# Nonparametric Bayesian Dictionary Learning for Analysis of Noisy and Incomplete Images 

Mingyuan Zhou,<br>Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708-0291 USA<br>Haojun Chen, Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708-0291 USA<br>John Paisley,<br>Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708-0291 USA<br>Lu Ren,<br>Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708-0291 USA

Lingbo Li,
Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708-0291 USA

Zhengming Xing,
Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708-0291 USA

David Dunson,
Department of Statistics, Duke University, Durham, NC 27708-0291 USA
Guillermo Sapiro, and
Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455 USA

Lawrence Carin [Senior Member, IEEE]
Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708-0291 USA


#### Abstract

Nonparametric Bayesian methods are considered for recovery of imagery based upon compressive, incomplete, and/or noisy measurements. A truncated beta-Bernoulli process is employed to infer an appropriate dictionary for the data under test and also for image recovery. In the context of compressive sensing, significant improvements in image recovery are manifested using learned dictionaries, relative to using standard orthonormal image expansions. The compressive-measurement projections are also optimized for the learned dictionary. Additionally, we consider simpler (incomplete) measurements, defined by measuring a subset of image pixels, uniformly selected at random. Spatial interrelationships within imagery are exploited through use


[^0]of the Dirichlet and probit stick-breaking processes. Several example results are presented, with comparisons to other methods in the literature.

## Index Terms

Bayesian nonparametrics; compressive sensing; dictionary learning; factor analysis; image denoising; image interpolation; sparse coding

## I. Introduction

## A. Sparseness and Dictionary Learning

Recently, there has been significant interest in sparse image representations, in the context of denoising and interpolation [1], [13], [24]-[26], [27], [28], [31], compressive sensing (CS) [5], [12], and classification [40]. All of these applications exploit the fact that images may be sparsely represented in an appropriate dictionary. Most of the denoising, interpolation, and CS literature assume "off-the-shelf" wavelet and DCT bases/dictionaries [21], but recent research has demonstrated the significant utility of learning an often overcomplete dictionary matched to the signals of interest (e.g., images) [1], [4], [12], [13], [24]-[26], [27], [28], [30], [31], [41].

Many of the existing methods for learning dictionaries are based on solving an optimization problem [1], [13], [24]-[26], [27], [28], in which one seeks to match the dictionary to the imagery of interest, while simultaneously encouraging a sparse representation. These methods have demonstrated state-of-the-art performance for denoising, superresolution, interpolation, and inpainting. However, many existing algorithms for implementing such ideas also have some restrictions. For example, one must often assume access to the noise/ residual variance, the size of the dictionary is set a priori or fixed via cross-validation-type techniques, and a single ("point") estimate is learned.

To mitigate the aforementioned limitations, dictionary learning has recently been cast as a factor-analysis problem, with the factor loading corresponding to the dictionary elements (atoms). Utilizing nonparametric Bayesian methods like the beta process (BP) [29], [37], [42] and the Indian buffet process (IBP) [18], [22], one may, for example, infer the number of factors (dictionary elements) needed to fit the data. Furthermore, one may place a prior on the noise or residual variance, with this inferred from the data [42]. An approximation to the full posterior may be manifested via Gibbs sampling, yielding an ensemble of dictionary representations. Recent research has demonstrated that an ensemble of representations can be better than a single expansion [14], with such an ensemble naturally manifested by statistical models of the type described here.

## B. Exploiting Structure and Compressive Measurements

In image analysis, there is often additional information that may be exploited when learning dictionaries, with this well suited for Bayesian priors. For example, most natural images may be segmented, and it is probable that dictionary usage will be similar for regions within a particular segment class. To address this idea, we extend the model by employing a probit stick-breaking process (PSBP), with this generalization of the Dirichlet process (DP) stickbreaking representation [34]. Related clustering techniques have proven successful in image processing [26]. The model clusters the image patches, with each cluster corresponding to a segment type; the PSBP encourages proximate and similar patches to be included within the same segment type, thereby performing image segmentation and dictionary learning simultaneously.

The principal focus of this paper is on applying hierarchical Bayesian algorithms to new compressive measurement techniques that have been recently developed. Specifically, we consider dictionary learning in the context of CS [5], [9], in which the measurements correspond to projections of typical image pixels. We consider dictionary learning performed "offline" based on representative (training) images, with the learned dictionary applied within CS image recovery. We also consider the case for which the underlying dictionary is simultaneously learned with inversion (reconstruction), with this related to "blind" CS [17]. Finally, we design the CS projection matrix to be matched to the learned dictionary (when this is done offline), and demonstrate, as in [12], that in practice, this yields performance gains relative to conventional random CS projection matrices.

While CS is of interest for its potential to reduce the number of required measurements, it has the disadvantage of requiring the development of new classes of cameras. Such cameras are revolutionary and interesting [11], [35], but there have been decades of previous research performed on development of pixel-based cameras, and it would be desirable if such cameras could be modified simply to perform compressive measurements. We demonstrate that one may perform compressive measurements of natural images by simply sampling pixels measured by a conventional camera, with the pixels uniformly selected at random. This is closely related to recent research on matrix completion [6], [23], [33], but here, we move beyond simple low-rank assumptions associated with such previous research.

## C. Contributions

This paper develops several hierarchical Bayesian models for learning dictionaries for analysis of imagery, with applications in denoising, interpolation, and CS. The inference is performed based on a Gibbs sampler, as is increasingly common in modern image processing [16]. Here, we demonstrate how generalizations of the beta-Bernoulli process allow one to infer the dictionary elements directly based on the underlying degraded image, without any a priori training data, while simultaneously inferring the noise statics and tolerating significant missingness in the imagery. This is achieved by exploiting the lowdimensional structure of most natural images, which implies that image patches may be represented in terms of a low-dimensional set of learned dictionary elements. Excellent results are realized in these applications, including for hyperspectral imagery, which has not been widely considered in these settings previously. We also demonstrate how the learned dictionary may be easily employed to define CS projection matrices that yield markedly improved image recovery, as compared with the much more typical random construction of such matrices.

