

MR Image Reconstruction Using Deep Density Priors

Kerem C. Tezcan¹, Christian F. Baumgartner¹, Roger Luechinger,
Klaas P. Pruessmann¹, and Ender Konukoglu

Abstract—Algorithms for magnetic resonance (MR) image reconstruction from undersampled measurements exploit prior information to compensate for missing k-space data. Deep learning (DL) provides a powerful framework for extracting such information from existing image datasets, through learning, and then using it for reconstruction. Leveraging this, recent methods employed DL to learn mappings from undersampled to fully sampled images using paired datasets, including undersampled and corresponding fully sampled images, integrating prior knowledge implicitly. In this letter, we propose an alternative approach that learns the probability distribution of fully sampled MR images using unsupervised DL, specifically variational autoencoders (VAE), and use this as an explicit prior term in reconstruction, completely decoupling the encoding operation from the prior. The resulting reconstruction algorithm enjoys a powerful image prior to compensate for missing k-space data without requiring paired datasets for training nor being prone to associated sensitivities, such as deviations in undersampling patterns used in training and test time or coil settings. We evaluated the proposed method with T1 weighted images from a publicly available dataset, multi-coil complex images acquired from healthy volunteers ($N = 8$), and images with white matter lesions. The proposed algorithm, using the VAE prior, produced visually high quality reconstructions and achieved low RMSE values, outperforming most of the alternative methods on the same dataset. On multi-coil complex data, the algorithm yielded accurate magnitude and phase reconstruction results. In the experiments on images with white matter lesions, the method faithfully reconstructed the lesions.

Index Terms—Reconstruction, MRI, prior probability, machine learning, deep learning, unsupervised learning, density estimation.

Manuscript received October 16, 2018; revised November 27, 2018; accepted December 8, 2018. Date of publication December 17, 2018; date of current version June 28, 2019. This work was supported by the Swiss National Science Foundation under Grant 205321_173016. (Corresponding author: Kerem C. Tezcan.)

K. C. Tezcan, C. F. Baumgartner, and E. Konukoglu are with the Computer Vision Laboratory, ETH Zürich, 8092 Zürich, Switzerland (e-mail: tezcan@vision.ee.ethz.ch; baumgartner@vision.ee.ethz.ch; ender.konukoglu@vision.ee.ethz.ch).

R. Luechinger and K. P. Pruessmann are with the Institute for Biomedical Engineering, ETH Zürich, 8092 Zürich, Switzerland, and also with the University of Zürich, 8006 Zürich, Switzerland (e-mail: luechinger@biomed.ee.ethz.ch; pruessmann@biomed.ee.ethz.ch).

This paper has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the author.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TMI.2018.2887072

I. INTRODUCTION

ACQUISITION time in magnetic resonance (MR) imaging is directly related to the number of samples acquired in k-space. For high quality images, a large number of samples, and therefore long acquisition times are necessary. Reducing acquisition time in a reliable manner is an important question in MR imaging and many methods exploiting different properties of the k-space [1] and hardware design [2] have found wide use in clinical practice.

The question also received considerable attention from the image and signal processing research communities, main focus being reconstruction methods from randomly or regularly undersampled k-space acquisitions. The random undersampling approach was primarily motivated by the compressed sensing framework [3]–[7], where the incoherence between k-space sampling and some sparsifying transform was exploited. Some works used hand-crafted sparsifying transformations, such as gradient in total variation [8] (TV) and wavelet decomposition [9], [10], while others used dictionary learning to determine the transformation from data [11], [12]. An alternative approach in [13] used block matching and 3D filtering, based on the BM3D denoising method [14], to tackle reconstruction as a denoising problem and exploit similarity of image patches to determine a sparsifying transform (BM3D-MRI). On the other hand, regular undersampling schemes and corresponding reconstruction algorithms were also extensively investigated [15]–[18]. The common aspect in all of these approaches, notably except BM3D-MRI, is that they invert underdetermined and ill-posed systems of equations by using good regularizers, which can be viewed as introducing explicit prior information on the expected structure that helps eliminate artifacts.

Recently, researchers started to employ deep neural networks (DNN) [19] for MR reconstruction [20]. Two main approaches have been explored. The first is to learn a mapping from undersampled images to the fully sampled versions. For a given undersampling pattern, [21]–[23] proposed to learn a feed-forward network and during test time fed the undersampled image to the network. Similarly, Han *et al.* [24] proposed to train a network to remove spiking artifacts in computed tomography images and modified the mapping for MR using domain adaptation. The feed-forward network strategy, however, could not guarantee data consistency at test time. To address this, Schlemper *et al.* [25] added

an explicit data-consistency. Here, a mapping block, based on Convolutional Neural Networks (CNN), is followed by a data consistency block and the dual block structure is cascaded. This strategy guaranteed consistency with measured data during test time while the CNNs perform de-aliasing. Alternatively, Wang *et al.* [26] proposed a model where a feed-forward mapping is learned and during test time its output is used as a regularization term alongside the data consistency and other regularizers. The aim was to make sure the final reconstruction did not deviate too much from the mapping output.

The second DL approach leverages networks to improve existing iterative reconstruction algorithms in terms of reconstruction quality. Yang *et al.* [27] showed that the iterations of the alternating direction method of multipliers (ADMM) algorithm can be unrolled as a multi-layer CNN. Instead of being fixed as in the original method, kernels and non-linear functions are parameterized and learned by the network, which improved reconstruction accuracy. A similar method was presented in [28] applying the technology developed in [29], which used the same principle as [27] for improving diffusion filtering. The same strategy of expressing an existing iterative method as a convolutional network was also used in [30] and [31] and in [32] using recurrent neural networks. Algorithms taking either of the DL approaches also integrate prior information on the structure of fully sampled images to account for missing k-space measurements but in an implicit manner, embedded in the trained weights of the feed-forward networks.

In this work, we employ a neural network as an explicit prior similar in essence to the non-DL based reconstruction methods proposed in the signal processing communities, the difference being the power of the prior model. Unsupervised learning with DNNs has been very successful in approximating probability distributions of high dimensional data, including images, from a set of samples. One such approach, of particular interest to the method proposed here, is the variational auto encoder (VAE) algorithm [33], [34]. Using VAEs, it is possible to approximate the distribution of MR image patches and likelihood of a previously unseen image. Furthermore, the approximate likelihood function is a network and therefore differentiable. These two aspects allow using a VAE as a prior model that can approximate distributions of large image patches, e.g. patches of 28×28 pixels, for iterative reconstruction.

We propose a novel probabilistic reconstruction method that uses priors learned via VAEs, which we term as Deep Density Prior (DDP) based reconstruction. We formulate a Bayesian model of the imaging process, including the prior and a data consistency term that embeds the encoding operation, and express DDP as the Maximum-A-Posteriori (MAP) estimation. Compared to non-DL based methods, the main difference of DDP is the powerful prior that can capture the distribution of large image patches. Compared to DL approaches, the fundamental differences of DDP are: (i) the explicit prior, which is trained to capture the structure of the fully sampled images as opposed to the implicit prior in feed-forward networks that is learned to reduce artifacts seen during training; and

(ii) the decoupling of the prior from the data consistency term. The latter difference leads to two theoretical advantages: (i) while previously proposed DL-based approaches required paired datasets for training, including undersampled and corresponding fully sampled images, DPP does not; and (ii) decoupling prior and data consistency terms eliminates possible sensitivities in accuracy to deviations in acquisition specifications between undersampled images used for training and encountered during test time, such as sampling parameters, coil settings and k-space trajectories, which for instance has been demonstrated to be the case in [21].

In the rest of the article, we first present the method and then show reconstruction results and comparisons with conventional approaches as well as recent DNN based methods.

II. METHODS

In the first two parts of this section, we provide a brief background on Bayesian formulation of the MR reconstruction problem and the VAE algorithm. We present our main technical contribution, learning a prior for MR patches and integrating it in the reconstruction problem, starting from Section II-C.

