

qMRI Diffuser

Quantitative T1 Mapping of the Brain Using a Denoising Diffusion Probabilistic Model

Wang, Shishuai; Ma, Hua; Hernandez-Tamames, Juan A.; Klein, Stefan; Poot, Dirk H.J.

DOI

[10.1007/978-3-031-72744-3_13](https://doi.org/10.1007/978-3-031-72744-3_13)

Publication date

2025

Document Version

Final published version

Published in

Deep Generative Models

Citation (APA)

Wang, S., Ma, H., Hernandez-Tamames, J. A., Klein, S., & Poot, D. H. J. (2025). qMRI Diffuser: Quantitative T1 Mapping of the Brain Using a Denoising Diffusion Probabilistic Model. In A. Mukhopadhyay, I. Oksuz, S. Engelhardt, D. Mehrof, & Y. Yuan (Eds.), *Deep Generative Models: 4th MICCAI Workshop, DGM4MICCAI 2024, Held in Conjunction with MICCAI 2024, Marrakesh, Morocco, October 10, 2024, Proceedings* (pp. 129-138). (Lecture Notes in Computer Science; Vol. 15224 LNCS). Springer. https://doi.org/10.1007/978-3-031-72744-3_13

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



qMRI Diffuser: Quantitative T1 Mapping of the Brain Using a Denoising Diffusion Probabilistic Model

Shishuai Wang¹(✉), Hua Ma¹, Juan A. Hernandez-Tamames^{1,2}, Stefan Klein¹,
and Dirk H. J. Poot¹

¹ Department of Radiology and Nuclear Medicine, Erasmus MC, Rotterdam,
The Netherlands

s.wang@erasmusmc.nl

² Department of Imaging Physics, TU Delft, Delft, The Netherlands

Abstract. Quantitative MRI (qMRI) offers significant advantages over weighted images by providing objective parameters related to tissue properties. Deep learning-based methods have demonstrated effectiveness in estimating quantitative maps from series of weighted images. In this study, we present qMRI Diffuser, a novel approach to qMRI utilising deep generative models. Specifically, we implemented denoising diffusion probabilistic models (DDPM) for T1 quantification in the brain, framing the estimation of quantitative maps as a conditional generation task. The proposed method is compared with the residual neural network (ResNet) and the recurrent inference machine (RIM) on both phantom and *in vivo* data. The results indicate that our method achieves improved accuracy and precision in parameter estimation, along with superior visual performance. Moreover, our method inherently incorporates stochasticity, enabling straightforward quantification of uncertainty. Hence, the proposed method holds significant promise for quantitative MR mapping.

Keywords: Quantitative MRI · Diffusion model · T1 mapping

1 Introduction

MRI is predominantly utilised for acquiring qualitative weighted images, which provide information regarding the contrast between various tissue types and are employed for visual assessment. However, compared with weighted images, quantitative MRI (qMRI) offers distinct advantages by measuring tissue properties, such as relaxation times and proton density. As the tissue properties are objective and intrinsic, they exhibit the promise as biomarkers for various diseases [4]. Furthermore, qMRI enables the direct comparisons of data acquired under diverse conditions, e.g., different patients, time periods, scanning protocols, and hardware configurations.

The generation of quantitative maps typically involves the acquisition of a series of weighted images with distinct settings (e.g. different inversion times),

from which the quantitative parameters can be estimated [14, 18]. Traditionally, the maximum likelihood estimator (MLE) [15] is used to estimate quantitative maps by iteratively fitting a physical signal model to the weighted images. However, MLE can suffer from high variability in scans with low signal-to-noise ratio (SNR) and requires substantial expertise for designing its regularisation term. Alternatively, convolutional neural networks (CNN) have been trained to directly translate weighted images into quantitative maps [3, 10, 19] and have proven effective. However, they suffer from low interpretability and may fail without notification. To leverage the advantages of both model-based and data-driven approaches, the recurrent inference machine (RIM) was proposed to iteratively reconstruct quantitative maps by training a recurrent neural network to predict the incremental update of each iteration step, achieving the state-of-the-art performance on T1 and T2 mapping tasks [17]. However, this approach becomes computationally expensive for tasks with complicated (e.g. Bloch simulation-based) physical signal models, as it requires evaluation of the derivative of the likelihood with respect to the tissue parameters.

