



# HHS Public Access

Author manuscript

*IEEE Trans Med Imaging*. Author manuscript; available in PMC 2023 February 02.

Published in final edited form as:

*IEEE Trans Med Imaging*. 2022 February ; 41(2): 279–291. doi:10.1109/TMI.2021.3112040.

## Bayesian Uncertainty Estimation of Learned Variational MRI Reconstruction

**Dominik Narnhofer,**

Institute of Computer Graphics and Vision, Graz University of Technology, Graz, Austria

**Alexander Effland,**

Institute of Applied Mathematics, University of Bonn, Bonn, Germany, and also with Silicon Austria Labs (TU Graz SAL DES Lab).

**Erich Kobler,**

Computer Science Department, Johannes Kepler University, Linz, Austria.

**Kerstin Hammernik,**

Technical University of Munich, Munich, Germany and Imperial College London, London, United Kingdom.

**Florian Knoll,**

NYU School of Medicine, New York, USA.

**Thomas Pock**

Institute of Computer Graphics and Vision, Graz University of Technology, Graz, Austria

### Abstract

Recent deep learning approaches focus on improving quantitative scores of dedicated benchmarks, and therefore only reduce the observation-related (aleatoric) uncertainty. However, the model-immanent (epistemic) uncertainty is less frequently systematically analyzed. In this work, we introduce a Bayesian variational framework to quantify the epistemic uncertainty. To this end, we solve the linear inverse problem of undersampled MRI reconstruction in a variational setting. The associated energy functional is composed of a data fidelity term and the total deep variation (TDV) as a learned parametric regularizer. To estimate the epistemic uncertainty we draw the parameters of the TDV regularizer from a multivariate Gaussian distribution, whose mean and covariance matrix are learned in a stochastic optimal control problem. In several numerical experiments, we demonstrate that our approach yields competitive results for undersampled MRI reconstruction. Moreover, we can accurately quantify the pixelwise epistemic uncertainty, which can serve radiologists as an additional resource to visualize reconstruction reliability.

### Keywords

magnetic resonance imaging; undersampled MRI; total deep variation; convolutional neural network; image reconstruction; epistemic uncertainty estimation; Bayes' theorem

## I. INTRODUCTION

A classical inverse problem related to magnetic resonance imaging (MRI) emerges from the undersampling of the raw data in Fourier domain (known as  $k$ -space) to reduce acquisition time. When directly applying the inverse Fourier transform, the quality of the resulting image is deteriorated by undersampling artifacts since in general the sampling rate does not satisfy the Nyquist–Shannon sampling theorem. Prominent approaches to reduce these artifacts incorporate parallel imaging [3] on the hardware side, or compressed sensing on the algorithmic side [4]. In further algorithmic approaches, the MRI undersampling problem is cast as an ill-posed inverse problem using a hand-crafted total variation-based regularizer [5]. In recent years, a variety of deep learning-based methods for general inverse problems have been proposed that can be adapted for undersampled MRI reconstruction, including deep artifact correction [6], learned unrolled optimization [7]–[9], or  $k$ -space interpolation learning [10]. We refer the interested reader to [11], [12] for an overview of existing methods and their applicability to MRI.

An established technique for solving ill-posed inverse problems are variational methods, in which the minimizer of an energy functional defines the restored output image. A probabilistic interpretation of variational methods is motivated by Bayes' theorem, which states that the posterior distribution  $p(x | z)$  of a reconstruction  $x$  and observed data  $z$  is proportional to  $p(z | x)p(x)$ . The maximum a posteriori (MAP) estimate [13] in a negative log-domain is the minimizer of the energy

$$\mathcal{E}(x, z) = \mathcal{D}(x, z) + \mathcal{R}(x) \quad (1)$$

among all  $x$ , where we define the data fidelity term as  $\mathcal{D}(x, z) \propto -\log(p(z | x))$  and the regularizer as  $\mathcal{R}(x) \propto -\log(p(x))$ . Deep learning has been successfully integrated in this approach in a variety of papers [9], [14], [15], in which the regularizer is learned from data. However, none of these publications addresses the quantification of the uncertainty in the model itself.

In general, two sources of uncertainties exist: aleatoric and epistemic. The former quantifies the uncertainty caused by observation-related errors, while the latter measures the inherent error of the model. Most of the aforementioned methods generate visually impressive reconstructions and thus reduce the aleatoric uncertainty, but the problem of quantifying the epistemic uncertainty is commonly not addressed. In practice, an accurate estimation of the epistemic uncertainty is vital for the identification of regions that cannot be reliably reconstructed such as hallucinated patterns, which could potentially result in a misdiagnosis. This problem has been addressed in a Bayesian setting by several approaches [16]–[19], but only very few are explicitly targeting uncertainty quantification for MRI. For instance, Schlemper et al. [20] use Monte Carlo-dropout and a heteroscedastic loss to estimate MRI reconstruction uncertainty in U-Net/DC-CNN based models. In contrast, Edupugantiet et al. [21] advocated a probabilistic variational autoencoder, which in combination with a Monte Carlo approach and Stein's unbiased risk estimator allows for a pixelwise uncertainty estimation.

There are two main contributions of this paper: First, we adapt the total deep variation (TDV) [15], [22] designed as a novel framework for general linear inverse problems to undersampled MRI reconstruction, and show that we achieve competitive results on the fastMRI data set [23], [24]. In detail, we apply our method to single and multi-coil undersampled MRI reconstruction, where in the latter case no coil sensitivities are used. Second, by roughly following the Bayes by Backprop framework [16] we pursue Bayesian inference and thus estimate the epistemic uncertainty in a pixelwise way (see Fig. 1 for an illustration). In detail, we draw the parameters of TDV from a multivariate Gaussian distribution, whose mean and covariance matrix are computed in an optimal control problem modeling training. By iteratively drawing the parameters from this distribution we can visualize the pixelwise standard deviation of the reconstructions, which measures the epistemic uncertainty. Ultimately, this visualization can aid clinical scientists to identify regions with potentially improper reconstructions.

## II. METHODS

In this section, we first recall the mathematical setting of undersampled MRI reconstruction. Then, we introduce the sampled optimal control problem for deterministic and stochastic MRI reconstruction, where the latter additionally allows for an estimation of the epistemic uncertainty.

### A. Magnetic Resonance Imaging

In what follows, we briefly recall the basic mathematical concepts of (undersampled) MRI. We refer the reader to [25] for further details.

Fully sampled raw data  $u \in \mathbb{C}^{nQ}$  are acquired in the Fourier domain commonly known as  $k$ -space with  $Q - 1$  measurement coils. Throughout this paper, the full resolution images and raw data are of size  $n = \text{width} \times \text{height}$  and are identified with vectors in  $\mathbb{C}^{nQ}$  (e.g.  $u \in \mathbb{C}^{nQ}$ ). If  $Q = 1$ , we refer to a single-coil, in all other cases to a multi-coil MRI reconstruction problem. Here, the associated uncorrupted data in image domain representing the ground truth are given by  $y = F^{-1}u \in \mathbb{C}^{nQ}$ , where  $F \in \mathbb{C}^{nQ \times nQ}$  denotes the channel-wise unitary matrix representation of the two-dimensional discrete Fourier transform and  $F^{-1}$  its inverse. In this case, the final root-sum-of-square image estimate  $Y$  for  $y$  with a resolution of  $n = \text{width} \times \text{height}$  is retrieved as

$$Y_i = \sqrt{\sum_{q=1}^Q |y_{i,q}|^2}$$

for  $i = 1, \dots, n$ , where  $y_{i,q}$  refers to the  $i^{\text{th}}$  pixel value of the  $q^{\text{th}}$  coil and  $|\cdot|$  denotes the absolute value or magnitude. Henceforth, we frequently denote the root-sum-of-square of an image by upper case letters. Acquiring the entirety of the Fourier space data (known as fully sampled MRI) results in long acquisition times and consequently a low patient throughput. To address this issue, a subset of the  $k$ -space data along lines defined by a certain sampling pattern is acquired. However, this approach violates the Nyquist–Shannon

theorem, which results in clearly visible backfolding artifacts. The aforementioned scheme is numerically realized by a downsampling operator  $M_R \in \mathbb{C}^{(nQ/R) \times nQ}$  representing  $R$ -fold Cartesian undersampling ( $R \in \mathbb{N}$ ), which only preserves  $\frac{1}{R}$  of the lines in frequency encoding direction. In this case, the linear forward operator is defined as  $A = M_R F \in \mathbb{C}^{(nQ/R) \times nQ}$ . Thus, the observations resulting from the forward formulation of the inverse problem are given by

$$z = Ay + v \in \mathbb{C}^{nQ/R}, \quad (2)$$

where  $v \in \mathbb{C}^{nQ/R}$  is additive noise.

## B. Deterministic MRI Reconstruction

The starting point of the proposed framework is a variant of the energy formulation (1). In this paper, we use the specific data fidelity term

$$\mathcal{D}(x, z) = \frac{1}{2} \|Ax - z\|_2^2.$$

The data-driven regularizer  $\mathcal{R}: \mathbb{R}^{2nQ} \times \Theta \rightarrow \mathbb{R}_0^+$  depends on the learned parameters  $\theta \in \Theta \subset \mathbb{R}^p$ , where  $\Theta$  is the space of admissible learned parameters. Note that we use the identification  $\mathbb{C} \cong \mathbb{R}^2$  to handle complex numbers. We emphasize that in our case the regularizer is *not* iteration-dependent, which implies that the learned parameters  $\theta$  are shared among all iterations leading to much fewer parameters compared to a scheme where each iteration has individual parameters. Throughout all numerical experiments, we use the total deep variation introduced in section II-D. Our approach is not exclusively designed for the total deep variation, which can consequently be replaced by any parametric regularizer.

