A Proper Strucutred Prior for Bayesian T_1 Mapping

Disi Lin[1*][0009−0001−9691−6042], Anders Garpebring[2][0000−0002−0532−232X], and Tommy Löfstedt[1][0000−0001−7119−7646]

[1] Department of Computing Science, Umeå University, Sweden
[2] Department of Diagnostics and Intervention, Umeå University, Sweden
*disil@cs.umu.se

Abstract. This work proposes a structured prior integrated within the Bayesian framework for variable flip angle T_1 mapping. The proposed structured prior combines total variation (TV) and ℓ_1 norm functions, and is proven to be a proper prior. The TV–ℓ_1 prior promotes sparsity in the spatial gradients of the parametric maps, resulting in smooth and coherent image reconstructions. Embedding the prior within the Bayesian framework enables uncertainty quantification for both T_1 and M_0 estimates. Posterior inference was performed using the No-U-Turn Sampler (NUTS). The proposed method is compared to maximum likelihood estimation and to alternative Bayesian models that employ uniform, Laplace, and bounded TV priors. The results show that the proposed method yields narrower probability density functions, indicating reduced uncertainty. The proposed method also achieves lower variance and exhibits a smaller negative bias, reflecting more stable estimates. Overall, the integration of TV and ℓ_1 functions in a prior within the Bayesian framework enhances spatial coherence in T_1 mapping and delivers improved uncertainty quantification, making it a promising tool for robust quantitative MRI parameter estimation.

Keywords: Bayesian inference · T_1 Mapping · Uncertainty quantification · Structured prior · Total variation.

1 Introduction

Magnetic resonance imaging (MRI) is a widely used medical imaging modality [20]. Compared to other imaging modalities, such as ultrasound or computed tomography (CT), MRI provides superior soft tissue contrast and eliminates the risks of exposure to ionizing radiation [18]. Despite its advantages, conventional MRI primarily provides qualitative images based on contrast differences, which limits its sensitivity to subtle or early pathological changes [12]. Quantitative MRI (qMRI), on the other hand, provides measurements in standardized physical units that allow for objective and reproducible tissue characterization, extending the capabilities of conventional MRI [9].

T_1 mapping is a qMRI technique that measures the longitudinal or spin-lattice relaxation time. T_1 mapping has, *e.g.*, been used to uncover pathological

processes in the brain [8] and myocardial tissue characterization [2]. Multiple T_1 mapping pulse sequences have been proposed [25], with inversion recovery widely regarded as the gold standard due to its accuracy [22]. Alternative methods have been proposed to improve efficiency, such as the Look-Locker (LL) technique [16], the Modified Look-Locker Inversion Recovery (MOLLI) sequence [17], and the Variable Flip Angle (VFA) method [4].

However, these T_1 mapping techniques are typically limited by long acquisition times and introduce systematic errors in the T_1 estimation [21]. An effective approach to accelerate qMRI is by undersampling the k-space, which results in underdetermined equation systems [23]. Compressed sensing that combines regularization, *e.g.*, based on low-rank constraint and/or total variation (TV), has been used to solve such ill-posed inverse problems. Zhang *et al.* [27] proposed a local low-rank constraint for VFA T_1 mapping, as well as multi-echo spin-echo T_2 mapping. Pandey *et al.* [19] proposed joint TV-based reconstruction for accelerated qMRI reconstruction. Le *et al.* [13] used recurrent neural networks with a cyclic, model-based loss for cardiac MOLLI T_1 mapping.

Uncertainty quantification is essential in medical image analysis, where estimated parameters often guide clinical decisions. The aforementioned methods follow a frequentist framework, estimating parameters via maximum likelihood—typically under Gaussian noise assumptions—reducing to (regularized) least-squares problems. While maximum likelihood estimation (MLE) offers computational efficiency and point estimates, it lacks measures of confidence or reliability, which are essential in risk-sensitive clinical contexts. In contrast, Bayesian inference provides a full posterior distribution over parameters and predictions, enabling quantification of confidence (credibility) and uncertainty. For instance, Beirinckx *et al.* [3] used a Bayesian approach with a TV prior for T_1 and T_2 mapping super-resolution. Löfstedt *et al.* [15] developed a Bayesian method based on a bounded TV prior for T_1 mapping that both reduced and estimated uncertainty in parameter maps. Huang *et al.* [11] proposed a Bayesian formulation that models the Wavelet coefficients of VFA-multi-echo images as Laplace distributed to achieve model-based T_1, T_2^*, and proton density map reconstructions.

