

RESEARCH ARTICLE

Magnetic Resonance in Medicine

Denoising and uncertainty estimation in parameter mapping with approximate Bayesian deep image priors

Max Hellström¹  | Tommy Löfstedt^{1,2}  | Anders Garpebring¹ 

¹Department of Radiation Sciences,
Umeå University, Umeå, Sweden

²Department of Computing Science,
Umeå University, Umeå, Sweden

Correspondence

Max Hellström, Department of Radiation
Sciences, Umeå University, Umeå,
Sweden.

Email: max.hellstrom@umu.se

Funding information

Cancerforskningsfonden i Norrland,
Grant/Award Number: AMP 18-912;
Lions Cancerforskningsfond,
Grant/Award Numbers: LP 18-2182, LP
22-2319; Region Västerbotten (Bas ALF),
Grant/Award Number: RV-970119; The
Swedish Research Council, Grant/Award
Number: 2019-0432

Abstract

Purpose: To mitigate the problem of noisy parameter maps with high uncertainties by casting parameter mapping as a denoising task based on Deep Image Priors.

Methods: We extend the concept of denoising with Deep Image Prior (DIP) into parameter mapping by treating the output of an image-generating network as a parametrization of tissue parameter maps. The method implicitly denoises the parameter mapping process by filtering low-level image features with an untrained convolutional neural network (CNN). Our implementation includes uncertainty estimation from Bernoulli approximate variational inference, implemented with MC dropout, which provides model uncertainty in each voxel of the denoised parameter maps. The method is modular, so the specifics of different applications (e.g., T1 mapping) separate into application-specific signal equation blocks. We evaluate the method on variable flip angle T1 mapping, multi-echo T2 mapping, and apparent diffusion coefficient mapping.

Results: We found that deep image prior adapts successfully to several applications in parameter mapping. In all evaluations, the method produces noise-reduced parameter maps with decreased uncertainty compared to conventional methods. The downsides of the proposed method are the long computational time and the introduction of some bias from the denoising prior.

Conclusion: DIP successfully denoise the parameter mapping process and applies to several applications with limited hyperparameter tuning. Further, it is easy to implement since DIP methods do not use network training data. Although time-consuming, uncertainty information from MC dropout makes the method more robust and provides useful information when properly calibrated.

KEYWORDS

deep image prior, denoising, parameter mapping, quantitative MRI, uncertainty estimation

1 | INTRODUCTION

In MRI, the image intensity depends on different tissue parameters of the imaged subject. The degree of dependence, and thus also the image contrast, can be altered to best suit the purpose of the imaging by selecting different scanner-parameter such as pulse repetition time (TR) and echo time (TE).

Quantitative MRI (QMRI) is a subset of MRI that involves quantifying these tissue parameters.¹ QMRI methods require a series of MRI images from which one isolates the tissue parameters using modelled relationships between the tissue parameters and the MRI signal. This process is often referred to as *parameter mapping* since the resulting tissue parameters are often presented as a two-dimensional heat map, displaying both the parameters' magnitude and spatial position.

There are several applications where tissue parameters provide helpful information. For instance, in oncology, apparent diffusion coefficient (ADC) estimation for prostate cancer staging,² and T1 mapping in dynamic contrast-enhanced-MRI for invasive breast cancer detection.³ In addition to T1 mapping, T2 mapping is also of interest in prostate cancer due to the provided ability to discriminate between cancer tissue, normal gland tissue, and prostate gland enlargement.⁴

The parameter map estimation is a postimaging procedure. Conventionally, the tissue parameters are obtained by least squares regression of the MRI magnitude signal at different user-defined scanner settings. Signal noise (e.g., from receiver coil readouts) is especially problematic in this process, which propagates to produce noisy parameter maps with high uncertainties in the parameter values. Therefore, there is great interest in improving QMRI by incorporating spatial denoising in the parameter estimation process. To this day, there are several published methods to achieve this. A common approach is implementing a Bayesian framework to denoise the parameter reconstruction with a spatial prior distribution. Examples are relaxation times estimation with Markov Random Fields,⁵ coherent region smoothness,^{6,7} and total variation priors.^{8,9} Other Bayesian methods address, for instance, intravoxel incoherent voxel modeling,^{10,11} field map reconstruction,¹² myelin water fraction estimation,¹³ and applications in dynamic contrast-enhanced-MRI.^{14–16} Other approaches address MRI signal denoising before the tissue parameter reconstruction, for example, with Marchenko–Pastur Principal Component Analysis¹⁷ or Beltrami Denoising.^{18–20} More modern data-driven deep learning approaches include mapping relaxation parameters with residual networks,^{21,22} frameworks with physical model constraints,^{23–26} supervised²⁷ and unsupervised²⁸ intravoxel incoherent motion estimation and

models that address uncertainty estimation in dynamic contrast-enhanced-MRI²⁹ and ADC mapping.³⁰

To some degree, any method that incorporates denoising will rely on specific prior knowledge rather than the data. Developing a framework thus involves both the selection/formulation of denoising prior and to what extent it should be trusted. The level of trust is usually selected by tuning model hyperparameters or regularization parameters, for example, size and distance of patches in nonlocal mean denoising and weight selection in total variation denoising.

In contrast to formulating an explicit image prior, Ulyanov et al.³¹ showed that denoising and other standard inverse problems could be addressed by Deep Image Priors (DIP), that is, by treating the *structure* of image generator network as an implicit image prior. DIP has been shown to correspond to a limiting case of a stationary Gaussian process.³² DIP introduces an alternative approach in-between explicitly formulated priors and data-driven machine learning algorithms. Low-level image features, for example, noise, are captured from untrained networks without prior training.

DIP for parameter mapping is interesting for several reasons, the most prominent of which are reports of successful results in other medical imaging fields, including PET reconstructions,³³ DWI denoising³⁴ and dynamic image reconstruction.³⁵ However, parameter mapping is, to our knowledge, unexplored. Further, the absence of prior training in DIP makes it directly available in parameter mapping without data collection for network training. Collecting network training data for parameter mapping is tedious since network training data must accurately cover all applications and use cases, including method-specific signal equations, scanner settings, and different anatomical regions, which are rarely available. Also, the absence of prior training prevents training-induced anatomical artifacts from entering the results, a common issue among network approaches. DIP is easy to customize for different applications since most of the framework is application-independent. Application specifics, that is, signal equations in parameter mapping, are easily handled so that a single generator network applies to several applications.

In its original implementation, DIP is sensitive to the values of hyperparameters, mainly the number of optimization steps, to achieve significant noise reduction. Also, the method does not present any information on model uncertainty. Gal et al.³⁶ presented a solution to this by using dropout³⁷ as a Bayesian approximation, later extended for convolutional neural networks as approximate Bernoulli variational inference with dropout network training.³⁸ This approximate Bayesian implementation, commonly called Monte Carlo (MC)

dropout, reduces the Bayesian implementation to simply performing dropout after every convolution layer when training a deterministic neural network. MC dropout is highly interesting in parameter mapping mainly due to its inherent robustness to overfitting and the added possibility of investigating model uncertainty. The combination of DIP and MC dropout has been addressed in magnitude image denoising³⁹ but has to our knowledge, not been applied in parameter mapping. This work addresses noisy parameter maps with high uncertainties in QMRI with the combination of DIP denoising and approximate Bernoulli variational inference with MC dropout. More formally, the purpose of this work was to

1. Investigate the possibility of implementing DIP denoising into different parameter mapping applications in QMRI, thus bridging the gap of unexplored applications in the current literature.
2. Investigate whether applying MC dropout can provide uncertainty information in DIP-enhanced parameter mapping.
3. Evaluate the proposed method regarding denoising properties, accuracy in the uncertainty estimate, computational requirements, and ease of use.

2 | METHODS

2.1 | Statistical model

Let $\beta \in \mathbb{R}^{P \times V}$ denote the collection of P tissue parameter maps with V voxels each. Parameter mapping, that is, the estimation of β , is based on the collection of an MRI dataset $\mathbf{y} \in \mathbb{R}^{M \times V}$ with M different scanner settings in the acquisition. The elements of \mathbf{y} , denoted $y_{m,v}$, are modeled as

$$y_{m,v} = s_m(\beta_{*,v}) + \epsilon_{m,v}, \quad (1)$$

$\forall m \in \{1, \dots, M\}$ and $\forall v \in \{1, \dots, V\}$, where $s_m(\beta_{*,v}) \in \mathbb{R}$ is the MRI signal equation, $\beta_{*,v} \in \mathbb{R}^P$ is the collection of all, denoted by index $*$ ¹, tissue parameters in the v th voxel, and $\epsilon_{m,v} \in \mathbb{R}$ describes noise in the measured signal. The noise in Equation (1) follows the Rice distribution for magnitude images but can be approximated as Gaussian at high signal-to-noise ratio.⁴⁰ By assuming that the model targets are independent and follow a Gaussian distribution, the data likelihood becomes

¹Throughout this paper, we will denote slices of high-dimensional objects with index $*$. For example, for the tissue parameters in $\beta \in \mathbb{R}^{P \times V}$, we denote all tissue parameters in the v th voxels $\beta_{*,v} \in \mathbb{R}^P$, and the entire map of the p th tissue parameter as $\beta_{p,*} \in \mathbb{R}^V$. Voxels are, for simplicity, indexed with a linear index $v \in \{1, 2, \dots, V\}$ in all objects.

$$p(\mathbf{y} | \beta, \sigma^2) = \prod_{m=1}^M \prod_{v=1}^V p(y_{m,v} | s_m(\beta_{*,v}), \sigma^2) \\ = (2\pi\sigma^2)^{-MV/2} \cdot \exp\left(-\frac{1}{2\sigma^2} \sum_{m=1}^M \sum_{v=1}^V (y_{m,v} - s_m(\beta_{*,v}))^2\right), \quad (2)$$

where σ^2 is the noise variance.