As a class of deep generative models, diffusion models have gained popularity by its impressive image generation capability. There is also a burgeoning interest in applying diffusion models within the realm of medical image analysis, encompassing various tasks such as image generation, reconstruction, registration, classification, segmentation, denoising and detection [11]. Inspired by the success of diffusion models, in this work, we present qMRI Diffuser and approach quantitative MR mapping as a generation task performed by a denoising diffusion probabilistic model (DDPM) [9] to take the advantage of diffusion models' image generation quality. To the best of our knowledge, this is the first use of diffusion model for qMRI. In the proposed approach, the generation of quantitative maps starts with a Gaussian noise, and the results are obtained by the gradual denoising using a deep learning model. To estimate the quantitative parameters (e.g. T1) corresponding to a given series of weighted images, we set the weighted images as the condition that guides the generation process. In addition, given the inherent stochastic nature of the proposed method, uncertainty can be readily quantified by repeating the inference process [5].

To investigate the advantages of our proposed method in qMRI, we compare it with two established deep learning-based methods: 1) ResNet as a baseline, and 2) the RIM-based method [17], in a T1 mapping task of the brain. To address the lack of training data, we generated a synthetic dataset for the training of the three methods, while the evaluation was conducted using real scan data acquired from a hardware phantom and a healthy volunteer. Compared with ResNet and RIM, we found that the proposed method exhibits better accuracy and precision as well as superior visual performance. Moreover, our proposed method provides meaningful uncertainty information without specific treatments, which is particularly useful. To summarise, the following contributions are made in this paper:

- To the best of our knowledge, this is the first diffusion model-based method developed for qMRI applications

- Improved visual performance and parameter estimation quality compared to other deep learning-based methods, as showcased by phantom and *in vivo* testing data
- Improved trustworthiness and interpretability of estimation results by the straightforward uncertainty quantification ability of the proposed method

2 Method

DDPM defines a forward diffusion process and the reverse process. Assuming x_0 represents the clean original data sample, the forward process gradually introduces Gaussian noise ϵ to x_0 through a scheduler over a series of T steps, thereby reshaping the original distribution of x_0 into a Gaussian distribution of x_T . The reverse process starts from x_T , which is Gaussian noise, and uses a model parametrised by θ to reverse the noise adding steps as [1, 9, 11]

$$p_{\theta}(x_{0:T} | y) = p(x_T) \prod_{t=1}^T p_{\theta}(x_{t-1} | x_t, y) \quad (1)$$

where y is the condition to guide the reverse (i.e. generation) process. Specifically, in our task, we set x_0 as the clean quantitative maps, where the number of channels corresponds to the number of quantitative parameters to estimate, and we use weighted images as the condition. The quantification task is then executed through a series of the reverse processes $p_{\theta}(x_{t-1} | x_t, y)$, where t starts from T and decreases to 1, ultimately producing quantitative maps corresponding to the respective conditioning weighted images. The training of the model is conducted by minimising the following objective function [9]

$$L = \mathbb{E}_{t, x_0, \epsilon} \left[\|\epsilon - \epsilon_{\theta}(x_t, y, t)\|^2 \right] \quad (2)$$

where $\epsilon_{\theta}(x_t, y, t)$ is trained to approximate ϵ given the time step t , the noisy sample x_t , and the condition y as inputs. Here, we provide y to the model by concatenating it with x_t in the channel dimension. The workflow of parameter mapping using the trained model is illustrated in Fig. 1.

3 Experiments

3.1 Scanning Protocol

As a proof of concept for the proposed method, we deployed a inversion recovery fast spin echo (IR-FSE) protocol on a GE M750 scanner with a 16-channel head and neck coil to acquire real scan data. For each scan, 30 axial slices were acquired with a voxel size of $1 \times 1 \times 2 \text{ mm}^3$. We used $\text{TR} = 10 \text{ s}$, $\text{TE} = 7.14 \text{ ms}$, $\text{ETL} = 64$, $\text{FOV} = [250, 250, 60] \text{ mm}$, and acquisition matrix = $[240, 240]$. The inversion times (TI) were set as $[0.05, 0.10, 0.25, 0.50, 0.85, 1.50, 2.50]$ seconds to get seven




Fig. 1. Illustration of qMRI Diffuser’s workflow. A U-Net model is trained to predict the noise added at a given time step. Weighted images acquired at different contrast states are set as the condition for the generation of quantitative maps. The number of channels in y and x_t should correspond to the number of weighted images used and quantitative parameters of interest respectively, and specifically here we set them as 7 and 3.

weighted images at different contrast states, which were used for quantitative parameter mapping.

With the selected protocol, there are three quantitative parameters to estimate: T1, proton density and inversion efficiency, with T1 being of the most interest. Although inversion efficiency is mostly related to radio-frequency (RF) and does not directly reflect tissue properties, it could confound the T1 estimation. Therefore, we are interested in estimating it as well.