TV regularization encourages smooth solutions by promoting sparsity in the spatial gradients of the coefficients, effectively favoring homogeneous parameter maps and sharp edges, and suppressing small, noisy fluctuations. In this work, we propose to combine both TV and the ℓ_1 norm to construct a proper prior. This approach retains the smoothness-promoting properties of TV while allowing full Bayesian inference with a well-defined posterior distribution. We applied this approach to VFA T_1 mapping within the Bayesian framework and evaluated the performance compared to MLE and alternative Bayesian models using uniform, Laplace, and bounded TV priors, defined in detail in Section 3.

2 Proposed Bayesian Model

VFA T_1 mapping quantifies T_1 values from repeated scans at different excitation flip angles [1]. The measured signal intensity S is a function of the longitudinal

relaxation time T_1, repetition time TR, flip angle α_i for $i = 1, 2, \ldots, I$, and the equilibrium longitudinal magnetization M_0, which is determined by the proton density and other factors [5]. The VFA signal model is defined as [4],

$$S_\theta(\alpha_i) = M_0 \frac{1 - e^{-\frac{\mathrm{TR}}{T_1}}}{1 - \cos(\alpha_i)\, e^{-\frac{\mathrm{TR}}{T_1}}} \sin \alpha_i, \tag{1}$$

where $\theta = (T_1, M_0)$ contains the parameters T_1 and M_0, which vary for tissues with different properties. However, residual unspoiled transverse magnetization may yield discrepancies between measured signal intensities and the theoretical values modeled in Equation (1) [1].

The measurement in the n-th pixel is thus modeled as,

$$y_n^{(i)} = S_{\theta_n}(\alpha_i) + \epsilon_n^{(i)}, \tag{2}$$

where the additive noise term $\epsilon_n^{(i)}$ is assumed to follow an independent zero-mean Normal distribution with variance σ^2, *i.e.*, $\epsilon_n^{(i)} \sim \mathcal{N}(0, \sigma^2)$. The parameters, $\theta_n = (T_1, M_0)$, represents the tissue properties in voxel n, with $n = 1, 2, \ldots, N$, for N the total number of voxels.

Assuming the measurements $y_n^{(i)} \sim \mathcal{N}(S_{\theta_n}(\alpha_i), \sigma^2)$ are independent, the probability of observing the data given the model parameters, *i.e.*, the likelihood function, is

$$p(\mathbf{y} \mid \boldsymbol{\theta}, \sigma^2) = (\sqrt{2\pi\sigma^2})^{-NI} \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^{I} \sum_{n=1}^{N} \left(y_n^{(i)} - S_{\theta_n}(\alpha_i) \right)^2 \right). \tag{3}$$

2.1 A Structured Prior Based on Total Variation

Neighboring voxels often correspond to the same underlying tissue, which suggests that the T_1 and M_0 maps are likely to exhibit spatial smoothness. To encourage this, we consider a structured prior based on TV. A TV prior can be defined using the Boltzmann distribution (Gibbs distribution) as

$$p_{\mathrm{TV}}(\mathbf{x}) = \frac{1}{Z} e^{-\lambda \mathrm{TV}(\mathbf{x})}, \tag{4}$$

where Z is a normalizing constant and $\mathbf{x}$ is an image. Specifically, this work considered the anisotropic TV [6], expressed as,

$$\mathrm{TV}(\mathbf{x}) = \|\boldsymbol{D}\bar{\mathbf{x}}\|_1,$$

where $\bar{\mathbf{x}} \in \mathbb{R}^N$ denotes a vectorized (flattened) version of $\mathbf{x}$, and $\boldsymbol{D}$ is a linear difference operator, that computes finite differences to approximate the spatial gradient in $\mathbf{x}$.