Maximizing the likelihood in Equation (2) coincides with solving a nonlinear least squares (NLLS) regression due to the assumption of Gaussian distributed noise. The signal equation is application-specific, for example, spoiled gradient echo (SPGR) for variable flip-angle (VFA) T1 mapping⁴¹ and multi echo-time spin-echo (SE) for T2 mapping.⁴²

In this work, we developed an improved parameter mapping method based on DIP image denoising.³¹ DIP performs image denoising by interpreting the output of a randomly initialized image generator as a parametrization of an image. The image generator is a neural network, defined as a function, f_θ , that maps a randomized input, \mathbf{z} , to an image, $f_\theta(\mathbf{z})$, of the same spatial shape. This method achieves denoising with the implicit bias of the generator instead of with an explicit prior defined in the space of images. We extend this concept into the domain of QMRI by instead interpreting the network output as a *parametrization of the tissue parameters* in parameter mapping applications. Given a QMRI magnitude signal dataset, \mathbf{y} , acquired with MRI signal, s , we define f_θ to be a parameter map generator and obtain denoised tissue parameter maps $\hat{\beta} \in \mathbb{R}^{P \times V}$ by solving the optimization

$$\hat{\theta} = \arg \min_{\theta} E(s_m(f_\theta(\mathbf{z})); \mathbf{y}), \quad (3a)$$

$$\hat{\beta} = f_{\hat{\theta}}(\mathbf{z}), \quad (3b)$$

where E is the mean squared error (MSE) loss function that compares the network output, $f_\theta(\mathbf{z}) \in \mathbb{R}^{P \times V}$, with the signal dataset, \mathbf{y} . This optimization problem is a two-step process in which the minimized weights, $\hat{\theta}$ is found by minimizing the loss function, and the network outputs the final parameter map estimate with $\hat{\theta}$ applied. Our main contribution is reinterpreting the neural network output for parameter mapping tasks, which involves mapping the output through the signal equation to estimate the network weights. To maximize the likelihood in Equation (2) is equivalent to minimizing its negative log-likelihood; this gives us the following MSE loss for DIP-enhanced parameter mapping

$$E(s_m(f_\theta(\mathbf{z})), \mathbf{y}) = \frac{1}{MV} \sum_{m=1}^M \sum_{v=1}^V \left\| y_{m,v} - s_m(f_\theta(\mathbf{z})_{*,v}) \right\|_2^2. \quad (4)$$

Note that to achieve denoising, the nonlinear optimization program in Equation (3a) needs to be terminated after some number of optimization steps to avoid the collapse of $\hat{\beta}$ to the maximum likelihood estimate.

We complement this denoising process with model uncertainty estimation. The image generator, f_θ , is interpreted as a BNN with Bernoulli approximate variational inference implemented using dropout after every convolution layer of f_θ .³⁸ The dropout layers randomly set elements of their inputs to zero, with probability p , that is, they add Bernoulli distributed multiplicative noise to the convolution layer outputs. The posterior in this model is approximated, and the posterior prediction is approximated by MC integration of N_s concurrent forward passes with dropout. With this approximation applied, the parameter estimation in Equation (3a) is changed to

$$\hat{\theta} = \arg \min_{\theta} E(s_m(f_\theta(\mathbf{z}, \epsilon)); \mathbf{y}), \quad (5a)$$

$$\hat{\beta} = \frac{1}{N_s} \sum_{n=1}^{N_s} f_{\hat{\theta}}(\mathbf{z}, \epsilon_n), \quad (5b)$$

where f_θ is redefined with the additional input ϵ , which denotes the collection of Bernoulli random variables required to implement dropout. The estimated parameter maps are now complemented with maps of estimated model uncertainty, $\sigma_{\text{est}} \in \mathbb{R}^{P \times V}$, with elements

$$\sigma_{\text{est},p,v} = \sqrt{\frac{1}{N_s} \sum_{n=1}^{N_s} (f_{\hat{\theta}}(\mathbf{z}, \epsilon_n)_{p,v} - \hat{\beta}_{p,v})^2}, \quad (6)$$

$\forall p \in \{1, \dots, P\}$ and $\forall v \in \{1, \dots, V\}$, which is the standard deviation over the N_s concurrent forward passes with dropout.

2.2 | Implementation

The proposed framework was implemented in Python (Python Software Foundation, <https://www.python.org/>) with the deep learning framework PyTorch². Our code³, is based on the original DIP example code and Laves et al.'s implementation of MC dropout.³⁹ The process (Algorithm 1) outputs denoised tissue parameter maps, $\hat{\beta}$, with estimated model uncertainty, σ_{est} , from a parameter mapping dataset \mathbf{y} . The process starts by optimizing the generator with the Adam stochastic gradient algorithm⁴³ with weight decay λ and learning rate γ . The

generator weights, θ , is updated until it reaches the final state θ_{N_b} after N_b optimization steps. The generator input, $\mathbf{z} \in \mathbb{R}^{N_c \times V}$, is N_c channels with V voxels of random uniform noise with standard deviation σ_{input} . \mathbf{z} is perturbed with noise-based regularization with SD σ_{reg} , just as in the original DIP implementation. After optimization, maps of tissue parameters and estimated model uncertainty are calculated from N_s MC dropout samples with fixed weights. The Bayesian approximation is enabled by drawing a new sample of the Bernoulli distributed random variables, ϵ , before passing the input \mathbf{z} through the network. The size of ϵ depends on the generator architecture and is omitted for simplicity in Algorithm 1. The application-specific signal equation s and the generator f are described in detail in Sections 2.2.1 and 2.2.2, respectively.

Algorithm 1. The proposed parameter mapping algorithm outputs denoised tissue parameter maps, $\hat{\beta}$, with associated model uncertainty, σ_{est} from a parameter mapping dataset \mathbf{y} . The hyperparameters of this method are the number of training iterations N_b , number of input channels, N_c , number of MC Dropout samples, N_s , dropout probability, p , input noise standard deviation, σ_{input} , and noise regularization SD, σ_{reg} . γ and λ are the learning rate and weight decay required by the Adam algorithm. Identity matrices with shape $N_c \times V$ are denoted $\mathbf{I}_{N_c \times V}$. The application-specific signal equation s and the generator f are described in detail in Section 2.2.1 and 2.2.2, respectively

procedure

```

 $\mathbf{z}_0 \sim \text{Unif}(\mathbf{0}, \sim \mathbf{I}_{N_c \times V} \cdot \sigma_{\text{input}}^2)$ 
for  $n \leftarrow 1$  to  $N_b$  do
     $\mathbf{z} \sim \mathcal{N}(\mathbf{z}_0, \mathbf{I}_{N_c \times V} \cdot \sigma_{\text{reg}}^2)$ 
     $\epsilon_n \sim \text{Bern}(p)$ 
     $\theta_n \leftarrow \text{Adam}\left[E\left(s_m(f_{\theta_{n-1}}(\mathbf{z}, \epsilon_n)), \mathbf{y}\right), \gamma, \lambda\right]$ 
end for
for  $n \leftarrow 1$  to  $N_s$  do
     $\epsilon_n \sim \text{Bern}(p)$ 
     $\beta^{(n)} \leftarrow f_{\theta_{N_b}}(\mathbf{z}, \epsilon_n)$ 
end for
for  $p \leftarrow 1$  to  $P$  do
    for  $v \leftarrow 1$  to  $V$  do
         $\hat{\beta}_{p,v} \leftarrow \frac{1}{N_s} \sum_{n=1}^{N_s} \beta_{p,v}^{(n)}$ 
         $\sigma_{\text{est},p,v} \leftarrow \sqrt{\frac{1}{N_s} \sum_{n=1}^{N_s} (\beta_{p,v}^{(n)} - \hat{\beta}_{p,v})^2}$ 
    end for
end for
end procedure

```

²<https://github.com/pytorch/pytorch>

³Publicly available at <https://github.com/MaxHellstrom/VI-DIP-QMRI>

We implemented Algorithm 1 with hyperparameters $N_c = 32$, $N_s = 1280$, $p = 0.1$, $\gamma = 3 \times 10^{-4}$, $\lambda = 1 \times 10^{-4}$, $\sigma_{\text{input}} = 1/10$, $\sigma_{\text{reg}} = 1/20$, and application-specific values for N_b .

