

Bayesian Nonparametric Dictionary Learning for Compressed Sensing MRI

Yue Huang, John Paisley, Qin Lin, Xinghao Ding, Xueyang Fu, and Xiao-Ping Zhang, *Senior Member, IEEE*

Abstract—We develop a Bayesian nonparametric model for reconstructing magnetic resonance images (MRIs) from highly undersampled k -space data. We perform dictionary learning as part of the image reconstruction process. To this end, we use the beta process as a nonparametric dictionary learning prior for representing an image patch as a sparse combination of dictionary elements. The size of the dictionary and patch-specific sparsity pattern are inferred from the data, in addition to other dictionary learning variables. Dictionary learning is performed directly on the compressed image, and so is tailored to the MRI being considered. In addition, we investigate a total variation penalty term in combination with the dictionary learning model, and show how the denoising property of dictionary learning removes dependence on regularization parameters in the noisy setting. We derive a stochastic optimization algorithm based on Markov chain Monte Carlo for the Bayesian model, and use the alternating direction method of multipliers for efficiently performing total variation minimization. We present empirical results on several MRI, which show that the proposed regularization framework can improve reconstruction accuracy over other methods.

Index Terms—Compressed sensing, magnetic resonance imaging, Bayesian nonparametrics, dictionary learning.

I. INTRODUCTION

MAGNETIC resonance imaging (MRI) is a widely used technique for visualizing the structure and functioning of the body. A limitation of MRI is its slow scan speed during data acquisition. Therefore, methods for accelerating the MRI process have been heavily researched. Recent advances in signal reconstruction from measurements sampled below the Nyquist rate, called compressed sensing (CS) [1], [2],

have had a major impact on MRI [3]. CS-MRI allows for significant undersampling in the Fourier measurement domain of MR images (called k -space), while still outputting a high-quality image reconstruction. While image reconstruction using this undersampled data is a case of an ill-posed inverse problem, compressed sensing theory has shown that it is possible to reconstruct a signal from significantly fewer measurements than mandated by traditional Nyquist sampling if the signal is sparse in a particular transform domain.

Motivated by the need to find a sparse domain for representing the MR signal, a large body of literature now exists on reconstructing MRI from significantly undersampled k -space data. Existing improvements in CS-MRI mostly focus on (i) seeking sparse domains for the image, such as contourlets [4], [5]; (ii) using approximations of the ℓ_0 norm for better reconstruction performance with fewer measurements, for example ℓ_1 , FOCUSS, ℓ_p quasi-norms with $0 < p < 1$, or using smooth functions to approximate the ℓ_0 norm [6], [7]; and (iii) accelerating image reconstruction through more efficient optimization techniques [8], [10], [29]. In this paper we present a modeling framework that is similarly motivated.

CS-MRI reconstruction algorithms tend to fall into two categories: Those which enforce sparsity directly within some image transform domain [3]–[8], [10]–[12], and those which enforce sparsity in some underlying latent representation of the image, such as an adaptive dictionary-based representation [9], [14]. Most CS-MRI reconstruction algorithms belong to the first category. For example Sparse MRI [3], the leading study in CS-MRI, performs MR image reconstruction by enforcing sparsity in both the wavelet domain and the total variation (TV) of the reconstructed image. Algorithms with image-level sparsity constraints such as Sparse MRI typically employ an off-the-shelf basis, which can usually capture only one feature of the image. For example, wavelets recover point-like features, while contourlets recover curve-like features. Since MR images contain a variety of underlying features, such as edges and textures, using a basis not adapted to the image can be considered a drawback of these algorithms.

Finding a sparse basis that is suited to the image at hand can benefit MR image reconstruction, since CS theory shows that the required number of measurements is linked to the sparsity of the signal in the selected transform domain. Using a standard basis not adapted to the image under consideration will likely not provide a representation that can compete in sparsity with an adapted basis. To this end, dictionary learning,

Manuscript received August 21, 2013; revised May 2, 2014 and July 25, 2014; accepted September 18, 2014. Date of publication September 24, 2014; date of current version October 13, 2014. This work was supported in part by the National Natural Science Foundation of China under Grant 30900328, Grant 61172179, Grant 61103121, and Grant 81301278, in part by the Fundamental Research Funds for the Central Universities under Grant 2011121051 and Grant 2013121023, and in part by the Natural Science Foundation of Fujian Province, China, under Grant 2012J05160. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Ali Bilgin. (*Corresponding author: Xinghao Ding.*)

Y. Huang, Q. Lin, X. Ding, and X. Fu are with the Department of Communications Engineering, Xiamen University, Xiamen 361005, China (e-mail: yhuang2010@xmu.edu.cn; fzqinwuye@gmail.com; dxh@xmu.edu.cn; fxy@stu.xmu.edu.cn).

J. Paisley is with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: jwp2128@columbia.edu).

X.-P. Zhang is with the Department of Electrical and Computer Engineering, Ryerson University, Toronto, ON M5B 2K3, Canada (e-mail: xzhang@ee.ryerson.ca).

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/TIP.2014.2360122

which falls in the second group of algorithms, learns a sparse basis on image subregions called patches that is adapted to the image class of interest. Recent studies in the image processing literature have shown that dictionary learning is an effective means for finding a sparse, patch-level representation of an image [19], [20], [25]. These algorithms learn a patch-level dictionary by exploiting structural similarities between patches extracted from images within a class of interest. Among these approaches, adaptive dictionary learning—where the dictionary is learned directly from the image being considered—based on patch-level sparsity constraints usually outperforms analytical dictionary approaches in denoising, super-resolution reconstruction, interpolation, inpainting, classification and other applications, since the adaptively learned dictionary suits the signal of interest [19]–[22].

Dictionary learning has previously been applied to CS-MRI to learn a sparse basis for reconstruction, see [14]. With these methods, parameters such as the dictionary size and patch sparsity are preset, and algorithms are considered that are non-Bayesian. In this paper, we consider a new dictionary learning algorithm for CS-MRI that is motivated by Bayesian nonparametric statistics. Specifically, we consider a nonparametric dictionary learning model called BPFA [23] that uses the beta process to learn the sparse representation necessary for CS-MRI reconstruction. The beta process is an effective prior for nonparametric learning of latent factor models; in this case the latent factors correspond to dictionary elements. While the dictionary size is therefore infinite in principle, through posterior inference the beta process learns a suitably compact dictionary in which the signal can be sparsely represented.

We organize the paper as follows. In Section II we review CS-MRI inversion methods and the beta process for dictionary learning. In Section III, we describe the proposed regularization framework and algorithm. We derive a Markov Chain Monte Carlo (MCMC) sampling algorithm for stochastic optimization of the dictionary variables in the objective function. In addition, we consider including a sparse total variation (TV) penalty, for which we perform efficient optimization using the alternating direction method of multipliers (ADMM). We then show the advantages of the proposed Bayesian nonparametric regularization framework on several CS-MRI problems in Section IV.

II. BACKGROUND AND RELATED WORK

We use the following notation: Let $x \in \mathbb{C}^N$ be a $\sqrt{N} \times \sqrt{N}$ MR image in vectorized form. Let $\mathcal{F}_u \in \mathbb{C}^{u \times N}$, $u \ll N$, be the undersampled Fourier encoding matrix and $y = \mathcal{F}_u x \in \mathbb{C}^u$ represent the sub-sampled set of k -space measurements. The goal is to estimate x from the small fraction of k -space measurements y . For dictionary learning, let R_i be the i^{th} patch extraction matrix. That is, R_i is a $P \times N$ matrix of all zeros except for a one in each row that extracts a vectorized $\sqrt{P} \times \sqrt{P}$ patch from the image, $R_i x \in \mathbb{C}^P$ for $i = 1, \dots, N$. We use overlapping image patches with a shift of one pixel and allow a patch to wrap around the image at the boundaries for mathematical convenience [15], [22]. All norms are extended to complex vectors when necessary, $\|a\|_p = (\sum_i |a_i|^p)^{1/p}$, where $|a_i|$ is the modulus of the complex number a_i .

A. Two Approaches to CS-MRI Inversion

We focus on single-channel CS-MRI inversion via optimizing an unconstrained function of the form

$$\arg \min_x h(x) + \frac{\lambda}{2} \|\mathcal{F}_u x - y\|_2^2, \quad (1)$$

where $\|\mathcal{F}_u x - y\|_2^2$ is a data fidelity term, $\lambda > 0$ is a parameter and $h(x)$ is a regularization function that controls properties of the image we want to reconstruct. As discussed in the introduction, the function h can take several forms, but tends to fall into one of two categories according to whether image-level or patch-level information is considered. We next review these two approaches.

1) *Image-Level Sparse Regularization*: CS-MRI with an image-level, or global regularization function $h_g(x)$ is one in which sparsity is enforced within a transform domain defined on the entire image. For example, in Sparse MRI [3] the regularization function is

$$h_g(x) = \|Wx\|_1 + \mu TV(x), \quad (2)$$

where W is the wavelet basis and $TV(x)$ is the total variation (spatial finite differences) of the image. Regularizing with this function requires that the image be sparse in the wavelet domain, as measured by the ℓ_1 norm of the wavelet coefficients $\|Wx\|_1$, which acts as a surrogate for ℓ_0 [1], [2]. The total variation term enforces homogeneity within the image by encouraging neighboring pixels to have similar values while allowing for sudden high frequency jumps at edges. The parameter $\mu > 0$ controls the trade-off between the two terms. A variety of other image-level regularization approaches have been proposed along these lines, see [4], [5], [7].