3.2 Datasets

The training dataset was synthesised following [17] and consists of paired quantitative maps and weighted images. The anatomy structure of quantitative maps were provided by 3D virtual discrete brain models from BrainWeb [6], and the parameter values were drawn from the realistic distributions of a given tissue type [2]. Subsequently, a series of weighted images were generated based on the synthetic quantitative maps and a forward signal model of the IR-FSE acquisition detailed in Sect. 3.1:

$$S_n = \left| PD \left(1 - Be^{-\frac{TIn}{T_1}} \right) \right| \quad (3)$$

where n is the index of weighed images, $S_n \in \{S_1, \dots, S_7\}$ represents the weighted image at inversion time TI_n . PD and B denotes proton density and inversion efficiency, respectively. In total 994 geometrically different axial slices were drawn from 14 brain models. For each slice, 4 realisations of quantitative parameter combination were generated, forming a training dataset with 3976 slices.

The testing dataset were acquired from: 1) a quantitative NIST/ISMRM phantom (CaliberMRI, Boulder, CO, USA), and 2) a healthy volunteer, with approval from the our institutional review board and informed consent from the volunteer obtained.

3.3 Network Training

The architecture and hyperparameters of ResNet and RIM were set as the same as in [17], except that we increased RIM’s number of inference steps from 6 to 9 to boost its performance. The qMRI Diffuser was deployed based on the open-source platform MONAI Generative Models [12]. Within the qMRI Diffuser, a U-Net with time step embedding [13, 16] was configured, comprising three levels with the number of channels set to 128, 256, and 256, respectively. The original DDPM noise scheduler was used [9], which means the expected range of quantitative maps is $[-1.0, 1.0]$. Hence, we scaled the training quantitative maps during pre-processing stage using $f(\mathbf{x}) = 2 \times \tanh(\mathbf{x}) - 1$. After training, the predicted quantitative maps were scaled back using the inverse of this function. All three methods were trained by using the same synthetic dataset with batch size of 8 and 100 epochs.

3.4 Evaluation

The phantom data and *in vivo* data were used to evaluate and compare ResNet, RIM and the proposed qMRI Diffuser. For qMRI Diffuser, we repeated the inference 10 times with different randomly initialised x_T and took the mean of the outputs as the final result. Meanwhile, the voxel-wise standard deviation of the 10 outputs can act as a measurement of uncertainty. The evaluation focused on the T1 and proton density (PD) channels, as inversion efficiency can not act as a biomarker. Visual evaluation was conducted on both phantom data and *in vivo* data, and T1 mapping results were also quantitatively evaluated using the phantom.

4 Results

Figure 2 shows a representative example of the T1 and PD estimation results for the phantom. With the used IR-FSE protocol, a typical artifact in weighted images is the blurring around the spheres. Similar visual effects are also observed in the T1 estimation from ResNet and PD estimations from all methods. For the T1 estimation from RIM, such artifacts did not directly cause blurring of the spheres but resulted in their distortion. All three methods exhibit blurring

around the phantom due to signal interference from the background, but the proposed method gives the clearest edge of the phantom. The PD estimations are visually comparable among all three methods, except that the estimation from ResNet is relatively lower in the water bolus.




Fig. 2. Example of T1 and PD estimation results in the phantom. A slice containing the $NiCl_2$ T1 spheres is presented.

The mean and standard deviation of T1 estimations in each T1-sphere are summarised in Table 1. Within the relevant T1 range typically found *in vivo* (i.e., from 400 to 2000 ms), the mean value of T1 estimations using qMRI Diffuser has the closest match to the ground truth values. Additionally, these estimations exhibit the smallest standard deviation in the ROIs. For T1 lower than 500 ms, qMRI Diffuser exhibits a bias towards values larger than the ground truth, with a standard deviation mostly smaller than that of RIM but larger than that of ResNet.

Figure 3 shows examples of T1 and PD estimations for *in vivo* data. In the T1 estimation, qMRI Diffuser shows the ability to preserve small anatomical structures, while blurriness around small structures are observed from both ResNet’s and RIM’s results. Within the brain mask, high uncertainty values given by qMRI Diffuser are observed in the high T1 region (above 2000 ms), while the uncertainty values in gray matter/white matter region is approximately homogeneous. The PD estimations from all methods are visually similar without outliers or noticeable artifacts, while RIM’s estimation is lower in the white matter and qMRI Diffuser’s estimation is slightly higher overall.