The distribution in Equation (4) results in high probabilities for sparse spatial changes in the model parameters. However, the normalizing constant, Z, diverges due to invariance in the direction of $\mathbb{1}$, *i.e.*, for constant images. To address this, we introduce an additional ℓ_1 norm to the TV function. We show that the resulting combination defines a proper prior distribution.

Lemma 1. *Let* TV_μ *be a modified TV function, such that*

$$\mathrm{TV}_\mu(\mathbf{x}) = \mathrm{TV}(\mathbf{x}) + \mu\|\bar{\mathbf{x}}\|_1, \tag{5}$$

where $\mu > 0$ *and* $\|\cdot\|_1$ *is the* ℓ_1 *norm, Then*

$$\int e^{-\lambda \mathrm{TV}_\mu(\mathbf{x})}\, d\bar{\mathbf{x}} < \infty,$$

for all $\lambda > 0$.

Proof. For any $\mathbf{x}$ we have that $\lambda\|\boldsymbol{D}\bar{\mathbf{x}}\|_1 \geq 0$, so

$$\begin{aligned}
\int e^{-\lambda \mathrm{TV}_\mu(\mathbf{x})}\, d\bar{\mathbf{x}} &= \int_{\mathbb{R}^N} e^{-\lambda\|\boldsymbol{D}\bar{\mathbf{x}}\|_1 - \lambda\mu\|\bar{\mathbf{x}}\|_1}\, d\bar{\mathbf{x}} \\
&\leq \int_{\mathbb{R}^N} e^{-\lambda\mu\|\bar{\mathbf{x}}\|_1}\, d\bar{\mathbf{x}} \\
&= \prod_{n=1}^{N} \int_{-\infty}^{\infty} e^{-\lambda\mu|x_n|}\, dx_n \\
&= \left(\frac{2}{\lambda\mu}\right)^N < \infty.
\end{aligned}$$

Lemma 1 guarantees the existence of a finite normalizing constant, which allows us to formulate the following theorem.

Theorem 1. *The modified prior,*

$$p_{\mathrm{TV}_\mu}(\mathbf{x}) = \frac{1}{Z} e^{-\lambda \mathrm{TV}_\mu(\mathbf{x})},$$

is a proper prior.

2.2 Posterior Distribution for VFA T_1 mapping

A Bayesian model combines a likelihood, such as that in Equation (3), with a prior distribution, that accounts for available prior information or any other prior beliefs. Specifically, a Bayesian model assigns a posterior distribution over the parameters, constructed using a likelihood and a prior, as

$$p(\boldsymbol{\theta} \mid \mathbf{y}) = \frac{p(\mathbf{y} \mid \boldsymbol{\theta})p(\boldsymbol{\theta})}{\int p(\mathbf{y} \mid \boldsymbol{\theta})p(\boldsymbol{\theta})d\boldsymbol{\theta}} \propto p(\mathbf{y} \mid \boldsymbol{\theta})p(\boldsymbol{\theta}), \tag{6}$$

where $p(\boldsymbol{\theta} \mid \mathbf{y})$, $p(\mathbf{y} \mid \boldsymbol{\theta})$ and $p(\boldsymbol{\theta})$ are the posterior distribution, the likelihood, and the prior distribution, respectively. The proportionality is in the parameters, $\boldsymbol{\theta}$. Since it is typically impossible or at least computationally intractable to compute the integral in the denominator, an efficient approach is to draw samples from the posterior using Markov chain Monte Carlo (MCMC) algorithms [10].