A. Bayesian Formulation of the MR Reconstruction Problem

An MR image is denoted as $\mathbf{m} \in \mathbb{C}^N$, where N is the number of pixels.¹ An imaging operation is given by an undersampling encoding operation $E = UFS$, where $S : \mathbb{C}^N \rightarrow \mathbb{C}^{N \times \gamma}$ is a sensitivity encoding operator. γ is the number of coils, $F : \mathbb{C}^{N \times \gamma} \rightarrow \mathbb{C}^{N \times \gamma}$ is the Fourier operator and $U : \mathbb{C}^{N \times \gamma} \rightarrow \mathbb{C}^{M \times \gamma}$ is an undersampling operator, with $M < N$. Let us also define $\mathbf{x} \in \mathbb{C}^P$ as an image patch of P pixels extracted from \mathbf{m} .

Assuming complex-valued, zero mean, normal distributed and uncorrelated additive noise, denoted as η , the acquired data $\mathbf{y} \in \mathbb{C}^{M \times \gamma}$ can be modelled as $\mathbf{y} = E\mathbf{m} + \eta$. Under this noise model the data likelihood becomes

$$p(\mathbf{y}|\mathbf{m}) = \mathcal{N}(\mathbf{y}|E\mathbf{m}, \sigma_\eta) = \frac{1}{(2\pi\sigma_\eta^2)^{M/2}} e^{-\frac{1}{2\sigma_\eta^2}(E\mathbf{m}-\mathbf{y})^H(E\mathbf{m}-\mathbf{y})}, \quad (1)$$

where H denotes the Hermitian transpose and σ_η is the standard deviation of the noise. In reality, the noise might have some correlation especially in multi-coil acquisitions, which can be corrected by prewhitening. Here, in order to keep things simple, we neglect the noise correlation. In reconstruction, the quantity of interest is the posterior distribution $p(\mathbf{m}|\mathbf{y})$, i.e. the probability of the image being \mathbf{m} given the k-space measurements. A common approach to model the reconstruction problem is to use the MAP estimation

$$\arg \max_{\mathbf{m}} p(\mathbf{m}|\mathbf{y}) = \arg \max_{\mathbf{m}} [p(\mathbf{y}|\mathbf{m})p(\mathbf{m})], \quad (2)$$

where we used the Bayes' theorem and dropped the term $p(\mathbf{y})$ since it does not depend on \mathbf{m} . $p(\mathbf{m})$ is called the prior term

¹In this work we focus on 2D imaging, however, the same techniques can be applied to 3D imaging and this extension will be a part of future work.

and represents the information one has about the fully sampled image before the data acquisition. Taking the log of both sides yields

$$\begin{aligned} \arg \max_{\mathbf{m}} \log p(\mathbf{m}|\mathbf{y}) &= \arg \max_{\mathbf{m}} [\log p(\mathbf{y}|\mathbf{m}) + \log p(\mathbf{m})] \\ &= \arg \max_{\mathbf{m}} \left[-\frac{1}{2\sigma_\eta} \|\mathbf{E}\mathbf{m} - \mathbf{y}\|_2^2 + \log p(\mathbf{m}) \right]. \end{aligned} \quad (3)$$

Taking the maximum (or equivalently taking the minimum of the negative of the expression), defining the constant $\lambda \triangleq 2\sigma_\eta$ and multiplying both terms with it recovers the conventional formulation of a reconstruction problem with a data consistency and a regularization term that is weighted by the trade-off parameter λ

$$\hat{\mathbf{m}} = \arg \min_{\mathbf{m}} \left[\|\mathbf{E}\mathbf{m} - \mathbf{y}\|_2^2 - \lambda \log p(\mathbf{m}) \right]. \quad (4)$$

In this work, we propose to estimate the prior term from examples of fully sampled images and approximate $-\log p(|\mathbf{x}|)$, i.e. the negative log prior of magnitude of image patches, with a neural network model. We train a VAE on patches extracted from fully sampled MR images to capture the distribution and use this prior for reconstruction. This allows us to utilize the prior independent of the sampling operation in contrast to the feed-forward mapping approach.

B. Learning the Data Distribution With VAEs

VAE is an unsupervised learning algorithm proposed to approximate high-dimensional data distributions [33], [34]. We introduce VAEs very briefly² and refer the reader to [33] for further details. VAE is a generic algorithm that can be applied to any signal but we focus on the magnitude image patches in our description.

The main goal of the VAE algorithm is to approximate the data distribution using a latent variable model and optimize its parameters for a given set of examples using variational approximation. The model is given as

$$p(|\mathbf{x}|) = \int_{\mathbf{z}} p(|\mathbf{x}|, \mathbf{z}) d\mathbf{z} = \int_{\mathbf{z}} p(|\mathbf{x}|\mathbf{z}) p(\mathbf{z}) d\mathbf{z}, \quad (5)$$

where $\mathbf{z} \in \mathbb{R}^L$ denotes the latent variable, $p(\mathbf{z})$ the prior over the \mathbf{z} 's and $L \ll P$. A known distribution is assumed for $p(\mathbf{z})$, e.g. unit Gaussian, and a parameterized $p(|\mathbf{x}|\mathbf{z})$ is optimized to maximize $\log p(|\mathbf{x}|)$ of observed samples. This modeling strategy is also taken in other probabilistic latent variable models, such as probabilistic principal component analysis [35]. The VAE model parameterizes $p(|\mathbf{x}|\mathbf{z})$ as a neural network whose set of parameters we denote with φ . To optimize $\log p(|\mathbf{x}|)$ for the given samples, the integral over \mathbf{z} needs to be evaluated and this is not feasible even for moderate L . Variational approximation uses an approximate distribution for the posterior $q(\mathbf{z}|\mathbf{x}) \approx p(\mathbf{z}|\mathbf{x})$ to address this problem. Using $q(\mathbf{z}|\mathbf{x})$, $\log p(|\mathbf{x}|)$ can be decomposed

into two terms [35]

$$\log p(|\mathbf{x}|) = \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} \left[\log \frac{p(|\mathbf{x}|, \mathbf{z})}{q(\mathbf{z}|\mathbf{x})} \right] + \text{D}_{\text{KL}} [q(\mathbf{z}|\mathbf{x}) || p(\mathbf{z}|\mathbf{x})]. \quad (6)$$

The first term is referred to as the evidence lower bound (ELBO) and the second term is the Kullback-Leibler divergence (KLD) between the approximate and true posteriors. The KLD term is intractable because the true posterior $p(\mathbf{z}|\mathbf{x})$ is intractable. It is, however, always larger than or equal to zero, which makes ELBO a lower bound for $\log p(|\mathbf{x}|)$. The strategy of VAE is to maximize the ELBO as a proxy to $\log p(|\mathbf{x}|)$.

Similar to $p(|\mathbf{x}|\mathbf{z})$, the VAE algorithm models $q(\mathbf{z}|\mathbf{x})$ as a separate neural network with parameters θ and during training optimizes both θ and φ to maximize the ELBO of the training samples. Rewriting the ELBO with $p(|\mathbf{x}|\mathbf{z})$, $p(\mathbf{z})$ and $q(\mathbf{z}|\mathbf{x})$, the optimization for training can be written as

$$\max_{\theta, \varphi} \sum_{n=1}^N \text{ELBO}(|\mathbf{x}^n|) = \max_{\theta, \varphi} \left[\sum_{n=1}^N \mathbb{E}_{q_\theta(\mathbf{z}|\mathbf{x}^n)} [\log p_\varphi(|\mathbf{x}^n|\mathbf{z})] - \text{D}_{\text{KL}} [q_\theta(\mathbf{z}|\mathbf{x}^n) || p(\mathbf{z})] \right], \quad (7)$$

where \mathbf{x}^n is the n^{th} training sample and we added the network parameters as subscript at the corresponding terms to indicate the dependence. Notice that the KLD term in Equation 7 is distinct from the one in Equation 6, it includes $p(\mathbf{z})$ instead of $p(\mathbf{z}|\mathbf{x})$, which makes it tractable.

