

Scan-Specific Self-Supervised Bayesian Deep Non-Linear Inversion for Undersampled MRI Reconstruction

Andrew P. Leynes^{ID}, Nikhil Deveshwar^{ID}, Srikantan S. Nagarajan^{ID}, *Fellow, IEEE*,
and Peder E. Z. Larson^{ID}, *Member, IEEE*

Abstract—Magnetic resonance imaging is subject to slow acquisition times due to the inherent limitations in data sampling. Recently, supervised deep learning has emerged as a promising technique for reconstructing sub-sampled MRI. However, supervised deep learning requires a large dataset of fully-sampled data. Although unsupervised or self-supervised deep learning methods have emerged to address the limitations of supervised deep learning approaches, they still require a database of images. In contrast, scan-specific deep learning methods learn and reconstruct using only the sub-sampled data from a single scan. Here, we introduce Scan-Specific Self-Supervised Bayesian Deep Non-Linear Inversion (DNLINV) that does not require an auto calibration scan region. DNLINV utilizes a Deep Image Prior-type generative modeling approach and relies on approximate Bayesian inference to regularize the deep convolutional neural network. We demonstrate our approach on several anatomies, contrasts, and sampling patterns and show improved performance over existing approaches in scan-specific calibrationless parallel imaging and compressed sensing.

Index Terms—Bayesian deep learning, self-supervised deep learning, deep image prior, undersampled MRI, calibrationless MRI.

I. INTRODUCTION

MAGNETIC resonance imaging is subject to slow acquisition times due to the inherent limitations in data sampling. The acquisition speed may be increased by either having gradients with higher slew rates, or by sub-sampling data below the Nyquist limit as in Parallel Imaging [1], [2] and

Manuscript received 9 December 2023; revised 23 January 2024; accepted 2 February 2024. Date of publication 9 February 2024; date of current version 3 June 2024. This work was supported in part by NIH under Grant R01AR078762 and Grant R01CA212148, in part by the University of California at San Francisco (UCSF) Graduate Research Mentorship Fellowship Award, and in part by the UCSF Cancer Center. (Corresponding author: Peder E. Z. Larson.)

This work involved human subjects or animals in its research. The authors confirm that all human/animal subject research procedures and protocols are exempt from review board approval.

The authors are with the Department of Radiology and Biomedical Imaging, University of California at San Francisco, San Francisco, CA 94158 USA, and also with the UC Berkeley—UC San Francisco Joint Graduate Program in Bioengineering, Berkeley, CA 94720 USA (e-mail: andudes@gmail.com; nikhil.deveshwar@ucsf.edu; sri@ucsf.edu; Peder.Larson@ucsf.edu).

Digital Object Identifier 10.1109/TMI.2024.3364911

Compressed Sensing [3]. Sub-sampling of data has been the focus of most research efforts as this allows for acceleration of scans without costly hardware changes.

Parallel imaging relies on the redundancy offered by multiple coils when sub-sampling k-space. By having multiple coils with their own local sensitivity profile, each coil image would result in different aliasing patterns. SENSE and SENSE-like approaches model the problem as a solution to a linear system of equations that unravels the aliasing artifacts to solve for the unaliased image with known coil sensitivity profiles [1], [4]. Alternatively, GRAPPA and GRAPPA-like approaches attempt to fill in the missing k-space regions using an auto calibration scan (ACS) [2]. At high acceleration rates, residual aliasing and noise amplification are encountered in these methods [5].

Compressed sensing MRI, meanwhile, relies on enforcing a sparse representation in some domain, typically using pseudo-random sub-sampling patterns that create incoherent aliasing artifacts in the sparsifying domain, and non-linear algorithms to fill in missing k-space data [6]. Compressed sensing can further accelerate MRI scans, although artifacts may be encountered that obscure anatomy and are difficult to identify [7].

Supervised deep learning has emerged as a promising technique for reconstructing sub-sampled MRI [8]. Supervised deep learning approaches learn to reconstruct images from sub-sampled data using a database of images consisting of matched pairs of sub-sampled data and fully-sampled data. These methods have shown possible acceleration factors that exceed what is possible in parallel imaging and compressed sensing [8]. However, supervised deep learning requires a large dataset of fully-sampled data.

There are scenarios where acquiring a large high-quality database of fully-sampled data may be impractical or even impossible. For example, simultaneously having ~ 1 mm spatial resolution at ~ 100 ms temporal resolution in 3D dynamic imaging is not possible with current hardware configurations and has required sophisticated reconstruction algorithms [9], [10]. Large high-quality databases are also challenging for emerging and SNR limited techniques such as non-proton imaging (e.g. ^2H , ^{23}Na , and ^{31}P), chemical exchange saturation transfer (CEST) MRI, and ultra-low field MRI. Finally, obtaining large datasets including rare disease

conditions are impractical due to limited availability of MRI scan data [11].

Recently, unsupervised and self-supervised deep learning have emerged to address the limitations of supervised deep learning approaches [12], [13], [14], [15], [16]. Rather than requiring fully-sampled images for supervision, these methods can utilize only sub-sampled data. Unsupervised approaches are designed to learn patterns, for example MRI priors, from the sub-sampled data, whereas self-supervised approaches explicitly create supervision signals from the sub-sampled data, which then can be used in the loss function and/or as validation data. Some recent approaches have utilized generative adversarial networks (GANs) [16], while others have directly used model-based approaches [12], [13], [14], [15]. These methods have been demonstrated to be superior to “classical” parallel imaging and compressed sensing and have competitive performance compared to fully supervised methods.

Within unsupervised and self-supervised deep learning approaches, scan-specific deep learning methods learn and reconstruct using only the sub-sampled data from a single scan ($N=1$). The major benefit of a scan-specific approach is that it removes the need for a training database and they operate at the lower bound of dataset size. (In practice, having additional data for testing and choose regularization parameters is often required [17], [18], similar to classical methods.) Scan-specific approaches promise to be more generalizable across hardware configurations, contrast mechanisms, pulse sequences, or data sampling schemes, much like classical reconstruction methods.

Current deep learning scan-specific methods including training a neural network to learn how to interpolate the sub-sampled regions of k-space from the fully-sampled auto-calibration scan region (like GRAPPA) [19], [20], [21], [22], [23]. A variation of Deep Decoder has also been used, either directly estimating individual coil images or sensitivity maps and a coil combined image [17].

Calibrationless methods allow for the greatest flexibility in k-space sampling by eliminating the need for an auto-calibration region, reducing sampling restrictions as well as providing more options for pulse sequence designs. This relies on the principles of jointly estimating image content and coil sensitivity profiles [24], [25], [26], and often low-rank matrix completion methods have been used to solve this reconstruction problem [27], [28], [29], [30]. Supervised deep learning [31], [32] and the deep image prior method [33] has been combined with joint estimation, although these primarily test on data with an auto-calibration region. Scan-specific methods using neural networks to interpolate sub-sampled k-space have also been extended to calibrationless sampling [22].

Of relevance to this work, Bayesian methods have shown promise integrated in deep learning-based MRI reconstruction methods, including within the loss functions to improve performance [34] as well as using diffusion models to provide estimates of reconstruction uncertainty [35].

We introduce Scan-Specific Self-Supervised Bayesian Deep Non-Linear Inversion (DNLIIVN) for under-sampled MRI reconstruction. Our contribution is two-fold. First, we utilize a network architecture that allows for the joint estimation

of image content and coil sensitivity profiles. Second, we use Bayesian deep learning to statistically regularize the network and reduce overfitting to noise. We demonstrate our approach on several anatomies, sampling patterns, and show superiority over existing approaches in scan-specific calibrationless parallel imaging and compressed sensing.

II. THEORY

A. Generative Modeling

We define x as the underlying true image. Then, using a SENSE signal model [1], the acquired coil-by-coil k-space data can be written as a linear equation:

$$y_i = SFC_i x + n | 1 \leq i \leq N \quad (1)$$

where N is the number of coils, y_i is the k-space data for the i -th coil, C_i is the coil sensitivity matrix for the i -th coil, F is the Fourier transform operator, S is the k-space sampling operator, and x is the underlying MR image. The coil sensitivities are acquired either through a calibration scan [1] or by estimating the sensitivities from an autocalibrating scan region [4]. The optimization problem then reduces to an inverse problem of obtaining x from the k-space measurements y_i .

Using sparsity-based regularization for reconstruction as in compressed sensing leads to the following optimization problem:

$$\hat{x} = \underset{x}{\operatorname{argmin}} \frac{1}{2} \|y - Ax\|_2^2 + \lambda_1 \|\Psi x\|_1 \quad (2)$$

$$y = [y_1 y_2 \dots y_N]^T \quad (3)$$

$$A = [SFC_1 \ SFC_2 \ \dots \ SFC_N]^T \quad (4)$$

where y is the complete measured k-space data from all coils, A is the forward operator, x is the image to be estimated, Ψ is a sparsifying transformation, and λ_1 is a regularization parameter. This enforces that the image x would be sparse in the transformed domain defined by Ψ . Common choices for Ψ are the Wavelet transform or the finite difference operator [3].