5 Discussion

According to the evaluation results, the proposed method trained purely on synthetic data generalises well to data from real scans, and shows resilience to

Table 1. Quantitative evaluation of T1 mapping (ms) on the phantom at room temperature. Ground truth (GT) values are from the manual (properties of T1 contrast spheres at 3T, 20 °C). The results for spheres with T1 values within the typical *in vivo* range are highlighted in bold.

GT	ResNet	RIM	Diffuser	GT	ResNet	RIM	Diffuser
1884 ± 30	1534 ± 72	1835 ± 39	1772 ± 36	175 ± 3	121 ± 13	186 ± 23	224 ± 14
1330 ± 20	4575 ± 1203	1365 ± 22	1353 ± 20	121 ± 2	124 ± 12	101 ± 31	176 ± 15
987 ± 14	2452 ± 418	966 ± 17	1003 ± 15	85 ± 1	89 ± 10	90 ± 140	161 ± 25
690 ± 10	1654 ± 415	683 ± 18	695 ± 16	60 ± 1	78 ± 11	1053 ± 62	139 ± 21
485 ± 7	617 ± 139	504 ± 15	503 ± 13	43 ± 1	60 ± 9	68 ± 98	116 ± 19
342 ± 5	175 ± 62	354 ± 16	370 ± 14	30 ± 1	66 ± 9	26 ± 18	100 ± 19
241 ± 3	139 ± 26	211 ± 18	263 ± 12	21 ± 1	117 ± 31	34 ± 41	118 ± 31

artifacts in the input weighted images. The test with phantom shows that qMRI Diffuser achieves the best visual results for T1 mapping by providing clean edges of spheres and uniform T1 values within the spheres, as well as higher accuracy and precision for T1 values in the range typically found *in vivo*. Though qMRI Diffuser’s performance for T1 values less than 500 ms is admittedly suboptimal, this could be attributed to the fact that such T1 values are rare in the synthetic human brain training dataset, and the fact that qMRI Diffuser is not physics model-based, which may limit the ability to extrapolate outside of the range of values seen during training. However, it could be improved by using a training dataset with a larger T1 range. The proposed method also provides superior visual performance on the *in vivo* data compared to ResNet and RIM, and the estimated T1 maps using qMRI Diffuser accurately depicted anatomical structures including the fine details.

One significant advantage of the proposed method is that it readily allows uncertainty quantification, which adds trustworthiness and interpretability to this method. For other methods, uncertainty quantification is either infeasible or dependent on separate operations (e.g. by adding dropout layers in the neural network [8]). The uncertainty quantification by using voxel-wise standard deviation of repeated inferences is meaningful, as the result reconstructed from a randomly initialised x_t can be viewed as a sample from the posterior distribution, given the training data the model has seen. It is noticeable that high-uncertainty region given by qMRI Diffuser overlaps with its error-prone areas (e.g. edge of spheres, high T1 regions), indicating the meaningfulness and usefulness of the uncertainty information provided by our method.

The proposed method effectively addresses the challenge of limited weighted images as demonstrated in the experiments, aligning with the objective of reduc-




Fig. 3. Example of T1 and PD estimation results on the *in vivo* data. The *in vivo* T1 maps are plotted following the colormap recommendations [7].

ing the scanning time to acquire quantitative maps. Compared to RIM which requires derivative computations through the signal model, the proposed method can be applied to other scanning protocols with more complicated signal models without introducing additional challenges or computational expenses as it does not require such computations, enhancing its usability across a wide range of scanning protocols.

6 Conclusion

We propose a novel quantitative MR mapping method, qMRI Diffuser, which frames the mapping process as a conditional generation task using a denoising diffusion probabilistic model (DDPM). The visual and quantitative evaluation on real scans demonstrates the superior performance of the proposed method, highlighting its potential in clinical use. Particularly, the proposed method allows readily uncertainty quantification. In conclusion, the proposed method could be a promising tool for qMRI.

Acknowledgments. This work was conducted within the “Trustworthy AI for MRI” ICAI lab within the project ROBUST, funded by the Dutch Research Council (NWO), GE Healthcare, and the Dutch Ministry of Economic Affairs and Climate Policy (EZK).

Disclosure of Interests. Shishuai Wang, Stefan Klein, Juan-Antonio Hernandez-Tamames and Dirk Poot received research grants from GE Healthcare. Hua Ma has no competing interests related to the scope of this article.