2.2.1 | Signal equation layers

We implement the signal equation as application-specific layers in the generator framework. Each application will require the network output to be mapped through the signal equation to compute the MSE loss in each optimization step; see Section 2.2. We have implemented parameter mapping for T1, T2, and ADC, based on the VFA SPGR, multi-echo spin echo and multiple b -value DWI acquisitions.¹ In all implementations, the generator outputs two tissue parameter maps, $f_\theta(\mathbf{z}, \epsilon) \in \mathbb{R}^{2 \times V}$, where the first channel contains the signal magnitude M , proportional to proton density, and the second channel contains the parameter of interest, that is, either T1, T2, or ADC. Strictly, the second channel equals the negative logarithm of the parameter of interest, for example, $f_\theta(\mathbf{z}, \epsilon) = -\log T1$ (omitted in the table for simplicity), due to better computational stability. The relevant scanner settings and tissue parameters are summarized in Table 1, together with the associated signal equation.

2.2.2 | Generator architecture

The implemented generator (Figure 1) utilizes a skip-connection encoder-decoder type CNN with MC dropout, which is a combination of the architecture provided by Ulyanov et al.³¹ for DIP denoising together with the MC dropout extension provided by Laves et al.³⁹ for Bayesian DIP denoising of magnitude images. The generator input, \mathbf{z} , is passed through a five-layer deep encoding-decoding process with added skip

connections. Lastly, a single convolution layer is added to create an output $f_\theta(\mathbf{z}, \epsilon) \in \mathbb{R}^{P \times V}$, that is, P parameter maps with V voxels each. These are then mapped through the application-specific signal equation layer s to update the network with MSE loss calculation. See Table 1 for specifics of different applications.

2.3 | Data

We used three separate data sources to evaluate the proposed method. One synthetic with known ground truth reference and two in vivo datasets with volunteering patients undergoing cancer treatment.

2.3.1 | Synthetic data

A single axial brain slice from the BrainWeb phantom,^{44–47} with tissue-specific masks for white matter (WM), grey matter (GM), and CSF, was selected for evaluation on synthetic data. Reference tissue parameter maps of PD, T1, and T2 (each with matrix size 256×256 and voxel size $0.98 \times 0.98 \times 2.00$ mm²) were generated from this slice using Hero Imaging (NONPI Medical AB; www.heroimaging.com). This source was used to compute two synthetic datasets:

- The PD and T1 were used to compute SPGR data with TR = 6.8 ms and flip angles FA = 2°, 4°, 11°, 13°, and 15°.
- The PD and T2 map was used to compute multi-echo spin echo data with TE = 50, 150, and 300 ms.

When creating the synthetic datasets, the PD reference tissue map was scaled to span the range [0, 1], that is, magnitude normalized, to decrease the need for application-specific hyperparameter tuning. Complex circular Gaussian noise was added to both datasets to

TABLE 1 Specifics of the three investigated applications, with MRI acquisition method, variable scanner setting, generator output interpretation, and associated signal equation.

		Generator output:		Application-specific signal equation
		$f_\theta(\mathbf{z}, \epsilon) \in \mathbb{R}^{2 \times V}$, where		
Acquisition method	Variable scanner setting	$f_\theta(\mathbf{z}, \epsilon)_{1,v} =$	$f_\theta(\mathbf{z}, \epsilon)_{2,v} =$	
VFA SPGR	flip angle, FA	M_v	$T1_v$	$M_v \sin \text{FA}_m \frac{1 - \exp(-\text{TR}/T1_v)}{1 - \cos \text{FA}_m \exp(-\text{TR}/T1_v)}$
Multi-echo SE	echo time, TE	M_v	$T2_v$	$M_v \exp(-\text{TE}_m/T2_v)$
DWI	b-value, b	M_v	ADC_v	$M_v \exp(-b_m \text{ADC}_v)$

Note: The network output two tissue parameter maps, $f_\theta(\mathbf{z}, \epsilon)_{1,*}$ and $f_\theta(\mathbf{z}, \epsilon)_{2,*}$, where the first contains the signal magnitude M , proportional to proton density, and the second contains the parameter of interest, that is, either T1, T2, or ADC.

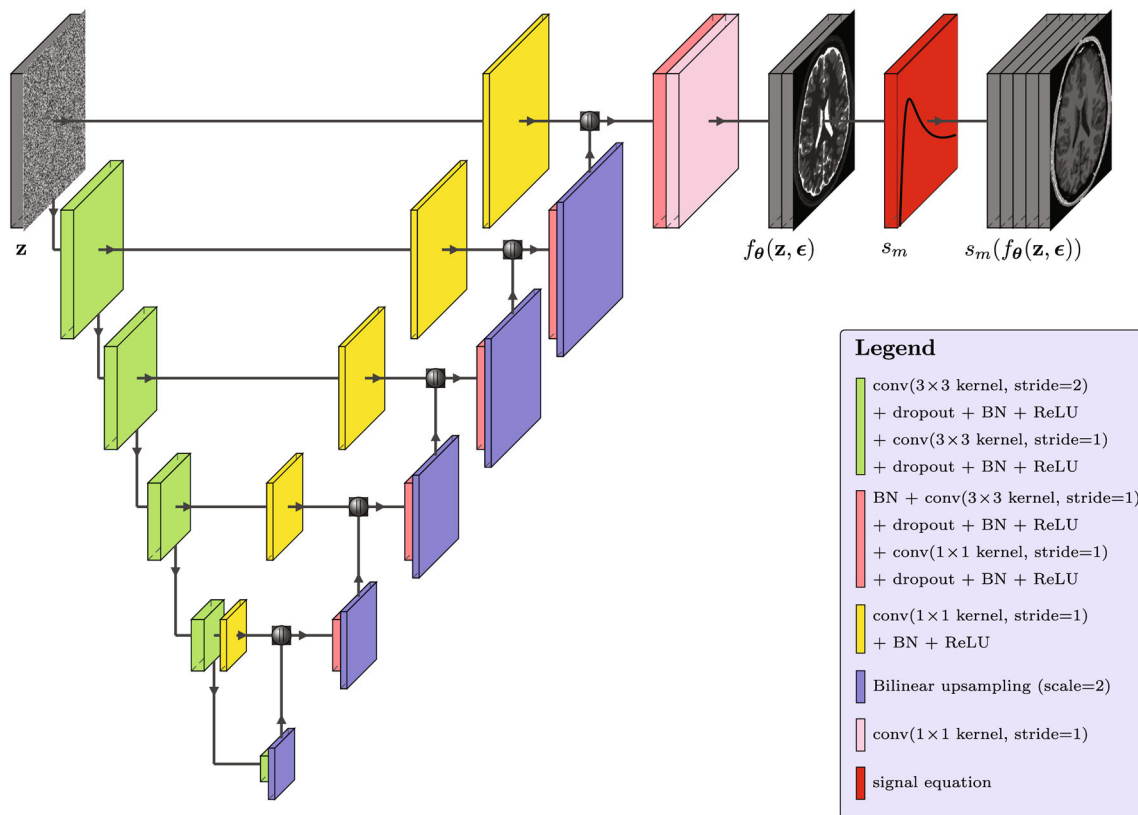


FIGURE 1 Illustration of the generator architecture. The generator input passes a five-level deep encoder-decoder process with added skip connections at each forward pass through the network. The encoding section consists of encoding blocks (green blocks with 128 output filters) which involve a factor 2 downsampling with strided convolutions. The decoding section consists of convolution blocks (red blocks with 128 output filters) and bilinear spatial upsampling (blue blocks). At each level, the decoding block is concatenated (black ball labelled ||) with a skip block (yellow blocks with four output filters). Lastly, a 1×1 convolution block (pink) transforms the output from 128 filters to 2, that is, to the two separate parameter maps, which are then mapped through the application-specific signal equation s (red) to enable MSE loss calculation. See section 2.2.1 for details.

simulate the impact of noise perturbations from Rice distributed noise. The noise SD was 0.01 (multi-echo data) and 0.02 (SPGR data) to simulate two different levels of the noise environment.

2.3.2 | In vivo data

This study used two in vivo datasets: one SPGR VFA study for T1 mapping and one DWI study for ADC mapping. Both datasets were collected with a PET/MRI scanner (SIGNA PET/MR 3T, GE Healthcare) after informed consent was obtained from the patients. This study was conducted by the principles embodied in the Declaration of Helsinki and was ethically approved by the national research ethics committee (nr 2019-02666).

- The SPGR dataset consists of eight patients scanned with a 3D SPGR sequence (seven male, one female, age span 39–75 years, mean age 52 years). Voxel

size, FAs, TR, and TE were set identical to the synthetic SPGR data (see Section 2.3.1 for details) and the voxel bandwidth was set to 488 Hz/voxel. For each patient, one single axial slice was selected for T1 estimation.