We can recast the compressed sensing model in a synthesis form that is equivalent to a generative model [36]:

$$\hat{z} = \underset{z}{\operatorname{argmin}} \frac{1}{2} \|y - A\Psi^{-1}z\|_2^2 + \lambda_1 \|z\|_1 \quad (5)$$

where instead of solving directly for x , we solve for sparse coefficients z and generate x from the known transformation Ψ^{-1} :

$$\hat{x} = \Psi^{-1}\hat{z} \quad (6)$$

In general, we can write this synthesis/generative modeling problem as follows:

$$\hat{z} = \underset{z}{\operatorname{argmin}} \frac{1}{2} \|y - Ag(z)\|_2^2 + R(z) \quad (7)$$

And produce the final image estimate as follows:

$$\hat{x} = g(\hat{z}) \quad (8)$$

where g is a function that generates an image x from a set of coefficients z , and R is a chosen regularization function that enforces the desired properties of z . This is the

form of compressed sensing with generative modeling with a previously determined generative model [37], which has been shown to provide MRI reconstructions that are robust to out-of-distribution shifts by using posterior sampling [38].

We may also want to estimate the generative function g in addition to estimating z , with a different chosen regularization function R also enforces desired properties of g , which has previously been applied to MRI reconstruction [39]:

$$\hat{z}, \hat{g} = \operatorname{argmin}_{z, g} \frac{1}{2} \|y - Ag(z)\|_2^2 + R(z, g) \quad (9)$$

Deep Image Prior (DIP) is a generative modeling technique that generates an image from a random noise vector using a learned convolutional neural network [40]. DIP depends on implicit regularization provided by the use of a CNN as the generative model, $g(z) = g_{CNN}(z; \theta)$, where g_{CNN} is a deep convolutional neural network, and can be described as a more specific case of Eq. 9:

$$\hat{\theta} = \operatorname{argmin}_{\theta} \frac{1}{2} \|y - Ag_{CNN}(z; \theta)\|_2^2 + R(\theta) \quad (10)$$

where z is a set of fixed random coefficients, θ are the parameters of the neural network, and R is a chosen regularization function that enforces the desired properties of θ . The parameters θ can be found using stochastic gradient descent using backpropagation [41] with an optimizer such as Adam [42]. The final image can be obtained similarly:

$$\hat{x} = g_{CNN}(z, \hat{\theta}) \quad (11)$$

B. Joint Estimation of the Image and Coil Sensitivities

In calibrationless parallel imaging and compressed sensing MRI, the lack of an autocalibration scan region means that the image and coil sensitivities in eq. (1) must be solved jointly. Thus, the problem becomes non-linear. The NLINV approach [24] formulates the problem as a low-rank matrix factorization problem:

$$\hat{x}, \hat{C} = \operatorname{argmin}_{x, C} \frac{1}{2} \|y - SFCx\|_2^2 + \lambda_2(\|x\|_2^2 + \|WC\|_2^2) \quad (12)$$

$$C = [C_1 \ C_2 \ \dots \ C_N]^T \quad (13)$$

where λ_2 is a regularization parameter and W is a weighting matrix penalizing high frequencies in the coil profiles. NLINV was improved by using multiple sets of coil sensitivity maps to account for model violations in ENLIVE [26], which was also shown to inherently promote low rank solutions.

For DNLINV, we reformulate the deep image prior approach by considering a new multi-task deep CNN architecture that simultaneously produces the image and coil sensitivities, shown in Figure 1. It consists of a U-net producing a coil-combined image output, denoted $g_{x_{CNN}}(z; \theta)$, with an upsampling pathway connected to the coarsest resolution feature space to produce coil sensitivities, denoted $g_{C_{CNN}}(z; \theta)$, with intended implicit low spatial frequency regularization. Note that these two outputs share the

same input coefficients, z , and many of the model weights, θ . The optimization problem for DNLINV is then:

$$\hat{\theta} = \operatorname{argmin}_{\theta} \frac{1}{2} \|y - SFg_{x_{CNN}}(z; \theta)g_{C_{CNN}}(z; \theta)\|_2^2 + R(\theta) \quad (14)$$

$$x = g_{x_{CNN}}(z; \theta), C = g_{C_{CNN}}(z; \theta) \quad (15)$$

This can also be solved using learning and inference procedures typical for deep neural networks.

C. Bayesian Inference

The DIP formulation (eq. 10) is known to overfit to noise [40], [43]. In fact, stochastic gradient descent is simply gradient descent and deep neural networks achieve zero loss under a scan-specific self-supervised regime (training $N = 1$) where cross-validation cannot be performed [44]. Other techniques such as splitting the measurements [13] could be used, but we chose to rely on Bayesian inference to combat the overfitting problem.

In contrast to regularized least-squares optimization (which is equivalent to *maximum a posteriori* (MAP) estimation with the appropriate log-prior corresponding to the regularization function R) performing Bayesian inference provides uncertainty estimates that has been demonstrated to reduce overfitting of the Deep Image Prior (DIP) model [43]. In particular, we utilize variational inference [45] and Monte-Carlo Dropout (MC-dropout) [46].

Reformulating the multicoil DIP problem (eq. 14) under Bayesian probabilistic modeling, we obtain the following:

$$y = SFg_{x_{CNN}}(z; \theta)g_{C_{CNN}}(z; \theta) + n \quad (16)$$

$$n \sim \text{Normal}(0, \Sigma_c) \quad (17)$$

$$p(\theta | y, z) = \frac{p(y, z | \theta) p(\theta)}{p(y, z)} \quad (18)$$

However, unlike deep image prior that only solves for the network parameters θ , we also perform inference for the input coefficients to the network z to solve for all unknowns in DNLINV:

$$p(\theta, z | y) = \frac{p(y | \theta, z) p(\theta, z)}{p(y)} \quad (19)$$

To use Bayesian inference we define the following priors:

$$p(z_j) = \text{Normal}(0, 1) \quad (20)$$

$$p(\theta_s) = \text{Normal}(0, \lambda) \quad (21)$$

$$p(z) = \prod_{j=1}^M p(z_j) \quad (22)$$

$$p(\theta) = \prod_{s=1}^S p(\theta_s) \quad (23)$$

where λ is the precision for the parameter θ_s from MC-dropout [46] and S is the total number of network parameters.

Since directly solving for the posterior is intractable, we utilize a mean-field approximation to the posterior as follows:

$$p(\theta, z | y) \approx q(\theta | y) q(z | y) \quad (24)$$

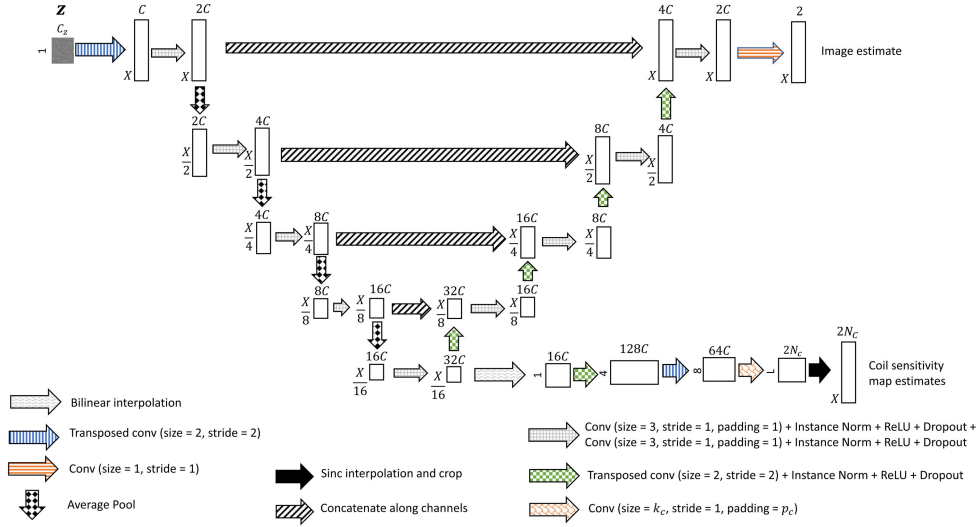


Fig. 1. Deep convolutional neural network architecture used for generative modeling. The spatial size is denoted at the bottom left of each block (e.g., X) and the number of channels is denoted at the top of each block (e.g., C). This network creates both image and coil sensitivity outputs which are described in Eqns. 14 and 15. The final image estimate and each coil image have 2 channels for real and imaginary dimensions, respectively.

Then, each of the approximate posteriors can be solved for independently.