2) *Patch-Level Sparse Regularization*: An alternative to the image-level sparsity constraint $h_g(x)$ is a patch-level, or local regularization function $h_l(x)$, which enforces that patches (square sub-regions of the image) have a sparse representation according to a dictionary. One possible general form of such a regularization function is,

$$h_l(x) = \sum_{i=1}^N \frac{\gamma}{2} \|R_i x - D\alpha_i\|_2^2 + f(\alpha_i, D), \quad (3)$$

where the dictionary matrix is $D \in \mathbb{C}^{P \times K}$ and α_i is a K -dimensional vector in \mathbb{R}^K . An important difference between $h_l(x)$ and $h_g(x)$ is the additional function $f(\alpha_i, D)$. While image-level sparsity constraints fall within a predefined transform domain, such as the wavelet basis, the sparse dictionary domain can be unknown for patch-level regularization and learned from data. The function f enforces sparsity by learning a D for which α_i is sparse.¹ For example, [9] uses K-SVD to learn D off-line, and then approximately optimize the objective function

$$\arg \min_{\alpha_{1:N}} \sum_{i=1}^N \|R_i x - D\alpha_i\|_2^2 \quad \text{subject to } \|\alpha_i\|_0 \leq T, \quad \forall i, \quad (4)$$

using orthogonal matching pursuits (OMP) [21]. In this case, the L_0 penalty on the additional parameters α_i make this a

¹The dependence of $h_l(x)$ on α and D is implied in our notation.

Algorithm 1 Dictionary Learning With BPFA

- 1) Construct a dictionary $D = [d_1, \dots, d_K]$, with

$$d_k \sim N(0, P^{-1}I_P), \quad k = 1, \dots, K.$$
- 2) Draw a probability $\pi_k \in [0, 1]$ for each d_k :

$$\pi_k \sim \text{Beta}(c\gamma/K, c(1 - \gamma/K)), \quad k = 1, \dots, K.$$
- 3) Draw precision values for noise and each weight

$$\gamma_\varepsilon \sim \text{Gam}(g_0, h_0), \quad \gamma_s \sim \text{Gam}(e_0, f_0).$$
- 4) For the i^{th} patch in x :
 - a) Draw the vector $s_i \sim N(0, \gamma_s^{-1}I_K)$.
 - b) Draw the binary vector z_i with $z_{ik} \sim \text{Bern}(\pi_k)$.
 - c) Define $\alpha_i = s_i \circ z_i$ by an element-wise product.
 - d) Sample noisy patch $R_i x \sim N(D\alpha_i, \gamma_\varepsilon^{-1}I_P)$.
- 5) Construct the image x as the average of all $R_i x$ that overlap on a given pixel.

non-convex problem. Using this definition of $h_l(x)$ in (1), a local optimal solution can be found by an alternating minimization procedure [32]: First solve the least squares solution for x using the current values of α_i and D , and then update α_i and D , or only α_i if D is learned off-line.

B. Dictionary Learning With Beta Process Factor Analysis

Typical dictionary learning approaches require a predefined dictionary size and, for each patch, the setting of either a sparsity level T , or an error threshold ϵ to determine how many dictionary elements are used. In both cases, if the settings do not agree with ground truth, the performance can significantly degrade. Instead, we consider a Bayesian nonparametric method called beta process factor analysis (BPFA) [23], which has been shown to successfully infer both of these values, as well as have competitive performance with algorithms in several application areas [23]–[26], and see [33]–[36] for related algorithms. The beta process is driven by an underlying Poisson process, and so its properties as a Bayesian nonparametric prior are well understood [27]. Originally used for survival analysis in the statistics literature, its use for latent factor modeling has been significantly increasing within the machine learning community [23]–[26], [28], [33]–[36].

1) *Generative Model*: We give the original hierarchical prior structure of the BPFA model in Algorithm 1, extending this to complex-valued dictionaries in Section III-A. With this approach, the model constructs a dictionary matrix $D \in \mathbb{R}^{P \times K}$ ($\mathbb{C}^{P \times K}$ below) of i.i.d. random variables, and assigns probability π_k to vector d_k . The parameters for these probabilities are set such that most of the π_k are expected to be small, with a few large. In Algorithm 1 we use an approximation to the beta process.² Under this parameterization, each patch $R_i x$ extracted from the image x is modeled as a sparse weighted combination of the dictionary elements, as determined by the element-wise product of $z_i \in \{0, 1\}^K$ with the Gaussian vector s_i . What makes the model nonparametric is that for

²For a finite $c > 0$ and $\gamma > 0$, the random measure $H = \sum_{k=1}^K \pi_k \delta_{d_k}$ converges weakly to a beta process as $K \rightarrow \infty$ [24], [27].

many values of k , the values of z_{ik} will equal zero for all i since π_k will be very small; the model learns the number of these unused dictionary elements and their index values from the data. Therefore, the value of K should be set to a large number that is more than the expected size of the dictionary. It can be shown that, under the assumptions of this prior, in the limit $K \rightarrow \infty$, the number of dictionary elements used by a patch is Poisson(γ) distributed and the total number of dictionary elements used by the data grows like $c\gamma \ln N$, where N is the number of patches [28]. The parameters of the model include $c, \gamma, e_0, f_0, g_0, h_0$ and K ; we discuss setting these values in Section IV.

2) *Relationship to K-SVD*: Another widely used dictionary learning method is K-SVD [20]. Though they are models for the same problem, BPFA and K-SVD have some significant differences that we briefly discuss. K-SVD learns the sparsity pattern of the coding vector α_i using the OMP algorithm [21] for each i . Holding the sparsity pattern fixed, it then updates each dictionary element and dimension of α jointly by a rank one approximation to the residual. Unlike BPFA, it learns as many dictionary elements as are given to it, so K should be set wisely. BPFA on the other hand automatically prunes unneeded elements, and updates the sparsity pattern by using the posterior distribution of a Bernoulli process, which is significantly different from OMP. It updates the weights and the dictionary from their Gaussian posteriors as well. Because of this probabilistic structure, we derive a sampling algorithm for these variables that takes advantage of marginalization, and naturally learns the auxiliary variables γ_ε and γ_s .

3) *Example Denoising Problem*: As we will see, the relationship of dictionary learning to CS-MRI is essentially as a denoising step. To this end, we briefly illustrate BPFA on a denoising problem. Denoising of an image using dictionary learning proceeds by first learning the dictionary representation of each patch, $R_i x \approx D\alpha_i$. The denoised reconstruction of x using BPFA is then $x_{\text{BPFA}} = \frac{1}{P} \sum_i R_i^T D\alpha_i$.

We show an example using 6×6 patches extracted from the noisy 512×512 image shown in Fig. 1(a). In Fig. 1(b) we show the resulting denoised image. For this problem we truncated the dictionary size to $K = 108$ and set all other model parameters to one. In Figs. 1(c) and 1(d) we show some statistics from dictionary learning. For example, Fig. 1(c) shows the values of π_k sorted, where we see that fewer than 100 elements are used by the data, many of which are very sparsely used. Fig. 1(d) shows the empirical distribution of the number of elements used per patch. We see the ability of the model to adapt the sparsity to the complexity of the patch. In Table I we show PSNR results for three noise variance levels. For K-SVD, we consider the case when the error parameter matches the ground truth, and when it mismatches it by a magnitude of five. As expected, when K-SVD does not have an appropriate parameter setting the performance suffers. BPFA on the other hand adaptively infers this value, which helps improve the denoising.

III. CS-MRI WITH BPFA AND TV PENALTY

We next present our approach for reconstructing single-channel MR images from highly undersampled k -space data.




Fig. 1. (a)-(b) An example of denoising by BPFA (image scaled to $[0,1]$). (c) Shows the final probabilities of the dictionary elements and (d) shows a distribution on the number of dictionary elements used per patch.