The basic hierarchical Bayesian architecture developed here serves as a foundation that may be flexibly extended to incorporate additional prior information. As examples, we show here how dictionary learning may be readily coupled with clustering, through the use of a DP [15]. We also incorporate spatial information within the image via a PSBP [32], with these extensions yielding significant advantages for the CS application. The basic modeling architecture may also exploit additional information manifested in the form of general covariates, as considered in a recent paper [43].

The remainder of the paper is organized as follows: In Section II, we review the classes of problems being considered. The beta-Bernoulli process is discussed in Section III, with relationships made with previous work in this area, including those based on the IBP. The Dirichlet and PSBPs are discussed in Section IV, and several example results are presented in Section V. Conclusions and a discussion of future work are provided in Section VI, and details of the inference equations are summarized in the Appendix.

## II. Problems Under Study

We consider data samples that may be expressed in the following form:

$$
\begin{equation*}
\boldsymbol{x}_{i}=\mathbf{D} w_{i}+\varepsilon_{i} \tag{1}
\end{equation*}
$$

where $\boldsymbol{x}_{i} \in \mathbb{R}^{P}, \boldsymbol{\varepsilon}_{i} \in \mathbb{R}^{P}$, and $\boldsymbol{w}_{i} \in \mathbb{R}^{K}$. The columns of matrix $\mathbf{D} \in \mathbb{R}^{P \times K}$ represent the $K$ components of a dictionary with which $\boldsymbol{x}_{i}$ is expanded. For our problem, the $\boldsymbol{x}_{i}$ will correspond to $B \times B$ (overlapped) pixel patches in an image [1], [13], [24], [25], [27], [42]. The set of vectors $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$ may be extracted from an image(s) of interest.

For the denoising problem, vectors $\boldsymbol{\varepsilon}_{i}$ may represent sensor noise, in addition to (ideally small) residual from representation of the underlying signal as $\mathbf{D} \boldsymbol{w}_{i}$. To perform denoising, we place restrictions on vectors $\boldsymbol{w}_{i}$, such that $\mathbf{D} \boldsymbol{w}_{i}$ by itself does not exactly represent $\boldsymbol{x}_{\boldsymbol{i}}$. A popular such restriction is that $\boldsymbol{w}_{i}$ should be sparse, motivated by the idea that any particular $\boldsymbol{x}_{i}$ may often be represented in terms of a small subset of representative dictionary elements, from the full dictionary defined by the columns of $\mathbf{D}$. There are several methods that have been recently developed to impose such a sparse representation, including $\ell_{1}$-based relaxation algorithms [24], [25], iterative algorithms [1], [13], and Bayesian methods [42]. One advantage of a Bayesian approach is that the noise/residual statistics may be nonstationary (with unknown noise statistics). Specifically, in addition to placing a sparseness-promoting prior on $\boldsymbol{w}_{i}$, we may also impose a prior on the components of $\boldsymbol{\varepsilon}_{i}$. From the estimated posterior density function on model parameters, each component of $\boldsymbol{\varepsilon}_{i}$, corresponding to the $i$ th $B \times B$ image patch, has its own variance. Given $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$, our goal may be to simultaneously infer $\mathbf{D}$ and $\left\{\boldsymbol{w}_{i}\right\}_{i=1, N}$ (and implicitly $\boldsymbol{\varepsilon}_{i}$ ), and then, the denoised version of $\boldsymbol{x}_{i}$ is represented as $\mathbf{D} \boldsymbol{w}_{i}$.

In many applications, the total number of pixels $N \cdot P$ may be large. However, it is well known that compression algorithms may be used on $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$ after the measurements have been performed, to significantly reduce the quantity of data that need be stored or communicated. This compression indicates that while the data dimensionality $N \cdot P$ may be large, the underlying information content may be relatively low. This has motivated the field of CS [5], [9], [11], [35], in which the total number of measurements performed may be much less than $N \cdot P$. Toward this end, researchers have proposed projection measurements of the following form:

$$
\begin{equation*}
y_{i}=\sum x_{i} \tag{2}
\end{equation*}
$$

where $\Sigma \in \mathbb{R}^{n \times P}$ and $\boldsymbol{y}_{i} \in \mathbb{R}^{n}$, ideally with $n \ll P$. The projection matrix $\Sigma$ has traditionally been randomly constituted [5], [9], with a binary or real alphabet (and $\Sigma$ may be also a function of the specific patch, and generalized as $\Sigma_{i}$ ). It is desirable that matrices $\Sigma$ and $\mathbf{D}$ be as incoherent as possible.

The recovery of $\boldsymbol{x}_{i}$ from $\boldsymbol{y}_{i}$ is an ill-posed problem, unless restrictions are placed on $\boldsymbol{x}_{i}$. We may exploit the same class of restrictions used in the denoising problem; specifically, the observed data satisfy $\boldsymbol{y}_{i}=\boldsymbol{\Phi} \boldsymbol{w}_{i}+\boldsymbol{\nu}_{i}$, with $\boldsymbol{\Phi}=\Sigma \mathrm{D}$ and $\boldsymbol{v}_{i}=\Sigma \boldsymbol{\varepsilon}_{i}$, and with sparse $\boldsymbol{w}_{i}$. Note that the sparseness constraint implies that $\left\{\boldsymbol{w}_{i}\right\}_{i=1, N}$ (and hence, $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$ ) occupy distinct subspaces of $\mathbb{R}^{P}$, selected from the overall linear subspace defined by the columns of $\Phi$.