The networks $q_\theta(\mathbf{z}|\mathbf{x})$ and $p_\varphi(|\mathbf{x}|\mathbf{z})$ are typically called the encoder and the decoder, respectively. The former takes a data sample $|\mathbf{x}|$ and encodes it into a posterior distribution in the latent space with network parameters θ . If the posterior distribution $q_\theta(\mathbf{z}|\mathbf{x})$ is modelled as a Gaussian, then the encoder outputs a mean and a covariance matrix for \mathbf{z} depending on $|\mathbf{x}|$. The decoder network on the other hand, takes a latent vector \mathbf{z} and maps it to a conditional distribution of the data given \mathbf{z} . During training, \mathbf{z} vectors are sampled from $q_\theta(\mathbf{z}|\mathbf{x})$ to evaluate the expectations in 7. In this work, we use the original VAE design [33] except for the data likelihood, for which we use a multi-modal Gaussian $p_\varphi(|\mathbf{x}|\mathbf{z}) = N(|\mathbf{x}|, \boldsymbol{\mu}_\varphi(\mathbf{z}), \boldsymbol{\Sigma}_\varphi(\mathbf{z}))$ with a diagonal covariance matrix, similar to [36]. We note that the Gaussian distribution here is different from the data likelihood given in Equation 1. The Gaussian in Equation 1 models the complex valued observation noise whereas the one here models conditional distribution of $|\mathbf{x}|$ given \mathbf{z} in the essence of compound probability distributions. We provide further network design details in the supplementary materials.

C. Deep Density Prior (DDP) Reconstruction Model

Once the VAE model is trained we can integrate the prior within a Bayesian formulation of the reconstruction problem as given in Equation 3. We make two key observations to achieve this. First, given by the theory, the ELBO($|\mathbf{x}|$) can be used as a proxy to the true distribution $\log p(|\mathbf{x}|)$. So, an approximate log likelihood of a magnitude image patch $|\mathbf{x}|$ can be obtained

²Brief explanation is due to space restrictions

by evaluating $\text{ELBO}(|\mathbf{x}|)$

$$\text{ELBO}(|\mathbf{x}|) = \mathbb{E}_{q_{\theta^*}(\mathbf{z}|\mathbf{x})} \left[\log p_{\phi^*}(|\mathbf{x}|\mathbf{z}) + \log \frac{p(\mathbf{z})}{q_{\theta^*}(\mathbf{z}|\mathbf{x})} \right], \quad (8)$$

where θ^* and ϕ^* are the optimal VAE parameters learned during training. The approximate log-likelihood allows us to formulate the proposed reconstruction model as the following MAP estimation problem

$$\arg \min_{\mathbf{m}} \left[\|\mathbf{E}\mathbf{m} - \mathbf{y}\|_2^2 - \sum_{\mathbf{x}_r \in \Omega(\mathbf{m})} \text{ELBO}(|\mathbf{x}_r|) \right], \quad (9)$$

where $\Omega(\mathbf{m})$ denotes a set of (overlapping) patches covering the image \mathbf{m} and $|\mathbf{x}_r|$ is the magnitude of the r^{th} image patch. Note that this approach assumes independence between different patches, ignoring statistical dependencies between them. It would be possible to extend the model to achieve this, which is left for future work.

Since an exact computation of the ELBO term requires evaluating the expectation with respect to $q(\mathbf{z}|\mathbf{x})$, which is computationally not feasible, we use a Monte Carlo sampling approach to calculate the ELBO as follows

$$\text{ELBO}(|\mathbf{x}|) \approx \frac{1}{J} \sum_{j=1}^J \log p(|\mathbf{x}|\mathbf{z}^j) + \log \frac{p(\mathbf{z}^j)}{q(\mathbf{z}^j|\mathbf{x})}, \quad (10)$$

$$\mathbf{z}^j \sim q(\mathbf{z}|\mathbf{x}).$$

Here J represents the number of Monte-Carlo samples.

Plugging the ELBO approximation into Equation 9, we obtain the formulation of the proposed DDP reconstruction problem

$$\arg \min_{\mathbf{m}} \|\mathbf{E}\mathbf{m} - \mathbf{y}\|_2^2 - \sum_{\mathbf{x}_r \in \Omega(\mathbf{m})} \left[\frac{1}{J} \sum_{j=1}^J \log p(|\mathbf{x}_r|\mathbf{z}^j) + \log \frac{p(\mathbf{z}^j)}{q(\mathbf{z}^j|\mathbf{x}_r|)} \right], \quad (11)$$

where $\mathbf{z}^j \sim q(\mathbf{z}|\mathbf{x}_r|)$, the first term is the usual data term and the second term within the summation is the regularization term that arises from the learned prior.

Our second key observation is that the approximation in Equation 10 is differentiable since each term is defined through networks that are themselves differentiable. This is the critical aspect that allows integrating the trained VAE as a prior into an iterative reconstruction algorithm. We can compute the total derivative of the prior term with respect to each image patch as follows

$$\mathcal{R}(|\mathbf{x}|, \mathbf{z}^j) \triangleq \log p(|\mathbf{x}|\mathbf{z}^j) + \log \frac{p(\mathbf{z}^j)}{q(\mathbf{z}^j|\mathbf{x})}$$

$$\frac{d}{d\mathbf{x}} \left[\frac{1}{J} \sum_{j=1}^J \mathcal{R}(|\mathbf{x}|, \mathbf{z}^j) \right] = \frac{1}{J} \sum_{j=1}^J \frac{d}{d\mathbf{x}} \mathcal{R}(|\mathbf{x}|, \mathbf{z}^j)$$

$$= \frac{\mathbf{x}}{|\mathbf{x}|} \left[\frac{1}{J} \sum_{j=1}^J \frac{\partial}{\partial |\mathbf{x}|} \mathcal{R}(|\mathbf{x}|, \mathbf{z}^j) + \frac{\partial}{\partial \mathbf{z}^j} \mathcal{R}(|\mathbf{x}|, \mathbf{z}^j) \frac{d\mathbf{z}^j}{d|\mathbf{x}|} \right], \quad (12)$$

where we defined $\mathcal{R}(|\mathbf{x}|, \mathbf{z}^j)$ for notational simplicity. The second term in the last line is due to the dependency of the samples \mathbf{z}^j on \mathbf{x} and $\mathbf{x}/|\mathbf{x}|$ is due to taking the derivative of the magnitude with respect to the image patch.

D. Optimization Using Projection Onto Convex Sets

We solve the DDP optimization problem given in Equation 11 using the projection onto convex sets (POCS) algorithm [37], specifically using the formulation in [38]. POCS is an iterative minimization process, where the solution variable is projected sequentially onto different convex sets, each defined by one of the constraints in the problem.

The projection for the data consistency term is implemented using the method proposed in [37], which is $\mathcal{P}_{DC}\mathbf{m} = \mathbf{m} - E^H(\mathbf{E}\mathbf{m} - \mathbf{y})$. When there are multiple coils, this projection implements SENSE reconstruction.³ Since we do not have a projection operator for the prior term, we approximate it by several gradient ascent steps with a small step size α as in [38]. We use the final image at the end of the ascent steps as the projected image patch. We define the prior projection with the following steps: i) create a set of patches $\{\mathbf{x}_r^t\} = \Omega(\mathbf{m}^t)$ from the image \mathbf{m}^t at iteration t , ii) obtain the derivatives for each of these patches using Equation 12, which have the same size as the patches themselves, iii) combine the derivatives of the patches to form a derivative image by averaging the values where the patches overlap, iv) update the image using the derivative image, v) repeat this K times. Notice that the set defined by the prior projection is not necessarily convex in $|\mathbf{x}|$, however we have not encountered any problems in convergence during our experiments. To reduce edge effects resulting from patchwise projections, we use four sets of overlapping patches.