TABLE I

PEAK SIGNAL-TO-NOISE RATIO (PSNR) FOR IMAGE DENOISED BY BPFA. COMPARED WITH K-SVD USING CORRECT (MATCH) AND INCORRECT (MISMATCH) NOISE PARAMETER

σ	K-SVD denoising (PSNR)		BPFA denoising (PSNR)	
	Match	Mismatch	Results	Learned noise
20/255	32.28	28.94	32.88	20.43/255
25/255	31.08	28.60	31.81	25.46/255
30/255	29.99	28.35	30.94	30.47/255

In reference to the discussion in Section II, we consider a sparsity constraint of the form

$$\arg \min_{x, \varphi} \lambda_g h_g(x) + h_l(x) + \frac{\lambda}{2} \|\mathcal{F}_u x - y\|_2^2,$$

$$h_g(x) := TV(x), \quad h_l(x) := \sum_{i=1}^N \frac{\gamma_\varepsilon}{2} \|R_i x - D \alpha_i\|_2^2 + f(\varphi_i). \quad (5)$$

For the local regularization function $h_l(x)$ we use BPFA as given in Algorithm 1 in Section II-B. The parameters to be optimized for this penalty are contained in the set $\varphi_i = \{D, s_i, z_i, \gamma_\varepsilon, \gamma_s, \pi\}$, and are defined in Algorithm 1. We note that only s_i and z_i vary in i , while the rest are shared by all patches. The regularization term γ_ε is a model variable that corresponds to an inverse variance parameter of the multivariate Gaussian likelihood. This likelihood is equivalently viewed as the squared error penalty term in $h_l(x)$ in (5). This term acts as the sparse basis for the image and also aids in producing a denoised reconstruction, as discussed in Sections II-B, III-B and IV-B. For the global regularization function $h_g(x)$ we use the total variation of the image. This term encourages homogeneity within contiguous regions of the image, while still allowing for sharp jumps in pixel value at edges due to the underlying ℓ_1 penalty. The regularization parameters λ_g , γ_ε and λ control the trade-off between the terms

in this optimization. Since we sample a new value of γ_ε with each iteration of the algorithm discussed shortly, this trade-off is adaptively changing.

For the total variation penalty $TV(x)$ we use the isotropic TV model. Let ψ_i be the $2 \times N$ difference operator for pixel i . Each row of ψ_i contains a 1 centered on pixel i , and the first row also has a -1 on the pixel directly above pixel i , while the second has a -1 corresponding to the pixel to the right, and zeros elsewhere. Let $\Psi = [\psi_1^T, \dots, \psi_N^T]^T$ be the resulting $2N \times N$ difference matrix for the entire image. The TV coefficients are $\beta = \Psi x \in \mathbb{C}^{2N}$, and the isotropic TV penalty is $TV(x) = \sum_i \|\psi_i x\|_2 = \sum_i \sqrt{|\beta|_{2i-1}^2 + |\beta|_{2i}^2}$, where i ranges over the pixels in the MR image. For optimization we use the alternating direction method of multipliers (ADMM) [30], [31]. ADMM works by performing dual ascent on the augmented Lagrangian objective function introduced for the total variation coefficients. For completeness, we give a brief review of ADMM in the appendix.

A. Algorithm

We present an algorithm for finding a local optimal solution to the non-convex objective function given in (5). We can write this objective as

$$L(x, \varphi) = \lambda_g \sum_i \|\psi_i x\|_2 + \sum_i \frac{\gamma_\varepsilon}{2} \|R_i x - D \alpha_i\|_2^2 + \sum_i f(\varphi_i) + \frac{\lambda}{2} \|\mathcal{F}_u x - y\|_2^2. \quad (6)$$

We seek to minimize this function with respect to x and the dictionary learning variables $\varphi_i = \{D, s_i, z_i, \gamma_\varepsilon, \gamma_s, \pi\}$.

Our first step is to put the objective into a more suitable form. We begin by defining the TV coefficients for the i^{th} pixel as $\beta_i := [\beta_{2i-1} \ \beta_{2i}]^T = \psi_i x$. We introduce the vector of Lagrange multipliers η_i , and then split β_i from $\psi_i x$ by relaxing the equality via an augmented Lagrangian. This results in the objective function

$$L(x, \beta, \eta, \varphi) = \sum_{i=1}^N \lambda_g \|\beta_i\|_2 + \eta_i^T (\psi_i x - \beta_i) + \frac{\rho}{2} \|\psi_i x - \beta_i\|_2^2 + \sum_{i=1}^N \frac{\gamma_\varepsilon}{2} \|R_i x - D \alpha_i\|_2^2 + f(\varphi_i) + \frac{\lambda}{2} \|\mathcal{F}_u x - y\|_2^2. \quad (7)$$

From the ADMM theory [32], this objective will have (local) optimal values β_i^* and x^* with $\beta_i^* = \psi_i x^*$, and so the equality constraints will be satisfied [31].³ Optimizing this function can be split into three separate sub-problems: one for TV, one for BPFA and one for updating the reconstruction x . Following the discussion of ADMM in the appendix, we define $u_i = (1/\rho)\eta_i$ and complete the square in the first line of (7). We then cycle through the following three sub-problems,

$$(P1) \quad \beta'_i = \arg \min_{\beta} \lambda_g \|\beta\|_2 + \frac{\rho}{2} \|\psi_i x - \beta + u_i\|_2^2,$$

$$u'_i = u_i + \psi_i x - \beta'_i, \quad i = 1, \dots, N,$$

³For a fixed $D, \alpha_{1:N}$ and x the solution is also globally optimal.

Algorithm 2 Outline of Algorithm

Input: y – Undersampled k -space data
Output: x – Reconstructed MR image
Initialize: $x = \mathcal{F}_u^H y$ and each $u_i = 0$. Sample D from prior.
Step 1. $P1$: Optimize each β_i via shrinkage.
Step 2. Update Lagrange multiplier vectors u_i .
Step 3. $P2$: Gibbs sample BPPA variables once.
Step 4. $P3$: Solve for x using Fourier domain.
if not converged then return to Step 1.

$$(P2) \quad \phi' = \arg \min_{\phi} \sum_i \frac{\gamma_\varepsilon}{2} \|R_i x - D \alpha_i\|_2^2 + f(\phi_i),$$

$$(P3) \quad x' = \arg \min_x \sum_i \frac{\rho}{2} \|\psi_i x - \beta'_i + u'_i\|_2^2 + \sum_i \frac{\gamma'_\varepsilon}{2} \|R_i x - D' \alpha'_i\|_2^2 + \frac{\lambda}{2} \|\mathcal{F}_u x - y\|_2^2.$$

Solutions for sub-problems $P1$ and $P3$ are globally optimal (conditioned on the most recent values of all other parameters). We cannot solve $P2$ analytically since the optimal values for the set of all BPPA variables do not have a closed form solution. Our approach for $P2$ is to use stochastic optimization by Gibbs sampling each variable of BPPA conditioned on current values of all other variables. We next present the updates for each sub-problem. We give an outline in Algorithm 2.

1) *Algorithm for P1 (Total Variation)*: We can solve for β_i exactly for each pixel $i = 1, \dots, N$ by using a generalized shrinkage operation [31],

$$\beta'_i = \max \left\{ \|\psi_i x + u_i\|_2 - \frac{\lambda_g}{\rho}, 0 \right\} \cdot \frac{\psi_i x + u_i}{\|\psi_i x + u_i\|_2}. \quad (8)$$

We recall that β_i corresponds to the 2D TV coefficients for pixel i , with differences in one direction vertically and horizontally. We then update the corresponding Lagrange multiplier, $u'_i = u_i + \psi_i x - \beta'_i$.

2) *Algorithm for P2 (BPPA)*: We update the parameters of BPPA using Gibbs sampling. We are therefore stochastically optimizing (7), but only for this sub-problem. With reference to Algorithm 1, the $P2$ sub-problem entails sampling new values for the complex dictionary D , the binary vectors z_i and real-valued weights s_i (with which we construct $\alpha_i = s_i \circ z_i$ through the element-wise product), the precisions γ_ε and γ_s , and the probabilities $\pi_{1:K}$, with π_k giving the probability that $z_{ik} = 1$. In principle, there is no limit to the number of samples that can be made, with the final sample giving the updates used in the other sub-problems. We found that a single sample is sufficient in practice and leads to a faster algorithm. We describe the sampling procedure below.

a) *Sample dictionary D* : We define the $P \times N$ matrix $X = [R_1 x, \dots, R_N x]$, which is a complex matrix of all vectorized patches extracted from the image x . We also define the $K \times N$ matrix $\alpha = [\alpha_1, \dots, \alpha_N]$ containing the dictionary weight coefficients for the corresponding columns in X such that $D\alpha$ is an approximation of X to which we add noise from a circularly-symmetric complex normal distribution.

The update for the dictionary D is

$$D = X\alpha^T (\alpha\alpha^T + (P/\gamma_\varepsilon)I_P)^{-1} + E, \quad (9)$$

$$E_{p,:} \stackrel{\text{ind}}{\sim} \mathcal{CN}(0, (\gamma_\varepsilon \alpha\alpha^T + P I_P)^{-1}), \quad p = 1, \dots, P,$$

where $E_{p,:}$ is the p^{th} row of E . To sample this, we can first draw $E_{p,:}$ from a multivariate Gaussian distribution with this covariance structure, followed by an i.i.d. uniform rotation of each value in the complex plane. We note that the first term in Equation (9) is the ℓ_2 -regularized least squares solution for D . The addition of correlated Gaussian noise in the complex plane generates the sample from the conditional posterior of D . Since both the number of pixels and γ_ε will tend to be very large, the variance of the noise is small and the mean term dominates the update for D .

b) *Sample sparse coding α_i* : Sampling α_i entails sampling s_{ik} and z_{ik} for each k . We sample these values using block sampling. We recall that to block sample two variables from their joint conditional posterior distribution, $(s, z) \sim p(s, z| -)$, one can first sample z from the marginal distribution, $z \sim p(z| -)$, and then sample $s|z \sim p(s|z, -)$ from the conditional distribution. The other sampling direction is possible as well, but for our problem sampling $z \rightarrow s|z$ is more efficient for finding a mode of the objective function.