- The DWI dataset consists of eight patients from the PAMP research project⁴ (PSMA, Acetate and Multiparametric MRI for Prostate cancer). Each patient was scanned with a diagnostic multiparametric pelvic protocol that included diffusion-weighted EPI acquisitions with b -value 0, 200, and 1000 s/mm² with 16 axial slices of 160×80 voxels. The voxel bandwidth was set to 1953 Hz/voxel. For each patient, the DWI data from a single axial slice was selected from the mid-prostate region for ADC estimation.

⁴<https://www.umu.se/en/research/projects/pamp-psma-acetate-and-multiparametric-mri-for-prostate-cancer/>

The in vivo datasets were magnitude normalized by dividing each voxel value with the maximum value, that is, $\mathbf{y}_{v,m} \rightarrow \mathbf{y}_{v,m} / \max(\mathbf{y})$ so that the signal data spanned the range [0, 1]. In all datasets, voxels with insufficient signal magnitude, for example, outside of the patient volume, were excluded by applying a binary mask.

2.4 | Analysis

2.4.1 | Parameter mapping

We analyzed the proposed method with synthetic and in-vivo T1 mapping, synthetic T2 mapping, and in vivo ADC mapping. For each task, we estimated the tissue parameter map, $\hat{\beta}$, and its associated model uncertainty map, σ_{est} , according to Algorithm 1. For reference, we conducted the same parameter mapping tasks with a conventional NLLS estimator, calculated with the SciPy least-squares function⁵ using the Trust Region Reflective algorithm, which in this case returns the solution to

$$\underset{\beta_{*,v} \in \mathcal{R}}{\text{minimize}} \sum_{m=1}^M ((y_{m,v} - s_m(\beta_{*,v}))^2, \quad (7)$$

where $\mathcal{R} \subset \mathbb{R}^2$ is the region limited by the bounds of the least-squares algorithm. We denote this method NONE, that is, with no implemented denoising. With these estimates, we computed the Cramér–Rao Bound (CRB) variance σ_{CRB}^2 , which equals the reciprocal of the Fisher information matrix, I . In this case of vector signal parameters estimated in approximate Gaussian noise,⁴⁸ the Fisher elements in the v th voxel equal

$$I[\beta_{*,v}]_{ij} = \frac{1}{\sigma^2} \sum_{m=1}^M \left[\frac{\partial s_m(\beta_{*,v})}{\partial \beta_{i,v}} \frac{\partial s_m(\beta_{*,v})}{\partial \beta_{j,v}} \right], \quad (8)$$

We complement this with more detailed experiments with synthetic T1 and T2 mapping with available maps of ground truth reference, $\beta^{(\text{ref})}$. We apply several denoising methods to the input signal prior to computing the NLLS with Equation (7). From the restoration module in scikit-image⁶, we implement wavelet denoising (WAVE), total-variation (TV), and nonlocal means (NLM). We also apply *block-matching and 3D filtering* (BM3D)^{49,50} with the python package BM3D⁷ and deep decoder (DD)

⁵https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.least_squares.html

⁶<https://scikit-image.org/docs/stable/api/skimage.restoration.html#skimage.restoration> hyperparameters: sigma=0.05 (WAVE), weight=0.02 (TV), patch_size=7, patch_distance=11, and h=0.02 (NLM)

⁷<https://pypi.org/project/bm3d/> implemented with hyperparameters noise_type=gw, noise_std=0.02/0.01 (T1/T2)

denoising⁸, the latter of which is an under-parameterized nonconvolutional alternative to DIP.

Hyperparameters were selected to achieve significant noise reduction in WM with a minimal blurring of spatial features.

With all of these methods, including the proposed method (labeled OURS), we simulated the actual uncertainty by repeating the same parameter mapping task $N_r = 100$ times with identical synthetic data but unique noise realization for each run. This enabled the calculation of bias and reference uncertainty maps with elements,

$$\text{Bias}(\hat{\beta}_{p,v}) = \frac{1}{N_r} \sum_{n=1}^{N_r} \hat{\beta}_{p,v}^{(n)} - \beta_{p,v}^{(\text{ref})}, \quad (9)$$

$$\sigma_{\text{ref}}(\hat{\beta}_{p,v}) = \sqrt{\frac{1}{N_r} \sum_{n=1}^{N_r} \left(\hat{\beta}_{p,v}^{(n)} - \frac{1}{N_r} \sum_{n=1}^{N_r} \hat{\beta}_{p,v}^{(n)} \right)^2}, \quad (10)$$

as well as a frequency mean of the model uncertainty with OURS,

$$\bar{\sigma}_{\text{est},p,v} = \frac{1}{N_r} \sum_{n=1}^{N_r} (\sigma_{\text{est},p,v})^{(n)}. \quad (11)$$

Equation (9) gives a direct measure of the precision of the parameter estimate, $\hat{\beta}$, while the differences between the outcomes in Equations (10) and (11) quantify the precision of the model uncertainty estimate, σ_{est} . We present the results of Equations (9)–(11) as heatmaps and ROI analysis.

2.4.2 | Uncertainty calibration

The model uncertainty with MC dropout is not inherently calibrated and can contain errors in data magnitude or scale differently for different datasets.⁵¹ We investigated calibrating σ_{est} to decrease the difference in model uncertainty and frequency outcome in Equation (10). For this, let $\mathbf{A} \in \mathbb{R}^{P \times V}$ be data magnitude maps of the estimated tissue parameters, with elements,

$$\mathbf{A}_{p,v} = \hat{\beta}_{p,v} / \max(\hat{\beta}_{p,*}), \quad (12)$$

such that \mathbf{A} is normalized to the range [0, 1] for tissue parameter p . Since the generator outputs the tissue parameter of interest on a logarithmic scale (see Section 2.2.1), we apply the correction in Equation (12) to the model

⁸https://github.com/reinhardt/supplement_deep_decoder

implemented with hyperparameters $k = 512$, numit = 10 000, rn = 0.0

uncertainty as

$$\sigma_{\text{est},2,v} \leftarrow \frac{\sigma_{\text{est},2,v}}{\log A_{2,v} - \min(\log A_{2,*})}, \quad (13)$$

for all voxels $v \in \{1, \dots, V\}$ of the model uncertainty.

2.4.3 | MC dropout in DIP

For completeness, we also include DIP-enhanced parameter mapping without the Bayesian approximation by disabling MC dropout. For this, we conduct T1 mapping with several different training iterations, N_b , both with and without MC dropout applied. We disable MC dropout by setting the dropout probability to zero so that concurrent forward passes return identical outputs. This comparison highlights the difference between Equations (3) and (5), in terms of the trade-off between convergence rate, computational cost, and the ability to avoid overfitting to the noise in the final result.

3 | RESULTS

The resulting tissue parameter maps and associated calibrated model uncertainty in all tested applications are presented in Figure 2. Comparing A and B, we conclude that the proposed method successfully incorporates denoising for all tested applications. Further, the model uncertainty produces uncertainty estimates similar to the CRB but smoother and with lower magnitude (Figure 2C,D). A detailed zoom of the T1 results with synthetic data is presented in Figure 3, which highlights the performance of using the proposed method compared to other denoising methods. Results associated with the rest of the patients included in this study are presented as Figures S1 and S2. The bias and uncertainty analysis are presented as heatmaps in Figures 4 and 5, complemented with ROI analysis for different tissues in Table 2. All denoising methods introduce some bias compared to NLLS without denoising (NONE). In T1 mapping, OURS introduces the least RMS bias in all ROIs. Equivalently for T2 mapping, OURS is the second lowest with a maximum deviation of 1.1% from the best performer. Note the increased bias with OURS in the border region of the imaged volume in T2 mapping (Figure 4G2). The simulated reference uncertainty (Equation 10) is presented in Figure 5. OURS seem visually most similar to TV, with significant uncertainty reduction in all tissues. Regarding ROI mean value, OURS deviates less than 1% from the best performer in all ROIs and applications; see Table 2. The model uncertainty from OURS, both with and without calibration, is appended