Using the modified TV_μ prior and a half-normal prior with a hyper-parameter σ_ϵ over σ, we define the posterior as,

$$\begin{aligned} p_{\mathrm{TV}_\mu}(\boldsymbol{\theta}, \sigma \mid \mathbf{y}) &\propto p(\mathbf{y} \mid \boldsymbol{\theta}, \sigma) p(\boldsymbol{\theta}, \sigma) = p(\mathbf{y} \mid \boldsymbol{\theta}, \sigma) p_{\mathrm{TV}_\mu}(\boldsymbol{\theta}) p(\sigma) \\ &= \left(\sigma\sqrt{2\pi}\right)^{-NI} \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^{I} \sum_{n=1}^{N} \left(y_n^{(i)} - S_{\theta_n}(\alpha_i)\right)^2\right) \\ &\quad \cdot \prod_{\boldsymbol{\theta} \in (T_1, M_0)} \frac{1}{Z_{\boldsymbol{\theta}}} e^{-\lambda \mathrm{TV}_\mu(\boldsymbol{\theta})} \cdot \frac{\sqrt{2}}{\sigma_\epsilon \sqrt{\pi}} \exp\left(-\frac{\sigma^2}{2\sigma_\epsilon^2}\right), \end{aligned}$$

and thus sample from this posterior distribution using MCMC.

3 Experiments

We evaluated the TV_μ prior for VFA T_1 mapping within a Bayesian framework, comparing its performance to that of MLE and to alternative Bayesian models using uniform, Laplace, and bounded TV priors. The hyperparameters λ and μ in the TV_μ prior were selected either by assigning hyperpriors—referred to as the Hierarchical TV_μ model—or through hyperparameter tuning, denoted simply as the TV_μ model.

When using hyper-priors, the λ, μ should be non-negative and were therefore both assigned to be $\Gamma(3, 1)$ distributed. When the hyper-parameters were found using a grid search, λ and μ were both independently explored over a predefined set of values, namely over $\{10^{-3}, 10^{-2}, 10^{-1}, 1, 10, 10^2, 10^3\}$. During the search, the best parameter combination was selected as the one that maximized the widely applicable information criterion (WAIC) [26], which estimates the expected log point-wise predictive density.

MLE formulates the estimation as a voxel-wise optimization problem that minimizes the least squares error (over a feasible parameter set), expressed as,

$$\begin{aligned} \hat{\boldsymbol{\theta}}_{\mathrm{ML}} = \arg\max_{\boldsymbol{\theta}} \quad & \sum_{i=1}^{I} \sum_{n=1}^{N} -\frac{1}{2\sigma^2} \left(y_n^{(i)} - S_{\theta_n}(\alpha_i)\right)^2 - NI \log(\sigma\sqrt{2\pi}) \\ \text{subject to} \quad & 0 \le \theta_n \le 50, \quad \forall\, n \end{aligned} \tag{7}$$

which was solved using the L-BFGS-B algorithm [14]. The baseline priors used in the experiments were Uniform$(0, 50)$, Laplace$(4, 1)$, and a bounded TV prior. Specifically, the uniform prior serves as a noninformative baseline, assigning equal probability to all T_1 values within $[0, 50]$ seconds—a conservative range covering physiologically plausible values. The Laplace prior acts as a shrinkage prior and keeps values from growing unrealistically. The bounded TV prior was introduced in [15] and is defined as

$$p_{b\mathrm{TV}}(\mathbf{x}) = p_{\mathrm{TV}}(\mathbf{x})\, p_{\mathrm{uni}}(\mathbf{x} \mid L, H), \quad \text{with} \quad L = 0 \text{ and } H = 50,$$

where the uniform prior, $p_{\mathrm{uni}}(\mathbf{x} \mid L, H)$, imposes hard bounds on the parameters' range. The hyper-parameter λ included in the bounded TV prior was selected through hyper-parameter search over the same set of values as for the TV_μ prior.

The baselines and the proposed model with the TV_μ prior were developed using PyMC 5.15.0[3]. Samples were drawn from the posterior distributions using the No-U-Turn Sampler (NUTS) algorithm [10]. Four independent Markov chains were run. Convergence was evaluated by inspecting trace plots and calculating the $\hat{R}$ statistic, which assesses the consistency between and within chains to confirm reliable mixing and sampling [7]. Specifically, each chain generates 2,000 samples after discarding 2,000 burn-in iterations.