We define $r_{i,-k}$ to be the residual error in approximating the i^{th} patch with the current values from BPPA minus the k^{th} dictionary element, $r_{i,-k} = R_i x - \sum_{j \neq k} (s_{ij} z_{ij}) d_j$. We then sample z_{ik} from its conditional posterior Bernoulli distribution $z_{ik} \sim p_{ik} \delta_1 + q_{ik} \delta_0$, where following a simplification,

$$p_{ik} \propto \pi_k \left(1 + (\gamma_\varepsilon/\gamma_s) d_k^H d_k \right)^{-\frac{1}{2}} \quad (10)$$

$$\times \exp \left\{ \frac{\gamma_\varepsilon}{2} \text{Re}(d_k^H r_{i,-k})^2 / (\gamma_s/\gamma_\varepsilon + d_k^H d_k) \right\},$$

$$q_{ik} \propto 1 - \pi_k. \quad (11)$$

The symbol H denotes the conjugate transpose. The probabilities can be obtained by dividing both of these terms by their sum. We observe that the probability that $z_{ik} = 1$ takes into account how well dictionary element d_k correlates with the residual $r_{i,-k}$. After sampling z_{ik} we sample the corresponding weight s_{ik} from its conditional posterior Gaussian distribution,

$$s_{ik}|z_{ik} \sim N \left(z_{ik} \frac{\text{Re}(d_k^H r_{i,-k})}{\gamma_s/\gamma_\varepsilon + d_k^H d_k}, \frac{1}{\gamma_s + \gamma_\varepsilon z_{ik} d_k^H d_k} \right). \quad (12)$$

When $z_{ik} = 1$, the mean of s_{ik} is the regularized least squares solution and the variance will be small if γ_ε is large. When $z_{ik} = 0$, s_{ik} can be sampled from the prior, but does not factor in the model in this case.

c) *Sample γ_ε and γ_s* : We next sample from the conditional gamma posterior distribution of the noise precision and weight precision,

$$\gamma_\varepsilon \sim \text{Gam} \left(g_0 + \frac{1}{2} P N, h_0 + \frac{1}{2} \sum_i \|R_i x - D \alpha_i\|_2^2 \right), \quad (13)$$

$$\gamma_s \sim \text{Gam} \left(e_0 + \frac{1}{2} \sum_{i,k} z_{ik}, f_0 + \frac{1}{2} \sum_{i,k} z_{ik} s_{ik}^2 \right). \quad (14)$$

The expected value of each variable is the first term of the distribution divided by the second, which is close to the inverse of the average empirical error for γ_ε .

d) *Sample π_k* : Sample each π_k from its conditional beta posterior distribution,

$$\pi_k \sim \text{Beta} \left(a_0 + \sum_{i=1}^N z_{ik}, b_0 + \sum_{i=1}^N (1 - z_{ik}) \right). \quad (15)$$

The parameters to the beta distribution include counts of how many times dictionary element d_k was used by a patch.

3) *Algorithm for P3 (MRI Reconstruction)*: The final sub-problem is to reconstruct the image x . Our approach takes advantage of the Fourier domain similar to other methods, see [14], [30]. The corresponding objective function is

$$x' = \arg \min_x \sum_{i=1}^N \frac{\rho}{2} \|\psi_i x - \beta_i + u_i\|_2^2 + \sum_{i=1}^N \frac{\gamma_\varepsilon}{2} \|R_i x - D\alpha_i\|_2^2 + \frac{\lambda}{2} \|\mathcal{F}_u x - y\|_2^2.$$

Since this is a least squares problem, x has a closed form solution that satisfies

$$\begin{aligned} & \left(\rho \Psi^T \Psi + \gamma_\varepsilon \sum_i R_i^T R_i + \lambda \mathcal{F}_u^H \mathcal{F}_u \right) x \\ & = \rho \Psi^T (\beta - u) + \gamma_\varepsilon P x_{\text{BPFA}} + \lambda \mathcal{F}_u^H y. \end{aligned} \quad (16)$$

We recall that Ψ is the matrix of stacked ψ_i . The vector β is also obtained by stacking each β_i and u is the vector formed by stacking u_i . The vector x_{BPFA} is the denoised reconstruction from BPFA using the current D and $\alpha_{1:N}$, which results from the definition $x_{\text{BPFA}} = \frac{1}{P} \sum_i R_i^T D\alpha_i$.

We observe that inverting the left $N \times N$ matrix is computationally prohibitive since N is the number of pixels in the image. Fortunately, given the form of the matrix in Equation (16) we can use the procedure described in [14] and simplify the problem by working in the Fourier domain. This allows for element-wise updates in k -space, followed by an inverse Fourier transform. We represent x as $x = \mathcal{F}^H \theta$, where θ is the Fourier transform of x . We then take the Fourier transform of each side of Equation (16) to give

$$\begin{aligned} & \mathcal{F} \left(\rho \Psi^T \Psi + \gamma_\varepsilon \sum_i R_i^T R_i + \lambda \mathcal{F}_u^H \mathcal{F}_u \right) \mathcal{F}^H \theta \\ & = \rho \mathcal{F} \Psi^T (\beta - u) + \gamma_\varepsilon \mathcal{F} P x_{\text{BPFA}} + \lambda \mathcal{F} \mathcal{F}_u^H y. \end{aligned} \quad (17)$$

The left-hand matrix simplifies to a diagonal matrix,

$$\begin{aligned} & \mathcal{F} \left(\rho \Psi^T \Psi + \gamma_\varepsilon \sum_i R_i^T R_i + \lambda \mathcal{F}_u^H \mathcal{F}_u \right) \mathcal{F}^H \\ & = \rho \Lambda + \gamma_\varepsilon P I_N + \lambda I_N^u. \end{aligned} \quad (18)$$

Term-by-term this results as follows: The product of the finite difference operator matrix Ψ with itself yields a circulant matrix, which has the rows of the Fourier matrix \mathcal{F} as its eigenvectors and eigenvalues equal to $\Lambda = \mathcal{F} \Psi^T \Psi \mathcal{F}^H$. The matrix $R_i^T R_i$ is a matrix of all zeros, except for ones on the diagonal entries that correspond to the indices of x associated with the i^{th} patch. Since each pixel appears in P patches, the sum over i gives $P I_N$, and the Fourier product cancels. The final diagonal matrix I_N^u also contains all zeros, except for ones along the diagonal corresponding to the indices in k -space that are measured, which results from $\mathcal{F} \mathcal{F}_u^H \mathcal{F}_u \mathcal{F}^H$.




Fig. 2. The three masks considered for a given sampling percentage. (a) Random 25%. (b) Cartesian 30%. (c) Radial 25%.

Since the left matrix is diagonal we can perform element-wise updating of the Fourier coefficients θ ,

$$\theta_i = \frac{\rho \mathcal{F}_i \Psi^T (\beta - u) + \gamma_\varepsilon P \mathcal{F}_i x_{\text{BPFA}} + \lambda \mathcal{F}_i \mathcal{F}_u^H y}{\rho \Lambda_{ii} + \gamma_\varepsilon P + \lambda \mathcal{F}_i \mathcal{F}_u^H \mathbf{1}}. \quad (19)$$

We observe that the rightmost term in the numerator and denominator equals zero if i is not a measured k -space location. We invert θ via the inverse Fourier transform \mathcal{F}^H to obtain the reconstructed MR image x' .

B. Discussion on λ

In noise-free compressed sensing, the fidelity term λ can tend to infinity giving an equality constraint for the measured k -space values [1]. However, when y is noisy the setting of λ is critical for most CS-MRI algorithms since this parameter controls the level of denoising in the reconstructed image. We note that a feature of dictionary learning CS-MRI approaches is that λ can still be set to a very large value, and so parameter selection isn't necessary here. This is because a denoised version of the image is obtained through dictionary learning (x_{BPFA} in this paper) and can be taken as the denoised reconstruction. In Equation (19), we observe that by setting λ to a large value, we are effectively fixing the measured k -space values and using the k -space projection of BPFA and TV to fill in the missing values. The reconstruction x will be noisy, but have artifacts due to sub-sampling removed. The output image x_{BPFA} is a denoised version of x using BPFA in essentially the same manner as in Section II-B3. Therefore, the quality of our algorithm depends largely on the quality of BPFA as an image denoising algorithm [25]. We show examples of this using synthetic and clinical data in Sections IV-B and IV-E.

IV. EXPERIMENTAL RESULTS

We evaluate the proposed algorithm on real-valued and complex-valued MRI, and on a synthetic phantom. We consider three sampling masks: 2D random sampling, Cartesian sampling with random phase encodes (1D random), and pseudo radial sampling.⁴ We show an example of each mask in Fig. 2. We consider a variety of sampling rates for each mask. As a performance measure we use PSNR, and also consider SSIM [37]. We compare with three other algorithms: Sparse MRI [3],⁵ which as discussed above is a

⁴We used codes referenced in [3], [8], and [10] to generate these masks.