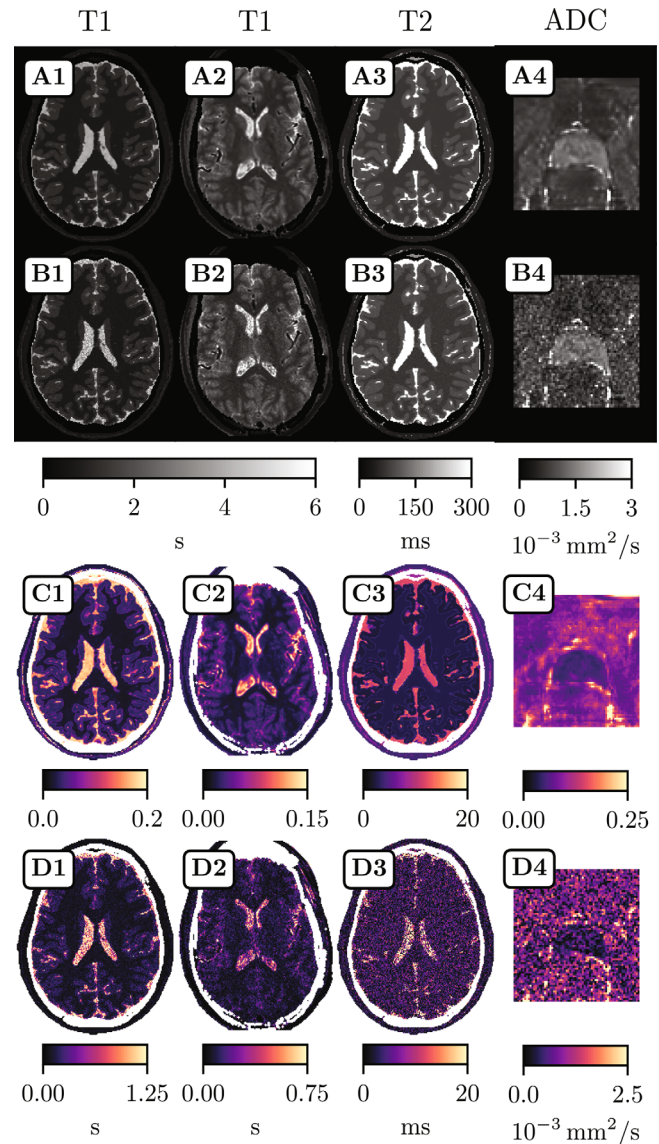


FIGURE 2 Parameter mapping results with the proposed method (row A) with associated calibrated model uncertainty (row C) compared to a conventional nonlinear least squares estimator without applied denoising (NONE, row B) with associated Cramér-Rao Bound standard deviation (row D). Each column shows a separate implementation: Brain T1 mapping with synthetic (col 1) and in-vivo (col 2) variable flip-angle data, brain T2 mapping with synthetic multi-echo spin echo (col 3), and prostate apparent diffusion coefficient mapping with in vivo DWI data (col 4). The ADC (ADC) maps in column 4 have been clipped to 70×70 voxels to highlight the region encompassing the prostate gland.

to the same plot for comparison. These metrics are also presented for different ROIs in Table 2.

The effect of disabling the Bayesian approximation is presented in Figure 6. The proposed MC dropout implementation is compared to the same model without MC dropout (labeled *non-bayesian*). OURS is significantly less prone to reconstruct noise in the input data

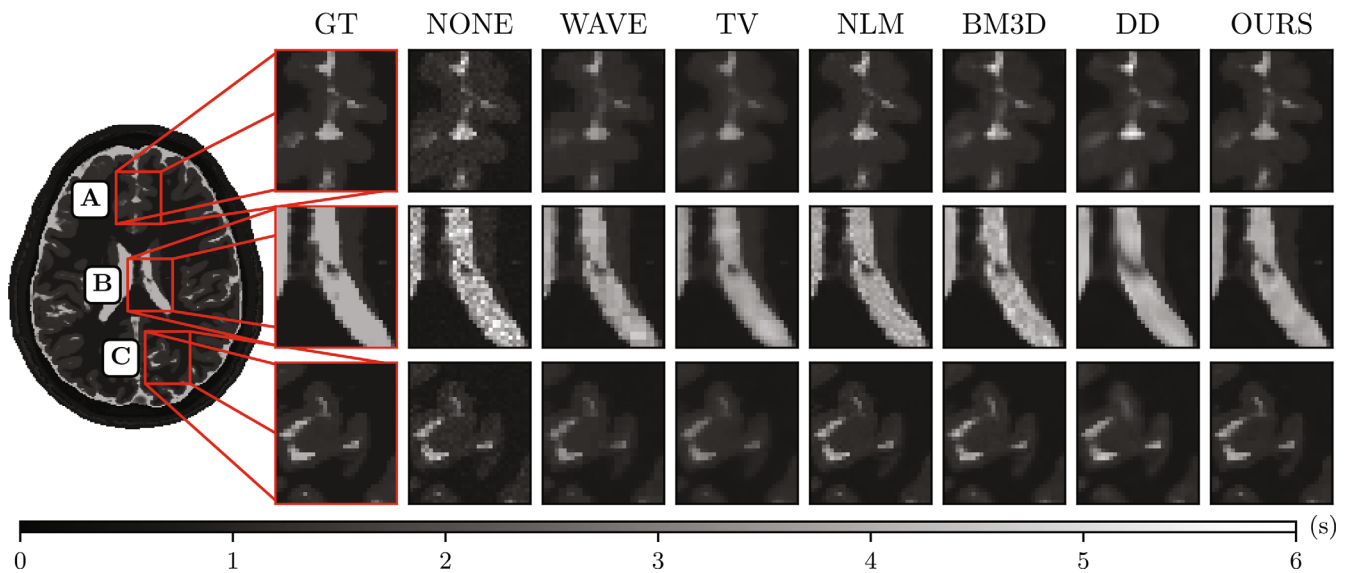


FIGURE 3 A detailed zoomed-in view of T1 mapping results with synthetic data. The proposed method (OURS, same as **A1** in Figure 2) is compared with the ground-truth reference (GT), and nonlinear least squares estimates with different methods of signal denoising: no denoising (NONE, same as **B1** in Figure 2), wavelet (WAVE), total variation (TV), nonlocal means (NLM), block-matching and three-dimensional filtering (BM3D), and deep decoder (DD). Each zoom-in region displays a rectangular patch of 70×70 voxels.

(i.e., overfitting), which facilitates the selection of appropriate hyperparameters. However, this comes at the price of a decrease in the convergence rate, which increases the computational cost.

All computations were conducted with Nvidia's parallel computing platform CUDA with a single RTX 2080 Ti GPU. The time consumption for single slice estimation with this setup is presented in Table 3, together with the application- and data-specific parameters on the number of voxels and the selected number of optimization steps.

4 | DISCUSSION

We have reformulated parameter mapping as a DIP denoising task with associated model uncertainty estimation.

Our main findings are that DIP adapts successfully to several applications of parameter mapping with only limited hyperparameter tuning, and that MC dropout mitigates the burden of overfitting noise. However, this is the most time-consuming aspect of the proposed method. Our code is generalized to handle several applications with a single generator by only altering the signal equation block. This modular approach is easy to expand to more applications. Also, the associated uncertainty estimate provides valuable information when adequately calibrated.

The main disadvantage of the proposed method is the high computational cost and the need for hyperparameter selection.

The proposed method is computationally heavy, with long estimation times in all applications (OURS in Table 3). T1 mapping was almost six times slower than ADC mapping, which shows that time consumption scales with data size (35k voxels vs. 11k voxels). This is a clear drawback in clinical settings where large matrix sizes with multiple slices are often required. Non-Bayesian DIP is much faster and requires only 10k iterations (8 min) to output similar results to OURS at 50k steps (Figure 6), that is, most of the computational cost originated from MC dropout. We did not include model backtracking to prevent unsatisfactory gradient updates. Backtracking would, however, make non-Bayesian DIP even faster and further increase the difference between the two methods.

Network approaches are conventionally time-consuming during training but fast in production. DIP is unique since DIP generators only use images from a single imaging occasion instead of large training sets with representative examples. This feature makes DIP easy to implement but slow in production since the training (or optimization) needs to be repeated for each imaging occasion. Also, this makes DIP methods less likely to do something unpredictable, like to hallucinate when encountering something unusual such as an artifact. This could be beneficial; however, trained models could possibly also understand that the artifacts should be removed if adequately trained, which one cannot anticipate for DIP-based methods.