With the data consistency and prior projections defined as above, one step of reconstruction within the POCS framework becomes

$$\mathbf{m}^{t+1} = \mathcal{P}_{DC}\mathcal{P}_{prior}\mathbf{m}^t. \quad (13)$$

The prior term in the DDP method does not explicitly provide information on the phase, therefore, reconstruction of the phase is driven by the data consistency projection in the update equation above. For acquisition with multiple coils, the reconstruction method recovers the phase without any modification. For single coil acquisitions, to account for the less amount of information we use an additional projection \mathcal{P}_{phase} in the update equation. A reconstruction iteration is then given as $\mathbf{m}^{t+1} = \mathcal{P}_{DC}\mathcal{P}_{phase}\mathcal{P}_{prior}\mathbf{m}^t$. We use the $\|C \exp(i\angle \mathbf{m})\|$ term [39] as the regularization with C as the finite difference operator, to prefer smooth phase images. We implement \mathcal{P}_{phase} as taking 10 steps with a step size of 0.1 in the negative gradient direction of the regularization term. It is however possible to change this to other constraints on the phase depending on the application, such as a zero-divergence constraint for phase contrast flow imaging reconstruction [40]. It is also possible to extend the

³For optimal signal-to-noise ratio, one needs to account for the noise covariance as well, which we ignore in this work for simplicity.

Algorithm 1 Deep Density Prior (DDP) Reconstruction Using POCS. See Text for a More Detailed Explanation

```

1:  $y$ : undersampled k-space data
2:  $E$ : undersampling encoding operator
3: VAE: the trained VAE
4: procedure DDPRECON( $y$ ,  $E$ , VAE)
5:    $\mathbf{m}^0 \leftarrow E^H y$   $\triangleright$  initialize with the zero-filled image
6:   for  $t = 0 : T - 1$  do  $\triangleright$  main loop: POCS iterations
7:      $\mathbf{m}^{t,0} \leftarrow \mathbf{m}^t$ 
8:     for  $k = 0 : K - 1$  do  $\triangleright$  inner loop: iterations for the
       prior projection  $\mathcal{P}_{prior}$ 
9:        $\{\mathbf{x}_r^{t,k}\} \leftarrow \text{image2patches}(\mathbf{m}^{t,k})$   $\triangleright$  creates a set of
       patches covering the image
10:      for  $r = 1 : \text{no of patches}$  do  $\triangleright$  loop over all the
        patches in  $\{\mathbf{x}_r^t\}$ 
11:         $\mathbf{g}_r \leftarrow \frac{d}{dx} \left[ \frac{1}{J} \sum_{j=1}^J \mathcal{R}(|\mathbf{x}|, \mathbf{z}^j) \right]_{\mathbf{x}=\mathbf{x}_r^{t,k}}$   $\triangleright$ 
        calculate the derivative acc. to Eq. 12
12:      end for
13:       $\mathbf{g} \leftarrow \text{patches2image}(\{\mathbf{g}_r\})$   $\triangleright \mathcal{P}_{prior} \mathbf{m}^t$ 
14:       $\mathbf{m}^{t,k+1} \leftarrow \mathbf{m}^{t,k} + \alpha \cdot \mathbf{g}$ 
15:       $\mathbf{m}^{t,k+1} \leftarrow \mathcal{P}_{phase} \mathbf{m}^{t,k+1}$   $\triangleright$  (optional)
16:    end for
17:     $\mathbf{m}^{t+1} \leftarrow \mathbf{m}^{t,K} - E^H (E \mathbf{m}^{t,K} - y)$   $\triangleright \mathcal{P}_{DC} \mathbf{m}^t$ 
18:  end for
19:  return  $\mathbf{m}^T$   $\triangleright$  Resulting reconstruction
20: end procedure

```

VAE model and train a prior for complex image patches with an appropriate training set.

We apply T POCS steps to complete the reconstruction. Algorithm 1 provides a summary of the reconstruction procedure.

E. Experimental Setup

1) MR Image Data: We used structural images from three different data sources to demonstrate the proposed algorithm.

First, we used images from the Human Connectome Project (HCP) data set [41] (see <https://www.humanconnectome.org/study/hcp-young-adult/document/500-subjects-data-release>). The high quality and large number of images from the HCP dataset are ideal for learning priors with the VAE model. We took 2D slices from the T1 weighted 3D MPRAGE images from 158 subjects (790 images in total) to train the prior VAE model. We normalized the training images by mapping their 99th intensity percentile to 1 per image slice. The VAE model was trained for 200k iterations with a batch size of 50.

Second, to verify that the proposed reconstruction method can be used on a domain that is different from the one the prior is trained on, we experimented with two slices from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) data set (for up-to-date information, see www.adni-info.org). The images were selected from subjects with Alzheimer’s disease and who have visible white matter lesions. Images with lesions allowed us to also test whether the proposed method will be able to faithfully reconstruct such lesions. We extracted the

central slices that showed the largest lesions from these images and further cropped the FOV to 168×224 to remove the empty regions in the images to accelerate computations.

Lastly, we acquired images of 8 healthy volunteers after written informed consent and according to the applicable ethics approval, using a similar protocol to the HCP and ADNI datasets with a 3T Philips Ingenia scanner. The FOV was planned according to the volunteer and a head coil array with 16 elements for receiving was used. We acquired fully sampled complex k-space data for retrospective undersampling study, measuring the coil sensitivities using the standard Philips SENSE reference scan (for only 6 volunteers, due to technical issues) and also estimating them using ESPIRiT [42] (Code: <http://people.eecs.berkeley.edu/~mlustig/Software.html>) with default parameters to obtain the autocalibrated coil maps.

2) Setup and Evaluation: We used images from 17 HCP subjects (separate from training subjects), the ADNI images and the acquired data from 8 subjects in our evaluation. We retrospectively undersampled the test images in k-space, reconstructed them back and compared the results with the original images. For the acquired data, the fully sampled k-space data was present while for HCP and ADNI images, we used Fourier transform to compute the fully sampled k-space data. We experimented with varying undersampling (US) ratios, which we denote with R when presenting results.

We used Cartesian US in one dimension while fully sampling in the other dimension, corresponding to phase-encoding and readout directions, respectively. We present an example of US patterns in Figure 1. We generated the patterns by randomly sampling a 1D Gaussian distribution along the phase-encoding dimension. We randomly drew many sampling patterns from the Gaussian distribution and selected the ones with the best peak-to-side ratio. In addition, we added the central 15 profiles to these selected patterns to fully sample the low-frequency components. We used 2, 3, 4 and 5 for net US ratios (including the fully sampled center). In addition, we used 4-fold radial US with non-uniform fast Fourier transform (NUFFT) [43]. We used the implementation given in [44] with a conjugate gradient solver. For reconstruction in this case, we only changed the fast Fourier transform (FFT) to NUFFT in the data consistency term and did not modify the prior projection.

At reconstruction time the undersampled images need to be in the intensity range for which the DDP prior is trained. To meet this requirement, we normalized the undersampled images to their 99th percentile before reconstruction with DDP. Notice that this may not be the same value as the 99th percentile of the fully sampled images, which would not be available in a real setting for test images.

While assessing our proposed DDP method, we generated a new random US pattern for each test image to make sure our empirical analyses assess the effects of the variability of the US patterns. We reconstructed the test images from HCP and ADNI using 30 POCS iterations ($T = 30$) for $R = 2$ and 3 and 60 for $R = 4$ and 5, 10 iterations for the prior projection ($K = 10$) and the step size was set to $\alpha = 1e-4$, where convergence was observed. We initialize the POCS algorithm using the zero-filled image. When reconstructing

images from volunteers, we run 10 iterations of only data consistency projections initially and then turn on the prior projection. We also use lower number of iterations ($T = 5$ for both $R = 2$ and $R = 3$) to prevent divergence due to discrepancy between the true and used coil sensitivities. We implemented the coil combination as given in [45] at each data projection step.