In most applications of CS, $\mathbf{D}$ is assumed known, corresponding to an orthonormal basis (e.g., wavelets or a DCT) [5], [9], [21]. However, such bases are not necessarily well matched to natural imagery, and it is desirable to consider design of dictionaries $\mathbf{D}$ for this purpose [12]. One may even consider recovering $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$ from $\left\{\boldsymbol{y}_{i}\right\}_{i=1, N}$ while
simultaneously inferring $\mathbf{D}$. Thus, we again have a dictionary-learning problem, which may be coupled with optimization of CS matrix $\Sigma$, such that it is matched to $\mathbf{D}$ (defined by a low coherence between the rows of $\Sigma$ and columns of $\mathbf{D}$ [5], [9], [12], [21]).

## III. Sparse Factor Analysis With the Beta-Bernoulli Process

When presenting example results, we will consider three problems. For denoising, it is assumed that we measure $\boldsymbol{x}_{i}=\mathbf{D} \boldsymbol{w}_{i}+\boldsymbol{\varepsilon}_{i}$, where $\boldsymbol{\varepsilon}_{i}$ represents measurement noise and model error; for the compressive-sensing application, we observe $\boldsymbol{y}_{i}=\Sigma\left(\mathbf{D} \boldsymbol{w}_{i}+\boldsymbol{\varepsilon}_{i}\right)=\boldsymbol{\Phi} \boldsymbol{w}_{i}+\boldsymbol{v}_{i}$, with $\Phi=\Sigma \mathbf{D}$ and $\boldsymbol{\nu}_{i}=\Sigma \boldsymbol{\varepsilon}_{i}$, and finally, for the interpolation problem, we observe $Q_{\varphi}\left(\mathbf{D} \boldsymbol{w}_{i}+\boldsymbol{\varepsilon}_{i}\right)$, where $Q_{\varphi}\left(\boldsymbol{x}_{i}\right)$ is a vector of elements from $\boldsymbol{x}_{i}$ contained within set $\varphi$. For all three problems, our objective is to infer underlying signal $\mathbf{D} \boldsymbol{w}_{i}$, with $\boldsymbol{w}_{i}$ assumed sparse; we generally wish to simultaneously infer $\mathbf{D}$ and $\left\{\boldsymbol{w}_{i}\right\}_{i=1, N}$. To address each of these problems, we consider a statistical model for $\boldsymbol{x}_{i}=\mathbf{D} \boldsymbol{w}_{i}+\boldsymbol{\varepsilon}_{i}$, placing Bayesian priors on $\mathbf{D}, \boldsymbol{w}_{i}$, and $\boldsymbol{\varepsilon}_{i}$, the way the model is used is slightly modified for each specific application. For example, when considering interpolation, only the observed $Q_{\varphi}\left(\boldsymbol{x}_{i}\right)$ are used within the model likelihood function.

## A. Beta-Bernoulli Process for Active-Set Selection

Let binary vector $z_{i} \in\{0,1\}^{K}$ denote which of the $K$ columns of $\mathbf{D}$ are used for representation of $\boldsymbol{x}_{i}$ (active set); if a particular component of $\boldsymbol{z}_{i}$ is equal to one, then the corresponding column of $\mathbf{D}$ is used in the representation of $\boldsymbol{x}_{i}$. Hence, for data $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$, there is an associated set of latent binary vectors $\left\{z_{i}\right\}_{i=1, N}$, and the beta-Bernoulli process provides a convenient prior for these vectors [29], [37], [42]. Specifically, consider the following model:

$$
\begin{equation*}
\boldsymbol{z}_{i} \sim \prod_{k=1}^{K} \operatorname{Bernoulli}\left(\pi_{k}\right), \quad \pi \sim \prod_{k=1}^{K} \operatorname{Beta}(a / K, b(K-1) / K) \tag{3}
\end{equation*}
$$

where $\pi_{k}$ is the $k$ th component of $\pi$, and $a$ and $b$ are model parameters; the impact of these parameters on the model are discussed below. Note that the use of product notation in (3) is meant to denote that each component of $z_{i}$ and $\boldsymbol{\pi}$ are independently drawn from distributions of the same form.

Considering limit $K \rightarrow \infty$, and after integrating out $\boldsymbol{\pi}$, the draws of $\left\{z_{i}\right\}_{i=1, N}$ may be constituted as follows. For each $z_{j}$, draw $c_{i} \sim \operatorname{Poisson}(a l b+i-1)$ and define $C_{i}=\sum_{j=1}^{i} c_{j}$, with $C_{0}=0$. Let $z_{i k}$ represent the $k$ th component of $z_{i}$, and $z_{i k}=0$ for $k>C_{i}$. For $k=1$,
 times the $k$ th component of $\left\{z_{j}\right\}_{j=1, i-1}$ is one). For $k=C_{i-1}+1, \ldots, C_{i}$, we set $z_{i k}=1$. Note that as $a /(b+i-1)$ becomes small, with increasing $i$, it is probable that $c_{i}$ will be small. Hence, with increasing $i$, the number of new nonzero components of $z_{i}$ diminishes. Furthermore, as a consequence of $\operatorname{Bernoulli}\left(n_{i k} / b+i-1\right)$, when a particular component of vectors $\left\{z_{j}\right\}_{j=1, i-1}$ is frequently one, it is more probable that it will be one for subsequent $z_{j}$, $j \geq i$. When $b=1$, this construction for $\left\{z_{i}\right\}_{i=1, N}$ corresponds to the IBP [18].

Since $z_{i}$ defines which columns of $\mathbf{D}$ are used to represent $\boldsymbol{x}_{i}$, (3) imposes that it is probable that some columns of $\mathbf{D}$ are repeatedly used among set $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$, whereas other columns of D may be more specialized to particular $\boldsymbol{x}_{i}$. As demonstrated below, this has been found to be a good model when $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$ are patches of pixels extracted from natural images.

