl, Marion I. Menzel, Bjoern H. Menze, Mike Davies, and Pedro A. Gómez. "Compressive MRI quantification using convex spatiotemporal priors and deep encoder-decoder networks." *Medical image analysis* 69 (2021): 101945.
15. Kai Zhang, Yawei Li, Wangmeng Zuo, Lei Zhang, Luc Van Gool, and Radu Timofte. "Plug-and-play image restoration with deep denoiser prior." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, no. 10 (2021): 6360-6376.

Figures

Method	MAPE (%)		RMSE (ms)		NRMSE			SSIM		
	T1	T2	T1	T2	T1	T2	TSMI	T1	T2	TSMI
SVDMRF	16.98	140.86	79.39	36.07	0.5566	1.8710	3.8360	0.9221	0.7547	0.7219
LRTV	15.94	35.96	76.62	15.40	0.5349	0.7353	1.2377	0.9376	0.8898	0.7932
DRUNet	8.93	21.07	18.28	9.39	0.1285	0.4458	0.5633	0.9693	0.9588	0.8388
SCQ	9.58	23.67	16.22	9.16	0.1143	0.4354	–	0.9750	0.9674	–
MRF-IDDPM (Ours)	7.20	17.69	14.13	7.67	0.1002	0.3687	0.4142	0.9758	0.9711	0.8729

Table 1 Reconstruction metrics for TSMIs (where applicable) and tissue parameter maps (T1/T2) for the evaluated algorithms. The evaluation dataset consisted of 24 axial brain slices from two different subjects. MRF-IDDPM reconstructions correspond to the average of ten realizations, each one from a different χ_T . Best reconstruction metrics are highlighted in bold.

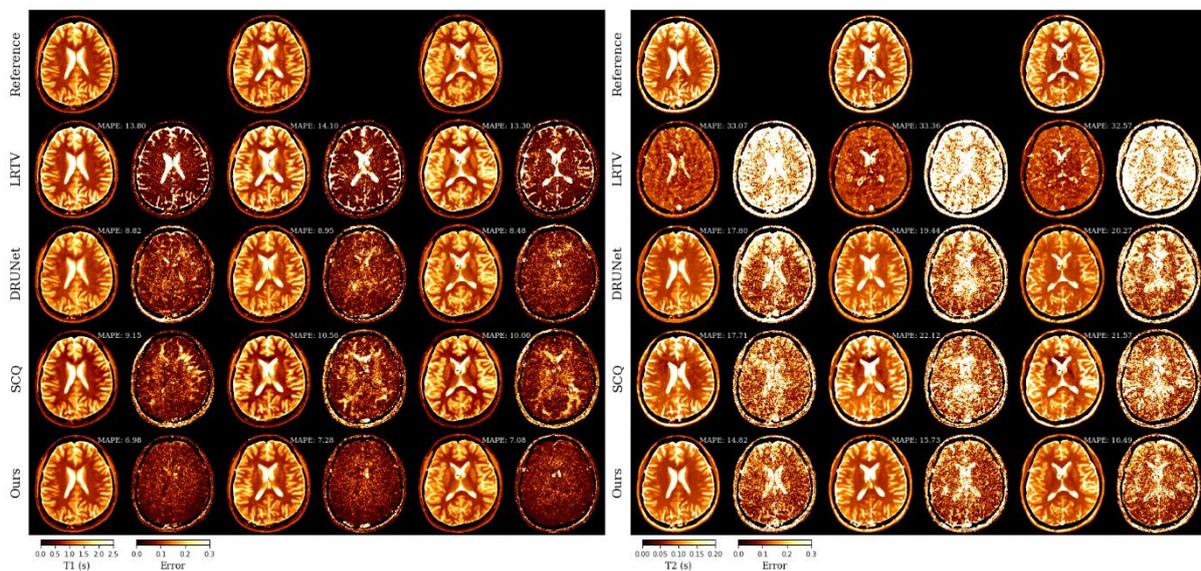


Figure 1. Reconstructed T1(left), T2(right) maps by our IDDPM method and baselines along with percentage error maps for three representative brain slices from the evaluation set. For our method, dictionary matching is used on the TSMIs computed from the pixel-wise average of 10 TSMI realizations of IDDPM.