⁵<http://www.eecs.berkeley.edu/~mlustig/Software.html>

combination of wavelets and total variation, DLMRI [14],⁶ which is a dictionary learning method based on K-SVD, and PBDW [15],⁷ which is patch-based method that uses directional wavelets and therefore places greater restrictions on the dictionary. We use the publicly available code for these algorithms indicated above and used the built-in parameter settings, or those indicated in the relevant papers. We also compare with the BPFa algorithm without using total variation by setting $\lambda_g = 0$.

A. Set-Up

For all images, we extract 6×6 patches where each pixel defines the upper left corner of a patch and wrap around the image at the boundaries; we investigate different patch sizes later to show that this is a reasonable size. We initialize x by zero-filling in k -space. We use a dictionary with $K = 108$ initial dictionary elements, recalling that the final number of dictionary elements will be smaller due to the sparse BPFa prior. If 108 is found to be too small, K can be increased with the result being a slower inference algorithm.⁸ We ran 1000 iterations and use the results of the last iteration.

For regularization parameters, we set the data fidelity term $\lambda = 10^{100}$. We are therefore effectively requiring equality with the measured values of k -space and allowing BPFa to fill in the missing values, as well as give a denoised reconstruction, as discussed in Section III-B and highlighted below in Sections IV-B and IV-E. After trying several values, we also found $\lambda_g = 10$ and $\rho = 1000$ to give good results. We set the BPFa hyperparameters as $c = \gamma = e_0 = f_0 = g_0 = h_0 = 1$. These settings result in a relatively non-informative prior given the amount of data we have. However, we note that our algorithm was robust to these values, since the data overwhelms these prior values when calculating posterior distributions.

B. Experiments on a GE Phantom

We consider a noisy synthetic example to highlight the advantage of dictionary learning for CS-MRI. In Fig. 3 we show results on a 256×256 GE phantom with additive noise having standard deviation $\sigma = 0.1$. In this experiment we use BPFa without TV to reconstruct the original image using 30% Cartesian sampling. We show the reconstruction using zero-filling in Fig. 3(a). Since $\lambda = 10^{100}$, we see in Fig. 3(b) that BPFa essentially helps reconstruct the underlying noisy image for x . However, using the denoising property of the BPFa model shown in Fig. 1, we obtain the denoised reconstruction of Fig. 3(c) by focusing on x_{BPFa} from Equation (16). This is in contrast with the best result we could obtain with TV in Fig. 3(d), which places the TV penalty on the reconstructed image. As discussed, for TV the setting of λ relative to λ_g is important. We set $\lambda = 1$ and swept through $\lambda_g \in (0, 0.15)$, showing the result with highest PSNR in Fig. 3(d). Similar to Fig. 1 we show statistics from the BPFa model in Figs. 3(e)-(g). We see that




Fig. 3. GE data with noise ($\sigma = 0.1$) and 30% Cartesian sampling. BPFa (b) reconstructs the original noisy image, and (c) denoises the reconstruction simultaneously. (d) TV denoises as part of the reconstruction. Also shown are the dictionary learning variables sorted by π_k . (e) the dictionary, (f) the distribution on the dictionary, π_k . (g) The normalized histogram of number of the dictionary elements used per patch.

roughly 80 dictionary elements were used (the unused noisy elements in Fig. 3(e) are draws from the prior). We note that 2.28 elements were used on average by a patch given that at least one was used, which discounts the black regions.

C. Experiments on Real-Valued (Synthetic) MRI

For our synthetic MRI experiments, we consider two publicly available real-valued 512×512 MRI⁹ of a shoulder and lumbar. We construct these problems by applying the relevant sampling mask to the projection of real-valued MRI

⁶<http://www.ifp.illinois.edu/~yoram/DLMRI-Lab/Documentation.html>

⁷http://www.quxiaobo.org/index_publications.html

⁸As discussed in Section II-B, in theory K can be infinitely large.

⁹www3.americanradiology.com/pls/web1/wwimggal.vmg/wwimggal.vmg

TABLE II

PSNR RESULTS FOR REAL-VALUED LUMBAR MRI AS FUNCTION OF SAMPLING PERCENTAGE AND MASK (CARTESIAN WITH RANDOM PHASE ENCODES, 2D RANDOM AND PSEUDO RADIAL)

Mask	Samp%	BPFA+TV	BPFA	DLMRI	SparseMRI	PBDW
<i>Cart.</i>	10	32.48	32.03	31.02	30.24	31.74
	20	36.07	35.84	33.92	33.44	35.19
	25	38.78	38.53	36.56	35.50	37.43
	30	41.08	40.12	38.87	35.57	39.23
	35	41.05	40.96	38.85	37.66	39.24
<i>Rand.</i>	10	42.82	40.81	38.25	25.87	37.09
	20	44.35	41.80	40.11	27.80	37.86
	25	48.11	47.09	43.51	37.22	43.65
	30	49.36	48.55	44.93	38.72	45.50
	35	50.20	49.19	45.87	41.70	46.85
<i>Rad.</i>	10	35.16	33.33	32.91	29.35	31.46
	20	41.69	41.18	38.38	35.69	38.01
	25	43.75	43.40	40.29	38.59	40.25
	30	45.22	44.95	41.86	37.37	42.11
	35	46.85	46.45	43.09	39.74	43.72

TABLE III

PSNR RESULTS FOR REAL-VALUED SHOULDER MRI AS FUNCTION OF SAMPLING PERCENTAGE AND MASK (CARTESIAN WITH RANDOM PHASE ENCODES, 2D RANDOM AND PSEUDO RADIAL)

Mask	Samp%	BPFA+TV	BPFA	DLMRI	SparseMRI	PBDW
<i>Cart.</i>	10	32.65	30.79	31.02	27.65	28.88
	20	36.96	35.77	34.52	30.64	32.10
	25	38.45	37.97	35.69	32.44	34.12
	30	41.43	41.22	38.11	34.26	36.73
	35	41.33	41.14	38.44	34.50	36.76
<i>Rand.</i>	10	41.00	39.96	38.18	30.72	36.48
	20	43.53	42.40	39.38	32.08	39.39
	25	45.43	45.44	42.58	40.81	41.31
	30	46.89	46.86	44.03	43.47	43.12
	35	47.95	47.87	45.01	44.89	44.45
<i>Rad.</i>	10	34.30	33.88	33.27	29.18	31.60
	20	39.41	39.47	38.06	35.50	36.38
	25	41.40	41.52	39.73	38.70	38.21
	30	43.14	43.45	41.20	39.98	40.30
	35	44.69	44.99	42.58	39.11	41.72

into k -space. Though using such real-valued MRI data may not reflect clinical reality, we include this idealized setting to provide a complete set of experiments similar to other papers [3], [14], [15]. We evaluate the performance of our algorithm using PSNR and compare with Sparse MRI [3], DLMRI [14] and PBDW [15]. Although the original data is real-valued, we learn complex dictionaries since the reconstructions are complex. We consider our algorithm with and without the total variation penalty, denoted BPFA+TV and BPFA, respectively.

We present the PSNR results for all sampling masks and rates in Tables II and III. From these values we see the competitive performance of the propose dictionary learning algorithm. We also see a slight improvement by the addition of the TV penalty. As expected, we observe that 2D random sampling produced the best results, followed by pseudo-radial sampling and Cartesian sampling, which is due to their decreasing level of incoherence, with greater incoherence producing artifacts that are more noise-like [3]. Since BPFA is good at denoising




Fig. 4. Absolute errors for 30% Cartesian sampling of synthetic lumbar MRI. (a) BPFA+TV. (b) PBDW. (c) DLMRI. (d) Sparse MRI.




Fig. 5. Absolute errors for 20% radial sampling of the shoulder MRI. (a) BPFA+TV. (b) PBDW. (c) DLMRI. (d) Sparse MRI.

images, the algorithm naturally performs well in this setting. In Figs. 4 and 5 we show the absolute value of the residuals of different algorithms using one experiment from each MRI. We see an improvement using the proposed method, which has more noise-like errors.