In what follows, we model the training process as a sampled optimal control problem [28]. To this end, let  $(y^i, z^i)_{i=1}^I \in \mathbb{C}^{nQ} \times \mathbb{C}^{nQ/R}$  be a collection of  $I$  pairs of uncorrupted data  $y^i$  in image domain and associated observed  $R$ -fold undersampled  $k$ -space data  $z^i$  for  $i = 1, \dots, I$ , where both are related by (2).

Next, we approximate the MAP estimator of  $\mathcal{G}$  (1) w.r.t.  $x$ . To this end, we use a proximal gradient scheme to increase numerical stability [34], which is equivalent to an explicit step in the regularizer and an implicit step in the data fidelity term. We recall that the proximal map of a function  $g$  with step size  $h > 0$  is defined as

$$\text{prox}_{hg}(\bar{x}) = \underset{x}{\text{argmin}} \frac{1}{2} \|\bar{x} - x\|_2^2 + hg(x). \quad (3)$$

Unrolling a proximal gradient scheme on (1), we obtain our model

$$x_{s+1} = \text{prox}_{\frac{T}{S}\mathcal{D}}\left(x_s - \frac{T}{S}\nabla_x \mathcal{R}(x_s, \theta)\right), \quad (4)$$

for  $s = 0, \dots, S-1$ . Here,  $S \in \mathbb{N}$  denotes a fixed number of iteration steps,  $T > 0$  is a learned scaling factor, and  $\nabla_x$  denotes the gradient with respect to the  $x$ -component. We define the initial state as  $x_0 = F^{-1}M_R^*z$  and the terminal state of the gradient descent  $x_S$  defines the output of our model. The considered proximal map exhibits the closed-form expression

$$\text{prox}_{\frac{T}{S}\mathcal{D}}(\bar{x}) = F^{-1}\left(\left(\text{Id} + \frac{T}{S}M_R^*M_R\right)^{-1}\left(F\bar{x} + \frac{T}{S}M_R^*z\right)\right),$$

for which we used  $A = M_R F$ . For a detailed computation we refer the reader to section V-A.

Following [14], [15], [22] we cast the training process as a discrete optimal control problem with control parameters  $T$  and  $\theta$ . Optimal control theory was introduced in the machine learning community to rigorously model the training process from a mathematical perspective in [28]. Intuitively, the control parameters, which coincide with the entity of learned parameters, determine the computed output by means of the state equation. During optimization, the control parameters are adjusted such that typically the generated output images are on average as close as possible to the respective ground truth images, where the discrepancy is quantified by the cost functional. In our case, the state equation is given by (4) with initial condition  $x_0 = F^{-1}M_R^*z$ . To define the associated cost functional, we denote by  $x_S(z, T, \theta)$  the terminal state of the state equation using the parameters  $T$  and  $\theta$  and the data  $z$ . Furthermore,  $X_S(z, T, \theta)$  defined as the root-sum-of-square of  $x_S(z, T, \theta)$  coincides with the reconstructed output image. We use the subsequent established loss functional

$$J(y, z, T, \theta) = \|X_S(z, T, \theta) - Y\|_1 + \tau(1 - \text{SSIM}(X_S(z, T, \theta), Y)) \quad (5)$$

for  $\tau > 0$ , which balances the  $\ell^1$ -norm and the SSIM score. Note that the loss functional only incorporates the difference of the magnitudes of the reconstruction  $X_S(z, T, \theta)$  and the target  $Y$ . In the cost functional given by

$$\inf_{T \in \mathbb{R}_0^+, \theta \in \Theta} \frac{1}{I} \sum_{i=1}^I J(y^i, z^i, T, \theta) \quad (6)$$

the discrepancy of the reconstructions and the targets among the entire data set is minimized. For further details we refer the reader to the literature mentioned above.

### C. Bayesian MRI Reconstruction

Inspired by [16], we estimate the epistemic uncertainty of the previous deterministic model by sampling the weights from a learned probability distribution. Here, we advocate the Gaussian distribution as a probability distribution for the parameters, which is justified by

the central limit theorem and has been discussed in several prior publications [29]–[31]. For instance, according to [32] a neural network with only a single layer and a parameter prior with bounded variance converges in the limit of the kernel size to a Gaussian process. For further examples of central limit type convergence estimates for neural networks we refer the reader to the aforementioned literature and the references therein. The second major advantage of this choice is the availability of a closed-form expression of the Kullback–Leibler divergence for Gaussian processes, which is crucial for the efficient proximal optimization scheme introduced below.

We draw the weights  $\theta$  of the regularizer from the multivariate Gaussian distribution  $\mathcal{N}(\mu, \Sigma)$  with a learned mean  $\mu \in \Theta \subset \mathbb{R}^p$  and covariance matrix  $\Sigma \in \mathbb{R}^{p \times p}$ . To decrease the amount of learnable parameter, we reparametrize  $\Sigma = LL^\top \in \mathbb{R}^{p \times p}$ , where  $L \in \mathbb{R}^{p \times p}$  is a learned lower triangular matrix with non-vanishing diagonal entries. In particular,  $\Sigma$  is always positive definite and symmetric. For simplicity, we assume that  $\Sigma$  admits a block diagonal structure, in which the diagonal can be decomposed into blocks of size  $o^2 \times o^2$ . Here, each block describes the covariance matrix of a single kernel of size  $o \times o$  of a CNN representing the regularizer. In particular, there is no correlation among the kernel weights of different kernels, i.e.  $o = 3$  throughout this work.

Realizations of parameters  $\theta$  can simply be computed by using the reparametrization  $\theta = \mu + Lz$  for  $z \sim \mathcal{N}(0, \text{Id})$ .

As a straightforward approach to model uncertainty, one could minimize (6) w.r.t.  $\mu$  and  $L$ . However, in this case a deterministic minimizer with  $\theta = \mu$  and  $\Sigma$  being the null matrix is retrieved. Thus, to enforce a certain level of uncertainty we include the Kullback–Leibler divergence KL in the loss functional [33]. We recall that KL for two multivariate probability distributions  $p_1$  and  $p_2$  with density functions  $f_1$  and  $f_2$  on a domain  $\Omega$  reads as

$$\text{KL}(p_1 \parallel p_2) = \int_{\Omega} f_1(x) \log \left( \frac{f_1(x)}{f_2(x)} \right) dx.$$

In particular, KL is non-negative and in general non-symmetric, and can be regarded as a discrepancy measure of two probability distributions. In the special case of multivariate Gaussian probability distributions  $p_1 = \mathcal{N}(\mu_1, \Sigma_1)$  and  $p_2 = \mathcal{N}(\mu_2, \Sigma_2)$ , the Kullback–Leibler divergence admits the closed-form expression

$$\begin{aligned} \text{KL}(p_1 \parallel p_2) &= \frac{1}{2} \left( \log \frac{|\Sigma_2|}{|\Sigma_1|} + \text{tr}(\Sigma_2^{-1} \Sigma_1) \right. \\ &\quad \left. + (\mu_2 - \mu_1)^\top \Sigma_2^{-1} (\mu_2 - \mu_1) - d \right). \end{aligned} \quad (7)$$

In this paper, we use the particular choice

$$p_1 = \mathcal{N}(\mu, \Sigma), \quad p_2 = \mathcal{N}(\mu, \alpha^{-1} \text{Id}),$$

where  $\mu$  and  $\Sigma = LL^\top$  are computed during optimization and  $\alpha > 0$  is an a priori given constant. This choice is motivated by the fact that the mean of the sampled weights should be determined in the optimal control problem while the constant  $\alpha$  is essential to control the level of uncertainty in the model.

Here, smaller values of  $\alpha$  enforce higher levels of uncertainty, and in the limit case  $\alpha \rightarrow \infty$  the deterministic model is retrieved.

Neglecting constants and scaling the Kullback–Leibler divergence with  $\beta = 0$  leads to the subsequent stochastic sampled optimal control problem

$$\begin{aligned} & \inf \left( \mathbb{E}_{\theta \sim \mathcal{N}(\mu, LL^\top)} \left[ \frac{1}{I} \sum_{i=1}^I J(y^i, z^i, T, \theta) \right] \right. \\ & \quad \left. + \beta (\alpha \text{tr}(LL^\top) - \log(\det(LL^\top))) \right) \\ & T \in \mathbb{R}_0^+, \mu \in \Theta, L \in \mathbb{R}^{p \times p} \text{ with } \det(L) \neq 0. \end{aligned} \quad (8)$$

Note that (8) coincides with the deterministic model if  $\beta = 0$ . Indeed, in this case the MAP estimate is retrieved which minimizes the functional  $J$  since no uncertainty is promoted. In summary,  $\alpha$  controls the covariance matrix of the Gaussian distribution to which the learned  $\Sigma$  should be close. The parameter  $\beta$  can be regarded as the strength of the penalization to enforce this constraint, and thus also controls the dynamics during optimization.