Although beyond the scope of this work, we emphasize that future studies should prioritize computational

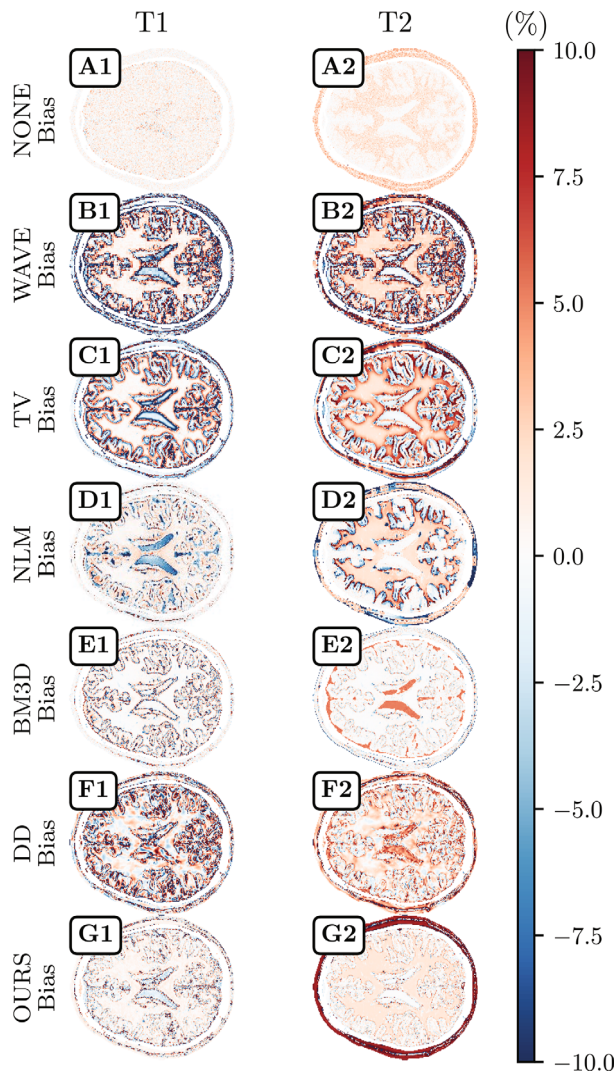


FIGURE 4 Estimated bias in parameter map estimates with an axial brain slice, calculated with Equation (9) for the case of T1 mapping (col 1), and T2 mapping (col 2). The proposed method (OURS) is compared with nonlinear least squares estimates with different methods of signal denoising: no denoising (NONE), wavelet (WAVE), total variation (TV), and nonlocal means (NLM), block-matching and three-dimensional filtering (BM3D), and deep decoder (DD). A region of interest (ROI) analysis of different tissues is presented in the bias columns of Table 2. All plots are normalized to the ground-truth reference parameter map and displayed as percentages.

cost since it is the most prominent limitation in clinical implementations.

Insufficient convergence in the optimization step (Equation 3 or 5) is concerning in applications where the image data has a meaningful interpretation, such as in QMRI parameter mapping. Stopping the optimization too early results in absent low-level image features (not necessarily noise). This can be seen in Figure 6 by comparing the initial output and a later stage. For

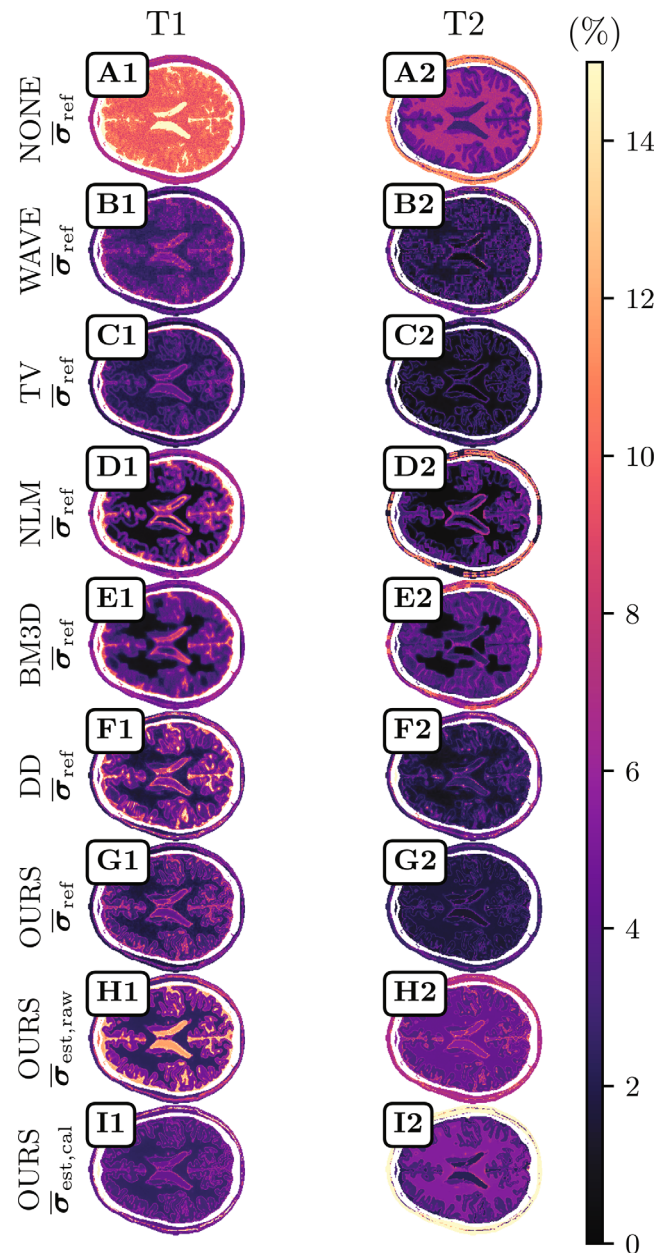


FIGURE 5 Simulated ground-truth (GT) reference uncertainty, σ_{ref} , in parameter map estimates with an axial brain slice, calculated with Equation (10) for the case of T1 mapping (col 1, row A–G) and T2 mapping (col 2, row A–G). The proposed method (OURS) is compared with nonlinear least squares estimates with different methods of signal denoising: no denoising (NONE), wavelet (WAVE), total variation (TV), nonlocal means (NLM), block-matching and three-dimensional filtering (BM3D), and deep decoder (DD). Additionally, the frequency mean of the uncertainty estimate from OURS, Equation (11), are presented both uncalibrated (row H) and calibrated according to Equation (13) (row I). The magnitude of the data magnitude errors from in the MC dropout implementation are examined by comparing G to H, and the precision of the proposed calibration is equivalently examined by comparing G to I. A region of interest (ROI) analysis of these metrics in different tissue is presented in the uncertainty column of Table 2. All plots are normalized to the GT parameter map and displayed as percentages.

TABLE 2 Region of interest (ROI) analysis of the bias and uncertainty results in Figures 4 and 5, separated in three separate ROIs encompassing either white matter (WM), grey matter (GM), or cerebrospinal fluid (CSF).

	Bias (RMS %)			$\sigma_{\text{ref}} (\%) / \sigma_{\text{est,raw}} (\%) / \sigma_{\text{est,cal}} (\%)$		
	WM	GM	CSF	WM	GM	CSF
T1 mapping						
NONE	1.09	1.18	1.92	9.7	10.46	17.48
WAVE	5.41	7.31	20.11	3.2	4.09	6.38
TV	4.34	5.26	14.6	2.39	3.17	5.73
NLM	2.47	3.53	6.32	2.1	4.47	10.56
BM3D	2.23	3.52	7.04	2.16	4.26	9.12
DD	4.31	6.85	12.42	3.11	5.18	10.2
OURS	1.79	2.77	4.92	2.77/3.22/3.19	4.18/5.98/3.84	6.11/12.62/4.74
T2 mapping						
NONE	1.58	0.65	0.30	7.03	4.64	2.93
WAVE	7.41	9.73	12.46	2.6	2.52	1.63
TV	6.13	3.69	3.68	1.52	1.86	1.4
NLM	4.81	4.03	0.94	1.68	3.09	2.67
BM3D	1.36	1.32	4.46	2.94	3.63	2.4
DD	3.07	2.63	5.51	1.87	2.42	3.18
OURS	2.02	1.6	2.02	1.94/4.54/5.55	2.18/4.42/3.83	2.11/6.08/2.63

Notes: The value in the bias column is presented as the root mean square (RMS) error of the bias metric in Equation (9), and the values in the uncertainty column are presented as the mean value of Equation (10). For OURS, we also append the mean model uncertainty, both uncalibrated ($\sigma_{\text{est,raw}}$, Equation 11) and calibrated ($\sigma_{\text{est,cal}}$, Equation 13) for region-specific analysis of the proposed calibration procedure.

example, see the isolated region of low T1 value in the ventricle region (Figure 3B), which is visible in the proposed method. This spatial feature is invisible when producing the same T1 map at $N_b = 10k$ steps (Figure 6).

To address this, we set the level of proposed hyperparameters to ensure proper convergence of all spatial image features, that is, until some noise appears in the final output.

For all tested applications and datasets, the results were obtained with identical network hyperparameters except the number of optimization steps, N_b , which were specifically selected for each application, mainly due to differences in the number of total voxels in the magnitude signal data which affects the convergence rate in the optimization process. The results suggest satisfactory denoising can be achieved in several applications with limited hyperparameter tuning.

The proposed method performs well when compared to other denoising methods, both standard methods (TV, WAVE, NLM), a more modern alternative (BM3D), and DD, an under-parametrized nonconvolutional deep learning alternative to DIP. OURS seems visually most similar

to TV, with significantly reduced noise and uncertainty in parameter mapping.