To solve for the approximate posterior of the network parameters $q(\theta|y)$, we use Monte-Carlo Dropout to convert the deep convolutional neural network into a Bayesian deep convolutional network. This is an implicit method where multiple models are trained but with different dropout masks to allow for estimation of $q(\theta|y)$.

To solve for $q(z|y)$, we derive the evidence lower bound (ELBO) [45], [47] and utilize the *reparameterization trick* [45], [47]. The ELBO cost function has the following general form:

$$\mathcal{L} = E_{q(z|y)q(\theta|y)} \left[\log \frac{p(y|z, \theta) p(z, \theta)}{q(z|y) q(\theta|y)} \right] \quad (25)$$

and we restrict $q(z)$ to be:

$$q(z_j) = \text{Normal}(\mu_{z_j}, \sigma_{z_j}), q(z) = \prod_{j=1}^J q(z_j) \quad (26)$$

This leads to the following cost function:

$$\begin{aligned} \mathcal{L} = & -\frac{N}{2} (\log 2\pi + \log \det(\Sigma_c)) \\ & - \frac{1}{2M} \sum_{m=1}^M \sum_{k=1}^K \left(y_k - S_k F g_C(z^{(m)}; \theta^{(m)}) \right. \\ & \times g_x(z^{(m)}; \theta^{(m)}) \left. \right)^T \\ & \Sigma_c^{-1} \left(y_k - S_k F g_C(z^{(m)}; \theta^{(m)}) g_x(z^{(m)}; \theta^{(m)}) \right) \\ & - 0.5 \sum_{j=1}^J \sum_{m=1}^M (z_j^{(m)})^2 + 0.5 \sum_{j=1}^J \log(\sigma_{z_j}^2) + 0.5 \quad (27) \end{aligned}$$

where y is the observed k-space data, $z^{(m)}$ is m -th sample of the input latent code to the network shown in Fig. 1, $\theta^{(m)}$ is the m -th sample of the network parameters, M is the number of Monte Carlo samples, Σ_c is the coil noise covariance matrix,

g_C is the coil sensitivity map outputs of the network, g_x is the image output of the network, σ_{z_j} is the standard deviation of the approximate posterior $q(z_j)$, J is the number of latent variables, K is the number of k-space sampling points, and N is the total number of k-space samples across all coils.

Finally, the optimization problem is:

$$\underset{\mu_{z_j}, \sigma_{z_j}, \theta, \Sigma_c}{\text{argmin}} \mathcal{L} \quad (28)$$

D. Implementation Details

The optimization problem in eq. 28 is solved using backpropagation stochastic gradient descent (SGD) using the Adam optimizer [42] with decoupled weight decay [48].

The nature of MC-Dropout is implicit: there is no explicit realization of $q(\theta|y)$ nor explicit cost function for it. Furthermore, we utilize decoupled weight decay [48] that places the gradient descent update corresponding to the Gaussian prior on the network parameters to the optimization update and thus does not appear in the cost function. Therefore, in the implementation of the optimization, we proceed using only \mathcal{L} as the cost function but each step of SGD is performed with a different dropout mask (as in MC-Dropout). The network parameters are updated at each SGD step including the decoupled weight decay updates.

The latent code distribution parameters, μ_{z_j}, σ_{z_j} , are updated via backpropagation using ELBO and the reparameterization described previously.

The coil noise covariance matrix Σ_c is decomposed to a lower triangular Cholesky factorization:

$$\Sigma_c = LL^T \quad (29)$$

and L is estimated instead to guarantee that Σ_c is positive semi-definite.

The algorithm requires determining the hyperparameters of learning rate, dropout rate, weight decay, and number of iterations. Thus far, these were empirically chosen based on

convergence behavior, PSNR/SSIM metrics, and image quality using a subset of test cases during algorithm development. We allowed for different hyperparameters to be used for each dataset described below (ENLIVE brain, ENLIVE knee, and fastMRI knee), but these were fixed for all samples within each datasets.

The computational complexity is like that of deep learning training. The optimization is based stochastic gradient descent (SGD) that has less than linear convergence. Using reparameterization gradient variational inference [45] reduces the variance of the gradient estimates. Having a larger number of Monte-Carlo samples per SGD step reduces the gradient estimate variance and accelerates convergence at the cost of requiring more memory per SGD step.

The final image and coil sensitivity estimates are then constructed through Monte-Carlo sampling of the expectation:

$$\hat{x} = E_{q(z)q(\theta)} [g_x(z; \theta)] \approx \frac{1}{M} \sum_{m=1}^M g_x(z^{(m)}; \theta^{(m)}) \quad (30)$$

$$\hat{C} = E_{q(z)q(\theta)} [g_C(z; \theta)] \approx \frac{1}{M} \sum_{m=1}^M g_C(z^{(m)}; \theta^{(m)}) \quad (31)$$

III. NUMERICAL EXPERIMENTS

In the spirit of reproducible research, we provide our code and scripts to reproduce all experiments in this paper at <https://www.github.com/LarsonLab/dnlinv>. All reconstruction parameters of the experiments can be found in the code and the results shown in this paper are immediately provided in the repository.

A. Evaluation Metrics

The evaluation metrics used were peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM). PSNR was calculated as:

$$PSNR = -20 \log I_{max} + 10 \log MSE \quad (32)$$

$$MSE = \sum_{r=1}^R (x_r - x_{ref,r})^2 \quad (33)$$

where x_r is the r -th voxel of the reconstructed sum-of-squares image, $x_{ref,r}$ is the r -th voxel of the reference fully-sampled sum-of-squares image, and R is the total number of voxels in an image, while SSIM was calculated in the same manner as in the fastMRI challenge [49].

B. 1d Denoising

As an initial test of the approach, including convergence and noise properties, we compare DIP and DNLINV, including an ablation study of the effect of the inference method (Table I) in a simple 1d denoising experiment. The data was simulated with a noise standard deviation of 0.01 and 0.1 with 8 random samples to emulate the multiple sampling provided by multiple coils. The noiseless and noisy signals are shown in Fig. 2.

Deep image prior was recommended to be used in conjunction with early stopping due to its overfitting

TABLE I

LIST OF INFERENCE METHODS FOR 1D DENOISING ABLATION EXPERIMENT. THIS STARTS WITH DEEP IMAGE PRIOR (*dip*), ADDS INFERENCE ON THE LATENT VECTOR ([45], Eq. 19, *dip_mc_inference*), ADDS MONTE CARLO DROPOUT ([46], *dip_mc_dropout*), COMBINES THE TWO (*dip_mc_z_mc_dropout*), AND INCORPORATES THE ELBO COST FUNCTION TO ESTIMATE POSTERIOR OF THE LATENT VECTOR (Eq. 27) WITH AND WITHOUT DROPOUT (*dnlinv_no_dropout*, *dnlinv_dropout*)

Inference method	Loss function	Inference on z	Inference on θ
Deep image prior (<i>dip</i>)	MSE	Fixed z	MAP θ
DIP MC- z (<i>dip_mc_inference</i>)	MSE	Draw from $p(z)$	MAP θ
DIP MC-Dropout (<i>dip_mc_dropout</i>)	MSE	Fixed z	Implicit $q(\theta)$
DIP MC- z MC-Dropout (<i>dip_mc_z_mc_dropout</i>)	MSE	Draw from $p(z)$	Implicit $q(\theta)$
DNLINV w/o MC-Dropout (<i>dnlinv_no_dropout</i>)	ELBO	Estimate $q(z)$	MAP θ
DNLINV (<i>dnlinv_dropout</i>)	ELBO	Estimate $q(z)$	Implicit $q(\theta)$

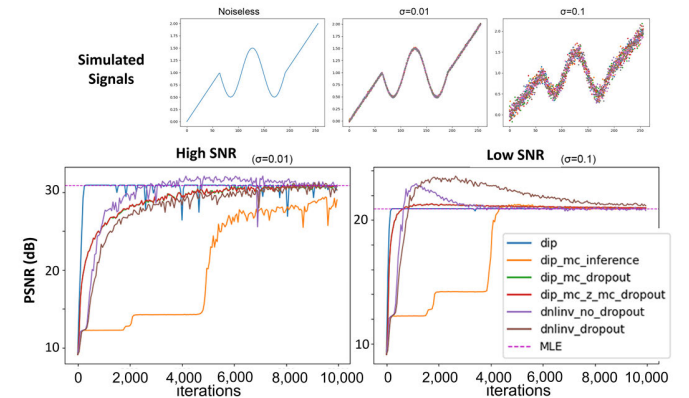


Fig. 2. 1d denoising ablation experiment results demonstrating the effect of inference method (Table I) on performance and convergence in both (bottom left) high SNR and (bottom right) low SNR regimes. The noiseless and noisy signals are shown in the top row. Peak signal-to-noise ratio (PSNR) is plotted over the number of optimization iterations. The optimal maximum likelihood estimator (MLE) PSNR is shown for comparison (dashed red line).

behavior [40], [43]. Thus, we also keep track of performance vs. number of optimization steps. In addition, we include the denoising performance the maximum likelihood estimator for Gaussian noise (mean of samples).