## B. Full Hierarchical Model

The hierarchical form of the model may now be expressed as
where $\boldsymbol{d}_{k}$ represents the $k$ th component (atom) of $\mathbf{D}, \odot$ represents the elementwise or Hadamard vector product, $\mathbf{I}_{P}\left(\mathbf{I}_{K}\right)$ represents a $P \times P(K \times K)$ identity matrix, and $\left\{z_{i}\right\}_{i=1, N}$ are drawn as in (3). Conjugate hyperpriors $\gamma_{s} \sim \operatorname{Gamma}(c, d)$ and $\gamma_{\varepsilon} \sim \operatorname{Gamma}(e, f)$ are also imposed. The construction in (4), and with the prior in (3) for $\left\{z_{i}\right\}_{i=1, N}$, is henceforth referred to as the BP factor analysis (BPFA) model. This model was first developed in [22], with a focus on general factor analysis; here, we apply and extend this construction for image-processing applications.

Note that we impose independent Gaussian priors for $\boldsymbol{d}_{k}$ and $\boldsymbol{s}_{i}$ and $\boldsymbol{\varepsilon}_{i}$ for modeling convenience (conjugacy of consecutive terms in the hierarchical model). However, the inferred posterior for these terms is generally not independent or Gaussian. The independent priors essentially impose prior information about the marginals of the posterior of each component, whereas the inferred posterior accounts for statistical dependence as reflected in the data.

To make connections of this model to more typical optimization-based approaches [24], [25], note that the negative logarithm of the posterior density function is

$$
\begin{gather*}
-\log p(\boldsymbol{\Theta} \mid \mathscr{D}, \mathscr{H})=\frac{\gamma_{\varepsilon}}{2} \sum_{i=1}^{N}\left\|\boldsymbol{x}_{i}-\mathbf{D}\left(\boldsymbol{s}_{i} \odot \boldsymbol{z}_{i}\right)\right\|_{2}^{2} \\
+\frac{P}{2} \sum_{k=1}^{K}\left\|\boldsymbol{d}_{k}\right\|_{2}^{2}+\frac{\gamma_{s}}{2} \sum_{i=1}^{N}\left\|\boldsymbol{s}_{i}\right\|_{2}^{2}  \tag{5}\\
\quad-\log f_{\text {Beta-Bern }}\left(\left\{\boldsymbol{z}_{i}\right\}_{i=1}^{N} ; \mathscr{H}\right) \\
\quad-\log \operatorname{Gammam}\left(\gamma_{\varepsilon} \mid \mathscr{H}\right) \\
-\log \operatorname{Gamma}\left(\gamma_{s} \mid \mathscr{H}\right)+\text { Const. }
\end{gather*}
$$

where $\Theta$ represents all unknown model parameters, $\mathcal{D}=\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}, f_{\text {Beta-Bem }}\left(\left\{\boldsymbol{z}_{i}\right\}_{i=1}^{N} ; \mathscr{H}\right)$ represents the beta-Bernoulli process prior in (3), and $\mathcal{H}$ represents model hyperparameters (i.e., $a, b, c, d, e$, and $f$ ). Therefore, the typical $\ell_{2}$ constraints [24], [25] on the dictionary elements $\boldsymbol{d}_{k}$ and on the nonzero weights $\boldsymbol{s}_{i}$ correspond here to the Gaussian priors employed in (4). However, rather than employing an $\ell_{1}$ (Laplacian prior) constraint [24], [25] to impose sparseness on $\boldsymbol{w}_{i}$, we employ the beta-Bernoulli process and $\boldsymbol{w}_{i}=s_{i} \odot z_{i}$. The betaBernoulli process imposes that binary $z_{i}$ should be sparse and that there should be a relatively consistent (re)use of dictionary elements across the image, thereby also imposing self-similarity. Furthermore, and perhaps most importantly, we do not constitute a point estimate, as one would do if a single $\Theta$ was sought to minimize (5). We rather estimate the full posterior density $p(\Theta \mid \mathcal{D} \mathcal{H})$, implemented via Gibbs sampling. A significant advantage of the hierarchical construction in (4) is that each Gibbs update equation is analytic, with detailed update equations provided in the Appendix. Note that consistent use of atoms is encouraged because the active sets are defined by binary vectors $\left\{z_{i}\right\}_{i=1, N}$, and these are all drawn from a shared probability vector $\pi$; this is distinct from drawing the active sets independent and identically distributed (i.i.d.) from a Laplacian prior. Furthermore, the beta-

Bernoulli prior imposes that many components of $\boldsymbol{w}_{i}$ are exactly zero, whereas with a Laplacian prior, many components are small but not exactly zero (hence, the former is analogous to $\ell_{0}$ regularization, with the latter closer to $\ell_{1}$ regularization).

## IV. Patch Clustering Via Dirichlet and PSBPs

## A. Motivation

In the model previously discussed, each patch $\boldsymbol{x}_{i}$ had a unique usage of dictionary atoms, defined by binary vector $z_{i}$, which selects columns of $\mathbf{D}$. One may wish to place further constraints on the model, thereby imposing a greater degree of statistical structure. For example, one may employ that the $\boldsymbol{x}_{i}$ cluster and that within each cluster of each of the associated $\boldsymbol{x}_{i}$ employ the same columns of $\mathbf{D}$. This is motivated by the idea that a natural image may be clustered into different types of textures or general image structure. However, rather than imposing that all $\boldsymbol{x}_{i}$ within a given cluster use exactly the same columns of $\mathbf{D}$, one may want to impose that all $\boldsymbol{x}_{i}$ within such a cluster share the same probability of dictionary usage, i.e., that all $\boldsymbol{x}_{i}$ within cluster $c$ share the same probability of using columns of $\mathbf{D}$, defined by $\boldsymbol{\pi}_{\mathcal{c}}$, rather than sharing a single $\boldsymbol{\pi}$ for all $\boldsymbol{x}_{i}$ (as in the original model above). Again, such clustering is motivated by the idea that natural images tend to segment into different textural or color forms. Below, we perform clustering in terms of vectors $\boldsymbol{\pi}_{\mathcal{c}}$, rather than explicit clustering of dictionary usage, which would entail cluster-dependent $\boldsymbol{z}_{c}$; the "softer" nature of the former clustering structure is employed to retain model flexibility while still encouraging sharing of parameters within clusters.