Next, we are concerned with the minimization of (8). First, we observe that (8) is actually composed of the nonconvex loss function

$$\mathcal{F}(T, \mu, L) := \mathbb{E}_{\theta \sim \mathcal{N}(\mu, LL^\top)} \left[ \frac{1}{I} \sum_{i=1}^I J(y^i, z^i, T, \theta) \right]$$

as well as the convex regularization term

$$f(L) := \beta (\alpha \text{tr}(LL^\top) - \log(\det(LL^\top))).$$

For minimizing the composite loss function (8), we again use a proximal gradient descent-based scheme. A proper optimization crucially relies on different step sizes  $h_u^k$  for each diagonal  $o^2 \times o^2$  block  $l_u$  of  $L$  due to the different magnitudes in the correlation matrix. Here, we used that the block diagonal structure of  $\Sigma$  translates to the corresponding structure in  $L$ , which admits a decomposition into  $o^2 \times o^2$  blocks. Hence, our update scheme is given by

$$\begin{pmatrix} T^{k+1} \\ \mu^{k+1} \\ l_u^{k+1} \end{pmatrix} = \begin{pmatrix} T^k - h_T^k \nabla_T \mathcal{F}(T^k, \mu^k, L^k) \\ \mu^k - h_\mu^k \nabla_\mu \mathcal{F}(T^k, \mu^k, L^k) \\ \text{prox}_{h_{l_u}^k} f(l_u^k - h_{l_u}^k \nabla_{L^k} \mathcal{F}(T^k, \mu^k, l_u^k)) \end{pmatrix}$$

for all  $u$ , where the iteration-dependent step sizes  $h_T^k, h_{\mu}^k, h_{l_u}^k > 0$  are adjusted by the ADAM optimizer [35]. On each block  $l_u$ , the proximal map is defined as

$$\text{prox}_{hf}(\bar{l}_u) = \underset{l_u}{\text{argmin}} \frac{1}{2h} \|l_u - \bar{l}_u\|_2^2 + f(l_u), \quad (9)$$

where the minimum is taken among all non-singular lower triangular matrices of size  $o^2 \times o^2$ . Given a regular lower triangular block matrix  $\bar{l} \in \mathbb{R}^{o^2 \times o^2}$ , the proximal map of  $f$  admits the closed-form expression

$$\text{prox}_{hf}(\bar{l})_{ab} = \begin{cases} \frac{\bar{l}_{aa} + \sqrt{\bar{l}_{aa}^2 + 8\beta h(1 + 2\alpha\beta h)}}{2(1 + 2\alpha\beta h)}, & a = b, \\ (1 + 2\alpha\beta h)^{-1} \bar{l}_{ab}, & a \neq b. \end{cases}$$

A detailed computation of this proximal map can be found in section V-B.

Finally, we stress that  $\beta$  determines the level of entropy inherent in the model. We define that the mean entropy  $\bar{H}$  as a measure of uncertainty in the model [36] for the convolutional kernels  $K_1$  and  $K_2$  of all residual blocks as

$$\bar{H}(\Sigma) = \frac{1}{2N_K} \sum_{i=1}^{N_K} \ln(2\pi \det(\Sigma_i)),$$

where  $N_K$  is the total number of stochastic convolutional kernels in the network and  $\Sigma_j$  is the collection of associated covariance matrices of each kernel.

#### D. Total Deep Variation

The data-driven TDV regularizer  $\mathcal{R}(x, \theta)$ , depending on the learned parameters  $\theta \in \Theta \subset \mathbb{R}^p$ , was originally proposed in [15], [22]. In detail,  $\mathcal{R}: \mathbb{R}^{2nQ} \times \Theta \rightarrow \mathbb{R}_0^+$  is computed by summing the pixelwise regularization energy  $r: \mathbb{R}^{2nQ} \times \Theta \rightarrow \mathbb{R}^n$ , i.e.  $\mathcal{R}(x, \theta) = \sum_{l=1}^n r(x, \theta)_l$ , which is defined as  $r(x, \theta) = w\psi(K_0x)$ . Note that we use the identification  $\mathbb{C} \cong \mathbb{R}^2$  to handle complex numbers. The building blocks of  $r$  are as follows:

- $K_0 \in \mathbb{R}^{nm \times 2nQ}$  is the matrix representation of a  $3 \times 3$  convolution kernel with  $m$  feature channels and zero-mean constraint, which enforces an invariance with respect to global shifts,
- $\psi: \mathbb{R}^{nm} \rightarrow \mathbb{R}^{nm}$  is a convolutional neural network (CNN) described below,
- $w \in \mathbb{R}^{n \times nm}$  is the matrix representation of a learned  $1 \times 1$  convolution layer.

Note that  $\theta$  encodes  $K_0$ , and all convolutional weights in  $\psi$  and  $w$ . The CNN  $\psi$  is composed of 3 macroblocks connected by skip connections (Fig. 2, second row), where each macroblock consists of 7 residual blocks (Fig. 2, third row). We remark that the



core architecture is inspired by a U-Net [26] with additional residual connections and a more sophisticated structure. Each residual block has two bias-free  $3 \times 3$  convolution layers  $K_1, K_2 \in \mathbb{R}^{nm \times nm}$  and a smooth log-student-t activation function  $\phi(x) = \frac{1}{2} \log(1 + x^2)$  as depicted in Fig. 2 (last row). This choice of the activation function is motivated by the pioneering work of Mumford and coworkers [27]. The down-/upsampling are realized by learned  $3 \times 3$  convolution layers and transposed convolutions, respectively, both with stride 2. For the stochastic setting, we henceforth assume that  $T$  is fixed and that  $K_0$ , the down- and upsampling operators, and  $w$  are always deterministic.

## E. Numerical Optimization and Training Data

To optimize the optimal control problems in the deterministic (6) and stochastic (8) regime, we use the ADAM algorithm [35] with a batch size of 8, momentum variables  $\beta_1 = 0.5$  and  $\beta_2 = 0.9$ , where the first and second moment estimates are reinitialized after 50000 parameter updates. The initial learning rate is  $10^{-4}$ , which is halved each 50000 iterations, and the total number of iterations is 120000 for  $Q = 1$  and 200000 for  $Q = 2$ . The memory consumption is reduced by randomly extracting patches of size  $96 \times 368$  in frequency encoding direction as advocated by [8]. To further stabilize the algorithm and increase training performance, we start with  $S = 2$  iterations, which is successively incremented by 1 after 7500 iterations. For the same reason, we retrain our model for  $R = 8$  starting from the terminal parameters for  $R = 4$  using 100000 iterations for  $Q = 1$  and 130000 for  $Q = 2$ .

In all experiments, we train our model with the data and the random downsampling operators of the single and multi-coil knee data of the fastMRI data set [23], for which two different acquisition protocols were used: PD (coronal proton density scans) and PD-FS (coronal proton density scans with fat saturation).

Moreover, the undersampling pattern defining  $M_R$  is created on-the-fly by uniformly sampling lines outside the fixed auto calibration area such that only  $\frac{1}{R}$  of the lines are preserved. The associated ground truth images are computed using the emulated single-coil methodology [37] in the single-coil case and the root-sum-of-squares reconstructions in the multi-coil setting, which is consistent with the methodology in [23]. We emphasize that no separate training for different modalities including the acquisition protocol, the field strength and the manufacturer is conducted.

Training and inference were performed on a 20 core 2.4GHz Intel Xeon machine equipped with a NVIDIA Titan V GPU. The entire training of a single example took roughly 10 days/14 days in the deterministic/stochastic case requiring 12GB of GPU memory. In both cases, inference takes 3 seconds per sample and requires 3.2GB of GPU memory.

## F. Stochastic Reconstruction

Next, we discuss how to retrieve estimates of the undersampled MRI reconstruction in the stochastic setting. First, we integrate the learned distribution  $p(\theta)$  of the parameters in the Bayesian formula by noting that  $p(z | x, \theta) = p(z | x)$  as follows:

$$p(x, \theta | z) = \frac{p(z | x)p(x | \theta)p(\theta)}{p(z)}.$$

By marginalizing over  $\theta$  we obtain

$$p(x | z) = \int_{\Theta} \frac{p(z | x)p(x | \theta)p(\theta)}{p(z)} d\theta.$$

In our case, a closed-form solution of the integral is not available due to the non-linearity in the regularizer. However, an approximation of the integral can be obtained by Monte–Carlo sampling

$$p(x | z) \approx \frac{1}{N} \sum_{i=1}^N \frac{p(z | x)p(x | \theta_i)}{p(z)}, \quad (10)$$

where  $\theta_i$  is randomly drawn from the probability distribution  $p(\theta)$ . We refer the reader to [1], [2] for the consistency of this approximation as well as the corresponding convergence rates. Furthermore, for each instance  $\theta_i$  we denote by  $x_S(z, T, \theta^i)$  the approximate optimal solution of the variational problem

$$\underset{x \in \mathbb{C}^{nQ}}{\operatorname{argmax}} p(z | x)p(x | \theta_i) \quad (11)$$

in this setting. Note that  $p(z)$  appearing in (10) does not affect the maximizer, that is why we omit this term in (11).

To retrieve estimates of the undersampled MRI reconstruction in the stochastic setting, we draw  $N \in \mathbb{N}$  instances,  $\underline{\theta}_N = (\theta_1, \dots, \theta_N) \in \Theta^N$  from  $\mathcal{N}(\mu, LL^\top)$ , where  $\mu$  and  $L$  are determined by (8). In a negative logarithmic domain, the maximization problem (11) is equivalent to

$$\underset{x \in \mathbb{C}^{nQ}}{\operatorname{argmin}} \mathcal{D}(x, z) + \mathcal{R}(x, \theta_i),$$

where we have identified the first factor with the data fidelity term and the second factor with the regularizer as above. As before, the approximate minimizer is denoted by  $x_S(z, T, \theta^i)$  and computed as in (4). Then, the average  $\bar{x}^N$  and the corresponding standard deviation  $\sigma^N$  of  $N$  independent

$$\bar{x}_S^N(z, T, \underline{\theta}_N) = \frac{1}{N} \sum_{i=1}^N x_S(z, T, \theta^i),$$

$$(\sigma_S^N(z, T, \underline{\theta}_N))_j^2 = \frac{1}{N} \sum_{i=1}^N \left( (x_S(z, T, \theta^i) - \bar{x}_S^N(z, T, \underline{\theta}_N))_j \right)^2$$

for each pixel  $j = 1, \dots, n$ . In particular,  $\bar{x}_S^N(z, T, \underline{\theta}_N)$  refers to the averaged output image. This approach is a special form of *posterior sampling* and summarized in Fig. 1. Finally, the root-sum-of-square reconstruction  $\bar{X}_S^N(z, T, \underline{\theta}_N)$  of  $\bar{x}_S^N(z, T, \underline{\theta}_N)$  and the corresponding standard deviation are given by

$$\bar{X}_S^N(z, T, \underline{\theta}_N) = \frac{1}{N} \sum_{i=1}^N X_S(z, T, \theta^i), \quad (12)$$

$$(\hat{\sigma}_S^N(z, T, \underline{\theta}_N))_j^2 = \frac{1}{N} \sum_{i=1}^N \left( (X_S(z, T, \theta^i) - \bar{X}_S^N(z, T, \underline{\theta}_N))_j \right)^2,$$

respectively.