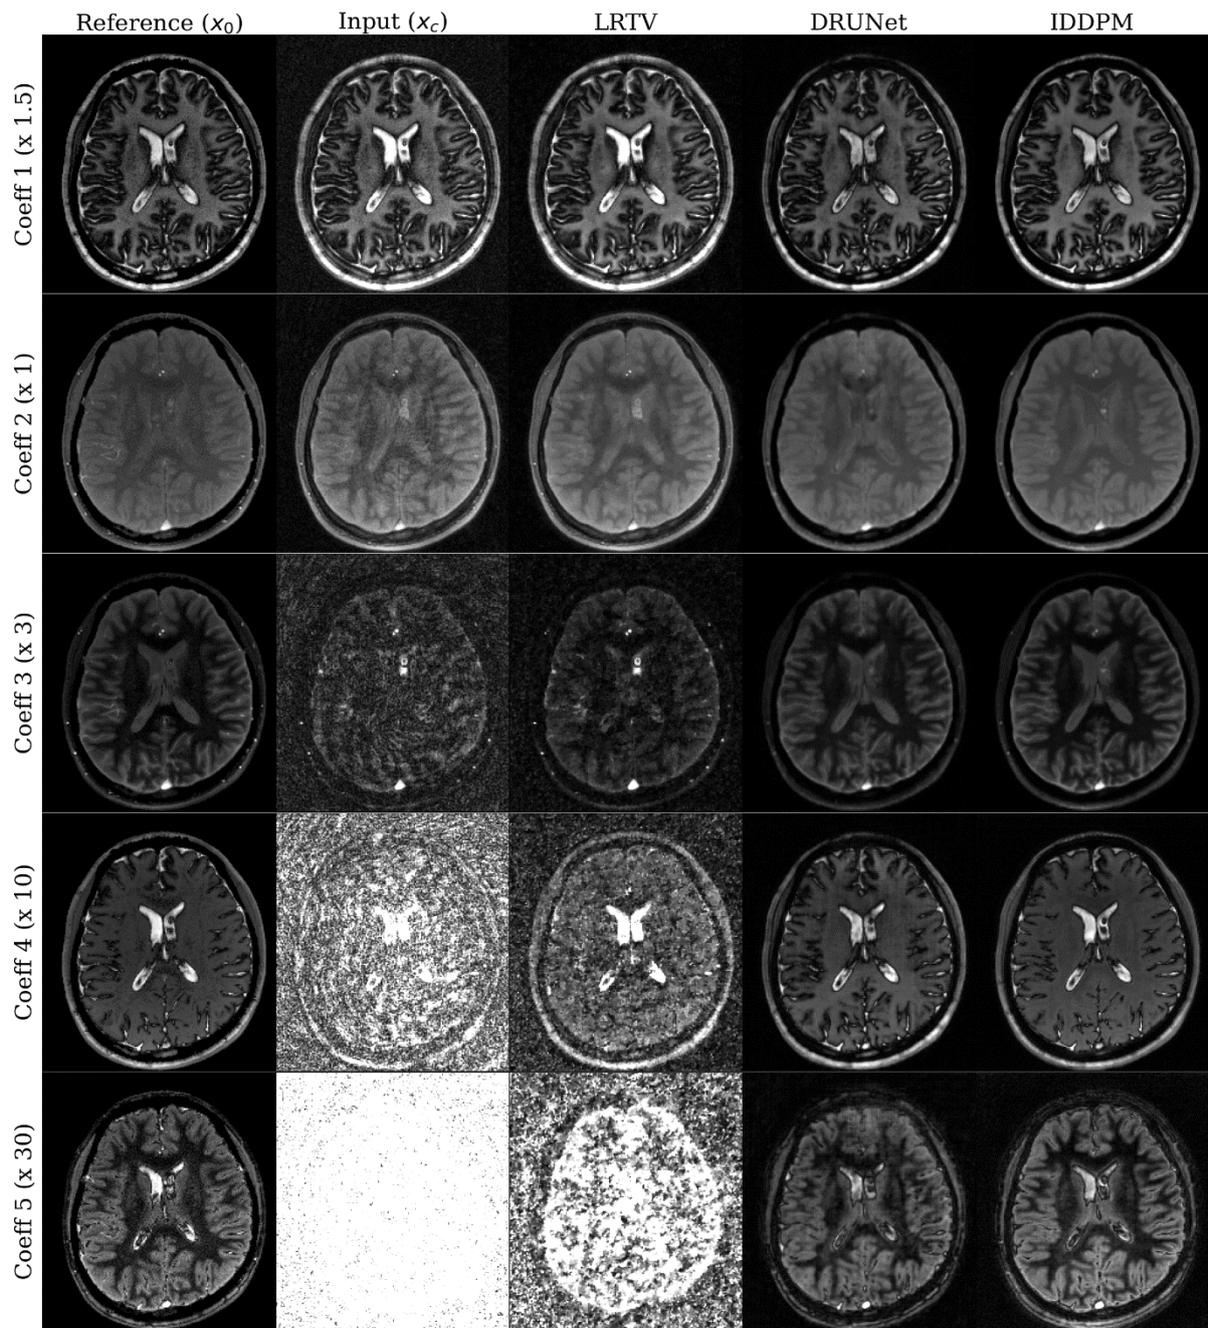


Figure 2. Magnitudes of the reference and reconstructed TSMIs (SVD compressed to 5 components from $l=200$) from a representative slice in the evaluation dataset, using different methods. IDDPM-MRF (also DRUNet) uses the gridding reconstruction from SVD-MRF as the input, generating restored TSMI outputs. Compare to baselines MRF-IDDPM notably improves the reconstruction of TSMI subspace components, especially in last/weakest SVD channels.

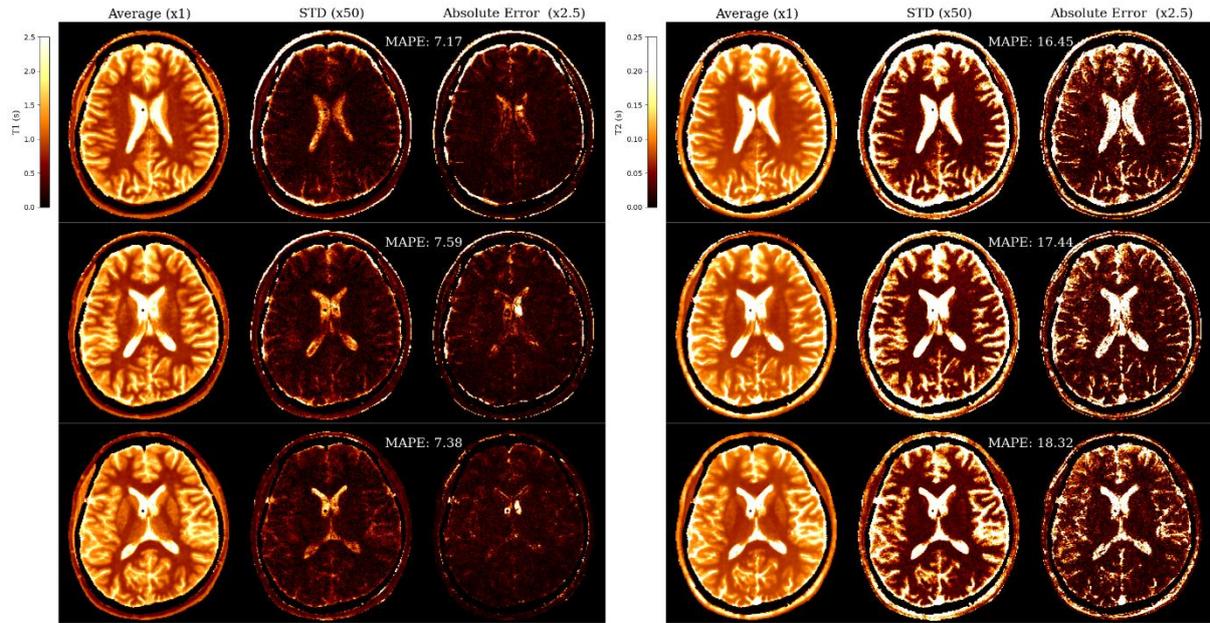


Figure 3 Uncertainty maps for T1 (left) T2 (right) reconstructions of three representative brain slices (rows) in evaluation set. Maps were obtained by ten runs of IDDPM-MRF with dictionary matching using the same condition x_c for different random x_r . Each subfigure shows the pixel-wise average (left), standard deviation (middle) and absolute error (right) between average and reference from Fig1. Higher uncertainties—correlated with greater errors—are mainly observed in tissues with longer relaxation times e.g. cerebrospinal fluid, while less so in white and grey matter areas.