C. SENSE Numerical Phantom Image Reconstruction

A SENSE numerical phantom experiment was performed where coil sensitivity profiles are known. An evenly spaced sampling scheme was used to create coherent aliasing artifacts and higher uncertainty in these regions. A birdcage coil configuration with 8 channels was simulated using the SigPy python package (v0.1.16) [50]. A Shepp-Logan phantom

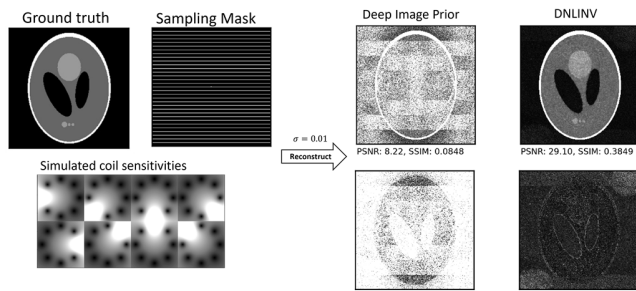


Fig. 3. SENSE phantom reconstruction experiment setup with 8x undersampling. 7500 iterations were used for the image and difference images (3x scaled) result shown.

was simulated with a maximum signal intensity of 1.0. Additive white Gaussian noise ($\sigma = 0.01$) was added to the image. Deep image prior and DNLINV reconstructions were compared.

D. In Vivo Retrospective Undersampling

Data from ENLIVE [26] and fastMRI [49] were used for retrospective undersampling experiments. The ENLIVE method was used for comparison with 2 sets of coil sensitivity maps. The number of iterations was selected individually for all methods and datasets, chosen to approximately maximize PSNR and SSIM.

1) *ENLIVE Data*: For ENLIVE data, we use both head and knee data under high acceleration factor parallel imaging and under calibrationless parallel imaging and compressed sensing sampling schemes.

Head data (3D FLASH, TR/TE = 11/4.9 ms, 3T, 32-channels) was retrospectively undersampled with a calibrationless variable-density poisson disk sampling pattern [3] at 4x, 7x, and 8.5x acceleration factors and were reconstructed using simultaneous autocalibrating and k-space estimation (SAKE) [27], ENLIVE, Deep Image Prior, and DNLINV. Additional CAIPIRINHA sampling [51] results including L2-ESPIRiT [4] are available in [52] and in our code repository.

Knee data (3D FSE, TR/TE = 1550/25 ms, 3T, 8-channels) was retrospectively undersampled with a calibrationless variable-density poisson disk sampling pattern [3] at 4x, 7x, and 8.5x acceleration factors and were reconstructed using SAKE, ENLIVE, Deep Image Prior, and DNLINV. Additional knee data used in ENLIVE were from the Stanford 3D fast-spin-echo dataset provided in mridata.org and were used to evaluate performance over a dataset.

2) *fastMRI Data*: We performed similar experiments using fastMRI knee data (2D TSE, TR/TE = 2200-3000/27-34 ms, 15-channel) [49]. Images were retrospectively undersampled with calibrationless and autocalibrated parallel imaging and compressed sensing sampling patterns (acceleration factors = [2, 3, 4]) were reconstructed with L1-ESPIRiT, SAKE, ENLIVE, DNLINV, and the pre-trained U-net fastMRI. A uniform distribution for selecting k-space lines was used for both sampling patterns.

E. In Vivo Ablation Study

Finally, we performed an ablation study using the knee ENLIVE data at 5x acceleration. Noise estimation is

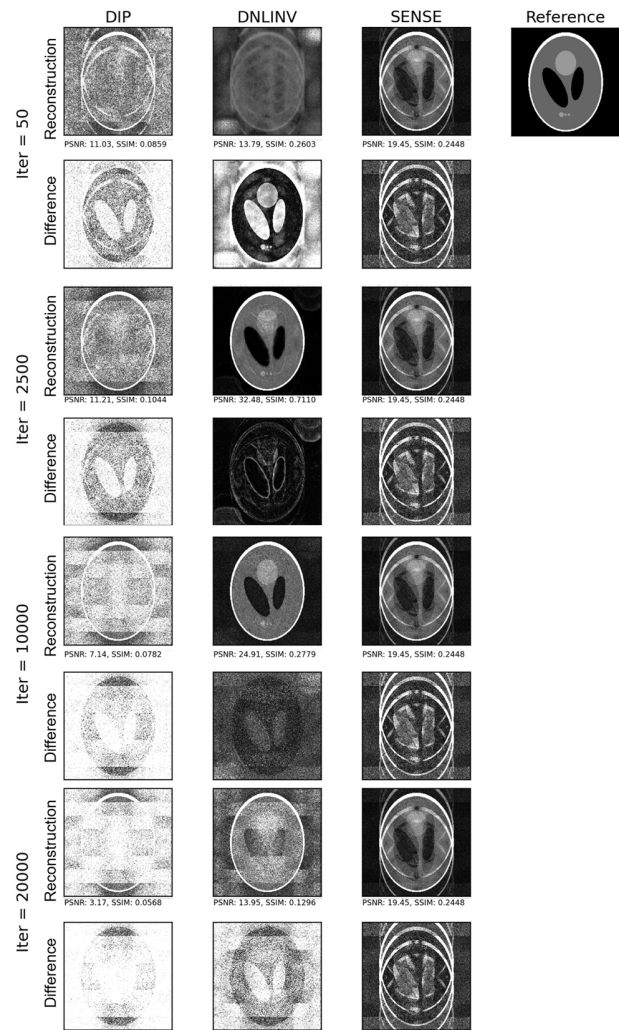


Fig. 4. Phantom image reconstructions at different iterations for Deep Image Prior (1st column), Deep Non-linear Inversion (DNLINV) (2nd column), and SENSE ($\lambda_{L2} = 0.005$) (3rd column). Under each image is the absolute difference image compared to the noiseless Shepp-Logan phantom with the intensities amplified by 3x to improve visibility. Deep Image Prior shows noise overfitting due to the uncertainty in the reconstruction that is characterized by the parallel imaging g-factor. DNLINV was able to avoid this noise overfitting in earlier iterations, although also eventually overfits.

performed in cases where acquiring noise samples is not done and we investigate whether it is needed for accurate reconstructions. Then, we tested how the use of non-linear activations (e.g., ReLU) in Deep image prior and DNLINV impacts the results. We gradually removed portions of DNLINV and assessed the performance impact. DNLINV was compared against DNLINV without noise estimation (keeping Σ constant), deep image prior, DIP MC-z, and DNLINV with linear activations (without ReLU).

IV. RESULTS

A. 1d Denoising Experiments

The 1d denoising experiments (Fig. 2) describe how the various components of the DNLINV method contribute to the overall results. They show that, in the high SNR