### III. NUMERICAL RESULTS

In this section, we present numerical results for single and multi-coil undersampled MRI reconstruction in the deterministic and stochastic setting. In all experiments, we set the initial lower triangular matrix  $L_0 = \sqrt{10^{-3}} \text{Id}$ ,  $\alpha = 10$ ,  $S = 15$  and  $N = 32$ , in all multi-coil results we have  $Q = 15$  coils.

#### A. MRI Reconstruction

Table I lists quantitative results for  $R \in \{4, 8\}$  of the initial zero filling, two state-of-the-art methods (U-Net [23] and iRim [38], values taken from the public leaderboard of the fastMRI challenge (<https://fastmri.org/leaderboards>) for both the deterministic and stochastic version of our approach. We stress that we jointly train our model for all contrasts without any further adaptations. In particular, we did not incorporate any metadata in the training process such as contrast levels, manufacturer or field strength. Moreover, our model exhibits an impressively low number of parameters compared to the competing methods, which have up to 300 times more parameters. Note that although the number of trainable parameters in the stochastic TDV is larger compared to the deterministic version, the number of sampled parameters  $\theta$  used for reconstruction is identical to the deterministic case. All MRI reconstructions are rescaled to the interval  $[0, 1]$  to allow for an easier comparison.

Fig. 3 depicts two prototypic ground truth images of the PD data in the single (first row) and multi-coil (second row) case, the corresponding zero filling results, the deterministic  $X_{15}$  and the mean  $\bar{X}_{15}^{32}$  with  $\beta = 10^{-4}$  for  $Q = 1$  and  $\beta = 7.5 \cdot 10^{-5}$  for  $Q = 15$  (see (12)), and the standard deviation  $\hat{\sigma}_{15}^{32}$  for an undersampling factor of  $R = 4$ . The associated entropy levels are  $\bar{H}(\Sigma) = -11.90$  and  $\bar{H}(\Sigma) = -13.68$ , respectively. In both the deterministic and

the stochastic reconstructions even fine details and structures are clearly visible, and the noise level is substantially reduced compared to the ground truth, which can be seen in the zoom with magnification factor 3. Note that hardly any visual difference is observed in both reconstructions. Clearly, the quality in the single-coil case is inferior to the multi-coil case. Moreover, large values of the standard deviation are concentrated in regions with clearly pronounced texture patterns, which are caused by the lack of data in high-frequency  $k$ -space regions. Thus, the standard deviation can be interpreted as a local measure for the epistemic uncertainty. Since the proximal operator is applied after the update of the regularizer, high values of the standard deviation can only be found in regions where data is unknown.

The  $k$ -space associated with the aforementioned singlecoil case is depicted in Fig. 4. In detail, the leftmost image visualizes the undersampling pattern resulting from the predefined Cartesian downsampling operator  $M_R$ , which yields the zero-filled observation (second image) when combined with the fully sampled raw data (third image), where we plot the magnitudes in a logarithmic scale. The fourth and the fifth image depict the mean  $\bar{x}_{15}^{32}$  and the standard deviation  $\sigma_{15}^{32}$  of the reconstruction in  $k$ -space. As a result, our proposed method accurately retrieves the central star-shaped structures of the  $k$ -space representing essential image features, although the undersampling pattern is still clearly visible. Moreover, the standard deviation peaks in the central star-shaped section when data is missing and thus empirically identifies regions with larger uncertainty.

Likewise, Fig. 5 depicts the corresponding results for PD-FS data and  $R = 4$  using the same entropy levels as before, all other parameters are the same as in the previous Fig. 3. We remark that the signal-to-noise ratio is smaller in PD-FS data than in PD data and thus the reconstructions have a tendency to include more noise and imperfections. The inferior quality compared to PD is also reflected in the higher average intensities of the standard deviations.

Fig. 6 shows the multi-coil reconstruction results for 8-fold undersampling and both data sets in the same arrangement as before, the entropy level is  $\bar{H}(\Sigma) = -31.41$ . As expected, the overall reconstruction quality is quantitatively and qualitatively inferior to the case  $R = 4$ . As before, the difference of the deterministic and the stochastic restored images is relatively small and the standard deviations properly identify regions with higher uncertainties. Finally, Fig. 7 depicts zooms of two different MRI reconstructions ( $R = 4$ , PD), in each row the ground truth, two realizations, the stochastic reconstruction and the standard deviation are visualized. The regions highlighted by the arrows indicate structures and patterns that differ among various samples. The variability of the single realizations can be interpreted as hallucinations, which are properly detected in the corresponding standard deviations. This empirically validates that our proposed method to measure the standard deviation actually quantifies the magnitude of the model-related uncertainty. Fig. 8 contains a visual comparison of our method with selected competitive methods from the fastMRI leader board. As a result, both E2EVN [42] and iRim [38] achieve slightly superior quantitative results at the expense of significantly more learnable parameters. In a qualitative comparison, we observe that our proposed method is capable of retrieving fine details, only

the signal of a few high-frequency patterns is lost. Finally, U-Net [23] results are inferior to the considered competitive methods – both quantitatively and qualitatively.

## B. Covariance Matrices

Fig. 9 contains triplets of color-coded covariance matrices of the convolution layers  $K_2$  in different macroblocks and residual blocks (using the abbreviations  $MB_i$  for  $i = 1, 2, 3$  and  $R_j$  for  $j = 1, 4, 7$ , respectively) in the multi-coil case with  $\beta = 7.5 \cdot 10^{-5}$ . Note that we use different scalings for positive and negative values among each residual block. Each visualized covariance matrix is the mean of the individual covariance matrices of each convolution block  $K_2$  appearing in  $\Sigma$ . The resulting covariance matrices are clearly diagonally dominant with a similar magnitude of the diagonal entries among each residual block, but different magnitudes among different residual blocks. Furthermore, most of the off-diagonal entries significantly differ from 0. As a result, the entries of the covariance matrices associated with the first residual block in each macroblock have a tendency to smaller values compared to the ones of the last residual block. Thus, the uncertainty of the network is primarily aggregated at latter residual blocks within each macroblock, which is in correspondence with error propagation theory: perturbations occurring shortly after the initial period have commonly a larger impact on a dynamical system than perturbations occurring later.

## C. Eigenfunction Analysis

Next, we perform a *nonlinear eigenfunction analysis* [39] following the approach in [14], [22] to heuristically identify local structures that are favorable in terms of energy. Classically, each pair  $(v, \lambda) \in \mathbb{C}^n \setminus \{0\} \times \mathbb{C}$  of eigen-function/eigenvalue for a given matrix  $A \in \mathbb{R}^{n \times n}$  solves  $Av = \lambda v$ , where the eigenvalue can be computed using the Rayleigh quotient  $\frac{v^* Av}{v^* v}$ . Nonlinear eigenfunctions  $v$  for the matrix  $\nabla_v \mathcal{R}(v, \theta)$  satisfy  $\nabla_v \mathcal{R}(v, \theta) = \Lambda(v)v$ , where the generalized Rayleigh quotient defining the corresponding *eigenvalues* is given by

$$\Lambda(v) = \frac{\langle \nabla_v \mathcal{R}(v, \theta), v \rangle}{\|v\|_2^2}.$$

Thus, nonlinear eigenfunctions  $v$  are minimizers of the variational problem

$$\min_{v \in \mathbb{C}^n Q} \frac{1}{2} \|\nabla_v \mathcal{R}(v, \theta) - \Lambda(v)v\|_2^2 \quad (13)$$

subject to a fixed initial image. We exploit Nesterov's projected gradient descent [40] for the optimization in (13). The nonlinear eigenfunctions locally reflect energetically minimal configurations and thus heuristically identify stable patterns that are favored by the model.

Fig. 10 depicts two pairs of reconstructed images  $X$  for the initialization and the corresponding root-sum-of-squares of the eigenfunctions in the deterministic singlecoil case. The resulting eigenfunctions predominantly exhibit piecewise smooth regions, where additional high-frequency stripe patterns and lines in the proximity of bone structures as

well as blood vessels are hallucinated. This behavior originates from two opposing effects: some backfolding artifacts caused by missing high-frequency components in  $k$ -space are removed in our approach, whereas certain high-frequency information are hallucinated.

#### D. Effects of Entropy Level and Averaging

In the final experiment, we analyze the effects of the entropy level and the averaging on the PSNR values. To this end, we draw 32 instances  $\underline{\theta}_{32} = (\theta_1, \dots, \theta_{32}) \in \Theta^{32}$  from  $\mathcal{N}(\mu, \Sigma)$ . Then, for different levels of the entropy enforced by different values of  $\beta$  we calculate the lower and upper bounds of the PSNR values of  $\bar{X}_s^N(z, T, \underline{\theta}_N)$  for  $1 \leq N \leq 32$ , where  $\underline{\theta}_N$  is any subset of  $\underline{\theta}_{32}$  with  $N$  elements. Fig. 11 depicts the resulting color-coded spans for  $N \in \{1, 4, 16, 32\}$  and five different levels of entropy (including the limiting case  $\bar{H}(\Sigma) = -\infty$  in the deterministic case). As a result, the PSNR curves monotonically decrease with higher levels of entropy, but even at the highest entropy level induced by  $\beta = 5 \cdot 10^{-4}$  we observe only a relatively small decrease in the PSNR value. Moreover, the spans of the different averaging processes clearly prove that higher values of  $N$  are beneficial, that is why an averaging among a larger number of realizations should be conducted whenever possible. Finally, we observe that the performance saturates with larger values of  $N$ .