A question when performing such clustering concerns the number of clusters needed, this motivating the use of nonparametric methods, like those considered in the next subsections. Additionally, since the aforementioned clustering is motivated by the segmentations characteristic of natural images, it is desirable to explicitly utilize the spatial location of each image patch, encouraging that patches $\boldsymbol{x}_{i}$ in a particular segment/cluster are spatially contiguous. This latter goal motivates use of a PSBP, as also detailed as follows.
B. DP

The DP [15] constitutes a popular means of performing non-parametric clustering. A random draw from a DP , i.e., $G \sim \mathrm{DP}\left(a G_{0}\right)$, with precision $a \in \mathbb{R}^{+}$and "base" measure $G_{0}$, may be constituted via the stick-breaking construction [34], i.e.,

$$
\begin{equation*}
G=\sum_{l=1}^{\infty} \beta_{l} \delta_{\theta_{l}^{*}}, \quad \theta_{l}^{*} \sim G_{0} \tag{6}
\end{equation*}
$$

where $\beta_{l}=V_{l} \prod_{h=1}^{l-1}\left(1-V_{h}\right)$ and $V_{h} \sim \operatorname{Beta}(1, a)$. The $\beta_{l}$ may be viewed as a sequence of fractional breaks from a "stick" of original length one, where the fraction of the stick broken off on break $I$ is $V_{l}$. The $\theta_{l}^{*}$ are model parameters, associated with the th data cluster. For our problem, it has proven effective to set $G_{0}=\prod_{k=1}^{K} \operatorname{Beta}(a / K, b(K-1) / K)$ analogous to (3), and hence, $G=\sum_{l=1}^{\infty} \beta_{l} \delta_{\pi_{l}^{*}} \cdot \pi_{l}^{*}$, drawn from $G_{0}$, corresponds to distinct probability vectors for using the $K$ dictionary elements (columns of $\mathbf{D}$ ). For sample $i$, we draw $\boldsymbol{\pi}_{i} \sim G$, and a separate sparse binary vector $\boldsymbol{z}_{i}$ is drawn for each sample $\boldsymbol{x}_{i}$, as $\boldsymbol{z}_{i} \sim \prod_{k=1}^{K} \operatorname{Bernoulli}\left(\pi_{i k}\right)$, with $\boldsymbol{\pi}_{i k}$ as the $k$ th component of $\boldsymbol{\pi}_{i}$. In practice, we truncate the infinite sum for the to $G$ to $N_{L}$
elements and impose $V_{N_{L}}=1$, such that $\sum_{l=1}^{N_{L}} \beta_{l}=1$. A (conjugate) gamma prior is placed on the DP parameter $a$.

## C. PSBP

The DP yields a clustering of $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$, but it does not account for our knowledge of the location of each patch within the image. It is natural to expect that if $\boldsymbol{x}_{i}$ and $\boldsymbol{x}_{i}{ }^{\prime}$ are proximate, then, they are likely to be constituted in terms of similar columns of $\mathbf{D} .{ }^{1}$ To impose this information, we employ the PSBP. A logistic stick-breaking process is discussed in detail in [32]. We employ the closely related probit version here because it may be easily implemented in a Gibbs sampler. We note that while the method in [32] is related to that discussed below, in [32], the concepts of learned dictionaries and beta-Bernoulli priors were not considered. Another related model, which employs a probit link function, is discussed in [7].

We augment the data as $\left\{\boldsymbol{x}_{i}, \boldsymbol{r}_{i}\right\}_{i=1, N}$, where $\boldsymbol{x}_{i}$ again represents pixel values from the $i t h$ image patch, and $\boldsymbol{r}_{i} \in \mathbb{R}^{2}$ represents the 2-D location of each patch. We wish to impose that proximate patches are more likely to be composed of the same or similar columns of $\mathbf{D}$. In the PSBP construction, all aspects of (4) are retained, except for the manner in which $z_{i}$ are constituted. Rather than drawing a single $K$-dimensional vector of probabilities $\boldsymbol{\pi}$, as in (3), we draw a library of such vectors, i.e.,

$$
\begin{equation*}
\pi_{l}^{*} \sim \prod_{k=1}^{K} \operatorname{Beta}(a / K, b(K-1) / K), \quad l=1, \ldots, N_{L} \tag{7}
\end{equation*}
$$

and each $\pi_{l}^{*}$ is associated with a particular segment in the image. One $\boldsymbol{\pi}_{i}$ is associated with location $\boldsymbol{r}_{i}$ and drawn

$$
\begin{equation*}
\pi_{i} \sim \sum_{l=1}^{N_{L}} \beta_{l}\left(\boldsymbol{r}_{i}\right) \delta_{\pi_{l}^{*}} \tag{8}
\end{equation*}
$$

with $\sum_{l=1}^{N_{L}} \beta_{l}\left(\boldsymbol{r}_{i}\right)=1$ for all $\boldsymbol{r}_{i}$, and $\delta_{\pi_{i}^{*}}$ represents a point measure concentrated at $\pi_{l}^{*}$. Once $\boldsymbol{\pi}_{i}$ is associated with a particular $\boldsymbol{x}_{i}$, the corresponding binary vector $\boldsymbol{z}_{i}$ is drawn as in the first line of (3). Note that the distinction between DP and PSBP is that, in the former, mixture

[^1]weights $\left\{\beta_{l}\right\}_{l=1, N_{L}}$ are independent of spatial position $\boldsymbol{r}$, whereas the latter explicitly utilizes $\boldsymbol{r}$ within $\left\{\beta_{k}(\boldsymbol{r})\right\}_{l=1, N_{L}}$ (and below, we impose that $\beta_{k}(\boldsymbol{r})$ smoothly changes with $\boldsymbol{r}$ ).