D. Experiments on Complex-Valued MRI

We also consider two clinically obtained complex-valued MRI: We use the T2-weighted brain MRI from [4], which is a 256×256 MRI of a healthy volunteer from a 3T Siemens Trio Tim MRI scanner using the T2-weighted turbo spin echo

TABLE IV

PSNR/SSIM RESULTS FOR COMPLEX-VALUED BRAIN MRI AS A FUNCTION OF SAMPLING PERCENTAGE. SAMPLING MASKS INCLUDE CARTESIAN SAMPLING WITH RANDOM PHASE ENCODES, 2D RANDOM SAMPLING AND PSEUDO RADIAL SAMPLING

Mask	Sample %	BPFA+TV	BPFA	DLMRI	Sparse MRI	PBDW	Zero-filling
<i>Cartesian</i>	25	35.62 / 0.951	34.86 / 0.948	29.90 / 0.812	25.29 / 0.696	34.69 / 0.935	24.13 / 0.591
	30	38.64 / 0.968	37.70 / 0.965	31.54 / 0.849	26.16 / 0.745	37.24 / 0.957	24.55 / 0.614
	35	39.36 / 0.972	38.87 / 0.971	32.35 / 0.863	27.35 / 0.795	37.90 / 0.963	24.94 / 0.616
	40	41.09 / 0.977	40.45 / 0.976	33.60 / 0.876	29.82 / 0.845	39.23 / 0.969	26.28 / 0.667
<i>Random</i>	10	31.57 / 0.923	31.24 / 0.920	29.38 / 0.821	24.85 / 0.756	31.15 / 0.921	23.23 / 0.536
	15	36.49 / 0.963	35.44 / 0.961	30.16 / 0.774	22.68 / 0.651	34.22 / 0.942	21.18 / 0.493
	20	38.83 / 0.962	38.38 / 0.964	31.62 / 0.804	26.28 / 0.672	36.29 / 0.960	23.52 / 0.504
	25	40.75 / 0.979	40.00 / 0.973	32.83 / 0.862	31.16 / 0.934	37.62 / 0.968	26.58 / 0.582
	30	42.70 / 0.984	42.24 / 0.984	34.09 / 0.887	31.90 / 0.965	39.38 / 0.976	27.67 / 0.630
<i>Radial</i>	10	30.76 / 0.914	30.68 / 0.914	27.78 / 0.680	19.79 / 0.482	30.78 / 0.886	19.06 / 0.367
	15	34.00 / 0.949	33.79 / 0.950	29.49 / 0.734	22.07 / 0.640	33.99 / 0.937	20.87 / 0.498
	20	36.92 / 0.967	36.60 / 0.967	30.78 / 0.768	24.22 / 0.739	36.34 / 0.958	22.57 / 0.537
	25	39.72 / 0.977	39.37 / 0.977	31.91 / 0.794	26.64 / 0.797	38.38 / 0.970	24.34 / 0.574
	30	41.81 / 0.982	41.54 / 0.982	32.77 / 0.807	28.20 / 0.827	39.74 / 0.975	25.43 / 0.600




Fig. 6. Reconstruction results for 25% pseudo radial sampling of a complex-valued MRI of the brain. (a) Original. (b) BPFA+TV. (c) BPFA. (d) PBDW. (e) DLMRI. (f) PSNR vs iteration. (g) BPFA+TV error. (h) BPFA error. (i) PBDW error. (j) DLMRI error.

sequence (TR/TE = 6100/99 ms, 220×220 mm field of view, 3 mm slice thickness). We also use an MRI scan of a lemon obtained from the Research Center of Magnetic Resonance and Medical Imaging at Xiamen University (TE = 32 ms, size = 256×256 , spin echo sequence, TR/TE = 10000/32 ms, FOV = 70×70 mm², 2-mm slice thickness). This MRI is from a 7T/160mm bore Varian MRI system (Agilent Technologies, Santa Clara, CA, USA) using a quadrature-coil probe.

For the brain MRI experiment we use both PSNR and SSIM as performance measures. We show these values in Table IV for each algorithm, sampling mask and sampling rate. As with the synthetic MRI, we see that our algorithm performs competitively with the state-of-the-art. We also see the significant improvement of all algorithms over zero-filling. Example reconstructions are shown for each MRI dataset in Figs. 6 and 7. Also in Fig. 7 are PSNR values for the lemon MRI. We see from the absolute error residuals for

these experiments that the BPFA algorithm learns a slightly finer detail structure compared with other algorithms, with the errors being more noise-like. We also show the PSNR of BPFA+TV and BPFA as a function of iteration. As is evident, the algorithm does not necessarily need all 1000 iterations, but performs competitively even in half that number.

E. Experiments in the Noisy Setting

The MRI we have considered thus far have been essentially noiseless. For some MRI machines this may be an unrealistic assumption. We continue our evaluation of noisy MRI begun with the toy GE phantom in Section IV-B by evaluating how our model performs on clinically obtained MRI with additive noise. We show BPFA results without TV to highlight the dictionary learning features, but note that results with TV provide a slight improvement in terms of PSNR and




Fig. 7. Reconstruction results for 35% 2D random sampling of a complex-valued MRI of a lemon. (a) Original. (b) BPFA+TV: PSNR = 39.64. (c) BPFA: PSNR = 38.21. (d) PBDW: PSNR = 37.89. (e) DLMRI: PSNR = 35.05. (f) PSNR vs iteration. (g) BPFA+TV error. (h) BPFA error. (i) PBDW error. (j) DLMRI error.

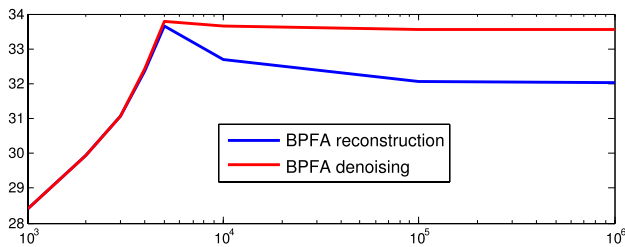


Fig. 8. PSNR vs λ in the noisy setting ($\sigma = 0.03$) for the complex-value brain MRI with 30% 2D random sampling.

SSIM. We again consider the brain MRI and use additive complex white Gaussian noise having standard deviation $\sigma = 0.01, 0.02, 0.03$. For all experiments we use the original noise-free MRI as the ground truth.

As discussed in Section III-B and illustrated in Section IV-B, dictionary learning allows us to consider two possible reconstructions: the actual reconstruction x , and the denoised BPFA reconstruction $x_{\text{BPFA}} = \frac{1}{P} \sum_i R_i^T D \alpha_i$. As detailed in these sections, as λ becomes larger the reconstruction will be noisier, but with the artifacts from sub-sampling removed. However, for all values of λ , x_{BPFA} produces a denoised version that essentially doesn't change. We see this clearly in Fig. 8, where we show the PSNR of each reconstruction as a function of λ . When λ is small, the performance degrades for both algorithms since too much smoothing is done by dictionary learning on x . As λ increases, both improve, but eventually the reconstruction of x degrades again because near equality to the noisy y is being more strictly enforced. The denoised reconstruction however levels off and does not degrade. We show PSNR values in Table V as a function of noise level.¹⁰

¹⁰We are working with a different scaling of the MRI than in [14] and made the appropriate modifications. Also, since DLMRI is a dictionary learning method it can output “ x_{KSVD} ”, though it was not originally motivated this way. Issues discussed in Sections II-B.2 and II-B.3 apply in this case.




Fig. 9. The denoising properties of dictionary learning on noisy complex-valued MRI with 35% Cartesian sampling and $\sigma = 0.03$. (a) Zero filling. (b) BPFA reconstruction (x). (c) BPFA denoising (x_{BPFA}). (d) DLMRI.

TABLE V
PSNR FOR 35% CARTESIAN SAMPLING OF COMPLEX-VALUED BRAIN MRI FOR VARIOUS NOISE STANDARD DEVIATIONS. ($\lambda = 10^{100}$)

Reconstruction method	$\sigma = 0$	$\sigma = 0.01$	$\sigma = 0.02$	$\sigma = 0.03$
BPFA–reconstruction	38.87	37.25	33.77	31.08
BPFA–denoising	37.99	37.19	34.43	32.39
DLMRI	32.35	32.12	31.61	30.65

Example reconstructions that parallel those given in Fig. 3 are also shown in Fig. 9. These results highlight the robustness of our approach to λ in the noisy setting, and we note that we

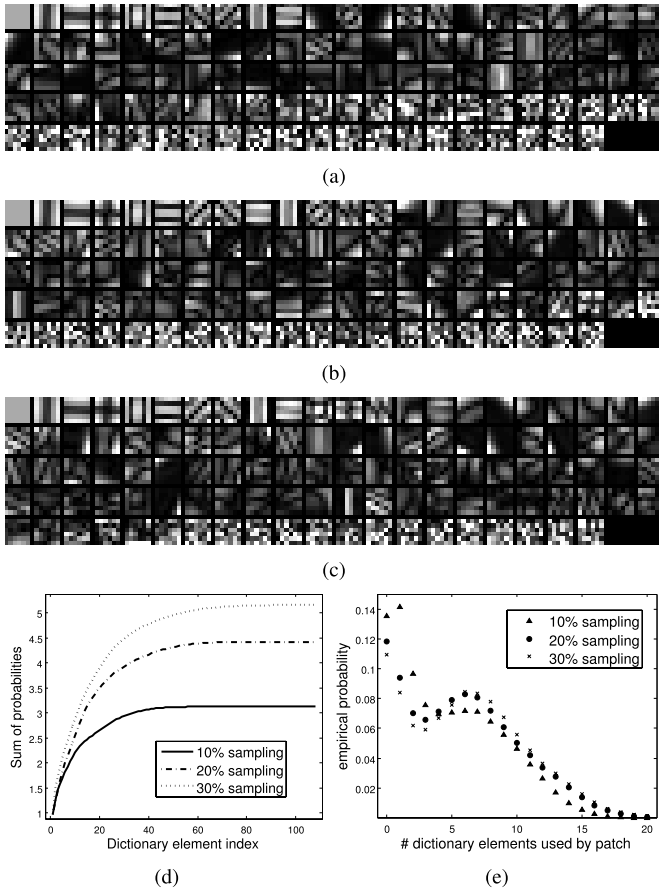


Fig. 10. Radial sampling for the Brain MRI. (a)-(c) The learned dictionary for various sampling rates. The noisy elements towards the end of each were unused and are samples from the prior. (d) The cumulative function of the sorted π_k from BPFA for each sampling rate. This gives information on sparsity and average usage of the dictionary. (e) The distribution on the number of elements used per patch for each sampling rate.

encountered no stability issues using extremely large values of λ .