#### E. Limitations

In the following we discuss potential limitations of our approach. Since in each iteration the network parameters have to be drawn from the learned distribution, the training takes longer in the stochastic compared to the deterministic case. Furthermore, to accurately quantify the uncertainty,  $N$  reconstructions have to be computed. This leads to an  $N$ -fold increased reconstruction time compared to the deterministic scheme. Higher levels of uncertainty result in a decrease of the PSNR score, as shown in Fig. 11. However, we believe that the advantage of having an estimate about the uncertainty outweighs the addressed limitations.

## IV. CONCLUSION

In this paper, we proposed a Bayesian framework for uncertainty quantification in single and multi-coil undersampled MRI reconstruction exploiting the total deep variation regularizer. To estimate the epistemic uncertainty, we introduced a stochastic optimal control problem, in which the weights of the regularizer are sampled from a learned multivariate Gaussian distribution.

With the proposed Bayesian framework, we can generate visually appealing reconstruction results alongside a pixelwise estimation of the epistemic uncertainty, which might aid medical scientists and clinicians to revise diagnoses based on structures clearly visible in the standard deviation plots.

## Acknowledgments

The authors acknowledge support from the ERC starting grant HOMOVIS (No. 640156), the US National Institute of Health (R01EB024532, P41 EB017183 and R21 EB027241), and the German Research Foundation under Germany's Excellence Strategy – EXC-2047/1 – 390685813 and – EXC2151 – 390873048.

## V.: APPENDIX

### A. Proximal map of the data fidelity

To derive a closed-form expression of  $\text{prox}_{\mathcal{D}}$ , we first recall that the proximal map of  $\mathcal{D}$  reads as

$$\text{prox}_{\frac{T}{S}\mathcal{D}}(\bar{x}) = \underset{x}{\text{argmin}} \underbrace{\frac{1}{2} \|x - \bar{x}\|_2^2 + \frac{T}{2S} \|Ax - z\|_2^2}_{=: G(x)}.$$

We note that

$$\nabla G(x) = x - \bar{x} + \frac{T}{S} A^* Ax - A^* z,$$

which implies that the first-order optimality condition for  $G$  is

$$\text{prox}_{\frac{T}{S}\mathcal{D}}(\bar{x}) = \left( \text{Id} + \frac{T}{S} A^* A \right)^{-1} (\bar{x} + A^* z).$$

Taking into account  $A = M_R F$  we can expand

$$\text{Id} + \frac{T}{S} A^* A = F^{-1} \left( \text{Id} + \frac{T}{S} M_R^* M_R \right) F,$$

where we exploited  $F^* = F^{-1}$  since  $F$  is unitary. Thus,

$$\text{prox}_{\frac{T}{S}\mathcal{D}}(\bar{x}) = F^{-1} \left( \text{Id} + \frac{T}{S} M_R^* M_R \right)^{-1} F \left( \bar{x} + \frac{T}{S} (F^{-1} M_R^* z) \right),$$

which ultimately leads to

$$\text{prox}_{\frac{T}{S}\mathcal{D}}(\bar{x}) = F^{-1} \left( \left( \text{Id} + \frac{T}{S} M_R^* M_R \right)^{-1} \left( F \bar{x} + \frac{T}{S} M_R^* z \right) \right).$$

Note that the inverse of the diagonal matrix can be computed very efficiently.

### B. Proximal map of Kullback–Leibler divergence

Next, we present a more detailed derivation of the proximal map eq. (9) for

$$f(L) := \beta \left( \text{atr}(LL^\top) - \log(\det(LL^\top)) \right)$$

appearing in (7). We first observe that

$$\text{tr}(LL^T) = \sum_{i,j} L_{ij}L_{ij} = \sum_{i,j} L_{ij}^2,$$

$$\det(LL^T) = \det(L)\det(L^T) = \det(L)^2 = \prod_i L_{ii}^2.$$

Recall that  $L$  admits a block diagonal structure, in which each block  $l_1, \dots, l_U \in \mathbb{R}^{o^2 \times o^2}$  is a regular lower triangular matrix. A straightforward computation reveals  $f(L) = \sum_{u=1}^U f(l_u)$ . Thus, we can restrict to a single block matrix  $l$  and rewrite  $f$  as follows:

$$f(l) = \alpha\beta \sum_{i,j=1}^p l_{ij}^2 - 2\beta \sum_{i=1}^p \log(l_{ii}).$$

The proximal map of the function  $f$  reads as

$$\text{prox}_{hf}(\bar{l}) = \underset{l}{\text{argmin}} \left\{ E(l) := \frac{1}{2h} \|l - \bar{l}\|_2^2 + f(l) \right\},$$

where the minimum is taken among all lower triangular and regular matrices. The gradient of  $E$  is given by

$$(\nabla E(l))_{ab} = \begin{cases} l_{aa}^2(1 + 2\alpha\beta h) - \bar{l}_{aa}l_{aa} - 2\alpha h, & a = b, \\ 2\alpha\beta h l_{ab} + l_{ab} - \bar{l}_{ab}, & a \neq b. \end{cases}$$

Thus, the optimization problem (9) can be optimized component-wise and results in a quadratic equation. Over-all, a closed-form expression for  $f$  reads as

$$\text{prox}_{hf}(\bar{l})_{ab} = \begin{cases} \frac{\bar{l}_{aa} + \sqrt{\bar{l}_{aa}^2 + 8\beta h(1 + 2\alpha\beta h)}}{2(1 + 2\alpha\beta h)}, & a = b, \\ (1 + 2\alpha\beta h)^{-1} \bar{l}_{ab}, & a \neq b. \end{cases}$$

The expression for the proximal map of  $f$  for the entire matrix  $L$  is given as the concatenation of the proximal maps for the individual blocks.

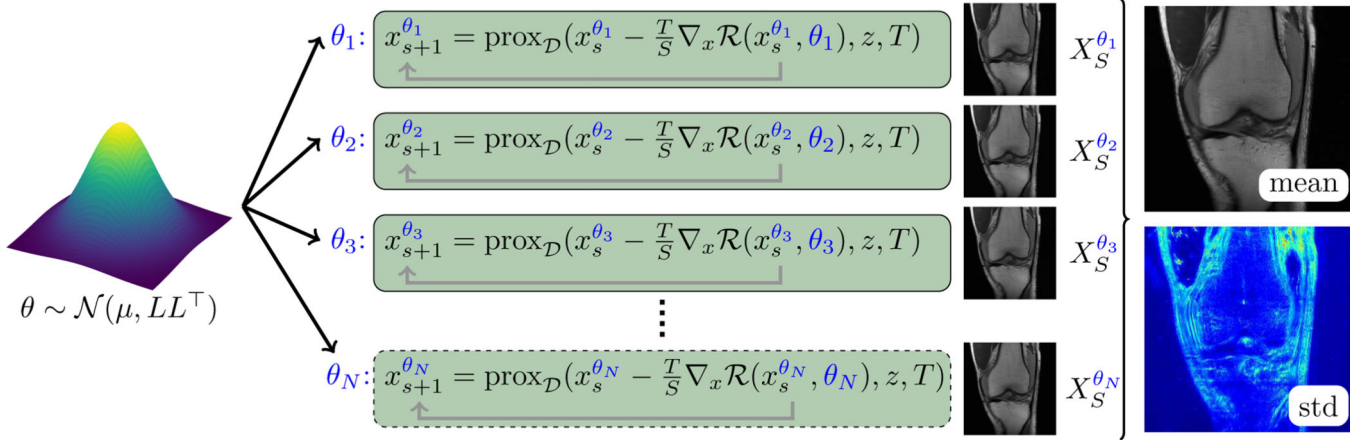
## REFERENCES

- [1]. Barbu A, Zhu S-C, Monte Carlo Methods. Springer Singapore, 2020.
- [2]. Caflisch RE, "Monte Carlo and quasi-Monte Carlo methods," Acta Numer., vol. 7, pp. 1–49, 1998.
- [3]. Pruessmann KP, Weiger M, Scheidegger MB, and Boesiger P, "SENSE: sensitivity encoding for fast MRI," Magn. Reson. Med, vol. 42, no. 5, pp. 952–962, Nov 1996.