The space-dependent weights are constructed as $\beta_{l}(\boldsymbol{r})=V_{l}(\boldsymbol{r}) \prod_{h=1}^{l-1}\left[1-V_{h}(\boldsymbol{r})\right]$, where $0<$ $V_{l}(\boldsymbol{r})<1$ constitute space-dependent probabilities. We set $V_{N_{L}}=1$, and for $l \leq N_{L}-1$, the $V_{I}$ are space-dependent probit functions, i.e.,

$$
\begin{equation*}
V_{l}(\boldsymbol{r})=\int_{-\infty}^{g_{l}(\boldsymbol{r})} d x \mathscr{N}(x \mid 0,1), \quad g_{l}(\boldsymbol{r})=\zeta_{l 0}+\sum_{i=1}^{N} \zeta_{l i} \mathscr{K}\left(\boldsymbol{r}, \boldsymbol{r}_{i} ; \psi_{l}\right) \tag{9}
\end{equation*}
$$

where $\mathcal{K}\left(\boldsymbol{r}, \boldsymbol{r}_{i} ; \psi_{l}\right)$ is a kernel characterized by parameter $\psi_{l}$ and $\left\{\zeta_{l i}\right\}_{i=0, N}$ is a sparse set of real numbers. To implement the sparseness on $\left\{\zeta_{l i}\right\}_{i=0, N}$, within prior $\zeta_{l i} \sim \mathscr{N}\left(0, \alpha_{l i}^{-1}\right)$, and (conjugate) $a_{l i} \sim \operatorname{Gamma}\left(a_{0}, b_{0}\right)$, with $\left(a_{0}, b_{0}\right)$ set to favor most $a_{l i}$ being large (if $a_{l i}$ is large, draw $\mathscr{N}\left(0, \alpha_{l i}^{-1}\right)$ is likely to be near zero, such that most $\left\{\zeta_{l i}\right\}_{i=0, N}$ are near zero). This sparseness-promoting construction is the same as that employed in the relevance vector machine (RVM) [39]. We here utilize radial basis function kernel $\mathcal{K}\left(\boldsymbol{r}, \boldsymbol{r}_{i} ; \psi_{I}\right)=\exp [-\|$ $\boldsymbol{r}_{i} \&$ minus $\left.; \boldsymbol{r} \|_{2} / \psi_{l}\right]$.

Each $g_{( }(\boldsymbol{r})$ is encouraged to only be defined by a small set of localized kernel functions, and via the probit link function $\int_{-\infty}^{g_{l}(\boldsymbol{r})} d x \mathscr{N}(x \mid 0,1)$, the probability $V_{l}(\boldsymbol{r})$ is characterized by localized segments over which the probability $V_{l}(\boldsymbol{r})$ is contiguous and smoothly varying.
The $V_{l}(\boldsymbol{r})$ constitute a space-dependent stick-breaking process. Since $V_{N_{L}}=1, \sum_{l=1}^{N_{L}} \beta_{l}(\boldsymbol{r})=1$ for all $\boldsymbol{r}$.

The PSBP model is relatively simple to implement within a Gibbs sampler. For example, as indicated above, sparseness on $\zeta_{I i}$ is imposed as in the RVM, and the probit link function is simply implemented within a Gibbs sampler (which is why it was selected, rather than a logistic link function). Finally, we define a finite set of possible kernel parameters $\left\{\psi_{j}\right\}_{j=1, N_{p}}$, and a multinomial prior is placed on these parameters, with the multinomial probability vector drawn from a Dirichlet distribution [32] (each of the $g_{( }(\boldsymbol{r})$ draws a kernel parameter from $\left\{\psi_{j}\right\}_{j=1, N_{p}}$ ). The model in (4), with the PSBP construction for $\left\{z_{i}\right\}_{i=1, N}$, is henceforth referred to as PSBP-BPFA.

## D. Discussion of Proposed Sparseness-Imposing Priors

The basic BPFA model is summarized in (4), and three related priors have been developed for sparse binary vectors $\left\{z_{i}\right\}_{i=1, N^{*}} 1$ ) the basic truncated beta-Bernoulli process in (3);2) a DP-based clustering of the underlying $\left\{\boldsymbol{\pi}_{i}\right\}_{i=1, N}$, and 3 ) a PSBP clustering of $\left\{\boldsymbol{\pi}_{i}\right\}_{i=1, N}$ that exploits knowledge of the location of the image patches. For 2 and $3, \boldsymbol{x}_{i}$ within a particular cluster have similar $z_{i}$, rather than exactly the same binary vector; we also considered the latter, but this worked less well in practice. As discussed further when presenting results, for denoising and interpolation, all three methods yield comparable performance. However, for CS, 2 and 3 yield marked improvements in image-recovery accuracy relative to 1 . In anticipation of these results, we provide a further discussion of the three priors on $\left\{z_{i}\right\}_{i=1, N}$ and on the three image-processing problems under consideration.