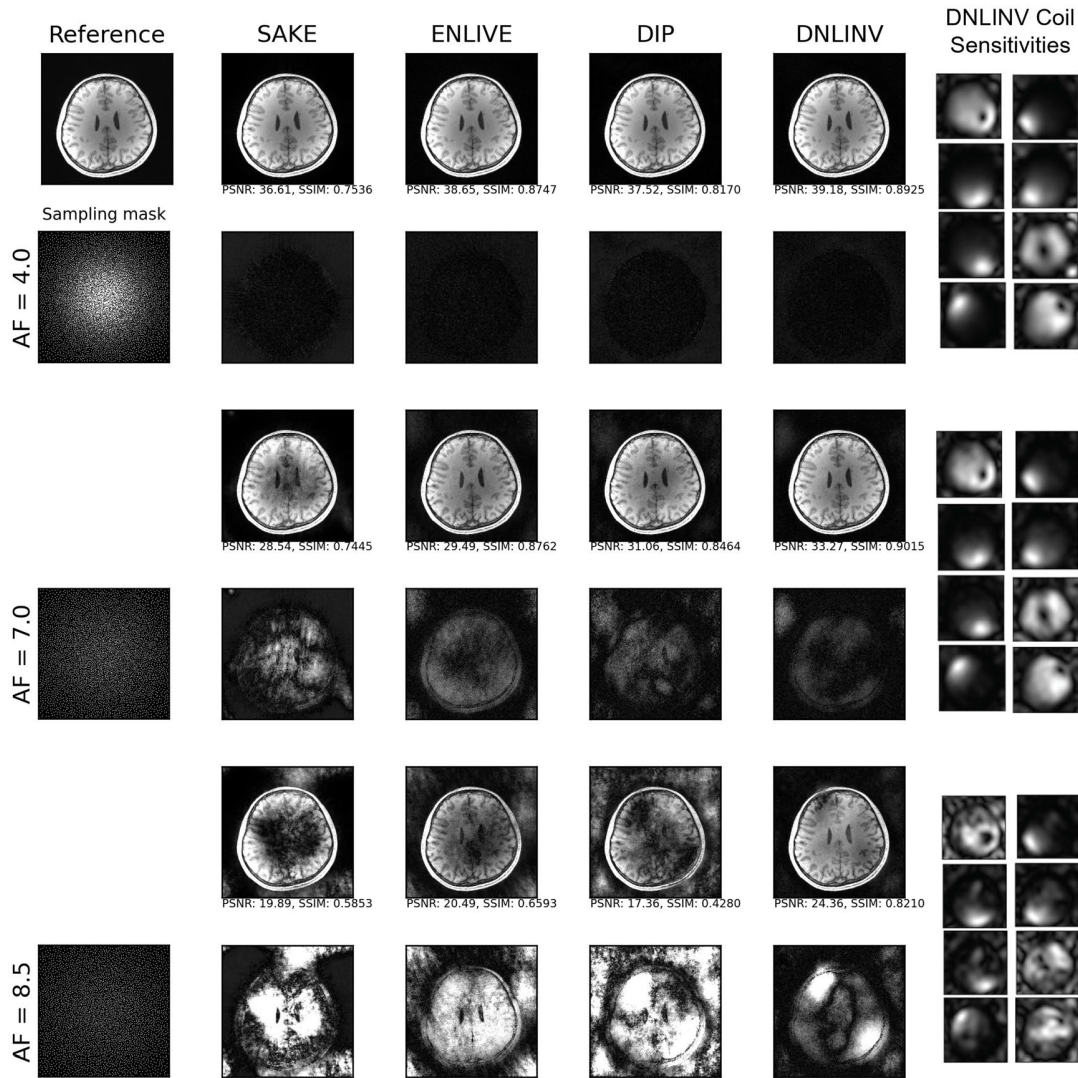


Fig. 5. Head calibrationless poisson disk retrospective undersampling results comparing SAKE, ENLIVE, DIP, and DNLINV. DNLINV estimated coil sensitivities are shown for 8 out of 32 channels with the highest SNR. The absolute difference images are shown below each reconstruction and the error intensities were amplified by 3x to improve visibility. While all methods perform qualitatively similar for the lowest acceleration factor (AF) = 4.0, DNLINV showed fewer artifacts and higher PSNR/SSIM at moderate to high AF.

regime, most all methods were able to achieve a similar PSNR, while differences in performance are found in the low SNR regime. Deep image prior by itself (“dip”) has the same performance as the maximum likelihood estimator for signals with Gaussian noise, but with extremely rapid convergence that is the result of the known DIP overfitting. The addition of drawing the latent variables from a prior distribution (“dip_mc_inference”) deteriorated performance, while instead adding MC-Dropout to DIP (“dip_mc_dropout” and “dip_mc_z_mc_dropout”) improved PSNR in the low SNR regime even when drawing latent variables from a prior distribution. Finally, creating DNLINV by incorporating the ELBO cost function to estimate posterior of the latent vector (Eq. 27) showed the best performance in the low SNR regime, and using dropout (“dnliv_dropout”) achieved the highest possible peak signal-to-noise ratio (PSNR). DNLINV also showed a broad range of iterations with high PSNR, indicating that early stopping can likely be used to achieve this performance.

B. SENSE Image Reconstruction Experiments

The main result for the SENSE image reconstruction experiment is shown on Fig. 3. Deep image prior was severely affected by the noise amplification (e.g. g-factor), which is expected where the parallel imaging problem is under-determined in this high undersampling (8x) regime. DNLINV, on the other hand, largely did not suffer from this noise amplification as was intended by using Bayesian inference.

Reconstructions at different number of optimization steps are shown in Fig. 4. Deep image prior immediately overfits to the noise at just 50 iterations whereas DNLINV slowly generated image content, similarly to the 1d denoising experiment. DNLINV took around 2500 iterations to reach the peak PSNR and SSIM (although residual aliasing artifacts remained) and took about 7500 iterations (in Fig. 3) to largely eliminate the aliasing artifacts. Also like in the 1d denoising experiments, DNLINV begins to overfit to noise at



Fig. 6. Knee 1d calibrationless retrospective undersampling comparing SAKE, ENLIVE, DNLINV, and the fastMRI pre-trained Unet reconstructions. The absolute difference image are shown below each reconstruction and the error intensities were amplified by 3x to improve visibility. Even at the lowest AF = 2.0, only DNLINV is the only method largely free of parallel imaging artifacts. At higher acceleration factors. While some aliasing artifacts appear at AF = 3.0 and 4.0 in DNLINV, the error is still much smaller compared to the other methods.

10,000 iterations and eventually overfits to the noise at 20,000 iterations.

C. In-Vivo Retrospective Undersampling Experiments

Figs. 5 and 6 show the main results for the calibrationless parallel imaging and compressed sensing sampling patterns. Aggregate PSNR and SSIM results over different datasets are provided in Figs. 7 and [52]. For the calibrationless 2D undersampling pattern (Fig. 5), DNLINV provided similar performance to prior reconstruction algorithms at low acceleration factors. At moderate acceleration factors, classical reconstructions (SAKE, ENLIVE) performed well but had increasing noise whereas DNLINV was able to reconstruct with similar quality to low acceleration factors. At high acceleration factors, classical reconstruction algorithms failed to properly reconstruct the images while DNLINV was still able to provide a good reconstruction. DNLINV also produced reasonable coil sensitivity profile maps without evidence of overfitting to data, although map quality degrades notably at AF=8.5. The aggregate results for 2D calibrationless sampling patterns shown in [52] demonstrate the improved performance in a dataset compared to ENLIVE and DIP.

Similar results are seen for the calibrationless 1D undersampling results in Fig. 6, although in this case DNLINV shows residual aliasing artifacts at moderate and high undersampling factors. For the aggregate dataset results in Fig. 7, DNLINV

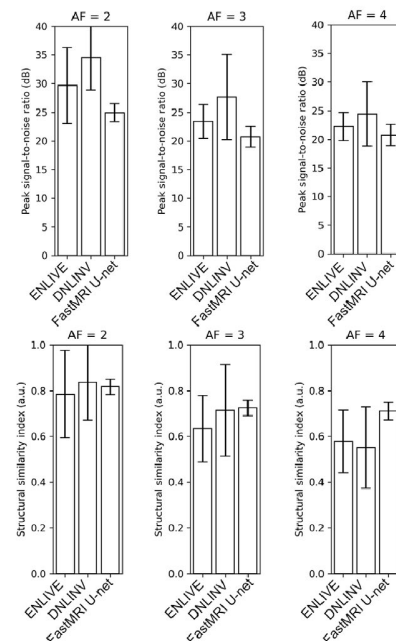


Fig. 7. Aggregate reconstruction results for a calibrationless sampling pattern from the validation set of the FastMRI 3T proton-density knee dataset (n=10). Error bars correspond to 1 standard deviation of the metrics.

had the highest PSNR but was 2nd only to the pre-trained fastMRI U-net in terms of SSIM.




Fig. 8. Knee 1d autocalibrated retrospective undersampling comparing ESPIRiT, ENLIVE, DNLINV, and the fastMRI pre-trained Unet reconstructions. The absolute difference images are shown below each reconstruction and the error intensities were amplified by 3x to improve visibility. With this sampling, ESPIRiT, ENLIVE, and DNLINV perform similarly, with noticeable aliasing artifacts at $AF = 3.0$ and 4.0 . FastMRI_Unet outperforms the other methods in this scenario, whereas DNLINV has more advantages with calibrationless sampling as shown in other figures.

Results from an autocalibrating undersampling pattern in Fig. 8 show that DNLINV does not have a notable advantage in this setting. Here, the U-Net performance is greatly improved when an ACS region is present, and performs better to DNLINV in this sampling environment that is closer to what the model was trained upon.

Figure 9 shows the ablation experiment, illustrating the value of each component of the DNLINV method. While DIP shows no obvious artifacts, it suffers from clear noise amplification. As in the 1D denoising example, adding MC inference slightly degrades performance and is no better than standard DIP. DNLINV suffers from artifacts or noise amplification without ReLU non-linear activations or noise estimation, while having these methods achieves the best performance. Similar performance patterns are shown in the estimated coil sensitivity maps in [52]. This data also shows results over various iterations of all methods and illustrate rapid overfitting to noise for DIP while DNLINV has a broader region of peak PSNR results similar to the 1d denoising experiment.