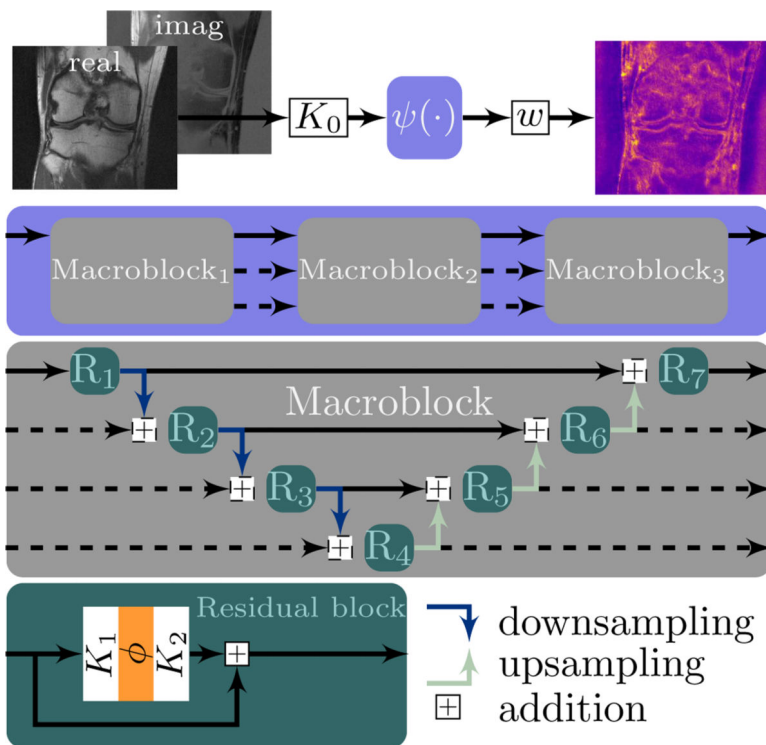


- [4]. Lustig M, Donoho D, and Pauly JM, “Sparse MRI: The application of compressed sensing for rapid MR imaging,” *Magn. Reson. Med.*, vol. 58, no. 6, pp. 1182–1195, Dec 2007, DOI: 10.1002/mrm.21391, [Online]. [PubMed: 17969013]
- [5]. Knoll F, Clason C, Bredies K, Uecker M, and Stollberger R, “Parallel imaging with nonlinear reconstruction using variational penalties,” *Magn. Reson. Med.*, vol. 67, no. 1, pp. 34–41, Jan 2012, DOI: 10.1002/mrm.22964, [Online]. [PubMed: 21710612]
- [6]. Ye JC, Han Y, and Cha E, “Deep convolutional framelets: a general deep learning framework for inverse problems,” *SIAM J. Imaging Sci.*, vol. 11, no. 2, pp. 991–1048, Jan 2018, DOI:10.1137/17M1141771, [Online].
- [7]. Aggarwal HK, Mani MP, and Jacob M, “MoDL: Model-Based Deep Learning Architecture for Inverse Problems,” *IEEE Trans. Med. Imaging*, vol. 38, no. 2, pp. 394–405, Feb 2019, DOI: 10.1109/TMI.2018.2865356, [Online]. [PubMed: 30106719]
- [8]. Schlemper J, Caballero J, Hajnal JV, Price AN, and Rueckert D, “A deep cascade of convolutional neural networks for dynamic mr image reconstruction,” *IEEE Trans. Med. Imaging*, vol. 37, no. 2, pp. 491–503, Feb 2018, DOI: 10.1109/TMI.2017.2760978, [Online]. [PubMed: 29035212]
- [9]. Hammernik K, Klatzer T, Kobler E, Recht MP, Sodickson DK, Pock T, and Knoll F, “Learning a variational network for reconstruction of accelerated MRI data, Learning a Variational Network for Reconstruction of Accelerated MRI Data,” *Magn. Reson. Med.*, vol. 79, no. 6, pp. 3055–3071, Jun 2018, DOI: 10.1002/mrm.26977, [Online]. [PubMed: 29115689]
- [10]. Akçakaya M, Moeller S, Weingärtner S, and Urbil K, “Scan-specific robust artificial-neural-networks for k-space interpolation (RAKI) reconstruction: Database-free deep learning for fast imaging,” *Magn. Reson. Med.*, vol. 81, no. 1, pp. 439–453, Jan 2019, DOI: 10.1002/mrm.27420, [Online]. [PubMed: 30277269]
- [11]. Lundervold AS and Lundervold A, “An overview of deep learning in medical imaging focusing on MRI,” *Zeitschrift für Medizinische Physik*, vol. 29, no. 2, pp. 102–127, May 2019, DOI: 10.1016/j.zemedi.2018.11.002, [Online]. [PubMed: 30553609]
- [12]. Knoll F, Hammernik K, Zhang C, Moeller S, Pock T, Sodickson DK, and Akçakaya M, “Deep-Learning Methods for Parallel Magnetic Resonance Imaging Reconstruction: A Survey of the Current Approaches, Trends, and Issues,” *IEEE Signal Process. Mag.*, vol. 37, no. 1, pp. 128–140, Jan 2020, DOI: 10.1109/MSP.2019.2950640, [Online]. [PubMed: 33758487]
- [13]. Murphy KP, *Machine learning: a probabilistic perspective*, ser. Adaptive Computation and Machine Learning. MIT Press, 2015.
- [14]. Effland A, Kobler E, Kunisch K, and Pock T, “Variational networks: an optimal control approach to early stopping variational methods for image restoration,” *J Math Imaging Vis.*, vol. 62, no. 3, pp. 396–416, Apr. 2020, DOI: 10.1007/s10851019-00926-8, [Online]. [PubMed: 32300264]
- [15]. Kobler E, Effland A, Kunisch K, and Pock T, “Total deep variation for linear inverse problems,” in *CVPR*, 2020.
- [16]. Blundell C, Cornebise J, Kavukcuoglu K, and Wierstra D, “Weight uncertainty in neural networks,” in *ICML*, vol. 37, 2015, pp. 1613–1622.
- [17]. Gal Y and Ghahramani Z, “Dropout as a Bayesian approximation: Representing model uncertainty in deep learning,” in *ICML*, 2016, pp. 1050–1059.
- [18]. Kendall A and Gal Y, “What uncertainties do we need in Bayesian deep learning for computer vision?” in *NIPS*. Curran Associates, Inc., 2017, pp. 5574–5584.
- [19]. Wenzel F, Roth K, Veeling BS, Wiatkowski J, Tran L, Mandt S, Snoek J, Salimans T, Jenatton R, and Nowozin S, “How good is the Bayes posterior in deep neural networks really?” *ICML*, 2020.
- [20]. Schlemper J, Castro DC, Bai W, Qin C, Oktay O, Duan J, Price AN, Hajnal J, and Rueckert D, “Bayesian deep learning for accelerated MR image reconstruction,” in *Machine Learning for Medical Image Reconstruction*, Knoll F, Maier A, and Rueckert D, Eds., Springer International Publishing, 2018, ch. 8, pp. 64–71.
- [21]. Edupuganti V, Mardani M, Vasanawala S, and Pauly J, “Uncertainty quantification in deep MRI reconstruction,” *IEEE Trans. Med. Imaging*, vol. 40, pp. 239–250, Jan. 2021, DOI: 10.1109/TMI.2020.3025065, [Online]. [PubMed: 32956045]
- [22]. Kobler E, Effland A, Kunisch K, and Pock T, “Total deep variation: A stable regularizer for inverse problems,” *arXiv*, 2020.


- [23]. Knoll F, Zbontar J, Sriram A, Muckley MJ, Bruno M, Defazio A, Parente M, Geras KJ, Katsnelson J, Chandarana H, Zhang Z, Drozdal M, Romero A, Rabbat M, Vincent P, Pinkerton J, Wang D, Yakubova N, Owens E, Zitnick LC, Recht MP, Sodickson DK, and Lui YW, “fastMRI: A publicly available raw k-space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning,” *Radiol. Artif. Intell.*, 2020.
- [24]. Knoll F, Murrell T, Sriram A, Yakubova N, Zbontar J, Rabbat M, Defazio A, Muckley MJ, Sodickson DK, Zitnick CL, and Recht MP, “Advancing machine learning for MR image reconstruction with an open competition: Overview of the 2019 fastMRI challenge,” *arXiv*, 2020.
- [25]. Brown RW, Cheng Y-CN, Haacke EM, Thompson MR, and Venkatesan R, *Magnetic resonance imaging: physical principles and sequence design*. John Wiley & Sons Ltd, Apr. 2014.
- [26]. Ronneberger O, Fischer P, and Brox T, “U-Net: Convolutional networks for biomedical image segmentation,” in *MICCAI*. Springer, 2015, pp. 234–241.
- [27]. Huang J and Mumford D, “Statistics of natural images and models,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, vol. 1, Jun. 1999, pp. 541–547.
- [28]. E W, Han J, and Li Q, “A mean-field optimal control formulation of deep learning,” *Res. Math. Sci.*, vol. 6, no. 1, pp. Paper No. 10, 41, Mar. 2019, DOI: 10.1007/s40687-018-0172-y, 6, [Online].
- [29]. Lee J, Bahri Y, Novak R, Schoenholz S, Pennington J, Sohl-Dickstein J “Deep Neural Networks as Gaussian Processes,” *ICLR*, vol. 6, 2018.
- [30]. Fortuin V, “Priors in Bayesian Deep Learning: A Review,” *arXiv*, 2021.
- [31]. Garriga-Alonso A, Rasmussen CE, Aitchison L, “Deep Convolutional Networks as shallow Gaussian Processes,” *ICLR*, vol. 7, 2019.
- [32]. Williams CK, “Computing with infinite networks,” *NIPS*, vol. 8, pp. 295–301, 1996.
- [33]. MacKay DJC, *Information theory, inference and learning algorithms*. Cambridge University Press, New York, 2003.
- [34]. Chambolle A and Pock T, “An introduction to continuous optimization for imaging,” *Acta Numerica*, vol. 25, pp. 161–319, May 2016, DOI: 10.1017/S096249291600009X, [Online]
- [35]. Kingma DP and Ba JL, “ADAM: a method for stochastic optimization,” in *ICLR*, 2015.
- [36]. Shannon CE, “A mathematical theory of communication,” *Bell Syst. Tech. J.*, vol. 27, pp. 379–423, 623–656, July, October 1948.
- [37]. Tygert M and Zbontar J, “Simulating single-coil MRI from the responses of multiple coils,” *arXiv*, 2018.
- [38]. Putzky P, Karkalousos D, Teuwen J, Miriakov N, Bakker B, Caan M, and Welling M, “i-RIM applied to the fastMRI challenge,” *arXiv*, 2020.
- [39]. Gilboa G, *Nonlinear eigenproblems in image processing and computer vision*, ser. *Advances in Computer Vision and Pattern Recognition*. Springer, Cham, 2018.
- [40]. Nesterov YE, “A method of solving a convex programming problem with convergence rate  $O(k^{-2})$ ,” *Dokl. Akad. Nauk SSSR*, vol. 269, no. 3, pp. 543–547, 1983.
- [41]. Hammernik K, Schlemper J, Qin C, Duan J, Summers RM, and Rueckert D, “ $\Sigma$ -net: Systematic evaluation of iterative deep neural networks for fast parallel MR image reconstruction,” *arXiv*, 2019.
- [42]. Sriram A, Zbontar J, Murrell T, Defazio A, Zitnick C, Yakubova N, Knoll F, and Johnson P “End-to-End Variational Networks for Accelerated MRI Reconstruction,” *MICCAI*, pp. 64–73, 2020.