The data used in the experiments were brain MR images simulated using Equations 1 and 2 with $TR = 6.8$ ms, $TE = 2.1$ ms, and flip angles $\alpha \in \{2°, 4°, 11°, 13°, 15°\}$. The ground truth values of T_1 and M_0 were generated from the BrainWeb phantom[4] using Hero Imaging[5]. The parameter maps generated had a matrix size of 256×256 and a voxel size of $0.98 \times 0.98 \times 2.00$ mm^3. The added noise was complex circular Gaussian noise following the work by Löfstedt *et al.* [15].

Results from different methods were evaluated by analyzing the estimated probability density functions (PDFs), as well as the bias and variance of the posterior samples. Specifically, PDFs were estimated voxel-wise using kernel density estimation with Gaussian kernels [24], based on 8,000 posterior samples per method. The uncertainty associated with each estimate was assessed by computing the credible interval spanning two standard deviations around the posterior mean. Bias and variance were also used for comparisons, and were estimated as,

$$\text{Bias} = \frac{1}{S}\sum_{s=1}^{S}(\mathbf{x}^s - \mathbf{x}), \tag{8}$$

$$\text{Varance} = \frac{1}{S-1}\sum_{s=1}^{S}\left(\mathbf{x}^s - \hat{\boldsymbol{\mu}}\right)^2, \quad \text{with} \quad \hat{\boldsymbol{\mu}} = \frac{1}{S}\sum_{s=1}^{S}\mathbf{x}^s, \tag{9}$$

where $\mathbf{x}$ denotes the ground truth and $\mathbf{x}^s$ is the s-th posterior sample.

4 Results and Discussion

This section presents the estimated PDFs of the T_1 values for six representative voxels, as illustrated in Figure 1, along with corresponding summary statistics reported in Table 1. The bias and variance for each method, computed using Equations 8 and 9, are visualized in Figure 2 as spatial maps and in Figure 3 as comparison metrics across all voxels.

Figure 1 illustrates that all Bayesian methods produce posterior PDFs that are generally centered around the ground truth (GT) T_1 values. However, the spread and peakedness of these distributions vary across methods. The proposed TV_μ model yields notably narrower PDFs, reflecting increased confidence and

[3] https://www.pymc.io
[4] https://brainweb.bic.mni.mcgill.ca/brainweb/
[5] https://www.heroimaging.com/

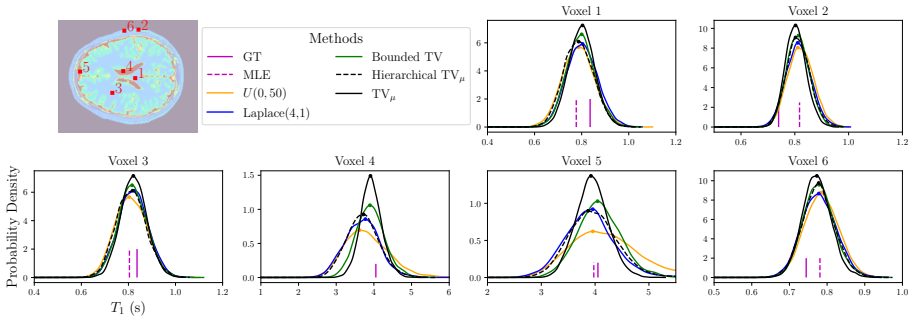

Fig. 1. PDFs of T_1 values at selected voxels (highlighted in red on the anatomical image, top-left). The estimated PDFs from different Bayesian methods are shown, along with the MLEs and corresponding GT T_1 values for comparison.