For the denoising and interpolation problems, we are provided with data $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$, albeit in the presence of noise and potentially with substantial missing pixels. However, for this problem, $N$ may be made quite large since we may consider all possible (overlapping) $B \times B$ patches. A given pixel (apart from near the edges of the image) is present in $B^{2}$ different
patches. Perhaps because we have such a large quantity of partially overlapping data, for denoising and interpolation, we have found that beta-Bernoulli process in (3) is sufficient for inferring the underlying relationships between the different data $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$ and processing these data collaboratively. However, the beta-Bernoulli construction does not explicitly segment the image, and therefore, an advantage of the PSBP-BPFA construction is that it yields comparable denoising and interpolation performance as (3) while also simultaneously yielding an effective image segmentation.

For the CS problem, we measure $\boldsymbol{y}_{i}=\Sigma \boldsymbol{x}_{i}$, and therefore, each of the $n$ measurements associated with each image patch $\left(\Sigma \in \mathbb{R}^{n \times P}\right)$ loses the original pixels in $\boldsymbol{x}_{i}$ (the projection matrix $\Sigma$ may also change with each patch, denoted by $\Sigma_{i}$ ). Therefore, for CS, one cannot consider all possible shifts of the patches as the patches are predefined and fixed in the CS measurement (in the denoising and interpolation problems, the patches are defined in the subsequent analysis). Therefore, for CS imposition of the clustering behavior, via DP or PSBP provides important information, yielding state-of-the-art CS-recovery results.

## E. Possible Extensions

The "Indian buffet franchise" and PSBP constructions considered above and in the results below draw the $K$-dimensional probability vectors $\pi_{l}^{*}$ independently. This implies that within the prior, we impose no statistical correlation between the components of vectors $\pi_{l}^{*}$ and $\pi_{l^{\prime}}^{*}$, for $l \neq l^{\prime}$. It may be desirable to impose such structure, imposing that there is a "global" probability of using particular dictionary elements, and the different mixture components within the DP/PSBP constructions correspond to specific draws from global statistics of dictionary usage. This will encourage the idea that there may be some "popular" dictionary elements that are shared across different mixture components (i.e., popular across different buffets in the franchise). One can impose this additional structure via a hierarchical BP (HBP) construction [37], related to the hierarchical DP (HDP) [36]. Briefly, in an HBP construction, one may draw the $\pi_{l}^{*}$ via the hierarchical construction, i.e.,

$$
\begin{gather*}
\pi_{l}^{*} \sim \prod_{k=1}^{K} \operatorname{Beta}\left(c \eta_{k}, c\left(1-\eta_{k}\right)\right), \\
\eta \sim \prod_{k=1}^{K} \operatorname{Beta}(a / K, b(K-1) / K) \tag{10}
\end{gather*}
$$

where $\eta_{k}$ is the $k$ th component of $\eta$. Vector $\boldsymbol{\eta}$ constitutes "global" probabilities of using each of the $K$ dictionary elements (across the "franchise"), and $\pi_{l}^{*}$ defines the probability of dictionary usage for the th buffet. This construction imposes statistical dependences among vectors $\left\{\pi_{l}^{*}\right\}$.

To simplify the presentation, in the following example results, we do not consider the HBP construction as good results have already been achieved with the (truncated) DP and PSBP models previously discussed. We note that in these analyses, the truncated versions of these models may actually help inference of statistical correlations among $\left\{\pi_{l}^{*}\right\}$ within the posterior (since the set $\left\{\pi_{l}^{*}\right\}$ is finite, and each vector $\pi_{l}^{*}$ is of finite length). If we actually considered the infinite limit on $K$ and on the number of mixture components, inference of such statistical relationships within the posterior may be undermined because specialized dictionary elements may be constituted across the different franchises, rather than encouraging sharing of highly similar dictionary elements.

While we do not focus on the HBP construction here, a recent paper has employed the HBP construction in related dictionary learning for image-processing applications, yielding very
encouraging results [43]. We therefore emphasize that the basic hierarchical Bayesian construction employed here is very flexible and may be extended in many ways to impose additional structure.

## V. Example Results

## A. Reproducible Research

The test results and the MATLAB code to reproduce them can be downloaded from http:// www.ee.duke.edu/~mz1/Results/BPFAImage/.

## B. Parameter Settings

For all BPFA, DP-BPFA, and PSBP-BPFA computations, the dictionary truncation level was set at $K=256$ or $K=512$ based on the size of the image. Not all $K$ dictionary elements are used in the model; the truncated beta-Bernoulli process infers the subset of dictionary elements employed to represent the data $\left\{x_{i}\right\}_{i=1, N}$. The larger the image, the more distinct types of structure are anticipated and, therefore, the more dictionary elements are likely to be employed; however, very similar results are obtained with $K=512$ in all examples, just with more dictionary elements are not employed for smaller images (therefore, to save computational resources, we set $K=256$ for the smaller images). The number of DP and PSBP sticks was set at $N_{L}=20$. The library of PSBP parameters is defined as in [32]; the PSBP kernel locations, i.e., $\left\{\boldsymbol{r}_{i}\right\}_{i=1, N}$, were situated on a uniformly sampled grid in each image dimension, situated at every fourth pixel in each direction (the results are insensitive to many related definitions of $\left\{\boldsymbol{r}_{i}\right\}_{i=1, N}$ ). The hyperparameters within the gamma distributions were set as $c=d=e=f=10^{-6}$, as is typically done in models of this type [39] (the same settings were used for the gamma prior for DP precision parameter $a$. The betadistribution parameters are set as $a=K$ and $b=1$ if random initialization is used or $a=K$ and $b=N / 8$ if a singular value decomposition (SVD) based initialization is used. None of these parameters have been optimized or tuned. When performing inference, all parameters are randomly initialized (as a draw from the associated prior) or based on the SVD of the image under test. The Gibbs samplers for the BPFA, DP-BPFA, and PSBP-BPFA have been found to mix and quickly converge, producing satisfactory results with as few as 20 iterations. The inferred images represent the average from the collection samples. All software was written in nonoptimized MATLAB. On a Dell Precision T3500 computer with a $2.4-\mathrm{GHz}$ central processing unit, for $N=148,836$ patches of size $8 \times 8 \times 3$ with $20 \%$ of the RGB pixels observed at random, the BPFA required about 2 min per Gibbs iteration (the DP version was comparable), and PSBP-BPFA required about 3 min per iteration. For the 106band hyperspectral imagery, which employed $N=428578$ patches of size $4 \times 84 \times 106$ with $2 \%$ of the voxels observed uniformly at random, each Gibbs iteration required about 15 min .