**Fig. 1.** Illustration of the stochastic MRI undersampling reconstruction model to calculate the epistemic uncertainty. Here,  $N$  instances of the model parameters  $\theta$  are drawn from  $\mathcal{N}(\mu, LL^T)$ , which lead to  $N$  output images  $X_S^{\theta_1}, \dots, X_S^{\theta_N}$ . The associated pixelwise mean and standard deviation are depicted on the right.

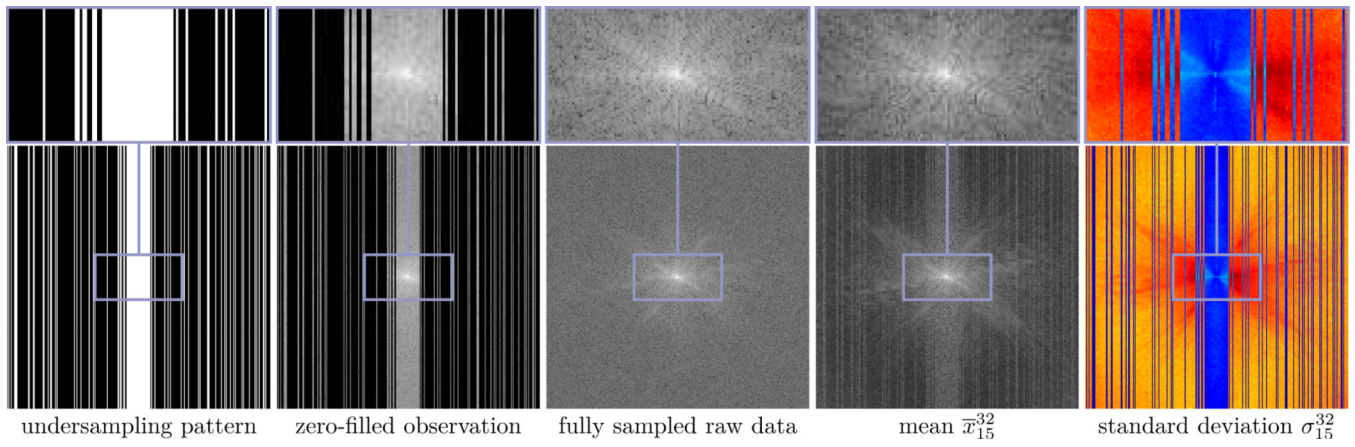



**Fig. 2.** The building blocks of the total deep variation with 3 macroblocks (Figure adapted from [15, Figure 1]). Complex data are transformed to a pixelwise energy as seen in the top right corner.

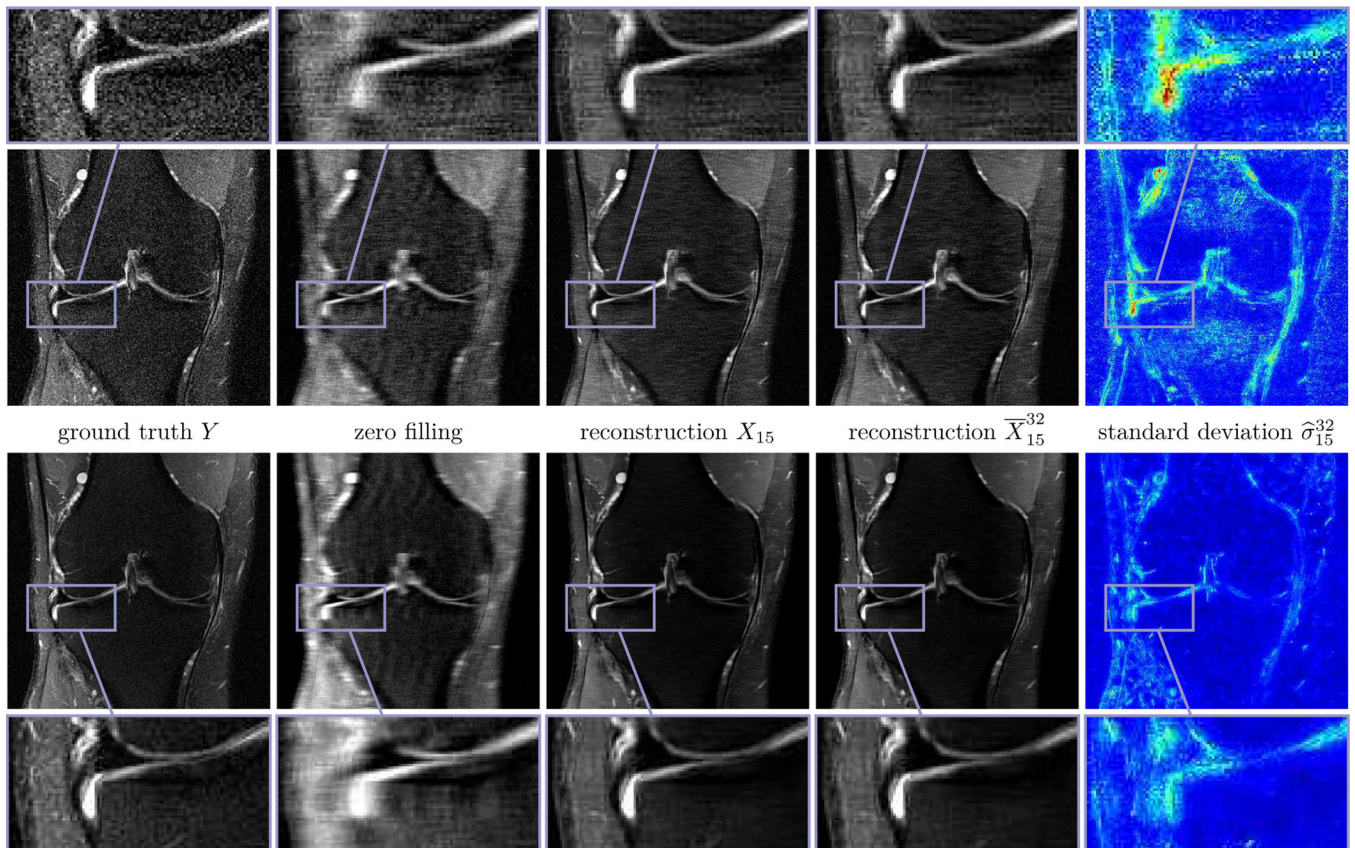



**Fig. 3.** Single (first row) and multi-coil (second row) MRI reconstruction results for PD data and  $R = 4$ . From left to right: ground truth images  $Y$ , zero filling, deterministic reconstructions  $X_{15}$ , stochastic reconstructions  $\bar{X}_{15}^{32}$  and standard deviations  $\hat{\sigma}_{15}^{32}$  (0  0.02).







**Fig. 4.** Visualization of the magnitude images in  $k$ -space (logarithmic scale) for  $R = 4$  in the single-coil case. From left to right: undersampling pattern, zero-filled observation, fully sampled raw data, mean  $\bar{x}_{15}^{32}$ , standard deviation  $\sigma_{15}^{32}$  ( $-14.1$    $-4.74$ ).





**Fig. 5.** Single (first row) and multi-coil (second row) MRI reconstruction results for PD-FS data and  $R = 4$ . From left to right: ground truth images  $Y$ , zero filling, deterministic reconstructions  $X_{15}$ , stochastic reconstructions  $\bar{X}_{15}^{32}$  and standard deviation  $\hat{\sigma}_{15}^{32}$  (0  0.035).