V. DISCUSSION

In this paper, we show DNLINV is more robust to noise than deep image prior and similar classical methods (e.g. SAKE, NLINV, ENLIVE) and thus can reconstruct data at high acceleration factors, especially for calibrationless parallel imaging and compressed sensing. Calibrationless sampling is especially difficult due to not only undersampling of the image content, but also due to the undersampling of the low frequency k-space locations that are crucial for accurate coil sensitivity estimation. The key innovations that enabled this scan-specific, self-supervised approach are: implicit regularization with DIP-type approach, addition of upsampling path to jointly estimate the image and coil sensitivity profiles that provides further implicit regularization, and the inclusion of probabilistic Bayesian inference in the reconstruction for noise robustness that plagued prior methods.

To the best of our knowledge, there is no other scan-specific nor self-supervised deep learning method that incorporates Bayesian deep learning in reconstruction, although Bayesian

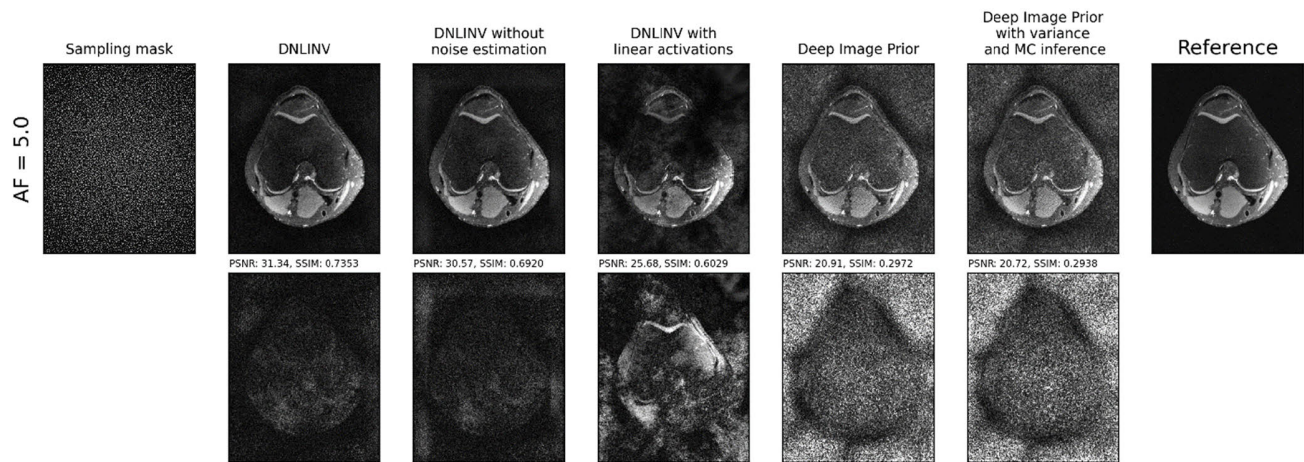


Fig. 9. Ablation experiment in a knee for calibrationless retrospective 2D undersampling. The absolute difference image compared to the reference are shown below each reconstruction and the error intensities were amplified by 3x to improve visibility. Coil sensitivity maps measured during each experiment are shown in [52].

deep learning has gained success in supervised learning [34], [35]. DNLINV builds off the success of DIP, and is the first of its kind to utilize Bayesian inference to minimize overfitting.

Other methods have utilized splitting the dataset into disjoint sets to separately learn the regularizer and the data fitting [13], use the autocalibration scan region for learning k-space interpolation [19], [20], [21], or rely on an architecture search to find an optimal network architecture [17].

The major component that makes DNLINV successful is the use of Bayesian evidence maximization [45], [53] using the ELBO cost function. The evidence maximization framework is what allows the network complexity to be better controlled. In DIP, only L_p regularization is available to control the network complexity and since the network is overparameterized, the network still overfits [40], [43]. We demonstrated in multiple experiments the impact that Bayesian evidence maximization can have on image reconstruction.

The next major component is the use of ReLU non-linear activations in the network. We demonstrated that using only linear layers lead to failure in image reconstruction (Fig. 9), showing the necessity of the non-linearities. The use of ReLU activations make the network like a multi-layer semi-negative convolutional sparse coding model [54].

The scan-specific and self-supervised nature of our approach allows for it to perform much like classical reconstruction algorithms and can be used for any hardware configuration, contrast, pulse sequence, or sampling pattern without any re-training. Supervised reconstruction algorithms have difficult handling test cases that are far from the training distribution. In the test example using the fastMRI data under the calibrationless context (Fig. 6), the pre-trained U-net produced significant artifacts. This is likely due to the U-net being trained on different sampling patterns that include an autocalibration region (Fig. 8). Dependence on sampling pattern and noise level is a limitation in current supervised learning approaches [55].

By eliminating the need of the autocalibration region, calibrationless parallel imaging and compressed sensing acquisitions allow for the greatest flexibility in k-space

sampling patterns. The optimal sampling strategy is an active area of research [56] and maybe coupled to the reconstruction approach, and DNLINV performs well across sampling patterns. Thus far, supervised deep learning methods have been demonstrated to reconstruct calibrationless parallel imaging and compressed sensing data [31], [32], [57], [58]. The advantage of a scan-specific approach such as DNLINV, DeepDecoder [17] or LORAKI [22] is that it can be used even in the most experimental of hardware configurations and pulse sequences.

Our results show that DNLINV is able to compete with supervised approaches but only up to a moderate acceleration factors, as shown in autocalibrated sampling results (Fig. 8, [52]). Supervised approaches have the benefit of learning from a database whereas DNLINV does not, and thus would be expected to have performance improvements in well-defined, more consistent circumstances.

A. Limitations

DNLINV still tends to overfit to noise over time, and requires early stopping, although this problem is not nearly as severe as it is for DIP. This is due to the variational approximations used, where mode-seeking and overfitting behavior is a well-known characteristic of variational Bayesian inference [45]. Using more sophisticated and computationally intensive inference methods such as through Langevin dynamics [43], expectation propagation [59], or ensemble methods [60] may be used to ultimately conquer the overfitting aspects, or by splitting up data to provide validation for early-stopping [61].

DNLINV also requires the choice of hyperparameters that are the learning rate, dropout rate, weight decay, and number of iterations. In this work, we chose the hyperparameters during algorithm development when reconstructing a small number of test cases. The learning rate, dropout rate, and weight decay were chosen by monitoring the convergence behavior. The number of iterations was selected to approximately maximize PSNR and SSIM, and is illustrated in detail in Supp. Figs. 2 and 9 [52]. Note that these results show DNLINV is relatively

insensitive to the choice of number of iterations. In this initial proof of concept, we allowed the hyperparameters to vary across each dataset (ENLIVE brain, ENLIVE knee, and fastMRI knee), as they have different SNR, resolution and contrast that could influence optimal hyperparameters. However, the hyperparameters were fixed for sample within a dataset. A more principled approach to choosing this method would simplify the use of DNINLV.

There are also large computational requirements for DNINLV. The Bayesian inference used has heavy compute requirements due to the Monte Carlo approximation of the expectations and integrals and the high noise of SGD. The reconstructions performed take 15 minutes to 2 hours on GPU depending on number of iterations required for convergence. Possible ways to accelerate this process is by pre-training on a database to initialize the network to a good configuration and thus require less iterations on reconstruction (i.e., transfer learning) [17]. Alternatively, meta-learning approaches may be used to learn an optimizer that also reduces the number of gradient steps needed (i.e., learning-to-learn [62]). Finally, the use of deep convolutional neural networks has a large memory requirement. Possible ways to save on memory is with the use of more compact neural network architectures (e.g., DCGAN/DeepDecoder [17], [63]) or by reducing the overall size of the network using an architecture search [17].