For model uncertainty calibration, we utilized a metric for ground-truth uncertainty: the actual spread in estimated parameter maps for different noise realizations (Equation 10). Comparing this to our uncertainty (Figure 5G,H) shows the extent of this issue, mainly errors in magnitude since all spatial features are still captured regardless. The suggested calibration procedure mitigates this issue by incorporating magnitude correction from the estimated tissue parameter map (Figure 5G,I, and uncertainty column in Table 2). The proposed calibration approach mitigates issues with incorrect uncertainty in regions of high data magnitude, for example, in CSF for T1 and T2. In T2 mapping, the calibration seems to fail at the border region and also to some degree in WM (Figure 5I2), possibly due to high estimation bias in those regions (Figure 4G2). Bias in those regions is also visible with NONE, that is, un-denoised NLLS, possibly relating this finding to insufficient TE selection. The computational cost is essentially unaffected since the calibration is merely a hyperparameters-free multiplication after training. We conclude that more extensive testing is required




FIGURE 6 Deep image prior (DIP)-enhanced parameter mapping with synthetic T1 variable flip-angle data (labeled y). T1 maps with nonlinear least squares without denoising (labeled NONE), the proposed method (labeled OURS), and non-Bayesian quantitative MRI DIP are presented for comparison. In both cases, our method successfully adapts DIP to parameter mapping. The two DIP implementations (Equations 3a and 5a) are presented at an array of different numbers of training iterations (N_b), which confirms that our method is less prone to overfitting noise due to the implementation of Monte Carlo dropout.

for use cases where uncertainty estimation is highly prioritized.

In this work, we aimed to investigate the possibility of addressing noise and uncertainties in parameter mapping by casting it as a DIP denoising task with MC dropout-produced uncertainty estimates. Based on our findings, we have identified the following possible extensions to improve the proposed method: automation of hyperparameter selection and the need for computational cost reduction. Automation of hyperparameter tuning is important to suppress the influence of subjective biases in the quantitative task of tissue parameter mapping. In particular, the number of optimization steps affects the denoising level in the estimated parameter map. Although not automatized, the proposed facilitates the selection of optimization steps due to the MC dropout implementation. MC dropout lowers the convergence rate, making finding a suitable number of steps easier since a larger range of optimization steps produces an acceptable outcome. To take this one step further would be to remove the

manual selection altogether, possibly by tracking convergence during the optimization step with some information criterion.

Addressing the computational cost and hyperparameter selection is an intertwined problem. For example, the generator architecture (Figure 1) depends on several parameters, such as U-Net depth and output filters in each convolution block. The selection of these parameters drastically affects the network size and computational load. We suggest to investigate a data-driven generator selection and to implement more effective pre-initialization of the network weights to reduce computational costs.

The experiments utilized a fixed standard architecture independent of the spatial size of the MRI data. This means that the spatial size of the most downsampled patch (i.e., in the bottleneck of the U-net) differs for different image sizes. Since this generator still produces satisfactory results for different image sizes, the number of parameters of the generator can likely be lowered to decrease computational cost.

TABLE 3 Computational time for the different parameter mapping estimators in this study.

Method	Number of voxels, V	Steps N_b	Time (min)
T1 synthetic:	25k		
OURS		50k	36
non-Bayesian DIP		10k	8
NONE/WAVE/TV		—	0.89/0.87/0.85
NLM/BM3D/DD		—	0.87/0.73/5.7
T2 synthetic:	25k		
OURS		30k	20
NONE/WAVE/TV		—	0.78/0.77/0.73
NLM/BM3D/DD		—	0.80/0.73/5.4
T1 in-vivo:	21k–25k		
OURS		50k	34–35
NONE		—	0.9–1.1
ADC in vivo:	11k		
OURS		15k	6
NONE		—	0.5

Notes: V is the number of active voxels in the selected slice (rounded to the nearest thousand), that is, the sum of all voxels minus voxels in background regions with nonrelevant MRI signal value. N_b is the number of optimization steps selected for DIP-network training. The non-Bayesian DIP approach severely overfitted at 50 000 (k) iterations. Therefore, we present the estimation time when stopping the process at 10k iterations, where it outputs a more reasonable level of noise reduction, see Figure 6.

One promising approach would be to investigate the DD architecture thoroughly. In contrast to the proposed over-parametrized DIP generator, DD is under-parametrized and, therefore, inherently efficient in suppressing overfitting when implemented with correct hyperparameters. Although we address overfitting with MC dropout, DD is still interesting, especially for reducing computational cost. The experiments show that DD signal denoising was significantly faster than the proposed method and the non-Bayesian DIP (see Table 3).

Regarding generator weight initialization, the experiments show that a significant portion of the optimization time is when the generator tries to find the correct output range (Figure 6, N_b up to 10k iteration). This fraction can likely be decreased by task-specific pre-initialization of the network weights.

Also, the estimation model used in this study was assumed Gaussian, but all experiments were conducted with either in vivo data or Ricean simulated noise in synthetic data. The Gaussian approximation model has very low bias at $\text{SNR} > 2$,⁴⁰ which was the case in the experiments. However, for very low signal-to-noise ratio applications, one should implement Ricean modeling to avoid biased results.

5 | CONCLUSIONS

We have implemented a DIP-based framework incorporating denoising and uncertainty estimation into the parameter mapping process. The implementation handles several applications with limited hyperparameter tuning. Although time-consuming, uncertainty information from MC dropout makes the method more robust and provides useful information when properly calibrated. We conclude that DIP successfully denoise the parameter mapping process and is easy to use since DIP methods do not use any network training data.

ACKNOWLEDGMENTS

This research was partially funded by grants from The Swedish Research Council (Grant No. 2019-0432), Region Västerbotten (Bas ALF, Grant No. RV-970119), Cancerforskningsfonden i Norrland (Grant No. AMP 18-912), and Lions Cancerforskningsfond (Grant No. LP 18-2182 and No. LP 22-2319).

CONFLICT OF INTEREST STATEMENT

Tommy Löfstedt and Anders Garpebring are co-owners of NONPI Medical AB, the developer of Hero Imaging—the medical image analysis software used in this work.