F. Dictionary Learning and Further Discussion

We investigate the model learned by BPFA. In Fig. 10 we show dictionary learning results learned by BPFA+TV for radial sampling of the complex Brain MRI. In the top portion, we show the dictionaries learned for 10%, 20% and 30% sampling. We see that they are similar in their shape, but the number of elements increases as the sampling percentage increases since more complex information about the image is contained in the k -space measurements. We again note that unused elements are represented by draws from the prior. In Fig. 10(d) we show the cumulative sum of the ordered π_k from BPFA. We can read off the average number of elements used per patch by looking at the right-most value. We see that more elements are used per patch as the fraction of observed k -space increases. We also see that for 10%, 20% and 30% sampling, roughly 70, 80 and 95, respectively, of the 108 total dictionary elements were significantly used, as indicated by the leveling off of these functions. This highlights the adaptive property of the nonparametric beta process prior. In Fig. 10(e) we show the empirical distribution on the number of dictionary elements used per patch for each sampling rate.

TABLE VI
PSNR AS A FUNCTION OF PATCH SIZE FOR A REAL-VALUED AND COMPLEX-VALUED BRAIN MRI WITH CARTESIAN SAMPLING

	4×4	5×5	6×6	7×7	8×8
Synthetic brain 25%	37.86	38.29	38.33	38.26	38.24
Complex brain 40%	40.53	40.84	41.09	41.11	41.15

TABLE VII
TOTAL RUNTIME IN MINUTES (SECONDS/ITERATION). WE RAN 1000 ITERATIONS OF BPFA, 100 OF DLMRI AND 10 OF SPARSE MRI

Sampling %	BPFA+TV	BPFA	DLMRI	Sparse MRI
10%	52.4 (3.15)	50.5 (3.03)	27.6 (16.5)	1.63 (9.78)
20%	51.3 (3.08)	49.5 (2.97)	38.3 (23.0)	1.59 (9.54)
30%	51.2 (3.07)	48.3 (2.90)	45.7 (27.4)	1.60 (9.60)

We see that there are two modes, one for the empty background and one for the foreground, and the second mode tends to increase as the sampling rate increases. The adaptability of this value to each patch is another characteristic of the beta process model.

We also performed an experiment with varying patch sizes and show our results in Table VI. We see that the results are not very sensitive to this setting and that comparisons using 6×6 patches are meaningful. We also compare the runtime for different algorithms in Table VII, showing both the total runtime of each algorithm and the per-iteration times using an Intel Xeon CPU E5-1620 at 3.60GHz, 16.0G ram. However, we note that we arguably ran more iterations than necessary for these algorithms; the BPFA algorithms generally produced high quality results in half the number of iterations, as did DLMRI (the authors of [14] recommend 20 iterations), while Sparse MRI uses 5 iterations as default and the performance didn't improve beyond 10 iterations. We note that the speed-up over DLMRI arises from the lack of the OMP algorithm, which in Matlab is much slower than our sparse coding update.¹¹ We note that inference for the BPFA model is easily parallelizable—as are the other dictionary learning algorithms—which can speed up processing time.

The proposed method has several advantages, which we believe leads to the improvement in performance. A significant advantage is the adaptive learning of the dictionary size and per-patch sparsity level using a nonparametric stochastic process that is naturally suited for this problem. Several other dictionary learning parameters such as the noise variance and the variances of the score weights are adjusted as well through a natural MCMC sampling approach. These benefits have been investigated in other applications of this model [25], and naturally translate here since CS-MRI with BPFA is closely related to image denoising as we have shown.

Another advantage of our model is the Markov Chain Monte Carlo inference algorithm itself. In highly non-convex Bayesian models (or similar models with a Bayesian interpretation), it is often observed by the statistics community that

¹¹BPFA is significantly faster than K-SVD in Matlab because it requires fewer loops. This difference may not be as large with other coding languages.

MCMC sampling can outperform deterministic methods, and rarely performs worse [38]. Given that BPFA is a Bayesian model, such sampling techniques are readily derived, as we showed in Section III-A.

V. CONCLUSION

We have presented an algorithm for CS-MRI reconstruction that uses Bayesian nonparametric dictionary learning. Our Bayesian approach uses a model called beta process factor analysis (BPFA) for *in situ* dictionary learning. Through this hierarchical generative structure, we can learn the dictionary size, sparsity pattern and additional regularization parameters. We also considered a total variation penalty term for additional constraints on image smoothness. We presented an optimization algorithm using the alternating direction method of multipliers (ADMM) and MCMC Gibbs sampling for all BPFA variables. Experimental results on real and complex-valued MRI showed that our proposed regularization framework compares favorably with other algorithms for various sampling trajectories and rates. We also showed the natural ability of dictionary learning to handle noisy MRI *without* dependence on the measurement fidelity parameter λ . To this end, we showed that the model can enforce a near equality constraint to the noisy measurements and use the dictionary learning result as a denoised output of the noisy MRI.

APPENDIX

We give a brief review of the ADMM algorithm [32]. We start with the convex optimization problem

$$\min_x \|Ax - b\|_2^2 + h(x), \quad (20)$$

where h is a non-smooth convex function, such as an ℓ_1 penalty. ADMM decouples the smooth squared error term from this penalty by introducing a second vector v such that

$$\min_x \|Ax - b\|_2^2 + h(v) \quad \text{subject to} \quad v = x. \quad (21)$$

This is followed by a relaxation of the equality $v = x$ via an augmented Lagrangian term

$$L(x, v, \eta) = \|Ax - b\|_2^2 + h(v) + \eta^T(x - v) + \frac{\rho}{2}\|x - v\|_2^2. \quad (22)$$

A minimax saddle point is found with the minimization taking place over both x and v and dual ascent for η .

Another way to write the objective in (22) is to define $u = (1/\rho)\eta$ and combine the last two terms. The result is an objective that can be optimized by cycling through the following updates for x , v and u ,

$$x' = \arg \min_x \|Ax - b\|_2^2 + \frac{\rho}{2}\|x - v + u\|_2^2, \quad (23)$$

$$v' = \arg \min_v h(v) + \frac{\rho}{2}\|x' - v + u\|_2^2, \quad (24)$$

$$u' = u + x' - v'. \quad (25)$$

This algorithm simplifies the optimization since the objective for x is quadratic and thus has a simple analytic solution, while the update for v is a proximity operator of h with penalty ρ , the difference being that v is not pre-multiplied

by a matrix as x is in (20). Such objective functions tend to be easier to optimize. For example when h is the TV penalty the solution for v is analytical.