Methods	Voxel 1	Voxel 2	Voxel 3	Voxel 4	Voxel 5	Voxel 6
GT	0.84	0.74	0.84	4.05	4.05	0.74
MLE	0.78	0.82	0.80	43.70	3.98	0.78
$U(0,50)$	0.79, 0.79 ± 0.14	0.81, 0.82 ± 0.10	0.80, 0.81 ± 0.13	3.62, 3.83 ± 1.20	3.96, 4.20 ± 1.33	0.78, 0.79 ± 0.09
Laplace(4, 1)	0.80, 0.80 ± 0.13	0.81, 0.82 ± 0.09	0.83, 0.82 ± 0.13	3.78, 3.69 ± 0.93	3.96, 3.91 ± 0.93	0.78, 0.78 ± 0.09
Bounded TV	0.80, 0.80 ± 0.12	0.81, 0.81 ± **0.08**	0.82, 0.82 ± 0.12	3.90, 3.91 ± 0.76	4.05, 4.08 ± 0.81	0.78, 0.78 ± **0.08**
Hierarchical TV$_\mu$	0.79, 0.79 ± 0.13	0.80, 0.81 ± 0.09	0.82, 0.81 ± 0.13	3.72, 3.71 ± 0.87	3.87, 3.97 ± 0.90	0.78, 0.78 ± 0.09
TV$_\mu$	0.80, 0.80 ± **0.11**	0.80, 0.80 ± **0.08**	0.82, 0.82 ± **0.11**	3.91, 3.93 ± **0.57**	3.92, 3.94 ± **0.60**	0.77, 0.77 ± **0.08**

Table 1. Quatitative comparison of results for each highlighted voxel in Figure 1. The table lists the GT T_1 values and the MLEs, alongside the results from different Bayesian methods, as indicated in the left column. For each Bayesian model, the table reports the posterior mode, the posterior mean, and the associated uncertainty expressed as two standard deviations around the mean, in the format "mode, mean ± 2σ". All T_1 values are reported in seconds (s).

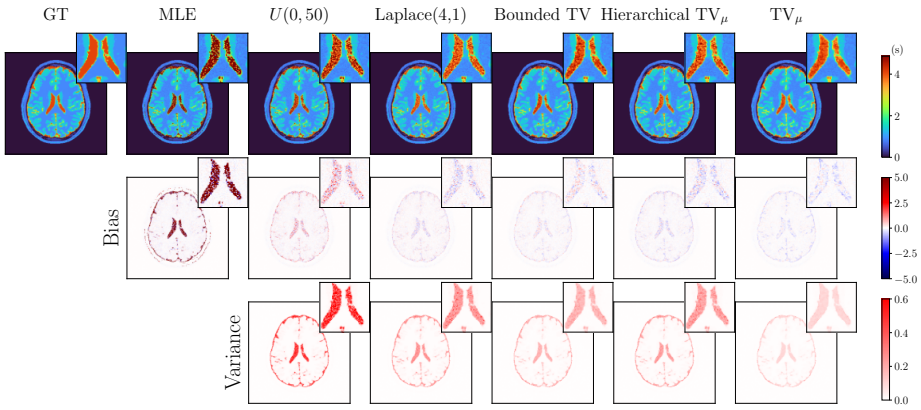

Fig. 2. Comparison of T_1, bias, and variance maps from different methods. The first row displays the GT T_1 map, the MLE, and the average T_1 maps obtained from each Bayesian model, as labeled at the top. The second row presents the corresponding bias maps, while the third row shows the variance maps. Each subplot includes a zoomed-in view of the central brain region in the top-right corner for detailed comparison.

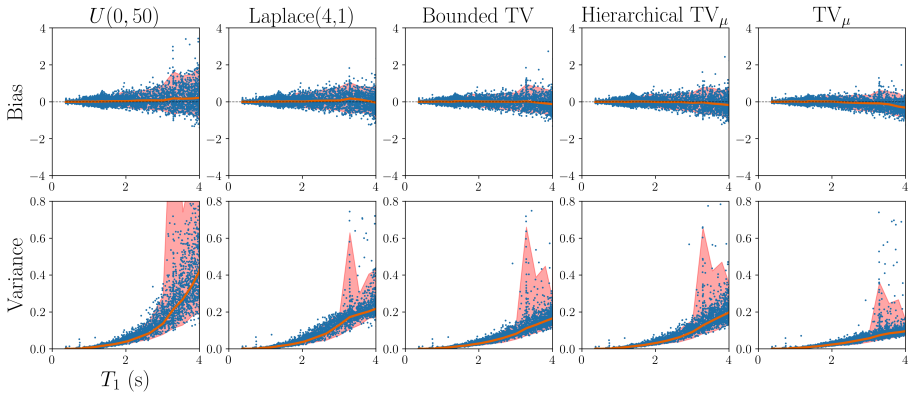

Fig. 3. Bias and variance analysis for the Bayesian methods. Unlike Figure 2, which presents spatially resolved maps, this figure illustrates the relationship between GT T_1 values and the corresponding bias or variance of the estimates. The horizontal axis shows the GT T_1 values, while the vertical axis indicates the bias/variance observed at each voxel. The orange curves represent the average bias/variance for a certain GT T_1 value. The red-shaded area contains 95% of the data points.