REFERENCES

- [1] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "SENSE: Sensitivity encoding for fast MRI," *Magn. Reson. Med.*, vol. 42, no. 5, pp. 952–962, Nov. 1999, doi: [10.1002/\(SICI\)1522-2594\(199911\)42:5<952::AID-MRM16>3.0.CO;2-S](https://doi.org/10.1002/(SICI)1522-2594(199911)42:5<952::AID-MRM16>3.0.CO;2-S).
- [2] M. A. Griswold et al., "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magn. Reson. Med.*, vol. 47, no. 6, pp. 1202–1210, 2002, doi: [10.1002/mrm.10171](https://doi.org/10.1002/mrm.10171).
- [3] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magn. Reson. Med.*, vol. 58, no. 6, pp. 1182–1195, 2007, doi: [10.1002/mrm.21391](https://doi.org/10.1002/mrm.21391).
- [4] M. Uecker et al., "ESPIRiT—An eigenvalue approach to autocalibrating parallel MRI: Where sense meets grappa," *Magn. Reson. Med.*, vol. 71, no. 3, pp. 990–1001, Mar. 2014, doi: [10.1002/mrm.24751](https://doi.org/10.1002/mrm.24751).
- [5] A. Deshmane, V. Gulani, M. A. Griswold, and N. Seiberlich, "Parallel MR imaging," *J. Magn. Reson. Imag.*, vol. 36, no. 1, pp. 55–72, 2012, doi: [10.1002/jmri.23639](https://doi.org/10.1002/jmri.23639).
- [6] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing MRI," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 72–82, Mar. 2008.
- [7] O. N. Jaspán, R. Fleysler, and M. L. Lipton, "Compressed sensing MRI: A review of the clinical literature," *Brit. J. Radiol.*, vol. 88, no. 1056, Dec. 2015, Art. no. 20150487, doi: [10.1259/bjr.20150487](https://doi.org/10.1259/bjr.20150487).
- [8] F. Knoll et al., "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](https://doi.org/10.1109/MSP.2019.2950640).
- [9] F. Ong et al., "Extreme MRI: Large-scale volumetric dynamic imaging from continuous non-gated acquisitions," *Magn. Reson. Med.*, vol. 84, no. 4, pp. 1763–1780, Oct. 2020, doi: [10.1002/mrm.28235](https://doi.org/10.1002/mrm.28235).
- [10] Z. Miller, A. Pirasteh, and K. M. Johnson, "Memory efficient model based deep learning reconstructions for high spatial resolution 3D Non-Cartesian acquisitions," 2022, *arXiv:2204.13862*, doi: [10.48550/arXiv.2204.13862](https://doi.org/10.48550/arXiv.2204.13862).
- [11] A. M. Rauschecker et al., "Artificial intelligence system approaching neuroradiologist-level differential diagnosis accuracy at brain MRI," *Radiology*, vol. 295, no. 3, pp. 626–637, Jun. 2020, doi: [10.1148/radiol.2020190283](https://doi.org/10.1148/radiol.2020190283).
- [12] O. Senouf, S. Vedula, T. Weiss, A. Bronstein, O. Michailovich, and M. Zibulevsky, "Self-supervised learning of inverse problem solvers in medical imaging," in *Proc. Domain Adaptation Represent. Transf. Med. Image Learn. Less Labels Imperfect Data, MICCAI Workshop*, 2019, pp. 111–119, doi: [10.1007/978-3-030-33391-1_13](https://doi.org/10.1007/978-3-030-33391-1_13).
- [13] B. Yaman, S. A. H. Hosseini, S. Moeller, J. Ellermann, K. Uğurbil, and M. Akçakaya, "Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data," *Magn. Reson. Med.*, vol. 84, no. 6, pp. 3172–3191, Dec. 2020, doi: [10.1002/mrm.28378](https://doi.org/10.1002/mrm.28378).
- [14] J. I. Tamir, S. X. Yu, and M. Lustig. (2019). *Unsupervised Deep Basis Pursuit: Learning Inverse Problems Without Ground-truth Data*. Accessed: Oct. 18, 2021. [Online]. Available: <https://openreview.net/forum?id=BylRn72cUH>
- [15] J. Liu, Y. Sun, C. Eldeniz, W. Gan, H. An, and U. S. Kamilov, "RARE: Image reconstruction using deep priors learned without groundtruth," *IEEE J. Sel. Topics Signal Process.*, vol. 14, no. 6, pp. 1088–1099, Oct. 2020, doi: [10.1109/JSTSP.2020.2998402](https://doi.org/10.1109/JSTSP.2020.2998402).
- [16] E. K. Cole, F. Ong, S. S. Vasanaawala, and J. M. Pauly, "Fast unsupervised MRI reconstruction without fully-sampled ground truth data using generative adversarial networks," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, vol. 21, pp. 3988–3997. Accessed: Oct. 18, 2021. [Online]. Available: https://openaccess.thecvf.com/content/ICCV2021W/LCI/html/Cole_Fast_Unsupervised_MRI_Reconstruction_Without_Fully-Sampled_Ground_Truth_Data_Using_ICCVW_2021_paper.html
- [17] M. Z. Darestani and R. Heckel, "Accelerated MRI with un-trained neural networks," *IEEE Trans. Comput. Imag.*, vol. 7, pp. 724–733, Jul. 2021, doi: [10.1109/TCI.2021.3097596](https://doi.org/10.1109/TCI.2021.3097596).
- [18] H. Gu, B. Yaman, K. Ugurbil, S. Moeller, and M. Akçakaya, "Compressed sensing MRI with 11-wavelet reconstruction revisited using modern data science tools," in *Proc. 43rd Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. (EMBC)*, Nov. 2021, pp. 3596–3600, doi: [10.1109/EMBC46164.2021.9630985](https://doi.org/10.1109/EMBC46164.2021.9630985).
- [19] M. Akçakaya, S. Moeller, S. Weingärtner, and K. Uğurbil, "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, 2019, doi: [10.1002/mrm.27420](https://doi.org/10.1002/mrm.27420).
- [20] S. A. H. Hosseini et al., "Accelerated coronary MRI with sRAKI: A database-free self-consistent neural network k-space reconstruction for arbitrary undersampling," *PLoS ONE*, vol. 15, no. 2, Feb. 2020, Art. no. e0229418, doi: [10.1371/journal.pone.0229418](https://doi.org/10.1371/journal.pone.0229418).
- [21] Y. Arefeen, O. Beker, J. Cho, H. Yu, E. Adalsteinsson, and B. Bilgic, "Scan-specific artifact reduction in k-space (SPARK) neural networks synergize with physics-based reconstruction to accelerate MRI," *Magn. Reson. Med.*, vol. 87, no. 2, pp. 764–780, Feb. 2022, doi: [10.1002/mrm.29036](https://doi.org/10.1002/mrm.29036).
- [22] T. Hyung Kim, P. Garg, and J. P. Haldar, "LORAKI: Autocalibrated recurrent neural networks for autoregressive MRI reconstruction in k-space," 2019, *arXiv:1904.09390*, doi: [10.48550/arXiv.1904.09390](https://doi.org/10.48550/arXiv.1904.09390).
- [23] J. P. Haldar, "Autocalibrated loraks for fast constrained MRI reconstruction," in *Proc. IEEE 12th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2015, pp. 910–913, doi: [10.1109/ISBI.2015.7164018](https://doi.org/10.1109/ISBI.2015.7164018).
- [24] M. Uecker, T. Hohage, K. T. Block, and J. Frahm, "Image reconstruction by regularized nonlinear inversion-joint estimation of coil sensitivities and image content," *Magn. Reson. Med.*, vol. 60, no. 3, pp. 674–682, Sep. 2008, doi: [10.1002/mrm.21691](https://doi.org/10.1002/mrm.21691).
- [25] L. Ying and J. Sheng, "Joint image reconstruction and sensitivity estimation in sense (JSENSE)," *Magn. Reson. Med.*, vol. 57, no. 6, pp. 1196–1202, 2007, doi: [10.1002/mrm.21245](https://doi.org/10.1002/mrm.21245).
- [26] H. C. M. Holme, S. Rosenzweig, F. Ong, R. N. Wilke, M. Lustig, and M. Uecker, "ENLIVE: An efficient nonlinear method for calibrationless and robust parallel imaging," *Sci. Rep.*, vol. 9, no. 1, Feb. 2019, Art. no. 1, doi: [10.1038/s41598-019-39888-7](https://doi.org/10.1038/s41598-019-39888-7).
- [27] P. J. Shin et al., "Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion: Calibrationless parallel imaging," *Magn. Reson. Med.*, vol. 72, no. 4, pp. 959–970, Oct. 2014, doi: [10.1002/mrm.24997](https://doi.org/10.1002/mrm.24997).
- [28] K. H. Jin, D. Lee, and J. C. Ye, "A general framework for compressed sensing and parallel MRI using annihilating filter based low-rank Hankel matrix," *IEEE Trans. Comput. Imag.*, vol. 2, no. 4, pp. 480–495, Dec. 2016, doi: [10.1109/TCI.2016.2601296](https://doi.org/10.1109/TCI.2016.2601296).
- [29] G. Ongie and M. Jacob, "Off-the-grid recovery of piecewise constant images from few Fourier samples," *SIAM J. Imag. Sci.*, vol. 9, no. 3, pp. 1004–1041, Jan. 2016, doi: [10.1137/15M1042280](https://doi.org/10.1137/15M1042280).