We compared the reconstruction with the fully sampled “ground truth” images using Root-Mean-Squared-Error (RMSE), Contrast-to-Noise-Ratio (CNR) and Contrast difference (CN) computed at the gray and white matter boundary in the brain. We use the FreeSurfer [46] to obtain the segmentations required for CNR and CN. Details are given in the supplementary materials. We present these results in Section III.

Additionally, we performed three experiments to test sensitivity of ADMM-Net to deviations in acquisition specifications between undersampled images used in training and test. Experiments in [21] provide evidence to this end for basic feed-forward networks. Authors showed that when k-space trajectories between training and test images differ, the feed-forward network’s performance decreased. ADMM-Net is notably different than a basic feed-forward network as it integrates explicit data consistency in the feed-forward architecture. In these experiments we test the null hypothesis that using a different US pattern in reconstruction than in training does not result in a decrease in the performance in the RMSE sense. When measuring RMSE, we used a mask to only measure the reconstruction errors in the brain tissue and skull ignoring artifacts in the background. In each experiment, we trained two networks using different US patterns and applied them on the same test images. In Experiment I, we used $R = 2$ and $R = 4$ for training and tested on $R = 2$. In Experiment II, we used the same networks as in I but tested on $R = 4$. Lastly, we trained networks with $R = 3$ Cartesian and $R = 3$ pseudo-radial patterns and tested on $R = 3$ Cartesian. We used the 17 test images for the evaluation and performed paired Wilcoxon signed-rank tests to assess the null hypothesis.

We further experimented with different latent space dimensions, patch sizes and signal-to-noise ratio in the measurements to better characterize the proposed algorithm. Due to space restrictions, we present these analyses in the supplementary materials.

3) Compared Methods: We implemented several methods to compare against our proposed approach. These are zero-filling reconstruction (ZF), total variation (TV) [3], dictionary learning (DLMRI) [11], ADMM-Net [27], BM3D-MRI [13], SIDWT [9], FDLCP [12] and PBDW [10]. Available implementations for the last three methods only ran on cropped FOV images (square image slices). To compare we used a corresponding DDP reconstruction on cropped FOV. We give the implementation details in the supplementary materials.

III. RESULTS

We start by showing visual reconstruction results. Figure 1 shows results for one of the test images from the HCP data




Fig. 1. Reconstruction results for $R = 3$. (a) shows the fully sampled image (FS) and the undersampling pattern (USP). (b-g) show reconstruction results and the error maps (intensities clipped to $(-0.3, 0.3)$) for the full FOV. (h-k) show the results for the cropped FOV. (l-v) show a zoomed region from the images above. (b-g) are produced with the undersampling pattern that was used to train the ADMM-Net for comparability. In (h-k), we also used the same pattern for undersampling for all methods for comparability.

for $R = 3$. The sampling pattern is also shown in the figure. Visual quality of reconstructions for different methods varied. TV and PBDW reconstructions had problems in restoring the structure and texture. Respective reconstructed images in Figure 1 appear cartoon-like and error images show higher error in regions where gray (GM) and white matter (WM) structures are intertwined. SIDWT did not complete dealiasing and ADMM-Net did not reconstruct the small GM island in the zoomed image. DDP, DLMRI and FDLCP perform well. We show four more randomly selected images from the test set in supplementary Figures S6 and S7 for $R = 2, 3$ as well as reconstruction results for $R = 2, 4, 5$ in supplementary Figures S8, S9 and S10.

We present the quantitative results for reconstruction accuracy in Table I. We show results for full FOV and cropped images separately as the numbers are not comparable. In terms of RMSE our DDP method performed best for the full FOV reconstructions (except for $R = 2$, where BM3D-MRI performed better) and second best against FDLCP for the cropped FOV setting. In terms of CNR, the proposed method performed equally or better than all the other methods for both full and cropped FOVs.

TABLE I

TABLE SUMMARIZING RESULTS FOR DIFFERENT RECONSTRUCTION QUALITY METRICS. NUMBERS INDICATE THE MEAN (AND STANDARD DEVIATION) OF THE ERROR METRIC FOR $N = 17$ TEST IMAGES. TOP GROUP ARE THE RESULTS FOR EXPERIMENTS WITH FULL FOV IMAGES AND BOTTOM GROUP ARE FOR CROPPED FOV IMAGES

	R=2			R=3			R=4			R=5		
	RMSE	CNR	CN	RMSE	CNR	CN	RMSE	CNR	CN	RMSE	CNR	CN
FS	-	0.48(0.10)	0.12(0.02)	-	0.48(0.10)	0.12(0.02)	-	0.48(0.10)	0.12(0.02)	-	0.48(0.10)	0.12(0.02)
Zero-fill	13.03(1.13)	0.40(0.09)	0.12(0.02)	21.15(1.36)	0.33(0.07)	0.09(0.02)	24.92(1.91)	0.31(0.06)	0.08(0.02)	27.36(1.79)	0.30(0.06)	0.08(0.02)
DDP	2.76(0.53)	0.48(0.11)	0.12(0.02)	4.25(0.61)	0.48(0.10)	0.12(0.02)	6.46(1.57)	0.46(0.11)	0.11(0.02)	11.13(2.39)	0.41(0.10)	0.10(0.02)
TV	3.87(0.47)	0.46(0.11)	0.12(0.02)	7.56(0.83)	0.40(0.10)	0.09(0.02)	11.40(1.39)	0.35(0.09)	0.08(0.02)	14.56(1.23)	0.31(0.08)	0.07(0.02)
DLMRI	4.48(0.52)	0.46(0.11)	0.12(0.02)	7.25(0.83)	0.40(0.10)	0.10(0.02)	10.72(1.31)	0.33(0.09)	0.08(0.02)	13.87(1.25)	0.30(0.08)	0.07(0.02)
ADMMNet	3.55(0.40)	0.48(0.11)	0.12(0.02)	7.06(0.52)	0.45(0.11)	0.11(0.02)	11.26(0.72)	0.36(0.09)	0.09(0.02)	13.05(0.70)	0.32(0.08)	0.08(0.02)
BM3D-MRI	1.92(0.36)	0.48(0.10)	0.12(0.02)	4.23(1.05)	0.46(0.10)	0.11(0.02)	8.08(2.48)	0.43(0.10)	0.10(0.02)	11.70(2.76)	0.38(0.09)	0.09(0.02)
FS	-	0.48(0.10)	0.12(0.02)	-	0.48(0.10)	0.12(0.02)	-	0.48(0.10)	0.12(0.02)	-	0.48(0.10)	0.12(0.02)
DDP	2.68(0.38)	0.48(0.10)	0.12(0.02)	4.61(1.12)	0.47(0.10)	0.11(0.02)	7.39(1.47)	0.45(0.10)	0.11(0.02)	13.00(3.01)	0.39(0.08)	0.10(0.02)
SIDWT	4.49(0.98)	0.45(0.11)	0.12(0.02)	9.42(1.62)	0.39(0.09)	0.12(0.02)	14.57(1.96)	0.33(0.08)	0.10(0.02)	18.76(2.80)	0.32(0.07)	0.10(0.02)
FDLCP	2.63(0.35)	0.48(0.10)	0.12(0.02)	4.35(0.87)	0.45(0.10)	0.14(0.03)	6.72(0.89)	0.41(0.10)	0.13(0.03)	9.62(1.48)	0.35(0.08)	0.11(0.02)
PBDW	3.24(0.38)	0.47(0.11)	0.12(0.02)	5.59(0.94)	0.44(0.10)	0.13(0.02)	8.51(0.98)	0.38(0.09)	0.12(0.02)	11.38(1.39)	0.34(0.08)	0.10(0.02)




Fig. 2. Reconstruction results for measured data. Rows show the magnitude and phase of (a) the fully sampled image, (b) the zero-filling image, (c) the DDP reconstruction (using ESPIRiT coil maps) and (d) the error (clipped to $(-0.3, 0.3)$). Upper two rows: A Cartesian US pattern ($R = 2$) was used to retrospectively undersample the k -space data. Lower two rows: A radial US pattern ($R = 4$) was used. In both cases the prior projection is the same, only the data consistency projection differs (using FFT and NUFFT, respectively). (a) FS (b) ZF (c) DDP (d) Error.