reduced uncertainty in the T_1 estimates compared to the baseline models. This qualitative observation is further substantiated by the quantitative results in Table 1, where the TV$_\mu$ model demonstrates a substantial improvement in estimation accuracy, effectively reducing both bias and credible interval widths across the selected voxels.

Figure 2 shows that the average T_1 maps obtained using the TV$_\mu$ prior appear smoother than those obtained with other priors. This observation is consistent with the bias maps. Additionally, the variance maps indicate a lower variance when using the TV$_\mu$ prior. Figure 3 confirms that samples from the TV$_\mu$ model exhibit a smaller negative bias, particularly at higher T_1 values, than other methods. Moreover, the proposed method consistently shows significantly lower variance. This trend is also evident in Table 1.

5 Conclusion

This work proposed a structured prior based on the combination of the TV and ℓ_1 functions. The proposed TV$_\mu$ prior was proven to be a proper prior and was used for VFA T_1 mapping in the Bayesian framework, which enabled uncertainty quantification in the reconstructions. The proposed method was compared to MLE and alternative Bayesian models. The results show that the proposed model provides smoother results and narrower credible intervals. As the evaluation was conducted exclusively on synthetic data, the results may not fully capture the complexity of clinical imaging scenarios. Future work will involve validating the method on real datasets and investigating the impact of hyperpriors on model performance.

Acknowledgments. This work was supported by the Swedish Research Council (Vetenskapsrådet; grant number 2021-04810); and Lion's Cancer Research Foundation in Northern Sweden (grant numbers LP 22-2319 and LP 24-2367).

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

References

1. Baudrexel, S., Nöth, U., Schüre, J.R., Deichmann, R.: T1 mapping with the variable flip angle technique: a simple correction for insufficient spoiling of transverse magnetization. Magnetic Resonance in Medicine **79**(6) (2018)
2. Becker, K.M., Schulz-Menger, J., Schaeffter, T., Kolbitsch, C.: Multi-parametric cardiac MRI for T1 mapping and cine imaging using iterative model-based image reconstruction. In: Proceedings of the International Society for Magnetic Resonance in Medicine (ISMRM 2017). vol. 25 (2017)
3. Beirinckx, Q., Jeurissen, B., Nicastro, M., Poot, D.H., Verhoye, M., den Dekker, A.J., Sijbers, J.: Model-based super-resolution reconstruction with joint motion estimation for improved quantitative MRI parameter mapping. Computerized Medical Imaging and Graphics **100** (2022)
4. Christensen, K.A., Grant, D.M., Schulman, E.M., Walling, C.: Optimal determination of relaxation times of Fourier transform nuclear magnetic resonance. determination of spin-lattice relaxation times in chemically polarized species. The Journal of Physical Chemistry **78**(19) (1974)
5. Deoni, S.C., Rutt, B.K., Peters, T.M.: Rapid combined T1 and T2 mapping using gradient recalled acquisition in the steady state. Magnetic Resonance in Medicine **49**(3) (2003)
6. Esedoḡlu, S., Osher, S.J.: Decomposition of images by the anisotropic Rudin-Osher-Fatemi model. Communications on Pure and Applied Mathematics **57**(12) (2004)
7. Gelman, A., Rubin, D.B.: Inference from iterative simulation using multiple sequences. Statistical Science **7**(4) (1992)
8. Gräfe, D., Frahm, J., Merkenschlager, A., Voit, D., Hirsch, F.W.: Quantitative T1 mapping of the normal brain from early infancy to adulthood. Pediatric Radiology **51** (2021)
9. Granziera, C., Wuerfel, J., Barkhof, F., Calabrese, M., De Stefano, N., Enzinger, C., Evangelou, N., Filippi, M., Geurts, J.J., Reich, D.S., et al.: Quantitative magnetic resonance imaging towards clinical application in multiple sclerosis. Brain **144**(5) (2021)
10. Hoffman, M.D., Gelman, A., et al.: The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. Journal of Machine Learning Research **15**(1) (2014)
11. Huang, S., Lah, J.J., Allen, J.W., Qiu, D.: Accelerated model-based T1, T2* and proton density mapping using a Bayesian approach with automatic hyperparameter estimation. Magnetic Resonance in Medicine **93**(2) (2025)
12. Lavrova, E., Lommers, E., Woodruff, H.C., Chatterjee, A., Maquet, P., Salmon, E., Lambin, P., Phillips, C.: Exploratory radiomic analysis of conventional vs. quantitative brain MRI: toward automatic diagnosis of early multiple sclerosis. Frontiers in Neuroscience **15** (2021)