ORCID

Max Hellström  <https://orcid.org/0000-0002-0200-6567>

Tommy Löfstedt  <https://orcid.org/0000-0001-7119-7646>

Anders Garpebring  <https://orcid.org/0000-0002-0532-232X>

REFERENCES

1. Cercignani M, Dowell NG, Tofts PS, eds. *Quantitative MRI of the Brain. Principles of Physical Measurement*. CRC Press; 2018.
2. Fütterer JJ. Multiparametric MRI in the detection of clinically significant prostate cancer. *Korean J Radiol*. 2017;18:597-606. doi:10.3348/kjr.2017.18.4.597
3. Tartar M, Comstock CE, Kipper MS, eds. Chapter 2 - evaluation of the symptomatic patient: diagnostic breast imaging. *Breast Cancer Imaging*. Mosby; 2008. doi:10.1016/B978-0-323-04677-0.50008-6
4. Mai J, Abubrig M, Lehmann T, et al. T2 mapping in prostate cancer. *Invest Radiol*. 2019;54:146-152. doi:10.1097/RLI.0000000000000520
5. Baseline F, Ferraioli G, Pascazio V. A Bayesian approach for relaxation times estimation in MRI. *Magn Reson Imaging*. 2016;34:312-325. doi:10.1016/j.mri.2015.10.020
6. Kumar D, Nguyen TD, Gauthier SA, Raj A. Bayesian algorithm using spatial priors for multiexponential T2 relaxometry from multiecho spin echo MRI. *Magn Reson Med*. 2012;68:1536-1543. doi:10.1002/mrm.24170
7. Raj A, Pandya S, Shen X, LoCastro E, Nguyen TD, Gauthier SA. Multi-compartment T2 relaxometry using a spatially constrained multi-Gaussian model. *PloS One*. 2014;9:e98391. doi:10.1371/journal.pone.0098391
8. 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. *Phys Med Biol*. 2020;65:225036. doi:10.1088/1361-6560/abb9f5
9. Beirinckx Q, Jeurissen B, Nicastrò M, et al. Model-based super-resolution reconstruction with joint motion estimation for improved quantitative MRI parameter mapping. *Comput Med Imaging Graph*. 2022;100:102071. doi:10.1016/j.compmedimag.2022.102071
10. While PT. A comparative simulation study of bayesian fitting approaches to intravoxel incoherent motion modeling in diffusion-weighted MRI. *Magn Reson Med*. 2017;78:2373-2387. doi:10.1002/mrm.26598
11. Orton MR, Collins DJ, Koh DM, Leach MO. Improved intravoxel incoherent motion analysis of diffusion weighted imaging by data driven Bayesian modeling. *Magn Reson Med*. 2014;71:411-420. doi:10.1002/mrm.24649
12. Baseline F, Ferraioli G, Shabou A. Field map reconstruction in magnetic resonance imaging using Bayesian estimation. *Sensors (Basel)*. 2010;10:266-279. doi:10.3390/s100100266
13. Bouhrara M, Spencer RG. Improved determination of the myelin water fraction in human brain using magnetic resonance imaging through Bayesian analysis of mcDESPOT. *Neuroimage*. 2016;127:456-471. doi:10.1016/j.neuroimage.2015.10.034
14. Schmid VJ, Whitcher B, Padhani AR, Taylor NJ, Yang GZ. Bayesian methods for pharmacokinetic models in dynamic contrast-enhanced magnetic resonance imaging. *IEEE Trans Med Imaging*. 2006;25:1627-1636. doi:10.1109/tmi.2006.884210
15. Kelm BM, Menze BH, Nix O, Zechmann CM, Hamprecht FA. Estimating kinetic parameter maps from dynamic contrast-enhanced MRI using spatial prior knowledge. *IEEE Trans Med Imaging*. 2009;28:1534-1547. doi:10.1109/TMI.2009.2019957
16. Sommer JC, Schmid VJ. Spatial two-tissue compartment model for dynamic contrast-enhanced magnetic resonance imaging. *J R Stat Soc Ser C Appl Stat*. 2014;63:695-713.
17. Grussu F, Battiston M, Veraart J, et al. Multi-parametric quantitative in vivo spinal cord MRI with unified signal readout and image denoising. *Neuroimage*. 2020;217:116884. doi:10.1016/j.neuroimage.2020.116884
18. Bustin A, Ferry P, Codreanu A, et al. Impact of denoising on precision and accuracy of saturation-recovery-based myocardial T1 mapping. *J Magn Reson Imaging*. 2017;46:1377-1388. doi:10.1002/jmri.25684
19. Nordio G, Bustin A, Henningsson M, et al. SASHA myocardial T1 mapping with high accuracy and improved precision [published correction appears in MAGMA. 2018 Oct 8]. *Magma*. 2019;32:281-289. doi:10.1007/s10334-018-0703-y
20. Nordio G, Bustin A, Odille F, et al. Faster 3D saturation-recovery based myocardial T1 mapping using a reduced number of saturation points and denoising. *PloS One*. 2020;15:e0221071. doi:10.1371/journal.pone.0221071
21. Fu Z, Mandava S, Keerthivasan MB, et al. A multi-scale residual network for accelerated radial MR parameter mapping. *Magn Reson Imaging*. 2020;73:152-162. doi:10.1016/j.mri.2020.08.013
22. Cai C, Wang C, Zeng Y, et al. Single-shot T2 mapping using overlapping-echo detachment planar imaging and a deep convolutional neural network. *Magn Reson Med*. 2018;80:2202-2214. doi:10.1002/mrm.27205
23. Liu F, Feng L, Kijowski R. MANTIS: model-augmented neural network with incoherent k-space sampling for efficient MR parameter mapping. *Magn Reson Med*. 2019;82:174-188. doi:10.1002/mrm.27707
24. Liu F, Kijowski R, Feng L, El Fakhri G. High-performance rapid MR parameter mapping using model-based deep adversarial learning. *Magn Reson Med*. 2020;74:152-160. doi:10.1016/j.mri.2020.09.021
25. Liu F, Kijowski R, El Fakhri G, Feng L. Magnetic resonance parameter mapping using model-guided self-supervised deep learning. *Magn Reson Med*. 2021;85:3211-3226. doi:10.1002/mrm.28659
26. Jun Y, Shin H, Eo T, Kim T, Hwang D. Deep model-based magnetic resonance parameter mapping network (DOPAMINE) for fast T1 mapping using variable flip angle method. *Med Image Anal*. 2021;70:102017. doi:10.1016/j.media.2021.102017
27. Mastropietro A, Procissi D, Scalco E, Rizzo G, Bertolino N. A supervised deep neural network approach with standardized targets for enhanced accuracy of IVIM parameter estimation from multi-SNR images. *NMR Biomed*. 2022;35:e4774. doi:10.1002/nbm.4774
28. Huang HM. An unsupervised convolutional neural network method for estimation of intravoxel incoherent motion parameters. *Phys Med Biol*. 2022;67:215018. doi:10.1088/1361-6560/ac9a1f

29. Bliesener Y, Acharya J, Nayak KS. Efficient DCE-MRI parameter and uncertainty estimation using a neural network. *IEEE Trans Med Imaging*. 2020;39:1712-1723. doi:10.1109/TMI.2019.2953901
30. Zormpas-Petridis K, Tunariu N, Collins DJ, Messiou C, Koh DM, Blackledge MD. Deep-learned estimation of uncertainty in measurements of apparent diffusion coefficient from whole-body diffusion-weighted MRI. *Comput Biol Med*. 2022;149:106091. doi:10.1016/j.combiomed.2022.106091
31. Ulyanov D, Vedaldi A, Lempitsky V. Deep image prior. *Int J Comput Vis*. 2020;128:1867-1888. doi:10.1007/s11263-020-01303-4
32. Cheng Z, Gadelha M, Maji S, Sheldon D. A Bayesian perspective on the deep image prior. Paper presented at: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR); 2019; Long Beach, CA:5438-5446. doi:10.1109/CVPR.2019.00559
33. Gong K, Catana C, Qi J, Li Q. PET image reconstruction using deep image prior. *IEEE Trans Med Imaging*. 2019;38:1655-1665. doi:10.1109/TMI.2018.2888491
34. Lin YC, Huang HM. Denoising of multi b-value diffusion-weighted MR images using deep image prior. *Phys Med Biol*. 2020;65:105003. doi:10.1088/1361-6560/ab8105
35. Yoo J, Jin KH, Gupta H, Yerly J, Stuber M, Unser M. Time-dependent deep image prior for dynamic MRI. *IEEE Trans Med Imaging*. 2021;40:3337-3348. doi:10.1109/TMI.2021.3084288
36. Gal Y, Ghahramani Z. Dropout as a Bayesian approximation: representing model uncertainty in deep learning. *arXiv:1506.02142* [stat.ML]. 2015.
37. Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R. Dropout: a simple way to prevent neural networks from overfitting. *J Mach Learn Res*. 2014;15:1929-1958.
38. Gal Y, Ghahramani Z. Bayesian convolutional neural networks with Bernoulli approximate variational inference. *arXiv:1506.02158* [stat.ML]. 2015.
39. Laves MH, Tölle M, Ortmaier T. Uncertainty estimation in medical image denoising with Bayesian deep image prior. *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Graphs in Biomedical Image Analysis. UNSURE GRAIL 2020. Lecture Notes in Computer Science*. Vol 12443. Springer; 2020. doi:10.1007/978-3-030-60365-6_9
40. Gudbjartsson H, Patz S. The Rician distribution of noisy MRI data. *Magn Reson Med*. 1995;34:910-914. doi:10.1002/mrm.1910340618
41. Fram EK, Herfkens RJ, Johnson GA, et al. Rapid calculation of T1 using variable flip angle gradient refocused imaging. *Magn Reson Imaging*. 1987;5:201-208. doi:10.1016/0730-725x(87)90021-x
42. Woermann FG, Barker GJ, Birnie KD, Meencke HJ, Duncan JS. Regional changes in hippocampal T2 relaxation and volume: a quantitative magnetic resonance imaging study of hippocampal sclerosis. *J Neurol Neurosurg Psychiatry*. 1998;65:656-664. doi:10.1136/jnnp.65.5.656
43. Kingma DP, Ba J. Adam: a method for stochastic optimization. *arXiv:1412.6980* [cs.LG]. 2014.
44. Kwan RKS, Evans AC, Pike GB. An extensible MRI simulator for post-processing evaluation. In: Heinz HK, Ron K, eds. *Visualization in Biomedical Computing*. Vol 135-140. Springer; 1996.
45. Cocosco CA, Kollokian V, Kwan RKS, Evans AC. BrainWeb: online Interface to a 3D MRI simulated brain database. *Neuroimage*. 1997;5:S425. Proceedings of 3rd International Conference on Functional Mapping of the Human Brain, Copenhagen, May 1997.
46. Collins DL, Zijdenbos AP, Kollokian V, et al. Design and construction of a realistic digital brain phantom. *IEEE Trans Med Imaging*. 1998;17:463-468. doi:10.1109/42.712135
47. Kwan RK, Evans AC, Pike GB. MRI simulation-based evaluation of image-processing and classification methods. *IEEE Trans Med Imaging*. 1999;18:1085-1097. doi:10.1109/42.816072
48. Kay SM. *Fundamentals of Statistical Signal Processing: Estimation Theory*. 1st ed. Pearson; 1993.
49. Dabov K, Foi A, Katkovnik V, Egiazarian K. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Trans Image Process*. 2007;16:2080-2095. doi:10.1109/tip.2007.901238
50. Makinen Y, Azzari L, Foi A. Collaborative filtering of correlated noise: exact transform-domain variance for improved shrinkage and patch matching. *IEEE Trans Image Process*. 2020;29:8339-8354. doi:10.1109/TIP.2020.3014721
51. Gal Y, Ghahramani Z. Dropout as a Bayesian approximation: appendix. *arXiv:1506.02157* [stat.ML]. 2015.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

Figure S1. ADC in-vivo results

Figure S2. T1 in-vivo results

How to cite this article: Hellström M, Löfstedt T, Garpebring A. Denoising and uncertainty estimation in parameter mapping with approximate Bayesian deep image priors. *Magn Reson Med*. 2023;90:2557-2571. doi: 10.1002/mrm.29823