In Figure 2 we show reconstruction results from retrospectively undersampled k -space data acquired within this study. First example is reconstruction of 2-fold Cartesian undersampling and the second is of 4-fold radial undersampling using NUFFT in the encoding operation. Our DDP method can reconstruct the original magnitude and phase fairly well, though the magnitude image was smoother than the fully sampled magnitude image in both cases. Quantitatively, the mean (std) RMSE for the Cartesian case for all subjects was 6.97%




Fig. 3. DDP reconstruction results for two images with white matter lesions due to Alzheimer's disease from the ADNI data set for $R = 2$. Images show (a) the original (FS), (b) zero-filling (ZF), (c) reconstructed (DDP) and (d) the error maps. Lesions are clearly visible in the reconstructed images as well. Error map values are clipped to $(-0.3, 0.3)$. Arrows denote lesions. (a) FS (b) ZF (c) DDP (d) Error.

(0.37) using ESPIRiT coil maps ($N = 8$) and 6.92% (0.58) using measured coil maps ($N = 6$) for $R = 2$. Mean (std) value for $R = 3$ was 10.35% (1.53) with ESPIRiT coil maps and 9.93% (1.82) with measured sensitivities. These values were higher than those for HCP data. In order to see the portion of the increase in RMSE due to the inaccuracies in the encoding operator and phase, we also reconstructed the images by setting the phase of the fully sampled image to zero before undersampling and expanding the image by known coil maps. In this case the mean RMSE reduced to 5.76% (0.56) ($R = 2$, $N = 8$). Furthermore, we also reconstructed the images using only data consistency projections, corresponding to SENSE reconstruction, to see the added value of the DDP projections. In this case the mean RMSE was 7.95% and 11.42% for ($R = 2, 3$, respectively, $N = 8$, ESPIRiT maps). We present more information on this point in the supplementary materials.

In Figure 3, we show DDP reconstructions for the ADNI images for $R = 2$. We used the VAE model that was trained on the HCP dataset, which had only healthy subjects, to reconstruct the images here. The reconstructed images recover GM

TABLE II

RECONSTRUCTION RESULTS FOR ADMM-NET, WHEN USING DIFFERENT US PATTERNS AT TRAINING AND TEST TIMES. STATISTICS WERE DERIVED FROM $N = 17$. R_{train} AND R_{test} DENOTE THE US FACTOR (PATTERN) USED IN TRAINING AND TESTING. IF NOT STATED, A CARTESIAN PATTERN (CART) WAS USED AND "RAD" DENOTES PSEUDO-RADIAL PATTERN. RMSE DIFF. REFERS TO THE N^{th} PERCENTILE OR MEDIAN VALUE OF THE DISTRIBUTION OF PAIRWISE RMSE DIFFERENCES. p-VALUES ARE CALCULATED WITH PAIRED WILCOXON SIGNED-RANK TEST. LAST ROW PRESENTS THE DDP RMSE VALUES FOR COMPARISON PURPOSES. THESE RMSE VALUES WERE CALCULATED WITH DISREGARDING THE ARTIFACTS IN THE BACKGROUND USING A BRAIN MASK

	Experiment I		Experiment II		Experiment III	
R_train	2	4	4	2	3 cart	3 rad
R_US	2	2	4	4	3 cart	3 cart
RMSE mean	3.17	4.11	10.55	10.70	6.49	6.67
RMSE std	0.39	0.44	0.68	0.66	0.49	0.51
RMSE diff. 25%	0.82		0.09		0.14	
RMSE diff. median	0.86		0.15		0.16	
RMSE diff. 75%	1.02		0.22		0.2	
p-value	0.0002		0.0009		0.0002	
RMSE DDP mean (std)	2.47 (0.44)		6.08 (1.45)		3.93 (0.56)	

and WM structures and edges faithfully. The WM lesions were also well reconstructed. The error maps do not indicate a specific increase in the lesion regions.

In addition to these results, we show the convergence of the POCS algorithm in Figure S3, and results for patch size, latent dimension and noise analyses in Figures S4 and S5 of the supplementary materials.

Lastly, we show results for the experiments assessing sensitivity to deviations in acquisition specifications between training and test images of ADMM-Net in Table II. All the differences were statistically significant at the 0.05 level according to the paired test, performance of the method decreased significantly when the training and test patterns/ratios differed. Performance differences in the Experiment I were particularly high.

IV. DISCUSSION

The reconstruction examples in Figure 1 and the quantitative results in Table I show that the proposed deep density prior reconstruction method produced highly accurate reconstructions both in terms of RMSE and CNR compared to the other methods. While BM3D-MRI and FDLCP could achieve slightly lower RMSE's, they both decreased CNR further than DDP, indicating that both methods sacrificed contrast to reduce RMSE, generating blurrier images. Leveraging the powerful prior over image patches, DDP was able to restore structures in the image faithfully while removing aliasing artifacts. Comparing methods that use explicit priors, DDP's performance demonstrate the huge potential in DL-based priors for MRI reconstruction. Compared to ADMM-Net, DDP's higher accuracy shows that it is a highly attractive alternative to feed-forward network approaches.

Experiments with images from the ADNI dataset and acquired k-space data showed further properties of DDP and also prompted some key questions for further research.

Results in Figure 2 show that the DDP method yielded high quality reconstructions for both the magnitude and phase images, though with higher mean RMSE values than those in the HCP test images. The fact that the RMSE decreased when the correct coil sensitivities and phase were used, suggest that some part of the error is due to the discrepancies between real and measured coil sensitivities. Hence, increasing accuracy of the coil maps would potentially decrease the error rate. Similarly, the method currently uses theoretical or no prior for the phase but it would benefit from better priors for this component as well.

The ADNI reconstruction results are encouraging in two respects. First, they show that the learned model does not blur out the lesions during the prior projection, which could have been expected since the training images did not contain any examples of lesions. We believe the lesions could be reconstructed because they are structurally not very different from healthy brain structures. In addition, the data consistency term makes sure the sampled information regarding the lesions is kept in the final image. However, the proposed method, similar to other reconstruction methods, requires further investigation as to where its limits lie, such as with bigger, brighter lesions. We believe for an optimal treatment of lesions, the training data set should include such images to provide the prior the capability to represent them.

Second, the proposed method performed reasonably well despite the domain difference between the training and test sets. Although the two data sets were acquired at the same field strength, they still differed in some acquisition protocol and imaging parameters. Their FOV and voxel resolution are different, and HCP was acquired with fat suppression while ADNI was not, which affects the image contrast. Our method is invariant to changes in FOV but not in scale. The lack of fat suppression also makes the dealiasing more challenging in the ADNI images since artifacts become more prominent. This results in very faint artifacts in the ventricles in the reconstructed images. Despite these differences, the quality of the reconstruction results indicate that the learned prior model could generalize to slightly different scales and similar but not identical imaging protocols. However, these experiments raised two key questions and exciting research directions regarding the sensitivity of the proposed method.

Even though the proposed method is not sensitive to variations in undersampling patterns or coil settings, when contrast or resolution of the acquired images differ substantially from training images used to learn the prior, we believe reconstruction quality will decrease. There are two interesting directions to remedy this issue. First is to use a different prior for different contrast and resolutions. Second is to further improve the reconstruction quality using appropriate domain adaptation methods. Integrating invariance to domain differences in the prior and even building a joint prior for multiple contrasts are interesting research directions. Similar sensitivity issues can also be expected when the underlying anatomy differs between the acquired image and the prior. In this scenario, the safest approach would be to train a prior for each anatomy but building a prior that is capable of representing multiple anatomies is also an interesting direction. We also note that

it is possible that all learning based methods suffer from the mentioned sensitivities [47].