13. Le, J.V., Mendes, J.K., McKibben, N., Wilson, B.D., Ibrahim, M., DiBella, E.V., Adluru, G.: Accelerated cardiac T1 mapping with recurrent networks and cyclic, model-based loss. Medical Physics **49**(11) (2022)
14. Liu, D.C., Nocedal, J.: On the limited memory BFGS method for large scale optimization. Mathematical Programming **45**(1) (1989)
15. Löfstedt, T., Hellström, M., Bylund, M., Garpebring, A.: Bayesian non-linear regression with spatial priors for noise reduction and error estimation in quantitative MRI with an application in T1 estimation. Physics in Medicine & Biology **65**(22) (2020)
16. Look, D.C., Locker, D.R.: Time saving in measurement of NMR and EPR relaxation times. Review of Scientific Instruments **41**(2) (1970)
17. Messroghli, D.R., Radjenovic, A., Kozerke, S., Higgins, D.M., Sivananthan, M.U., Ridgway, J.P.: Modified look-locker inversion recovery (MOLLI) for high-resolution T1 mapping of the heart. Magnetic Resonance in Medicine **52**(1) (2004)
18. Owrangi, A.M., Greer, P.B., Glide-Hurst, C.K.: MRI-only treatment planning: benefits and challenges. Physics in Medicine & Biology **63**(5) (2018)
19. Pandey, S., Snider, A.D., Moreno, W.A., Ravi, H., Bilgin, A., Raghunand, N.: Joint total variation-based reconstruction of multiparametric magnetic resonance images for mapping tissue types. NMR in Biomedicine **34**(12) (2021)
20. Poldrack, R.A., Mumford, J.A., Nichols, T.E.: Handbook of functional MRI data analysis. Cambridge University Press (2024)
21. Preibisch, C., Deichmann, R.: Influence of RF spoiling on the stability and accuracy of T1 mapping based on spoiled FLASH with varying flip angles. Magnetic Resonance in Medicine **61**(1) (2009)
22. Pykett, I., Rosen, B., Buonanno, F., Brady, T.: Measurement of spin-lattice relaxation times in nuclear magnetic resonance imaging. Physics in Medicine & Biology **28**(6) (1983)
23. 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 **90**(3) (2023)
24. Silverman, B.W.: Density estimation for statistics and data analysis. Routledge (2018)
25. Tirkes, T., Zhao, X., Lin, C., Stuckey, A.J., Li, L., Giri, S., Nickel, D.: Evaluation of variable flip angle, MOLLI, SASHA, and IR-SNAPSHOT pulse sequences for T_1 relaxometry and extracellular volume imaging of the pancreas and liver. Magnetic Resonance Materials in Physics, Biology and Medicine **32** (2019)
26. Watanabe, S., Opper, M.: Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. Journal of Machine Learning Research **11**(12) (2010)
27. Zhang, T., Pauly, J.M., Levesque, I.R.: Accelerating parameter mapping with a locally low rank constraint. Magnetic Resonance in Medicine **73**(2) (2015)