References

1. Batzolis, G., Stanczuk, J., Schönlieb, C.B., Etmann, C.: Conditional image generation with score-based diffusion models. arXiv preprint [arXiv:2111.13606](https://arxiv.org/abs/2111.13606) (2021)
2. Bojorquez, J.Z., Bricq, S., Acquitier, C., Brunotte, F., Walker, P.M., Lalande, A.: What are normal relaxation times of tissues at 3T? Magnetic Resonance Imaging **35**, 69–80 (2017)
3. Cai, C., Wang, C., Zeng, Y., Cai, S., Liang, D., Wu, Y., Chen, Z., Ding, X., Zhong, J.: Single-shot T2 mapping using overlapping-echo detachment planar imaging and a deep convolutional neural network. Magnetic Resonance in Medicine **80**(5), 2202–2214 (2018)
4. Cashmore, M.T., McCann, A.J., Wastling, S.J., McGrath, C., Thornton, J., Hall, M.G.: Clinical quantitative MRI and the need for metrology. The British Journal of Radiology **94**(1120), 20201215 (2021)
5. Chung, H., Ye, J.C.: Score-based diffusion models for accelerated mri. Medical Image Analysis **80**, 102479 (2022)
6. Cocosco, C.A.: Brainweb: Online interface to a 3D MRI simulated brain database. NeuroImage p. 425 (1997)
7. Fuderer, M., Crameri, F., deSouza, N., Baeßler, B., Gulani, V., Wang, M., de Boer, R., Cashmore, M., de Graaf, W., Keenan, K.E., Ma, D., Pirkil, C., Sollmann, N., Weingärtner, S., Mandija, S., Wichtmann, B., Golay, X.: Colormap recommendations for quantitative MRI. <https://www.cig-utrecht.org/blog/2023/11/27/colormap-relaxometry.html> (2023)
8. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In: Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48. p. 1050–1059. ICML’16, JMLR.org (2016)
9. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in Neural Information Processing Systems **33**, 6840–6851 (2020)
10. Jeelani, H., Yang, Y., Zhou, R., Kramer, C.M., Salerno, M., Weller, D.S.: A myocardial T1-mapping framework with recurrent and U-Net convolutional neural networks. In: 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI). pp. 1941–1944. IEEE (2020)
11. Kazerouni, A., Aghdam, E.K., Heidari, M., Azad, R., Fayyaz, M., Hacihaliloglu, I., Merhof, D.: Diffusion models in medical imaging: A comprehensive survey. Medical Image Analysis p. 102846 (2023)
12. Pinaya, W.H., Graham, M.S., Kerfoot, E., Tudosiu, P.D., Dafflon, J., Fernandez, V., Sanchez, P., Wolleb, J., da Costa, P.F., Patel, A., et al.: Generative AI for medical imaging: extending the monai framework. arXiv preprint [arXiv:2307.15208](https://arxiv.org/abs/2307.15208) (2023)

13. Pinaya, W.H., Tudosi, P.D., Dafflon, J., Da Costa, P.F., Fernandez, V., Nachev, P., Ourselin, S., Cardoso, M.J.: Brain imaging generation with latent diffusion models. In: MICCAI Workshop on Deep Generative Models. pp. 117–126. Springer (2022)
14. Pine, K., Kirilina, E.: Quantitative MRI and multiparameter mapping. In: *Advances in Magnetic Resonance Technology and Applications*, vol. 10, pp. 227–241. Elsevier (2023)
15. Ramos-Llorden, G., Den Dekker, A.J., Van Steenkiste, G., Jeurissen, B., Vanhevel, F., Van Audekerke, J., Verhoye, M., Sijbers, J.: A unified maximum likelihood framework for simultaneous motion and T1 estimation in quantitative MR T1 mapping. *IEEE Transactions on Medical Imaging* **36**(2), 433–446 (2016)
16. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10684–10695 (2022)
17. Sabidussi, E.R., Klein, S., Caan, M.W., Bazrafkan, S., den Dekker, A.J., Sijbers, J., Niessen, W.J., Poot, D.H.: Recurrent inference machines as inverse problem solvers for MR relaxometry. *Medical Image Analysis* **74**, 102220 (2021)
18. Shafieizargar, B., Byanju, R., Sijbers, J., Klein, S., den Dekker, A.J., Poot, D.H.: Systematic review of reconstruction techniques for accelerated quantitative MRI. *Magnetic Resonance in Medicine* (2023)
19. Shao, J., Ghodrati, V., Nguyen, K.L., Hu, P.: Fast and accurate calculation of myocardial T1 and T2 values using deep learning Bloch equation simulations (DeepBLESS). *Magnetic Resonance in Medicine* **84**(5), 2831–2845 (2020)