An underlying assumption of our model is the unit Gaussian prior for the latent space, for which different extensions such as Gaussian mixtures or graphical models were investigated in [48] and [49], which could in principle improve the representation capacity of the prior. Similarly, different density estimation methods are investigated, for instance using generative adversarial networks (GAN) [50], [51]. Ledig *et al.* [52] proposed to use the discriminator of a GAN to superresolve images. Shah and Hegde *et al.* [53] explain the drawbacks of this approach and similar to [54] suggest projecting images onto the span of the generator. A similar approach is given in [55]. The advantage of using VAEs is the simplicity of the inverse mapping. Determining the latent space representation of a given image is much more straightforward with VAEs, not requiring a separate optimization as GANs.

Another alternative approach to the VAE would be to use denoising autoencoders (DAE) [56]. Even though these might offer faster algorithms, the advantages of VAEs compared to DAEs are i) a more principled way of approximating the target distribution and ii) ease and rigor in approximating the likelihood of an image patch due to the variational inference mechanism that underlies VAEs.

One limitation of our method is the requirement of training in contrast to methods using fixed bases for regularization. Training of the prior relies on the availability of high quality data. Such data is available for commonly used sequences, such as structural T1w and T2w MRI. For other sequences, such as functional or diffusion MRI, construction of appropriate training set is of interest for future research.

Experiments with different configurations, which are presented in the supplementary materials, showed that the DDP reconstruction is not sensitive to patch size and latent space dimensions for a reasonable range. Furthermore, reconstruction quality is higher when patches larger than 12×12 are used. In addition, as expected, the performance degrades with decreasing SNR, however, with a low rate indicating robustness of the method to noise. This means the method is likely to perform fairly well in the regions with higher g-factor that have lower SNR due to parallel imaging.

In our current implementation a single evaluation of the derivative through the network takes around 9 seconds for $360 \times 28 \times 28$ patches (0.6 seconds per batch of 25 patches in parallel). The total run time is given by this multiplied by number of iterations and all the other operations including mainly the phase projection, data projection and application of the derivatives for these operations. This time can be reduced by optimizing the code, increasing the parallelization and changing the network to work with images directly rather than patches. Furthermore, we demonstrate our method on 2D slices with a single phase encoding/undersampling direction as a proof of concept. However the method can be extended to a 3D setting, either by doing slice-by-slice reconstruction with a VAE trained on the whole brain or by training a VAE with 3D patches.

Results presented in Table II demonstrate that ADMM-Net model is also sensitive to deviations in undersampling patterns

used in training and test images, similar to the method proposed in [21]. Other feed-forward networks may also show similar sensitivities however, to the best of our knowledge, such an analysis is rarely performed. In contrast, we emphasize that the proposed method, due to decoupling of the prior and the data consistency term, does not share the same sensitivity. It learns the prior distribution on fully-sampled images and can be used to reconstruct any sampling scheme faithful to the measured data without the need of retraining as long as the images are from the same domain.

V. CONCLUSION

In this paper we proposed a novel method termed DDP for MR reconstruction from undersampled k-space acquisitions. The method uses the VAE algorithm to learn the distribution of MR patches from fully sampled images, removing the sensitivity of the model to the sampling pattern. The model then uses this learned distribution as a probabilistic prior in a Bayesian reconstruction framework. We have shown that the reconstruction with the DDP approach yielded promising results for HCP and ADNI data sets as well as multi-channel k-space measurements in terms of visual quality and quantitative measures.

ACKNOWLEDGMENTS

The authors would like to acknowledge Prof. Sebastian Kozlerke for valuable discussions. They thank GyroTools for their MRRecon software. They thank NVIDIA for their GPU donation. They also thank the reviewers for comments/suggestions that substantially improved this work. They used images from the HCP and ADNI datasets.

REFERENCES

- [1] D. A. Feinberg, J. D. Hale, J. C. Watts, L. Kaufman, and A. Mark, "Halving mr imaging time by conjugation: Demonstration at 3.5 kg," *Radiology*, vol. 161, no. 2, pp. 527–531, 1986, doi: 10.1148/radiology.161.2.3763926.
- [2] 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, 1999.
- [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.
- [4] M. Guerin-kern, M. Häberlin, K. P. Pruessmann, and M. Unser, "A fast wavelet-based reconstruction method for magnetic resonance imaging," *IEEE Trans. Med. Imag.*, vol. 30, no. 9, pp. 1649–1660, Sep. 2011.
- [5] S. G. Lingala, Y. Hu, E. DiBella, and M. Jacob, "Accelerated dynamic MRI exploiting sparsity and low-rank structure: k-t SLR," *IEEE Trans. Med. Imag.*, vol. 30, no. 5, pp. 1042–1054, May 2011.
- [6] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [7] E. J. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Commun. Pure Appl. Math.*, vol. 59, no. 8, pp. 1207–1223, 2006.
- [8] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D, Nonlinear Phenomena*, vol. 60, nos. 1–4, pp. 259–268, 1992.
- [9] B. Ning, X. Qu, D. Guo, C. Hu, and Z. Chen, "Magnetic resonance image reconstruction using trained geometric directions in 2D redundant wavelets domain and non-convex optimization," *Magn. Reson. Imag.*, vol. 31, no. 9, pp. 1611–1622, 2013.
- [10] X. Qu *et al.*, "Undersampled MRI reconstruction with patch-based directional wavelets," *Magn. Reson. Imag.*, vol. 30, no. 7, pp. 964–977, Sep. 2012.