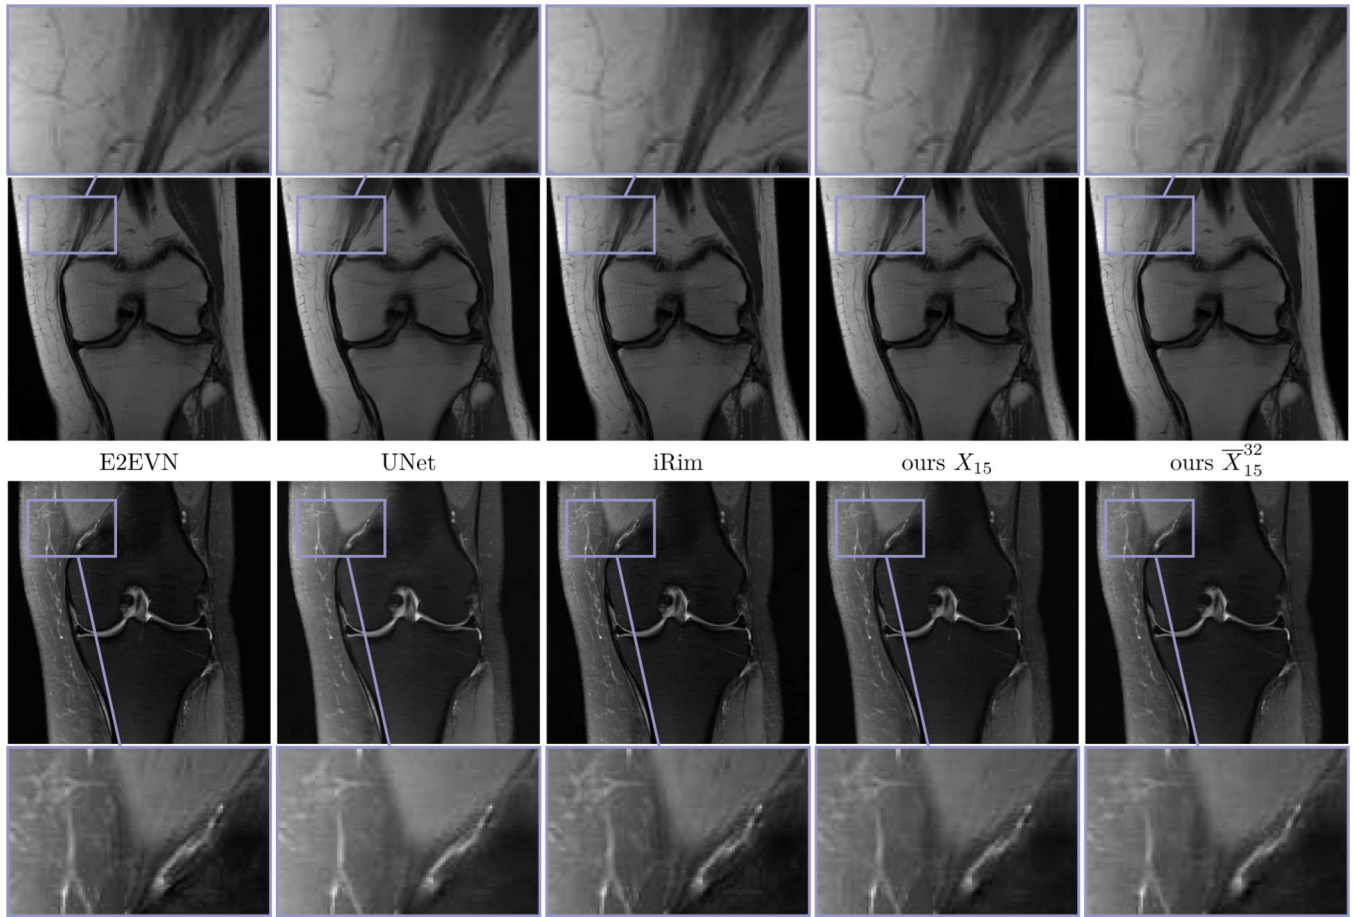


**Fig. 6.** Multi-coil MRI reconstruction results for PD (first row) and PD-FS (second row) data and  $R = 8$ . From left to right: ground truth images  $Y$ , zero filling, deterministic reconstructions  $X_{15}$ , stochastic reconstructions  $\bar{X}_{15}^{32}$  and standard deviation  $\hat{\sigma}_{15}^{32}$  (0  0.02).

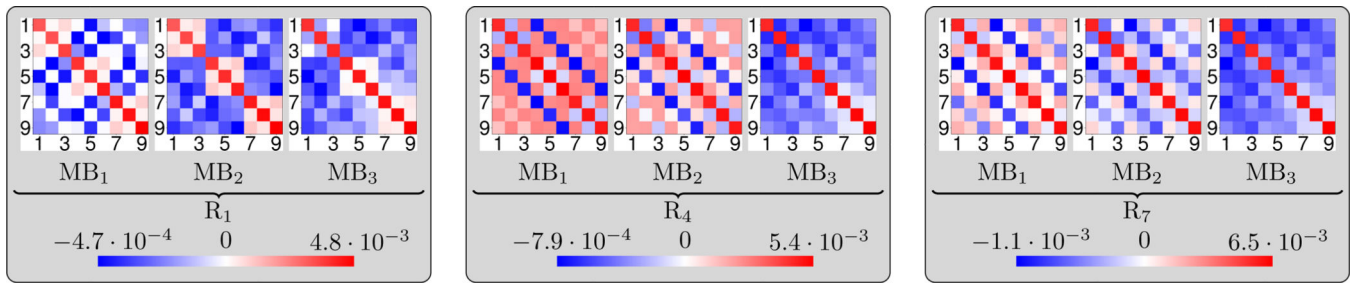




**Fig. 7.** Zooms of multi-coil MRI reconstruction ( $R = 4$ , PD). From left to right: ground truth, two distinct samples, stochastic reconstruction and standard deviation (0  0.03). The arrows highlight patterns that are only visible in distinct samples.



**Fig. 8.** Visual comparison of selected MRI reconstruction methods for PD (first row) and PD-FS (second row), both with  $R = 4$ . From left to right: E2EVN, UNet, iRim, TDV(deterministic), TDV(stochastic).




**Fig. 9.** From left to right: triplets of color-coded covariance matrices of the convolution layers  $\mathbf{K}_2$  in different macroblocks ( $\text{MB}_i$  for  $i = 1, 2, 3$ ) and residual blocks ( $\text{R}_j$  for  $j = 1, 4, 7$ ). Note that we use different scalings for positive and negative values among each residual block.

Author Manuscript

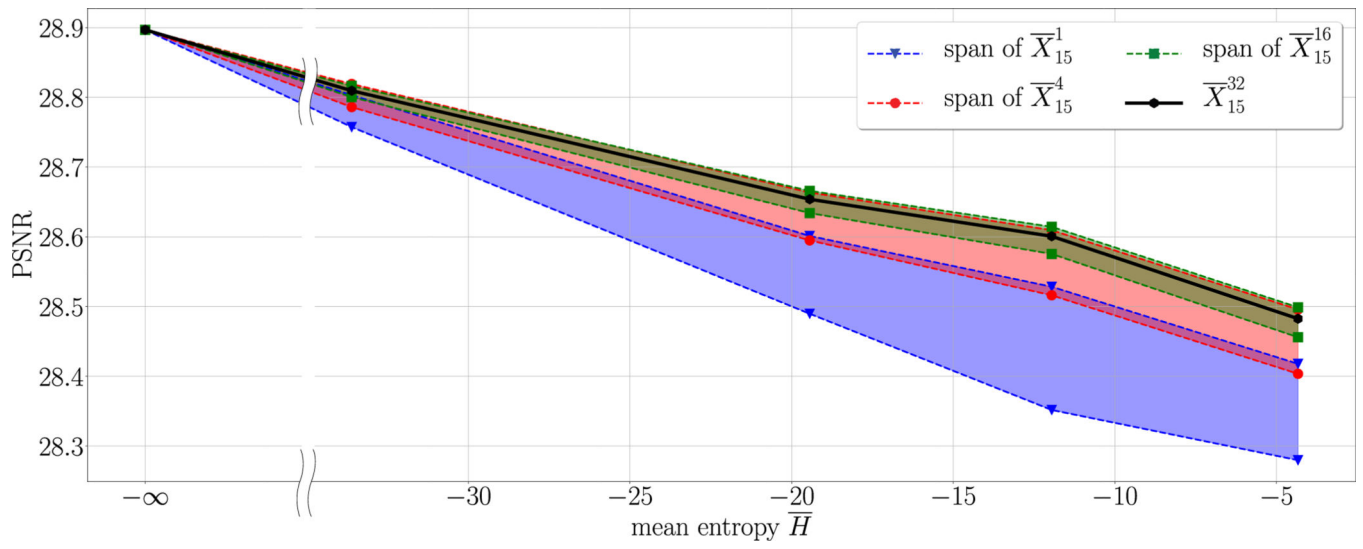
Author Manuscript

Author Manuscript

Author Manuscript



**Fig. 10.** Pairs of fully sampled initial images along with the corresponding eigenfunctions using PD (first pair) and PD-FS (second pair) as initialization.



**Fig. 11.** Dependency of the PSNR value on the entropy and the averaging.

QUANTITATIVE RESULTS FOR VARIOUS SINGLE AND MULTI-COIL MRI RECONSTRUCTION METHODS FOR  $R \in \{4,8\}$ .

TABLE I

Acquisition	Method	$R = 4$				$R = 8$				Parameters ( $\times 10^6$ )
		PSNR $\uparrow$	NMSE $\downarrow$	SSIM $\uparrow$	PSNR $\uparrow$	NMSE $\downarrow$	SSIM $\uparrow$	SSIM $\uparrow$		
single-coil	zero filling	30.5	0.0438	0.687	26.6	0.0839	0.543	—	—	
	U-Net [23]	32.2	0.032	0.754	29.5	0.048	0.651	214.16	214.16	
	$\Sigma$ -Net [41]	33.5	0.0279	0.777	n.a.	n.a.	n.a.	140.92	140.92	
	iRim [38]	33.7	0.0271	0.781	30.6	0.0419	0.687	275.25	275.25	
	TDV (deterministic)	33.8	0.0257	0.768	30.5	0.0407	0.665	2.21	2.21	
	TDV (stochastic)	33.5	0.0269	0.762	30.4	0.0412	0.662	9.95	9.95	
multi-coil	zero filling	32.0	0.0255	0.848	28.4	0.0549	0.778	—	—	
	U-Net [23]	35.9	0.0106	0.904	33.6	0.0171	0.858	214.16	214.16	
	$\Sigma$ -Net [41]	39.8	0.0051	0.928	36.7	0.0091	0.888	675.97	675.97	
	iRim [38]	39.6	0.0051	0.928	36.7	0.0091	0.888	329.67	329.67	
	E2EVN [42]	39.9	0.0049	0.930	36.9	0.0089	0.890	30.0	30.0	
	TDV (deterministic)	39.3	0.0054	0.923	35.9	0.0108	0.876	2.22	2.22	
TDV (stochastic)	38.9	0.0058	0.919	35.2	0.0123	0.867	9.97	9.97		