- [30] J. P. Haldar, "Low-rank modeling of local k-space neighborhoods (LORAKS) for constrained MRI," *IEEE Trans. Med. Imag.*, vol. 33, no. 3, pp. 668–681, Mar. 2014, doi: [10.1109/TMI.2013.2293974](https://doi.org/10.1109/TMI.2013.2293974).
- [31] N. Meng, Y. Yang, Z. Xu, and J. Sun, "A prior learning network for joint image and sensitivity estimation in parallel MR imaging," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, Cham, Switzerland: Springer, 2019, pp. 732–740, doi: [10.1007/978-3-030-32251-9_80](https://doi.org/10.1007/978-3-030-32251-9_80).
- [32] M. Arvinte, S. Vishwanath, A. H. Tewfik, and J. I. Tamir, "Deep J-Sense: Accelerated MRI reconstruction via unrolled alternating optimization," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, Cham, Switzerland: Springer, Sep. 2021, pp. 350–360, doi: [10.1007/978-3-030-87231-1_34](https://doi.org/10.1007/978-3-030-87231-1_34).
- [33] G. Luo, X. Wang, V. Roeloffs, Z. Tan, and M. Uecker, "Joint estimation of coil sensitivities and image content using a deep image prior," in *Proc. Presented Annu. Meeting (ISMRM)*, p. 280. Accessed: Apr. 21, 2023. [Online]. Available: <https://archive.ismrm.org/2021/0280.html>
- [34] G. Luo, N. Zhao, W. Jiang, E. S. Hui, and P. Cao, "MRI reconstruction using deep Bayesian estimation," *Magn. Reson. Med.*, vol. 84, no. 4, pp. 2246–2261, Oct. 2020, doi: [10.1002/mrm.28274](https://doi.org/10.1002/mrm.28274).
- [35] G. Luo, M. Blumenthal, M. Heide, and M. Uecker, "Bayesian MRI reconstruction with joint uncertainty estimation using diffusion models," *Magn. Reson. Med.*, vol. 90, no. 1, pp. 295–311, Jul. 2023, doi: [10.1002/mrm.29624](https://doi.org/10.1002/mrm.29624).
- [36] J. A. Fessler, "Optimization methods for magnetic resonance image reconstruction: Key models and optimization algorithms," *IEEE Signal Process. Mag.*, vol. 37, no. 1, pp. 33–40, Jan. 2020, doi: [10.1109/MSP.2019.2943645](https://doi.org/10.1109/MSP.2019.2943645).
- [37] G. Joseph, S. Kifle, and P. K. Varshney, "One-bit compressed sensing using generative models," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*, May 2020, pp. 537–546. [Online]. Available: <https://proceedings.mlr.press/v70/bora17a.html>
- [38] A. Jalal, M. Arvinte, G. Daras, E. Price, A. G. Dimakis, and J. Tamir, "Robust compressed sensing MRI with deep generative priors," in *Proc. Adv. Neural Inf. Process. Syst. Curran Associates, Inc.*, 2021, pp. 14938–14954. Accessed: Apr. 21, 2023. [Online]. Available: <https://proceedings.neurips.cc/paper/2021/hash/7d6044e95a16761171b130dcb476a43e-Abstract.html>
- [39] D. Narnhofer, K. Hammernik, F. Knoll, and T. Pock, "Inverse GANs for accelerated MRI reconstruction," in *Proc. Wavelets Sparsity 18th*, Sep. 2019, pp. 381–392, doi: [10.1117/12.2527753](https://doi.org/10.1117/12.2527753).
- [40] D. Ulyanov, A. Vadaldi, and V. Lempitsky, "Deep image prior," *Int. J. Comput. Vis.*, vol. 128, no. 7, pp. 1867–1888, 2020, doi: [10.1007/s11263-020-01303-4](https://doi.org/10.1007/s11263-020-01303-4).
- [41] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, no. 6088, pp. 533–536, Oct. 1986, doi: [10.1038/323533a0](https://doi.org/10.1038/323533a0).
- [42] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
- [43] Z. Cheng, M. Gadelha, S. Maji, and D. Sheldon, "A Bayesian perspective on the deep image prior," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2019, pp. 5438–5446, doi: [10.1109/CVPR.2019.00559](https://doi.org/10.1109/CVPR.2019.00559).
- [44] S. Du, J. Lee, H. Li, L. Wang, and X. Zhai, "Gradient descent finds global minima of deep neural networks," in *Proc. 36th Int. Conf. Mach. Learn.*, May pp. 1675–1685. Accessed: Oct. 18, 2021. [Online]. Available: <https://proceedings.mlr.press/v97/du19c.html>
- [45] C. Zhang, J. Bütetage, H. Kjellström, and S. Mandt, "Advances in variational inference," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 8, pp. 2008–2026, Aug. 2019, doi: [10.1109/TPAMI.2018.2889774](https://doi.org/10.1109/TPAMI.2018.2889774).
- [46] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," 2015, *arXiv:1506.02142*.
- [47] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," 2013, *arXiv:1312.6114*.
- [48] I. Loshchilov and F. Hutter. (2018). *Decoupled Weight Decay Regularization*. Accessed: Dec. 17, 2020. [Online]. Available: <https://openreview.net/forum?id=Bkg6RiCqY7>
- [49] J. Zbontar et al., "FastMRI: An open dataset and benchmarks for accelerated MRI," 2018, *arXiv:1811.08839*.
- [50] F. Ong and M. Lustig, "SigPy: A Python package for high performance iterative reconstruction," in *Proc. Int. Soc. Magn. Reson. Med. Montr. QC*, May 2019, p. 4819. [Online]. Available: <https://archive.ismrm.org/2019/4819.html>
- [51] K. Setsompop, B. A. Gagoski, J. R. Polimeni, T. Witzel, V. J. Wedeen, and L. L. Wald, "Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty," *Magn. Reson. Med.*, vol. 67, no. 5, pp. 1210–1224, May 2012, doi: [10.1002/mrm.23097](https://doi.org/10.1002/mrm.23097).
- [52] A. P. Leynes, N. Deveshwar, S. S. Nagarajan, and P. E. Z. Larson, "Scan-specific self-supervised Bayesian deep non-linear inversion for undersampled MRI reconstruction," 2022, *arXiv:2203.00361*, doi: [10.48550/arXiv.2203.00361](https://doi.org/10.48550/arXiv.2203.00361).
- [53] C. M. Bishop, *Pattern Recognition and Machine Learning in Information Science and Statistics*. New York: Springer, 2006.
- [54] G. Trigeorgis, K. Bousmalis, S. Zafeiriou, and B. W. Schuller, "A deep semi-NMF model for learning hidden representations," in *Proc. Int. Conf. Mach. Learn.*, 1692, Jun. 2014, p. 1700.
- [55] D. Liang, J. Cheng, Z. Ke, and L. Ying, "Deep magnetic resonance image reconstruction: Inverse problems meet neural networks," *IEEE Signal Process. Mag.*, vol. 37, no. 1, pp. 141–151, Jan. 2020, doi: [10.1109/MSP.2019.2950557](https://doi.org/10.1109/MSP.2019.2950557).
- [56] J. P. Haldar and K. Setsompop, "Linear predictability in magnetic resonance imaging reconstruction: Leveraging shift-invariant Fourier structure for faster and better imaging," *IEEE Signal Process. Mag.*, vol. 37, no. 1, pp. 69–82, Jan. 2020, doi: [10.1109/MSP.2019.2949570](https://doi.org/10.1109/MSP.2019.2949570).
- [57] A. Pramanik, H. K. Aggarwal, and M. Jacob, "Deep generalization of structured low-rank algorithms (Deep-SLR)," *IEEE Trans. Med. Imag.*, vol. 39, no. 12, pp. 4186–4197, Dec. 2020, doi: [10.1109/TMI.2020.3014581](https://doi.org/10.1109/TMI.2020.3014581).
- [58] A. Pramanik, H. Aggarwal, and M. Jacob, "Calibrationless parallel MRI using model based deep learning (C-MODL)," in *Proc. IEEE 17th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2020, pp. 1428–1431, doi: [10.1109/ISBI45749.2020.9098490](https://doi.org/10.1109/ISBI45749.2020.9098490).
- [59] J. Zhao, X. Liu, S. He, and S. Sun, "Probabilistic inference of Bayesian neural networks with generalized expectation propagation," *Neurocomputing*, vol. 412, pp. 392–398, Oct. 2020, doi: [10.1016/j.neucom.2020.06.060](https://doi.org/10.1016/j.neucom.2020.06.060).
- [60] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," in *Proc. Adv. Neural Inf. Process. Syst.*, pp. 6402–6413. Accessed: Sep. 25, 2019. [Online]. Available: <http://papers.nips.cc/paper/7219-simple-and-scalable-predictive-uncertainty-estimation-using-deep-ensembles.pdf>
- [61] B. Yaman, S. Amir Hossein Hosseini, and M. Akçakaya, "Zero-shot self-supervised learning for MRI reconstruction," 2021, *arXiv:2102.07737*.
- [62] M. Andrychowicz et al., "Learning to learn by gradient descent by gradient descent," in *Proc. Advances in Neural Information Processing Systems*, 2016, pp. 1–8. Accessed: Oct. 18, 2021. [Online]. Available: <https://proceedings.neurips.cc/paper/2016/hash/fb87582825f9d28a8d42c5e5e5e8b23d-Abstract.html>
- [63] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," 2015, *arXiv:1511.06434*.