- [11] S. Ravishanker and Y. Bresler, "Mr image reconstruction from highly undersampled k-space data by dictionary learning," *IEEE Trans. Med. Imag.*, vol. 30, no. 5, pp. 1028–1041, May 2010.
- [12] Z. Zhan, J.-F. Cai, D. Guo, Y. Liu, Z. Chen, and X. Qu, "Fast multiclass dictionaries learning with geometrical directions in MRI reconstruction," *IEEE Trans. Biomed. Eng.*, vol. 63, no. 9, pp. 1850–1861, Sep. 2016.
- [13] E. M. Eksioğlu, "Decoupled algorithm for MRI reconstruction using nonlocal block matching model: BM3D-MRI," *J. Math. Imag. Vis.*, vol. 56, no. 3, pp. 430–440, Nov. 2016.
- [14] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.
- [15] L. Wissmann *et al.*, "Analysis of spatiotemporal fidelity in quantitative 3D first-pass perfusion cardiovascular magnetic resonance," *J. Cardiovascular Magn. Reson.*, vol. 19, no. 1, p. 11, Jan. 2017.
- [16] K. P. Pruessmann, M. Weiger, P. Börner, and P. Boesiger, "Advances in sensitivity encoding with arbitrary K -space trajectories," *Magn. Reson. Med.*, vol. 46, no. 4, pp. 638–651, 2001.
- [17] H. Pedersen, S. Kozerke, S. Ringgaard, K. Nehrke, and Y. K. Won, " K - t PCA: Temporally constrained k - t BLAST reconstruction using principal component analysis," *Magn. Reson. Med.*, vol. 62, no. 3, pp. 706–716, 2009.
- [18] J. F. M. Schmidt, L. Wissmann, R. Manka, and S. Kozerke, "Iterative k - t principal component analysis with nonrigid motion correction for dynamic three-dimensional cardiac perfusion imaging," *Magn. Reson. Med.*, vol. 72, no. 1, pp. 68–79, 2014.
- [19] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, pp. 436–444, May 2015.
- [20] G. Litjens *et al.*, "A survey on deep learning in medical image analysis," *Med. Image Anal.*, vol. 42, pp. 60–88, Dec. 2017.
- [21] B. Zhu, J. Z. Liu, S. F. Cauley, B. R. Rosen, and M. S. Rosen, "Image reconstruction by domain-transform manifold learning," *Nature*, vol. 555, pp. 487–492, Mar. 2018.
- [22] K. Kwon, D. Kim, and H. Park, "A parallel mr imaging method using multilayer perceptron," *Med. Phys.*, vol. 44, no. 12, pp. 6209–6224, 2017.
- [23] D. Lee, J. Yoo, and J. C. Ye, "Deep residual learning for compressed sensing MRI," in *Proc. IEEE 14th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2017, pp. 15–18.
- [24] Y. S. Han, J. Yoo, and J. C. Ye, (Mar. 2017). "Deep learning with domain adaptation for accelerated projection-reconstruction MR." [Online]. Available: <https://arxiv.org/abs/1703.01135>
- [25] J. Schlemper, J. Caballero, J. V. Hajnal, A. N. Price, and D. Rueckert, *A Deep Cascade of Convolutional Neural Networks for MR Image Reconstruction*. Cham, Switzerland: Springer, 2017, pp. 647–658.
- [26] S. Wang *et al.*, "Accelerating magnetic resonance imaging via deep learning," in *Proc. IEEE 13th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2016, pp. 514–517.
- [27] Y. Yang, J. Sun, H. Li, and Z. Xu, "Deep ADMM-net for compressive sensing MRI," in *Advances in Neural Information Processing Systems*, vol. 29, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, Eds. Red Hook, NY, USA: Curran Associates, 2016, pp. 10–18.
- [28] K. Hammernik *et al.*, "Learning a variational network for reconstruction of accelerated MRI data," *Magn. Reson. Med.*, vol. 79, no. 6, pp. 3055–3071, 2017.
- [29] Y. Chen, W. Yu, and T. Pock. (2015). "On learning optimized reaction diffusion processes for effective image restoration." [Online]. Available: <https://arxiv.org/abs/1503.05768>
- [30] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, "Deep convolutional neural network for inverse problems in imaging," *IEEE Trans. Image Process.*, vol. 26, no. 9, pp. 4509–4522, Sep. 2017.
- [31] H. K. Aggarwal, M. P. Mani, and M. Jacob. (2017). "MoDL: Model based deep learning architecture for inverse problems." [Online]. Available: <https://arxiv.org/abs/1712.02862>
- [32] P. Putzky and M. Welling. (2017). "Recurrent inference machines for solving inverse problems." [Online]. Available: <https://arxiv.org/abs/1706.04008>
- [33] D. Kingma and M. Welling, "Auto-encoding variational Bayes," in *Proc. 2nd Int. Conf. Learn. Represent. (ICLR)*, 2014.
- [34] D. J. Rezende, S. Mohamed, and D. Wierstra, "Stochastic backpropagation and approximate inference in deep generative models," in *Proc. 31st Int. Conf. Mach. Learn. (Proceedings of Machine Learning Research)*, vol. 32, no. 2, E. P. Xing and T. Jebara, Eds. Beijing, China: PMLR, Jun. 2014, pp. 1278–1286.
- [35] C. M. Bishop, *Pattern Recognition and Machine Learning*. Berlin, Germany: Springer, 2006.
- [36] S. Yeung, A. Kannan, Y. Dauphin, and L. Fei-Fei. (2017). "Tackling over-pruning in variational autoencoders." [Online]. Available: <https://arxiv.org/abs/1706.03643>
- [37] A. A. Samsonov, E. G. Kholmovski, D. L. Parker, and C. R. Johnson, "POCSense: POCS-based reconstruction for sensitivity encoded magnetic resonance imaging," *Magn. Reson. Med.*, vol. 52, no. 6, pp. 1397–1406, 2004.
- [38] A. R. De Pierro and E. S. H. Neto, "From convex feasibility to convex constrained optimization using block action projection methods and underrelaxation," *Int. Trans. Oper. Res.*, vol. 16, no. 4, pp. 495–504, 2009.
- [39] F. Zhao, D. C. Noll, J.-F. Nielsen, and J. A. Fessler, "Separate magnitude and phase regularization via compressed sensing," *IEEE Trans. Med. Imag.*, vol. 31, no. 9, pp. 1713–1723, Sep. 2012.
- [40] C. Santelli, M. Loecherer, J. Busch, O. Wieben, T. Schaeffter, and S. Kozerke, "Accelerating 4D flow MRI by exploiting vector field divergence regularization," *Magn. Reson. Med.*, vol. 75, no. 1, pp. 115–125, 2016.
- [41] D. C. Van Essen *et al.*, "The human connectome project: A data acquisition perspective," *NeuroImage*, vol. 62, no. 4, pp. 2222–2231, 2012.
- [42] 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, 2014.
- [43] J. A. Fessler and B. P. Sutton, "Nonuniform fast Fourier transforms using min-max interpolation," *IEEE Trans. Signal Process.*, vol. 51, no. 2, pp. 560–574, Feb. 2003.
- [44] J.-M. Lin and H.-W. Chung, "Pynufft: Python non-uniform fast Fourier transform for MRI," in *Building Bridges in Medical Sciences*. Cambridge, U.K.: St John's College, 2017.
- [45] M. Bydder, D. J. Larkman, and J. V. Hajnal, "Combination of signals from array coils using image-based estimation of coil sensitivity profiles," *Magn. Reson. Med.*, vol. 47, no. 3, pp. 539–548, 2002.
- [46] M. Reuter, N. J. Schmansky, H. D. Rosas, and B. Fischl, "Within-subject template estimation for unbiased longitudinal image analysis," *NeuroImage*, vol. 61, no. 4, pp. 1402–1418, 2012.
- [47] F. Knoll, K. Hammernik, E. Kobler, T. Pock, M. P. Recht, and D. K. Sodickson, "Assessment of the generalization of learned image reconstruction and the potential for transfer learning," *Magn. Reson. Med.*, vol. 81, no. 1, pp. 116–128, 2019.
- [48] N. Dilokthanakul *et al.* (2016). "Deep unsupervised clustering with gaussian mixture variational autoencoders." [Online]. Available: <https://arxiv.org/abs/1611.02648>
- [49] M. J. Johnson, D. Duvenaud, A. B. Wiltschko, S. R. Datta, and R. P. Adams, "Composing graphical models with neural networks for structured representations and fast inference," in *Advances in Neural Information Processing Systems*, vol. 29, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, Eds. Red Hook, NY, USA: Curran Associates, 2016, pp. 2946–2954.
- [50] I. J. Goodfellow *et al.*, "Generative adversarial nets," in *Advances in Neural Information Processing Systems*, vol. 27, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, Eds. Red Hook, NY, USA: Curran Associates, 2014, pp. 2672–2680. [Online]. Available: <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>
- [51] J. H. R. Chang, C.-L. Li, B. Póczos, and B. V. K. V. Kumar, "One network to solve them all—Solving linear inverse problems using deep projection models," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 5889–5898.
- [52] C. Ledig *et al.*, "Photo-realistic single image super-resolution using a generative adversarial network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 105–114.
- [53] V. Shah and C. Hegde. (2018). "Solving linear inverse problems using GAN priors: An algorithm with provable guarantees." [Online]. Available: <https://arxiv.org/abs/1802.08406>
- [54] R. A. Yeh, C. Chen, T. Y. Lim, A. G. Schwing, M. Hasegawa-Johnson, and M. N. Do. (Jul. 2016). "Semantic image inpainting with deep generative models." [Online]. Available: <https://arxiv.org/abs/1607.07539>
- [55] M. Mardani, H. Monajemi, V. Pappayan, S. Vasanaawala, D. L. Donoho, and J. M. Pauly. (2017). "Recurrent generative adversarial networks for proximal learning and automated compressive image recovery." [Online]. Available: <https://arxiv.org/abs/1711.10046>
- [56] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol, "Extracting and composing robust features with denoising autoencoders," in *Proc. ACM 25th Int. Conf. Mach. Learn. (ICML)*, W. W. Cohen, A. McCallum, and S. T. Roweis, Eds. 2008, pp. 1096–1103.