ACKNOWLEDGMENT

Y. Huang and J. Paisley contributed equally to this work.

REFERENCES

- [1] E. J. Candés, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [2] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [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] X. Qu, W. Zhang, D. Guo, C. Cai, S. Cai, and Z. Chen, "Iterative thresholding compressed sensing MRI based on contourlet transform," *Inverse Problems Sci. Eng.*, vol. 18, no. 6, pp. 737–758, Jun. 2009.
- [5] X. Qu, X. Cao, D. Guo, C. Hu, and Z. Chen, "Combined sparsifying transforms for compressed sensing MRI," *Electron. Lett.*, vol. 46, no. 2, pp. 121–123, Jan. 2010.
- [6] J. Trzasko and A. Manduca, "Highly undersampled magnetic resonance image reconstruction via homotopic ℓ_0 -minimization," *IEEE Trans. Med. Imag.*, vol. 28, no. 1, pp. 106–121, Jan. 2009.
- [7] H. Jung, K. Sung, K. S. Nayak, E. Y. Kim, and J. C. Ye, "k-t FOCUSS: A general compressed sensing framework for high resolution dynamic MRI," *Magn. Reson. Med.*, vol. 61, no. 1, pp. 103–116, 2009.
- [8] J. Yang, Y. Zhang, and W. Yin, "A fast alternating direction method for TVL1-L2 signal reconstruction from partial Fourier data," *IEEE J. Sel. Topics Signal Process.*, vol. 4, no. 2, pp. 288–297, Apr. 2010.
- [9] Y. Chen, X. Ye, and F. Huang, "A novel method and fast algorithm for MR image reconstruction with significantly under-sampled data," *Inverse Problems Imag.*, vol. 4, no. 2, pp. 223–240, 2010.
- [10] J. Huang, S. Zhang, and D. Metaxas, "Efficient MR image reconstruction for compressed MR imaging," *Med. Image Anal.*, vol. 15, no. 5, pp. 670–679, 2011.
- [11] S. Ji, Y. Xue, and L. Carin, "Bayesian compressive sensing," *IEEE Trans. Signal Process.*, vol. 56, no. 6, pp. 2346–2356, Jun. 2008.
- [12] X. Ye, Y. Chen, W. Lin, and F. Huang, "Fast MR image reconstruction for partially parallel imaging with arbitrary k-space trajectories," *IEEE Trans. Med. Imag.*, vol. 30, no. 3, pp. 575–585, Mar. 2011.
- [13] M. Akçakaya *et al.*, "Low-dimensional-structure self-learning and thresholding: Regularization beyond compressed sensing for MRI reconstruction," *Magn. Reson. Med.*, vol. 66, no. 3, pp. 756–767, 2011.
- [14] 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 2011.
- [15] X. Qu *et al.*, "Undersampled MRI reconstruction with patch-based directional wavelets," *Magn. Reson. Imag.*, vol. 30, no. 7, pp. 964–977, 2012.
- [16] Z. Yang and M. Jacob, "Robust non-local regularization framework for motion compensated dynamic imaging without explicit motion estimation," in *Proc. 9th IEEE Int. Symp. Biomed. Imag.*, May 2012, pp. 1056–1059.
- [17] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, "Non-local sparse models for image restoration," in *Proc. 12th Int. Conf. Comput. Vis.*, Sep./Oct. 2009, pp. 2272–2279.
- [18] K. Dabov, A. Foi, V. Katkovich, 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.
- [19] K. Engan, S. O. Aase, and J. H. Husøy, "Method of optimal directions for frame design," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Mar. 1999, pp. 2443–2446.
- [20] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4311–4322, Nov. 2006.
- [21] Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proc. 27th Asilomar Conf. Signals, Syst. Comput.*, Nov. 1993, pp. 40–44.
- [22] M. Protter and M. Elad, "Image sequence denoising via sparse and redundant representations," *IEEE Trans. Image Process.*, vol. 18, no. 1, pp. 27–36, Jan. 2009.

- [23] J. Paisley and L. Carin, "Nonparametric factor analysis with beta process priors," in *Proc. 26th Annu. Int. Conf. Mach. Learn.*, 2009, pp. 777–784.
- [24] J. W. Paisley, D. M. Blei, and M. I. Jordan, "Stick-breaking beta processes and the Poisson process," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2012, pp. 850–858.
- [25] M. Zhou *et al.*, "Nonparametric Bayesian dictionary learning for analysis of noisy and incomplete images," *IEEE Trans. Image Process.*, vol. 21, no. 1, pp. 130–144, Jun. 2012.
- [26] X. Ding, L. He, and L. Carin, "Bayesian robust principal component analysis," *IEEE Trans. Image Process.*, vol. 20, no. 12, pp. 3419–3430, Dec. 2011.
- [27] N. L. Hjort, "Nonparametric Bayes estimators based on beta processes in models for life history data," *Ann. Statist.*, vol. 18, no. 3, pp. 1259–1294, 1990.
- [28] R. Thibaux and M. I. Jordan, "Hierarchical beta processes and the Indian buffet process," in *Proc. Int. Conf. Artif. Intell. Statist.*, San Juan, PR, USA, 2007, pp. 564–571.
- [29] D. Gabay and B. Mercier, "A dual algorithm for the solution of nonlinear variational problems via finite element approximation," *Comput. Math. Appl.*, vol. 2, no. 1, pp. 17–40, 1976.
- [30] W. Yin, S. Osher, D. Goldfarb, and J. Darbon, "Bregman iterative algorithms for ℓ_1 -minimization with applications to compressed sensing," *SIAM J. Imag. Sci.*, vol. 1, no. 1, pp. 143–168, 2008.
- [31] T. Goldstein and S. Osher, "The split Bregman method for L1-regularized problems," *SIAM J. Imag. Sci.*, vol. 2, no. 2, pp. 323–343, 2009.
- [32] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2010.
- [33] J. Paisley, L. Carin, and D. Blei, "Variational inference for stick-breaking beta process priors," in *Proc. 28th Int. Conf. Mach. Learn.*, Bellevue, WA, USA, pp. 889–896, 2011.
- [34] D. Knowles and Z. Ghahramani, "Infinite sparse factor analysis and infinite independent components analysis," in *Independent Component Analysis and Signal Separation*. New York, NY, USA: Springer-Verlag, 2007, pp. 381–388.
- [35] E. B. Fox, E. B. Sudderth, M. I. Jordan, and A. S. Willsky, "Sharing features among dynamical systems with beta processes," in *Advances in Neural Information Processing*. Vancouver, BC, Canada: Curran Associates, Inc., 2011.
- [36] T. Griffiths and Z. Ghahramani, "Infinite latent feature models and the Indian buffet process," in *Advances in Neural Information Processing Systems*. Cambridge, MA, USA: MIT Press, 2006, pp. 475–482.
- [37] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [38] C. Faes, J. T. Ormerod, and M. P. Wand, "Variational Bayesian inference for parametric and nonparametric regression with missing data," *J. Amer. Statist. Assoc.*, vol. 106, no. 495, pp. 959–971, 2011.




Yue Huang received the B.S. degree in electrical engineering from Xiamen University, Xiamen, China, in 2005, and the Ph.D. degree in biomedical engineering from Tsinghua University, Beijing, China, in 2010. Since 2010, she has been an Assistant Professor with the School of Information Science and Engineering, Xiamen University. Her main research interests include image processing, machine learning, and biomedical engineering.



John Paisley is currently an Assistant Professor with the Department of Electrical Engineering, Columbia University, New York, NY, USA. Prior to that, he was a Post-Doctoral Researcher with the Departments of Computer Science, University of California at Berkeley, Berkeley, CA, USA, and Princeton University, Princeton, NJ, USA. He received the B.S., M.S., and Ph.D. degrees in electrical and computer engineering from Duke University, Durham, NC, USA, in 2004, 2007, and 2010, respectively. His research is in the area of statistical machine learning

and focuses on probabilistic modeling and inference techniques, Bayesian nonparametric methods, and text and image processing.



Qin Lin is currently pursuing the degree with the Department of Communication Engineering, Xiamen University, Xiamen, China. His research interests include computer vision, machine learning, and data mining.



Xinghao Ding was born in Hefei, China, in 1977. He received the B.S. and Ph.D. degrees from the Department of Precision Instruments, Hefei University of Technology, Hefei, China, in 1998 and 2003, respectively.

He was a Post-Doctoral Researcher with the Department of Electrical and Computer Engineering, Duke University, Durham, NC, USA, from 2009 to 2011. Since 2011, he has been a Professor with the School of Information Science and Engineering, Xiamen University, Xiamen, China. His main

research interests include image processing, sparse signal representation, and machine learning.



Xueyang Fu is currently pursuing the degree with the Department of Communication Engineering, Xiamen University, Xiamen, China. His research interests include image processing, sparse representation, and machine learning.



Xiao-Ping Zhang (M'97–SM'02) received the B.S. and Ph.D. degrees in electronic engineering from Tsinghua University, Beijing, China, in 1992 and 1996, respectively, and the M.B.A. (Hons.) degree in finance, economics, and entrepreneurship from the Booth School of Business, University of Chicago, Chicago, IL, USA.

He has been with the Department of Electrical and Computer Engineering, Ryerson University, Toronto, ON, Canada, since Fall 2000, where he is currently a Professor and the Director of the Communication

and Signal Processing Applications Laboratory. He has served as the Program Director of Graduate Studies. He is cross-appointed with the Department of Finance, Ted Rogers School of Management, Ryerson University. Prior to joining Ryerson University, he was a Senior DSP Engineer with SAM Technology, Inc., San Francisco, CA, USA, and a Consultant with San Francisco Brain Research Institute, San Francisco. He held research and teaching positions at the Communication Research Laboratory, McMaster University, Hamilton, ON, USA, and was a Post-Doctoral Fellow with the Beckman Institute, University of Illinois at Urbana-Champaign, Champaign, IL, USA, and the University of Texas at San Antonio, San Antonio, TX, USA. His research interests include statistical signal processing, multimedia content analysis, sensor networks and electronic systems, computational intelligence, and applications in bioinformatics, finance, and marketing. He is a frequent Consultant for biotech companies and investment firms. He is the Co-Founder and CEO of EidoSearch, an Ontario-based company offering a content-based search and analysis engine for financial data.

Dr. Zhang is a registered Professional Engineer in Ontario, and a member of the Beta Gamma Sigma Honor Society. He is the General Chair of MMSP'15, Publicity Chair of ICME'06, and Program Chair of ICIC'05 and ICIC'10. He served as a Guest Editor of *Multimedia Tools and Applications* and the *International Journal of Semantic Computing*. He is a Tutorial Speaker of ACMMM2011, ISCAS2013, ICIP2013, and ICASSP2014. He is currently an Associate Editor of the *IEEE TRANSACTIONS ON SIGNAL PROCESSING*, the *IEEE TRANSACTIONS ON MULTIMEDIA*, the *IEEE SIGNAL PROCESSING LETTERS*, and the *Journal of Multimedia*.