## C. Denoising

The BPFA denoising algorithm is compared with the original KSVD [13] and for both grayscale and colored images. Newer denoising algorithms include block matching with 3-D filtering [8], the multiscale KSVD [28], and KSVD with the nonlocal mean constraints [26]. These algorithms assume that the noise variance is known, whereas the proposed model automatically infers, as part of the same model, the noise variance from the image under test. There are existing methods for estimation of the noise variance, as a preprocessing step, e.g., via wavelet shrinkage [10]. However, it was shown in [42] that the denoising accuracy of methods like that in [13] can be sensitive to small errors in the estimated variance, which are likely to occur in practice. Additionally, when doing joint denoising and image interpolation, as we consider below, methods like that in [10] may not be directly applied to estimate the noise variance, as there is a large fraction of missing data. Moreover, the BPFA,

DP-BPFA, and PSBP-BPFA models infer a potentially nonstationary noise variance, with a broad prior on the variance imposed by the gamma distribution.

In the denoising examples, we consider the BPFA model in (3); similar results are obtained via the DP-BPFA and PSBP-BPFA models discussed in Section IV. To infer the final signal value at each pixel, we average the associated pixel contribution from each of the overlapping patches in which it is included (this is true for all results presented below in which overlapping patches were used).

In Table I, we consider images from [13]. The proposed BPFA performs very similarly to KSVD. As one representative example of the model's ability to infer the noise variance, we consider the Lena image from Table I. The mean inferred noise standard deviations are 5.83, $10.59,15.53,20.48,25.44,50.46$, and 100.54 for images contaminated by noise with the respective standard deviations of $5,10,15,20,25,50$, and 100 . Each of these noise variances was automatically inferred using exactly the same model, with no changes to the gamma hyperparameters (whereas for the KSVD results, it was assumed that the noise variance was known exactly a priori).

In Table II, we present similar results, for denoising RGB images; the KSVD comparisons come from [27]. An example denoising result is shown in Fig. 1. As another example of the BPFA's ability to infer the underlying noise variance, for the castle image, the mean (automatically) inferred variances are $5.15,10.18,15.22$, and 25.23 for images with additive noise with true respective standard deviations of $5,10,15$, and 25 . The sensitivity of the KSVD algorithm to a mismatch between the assumed and true noise variances is shown in [42, Fig. 1], and the insensitivity of BPFA to changes in the noise variance and to requiring knowledge of the noise variance is deemed an important advantage.

It is also important to note that the gray-scale KSVD results in Table I were initialized using an overcomplete DCT dictionary, whereas the RGB KSVD results in Table II employed an extensive set of training imagery to learn dictionary $\mathbf{D}$ that was used to initialize the denoising computations. All BPFA, DP-BPFA, and PSBP-BPFA results employ no training data, with the dictionary initialized at random using draws from the prior or with the SVD of the data under test.

## D. Image Interpolation

For the initial interpolation examples, we consider standard RGB images, with $80 \%$ of the RGB pixels uniformly missing at random (the data under test are shown in Fig. 2). Results are first presented for the Castle and Mushroom images, with comparisons between the BPFA model in (3) and the PSBP-BPFA model discussed in Section IV. The difference between the two is that the former is a "bag-of-patches" model, whereas the latter accounts for the spatial locations of the patches. Furthermore, the PSBP-BPFA simultaneously performs image recovery and segmentation. The results are shown in Fig. 3, presenting the mean reconstructed images and inferred segmentations. Each color in the inferred segmentation represents one PSBP mixture component, and the figure shows the last Gibbs iteration (to avoid issues with label switching between Gibbs iterations). While the BPFA does not directly yield a segmentation, its peak signal-to-noise ratio (PSNR) results are comparable to those inferred by PSBP-BPFA, as summarized in Table III.

An important additional advantage of Bayesian models such as BPFA, DP-BPFA, and PSBP-BPFA is that they provide a measure of confidence in the accuracy of the inferred image. In Fig. 4, we plot the variance of the inferred error $\left\{\boldsymbol{\varepsilon}_{i}\right\}_{i=1, N}$, computed via the Gibbs collection samples.

To provide a more thorough examination of model performance, in Table IV, we present results for several well-studied gray-scale and RGB images, as a function of the fraction of pixels missing. All of these results are based upon BPFA, with DP-BPFA and PSBP-BPFA yielding similar results. Finally, in Table V, we perform interpolation and denoising simultaneously, again with no training data and without prior knowledge of the noise level (again, as previously discussed, it would be difficult to estimate the noise variance as a preprocessing step in this case using methods like that in [10] as, in this case, there is also a large number of missing pixels). An example result is shown in Fig. 5. To our knowledge, this is the first time that denoising and interpolation have been jointly performed while simultaneously inferring the noise statistics.

For all of the examples previously considered, for both gray-scale and RGB images, we also attempted a direct application of matrix completion based on the incomplete matrix $\mathbf{X} \in$ $\mathbb{R}^{P \times N}$, with columns defined by the image patches (i.e., for $N$ patches, with $P$ pixels in each, the incomplete matrix is of size $P \times N$, with the $i$ th column defined by the pixels in $x_{i}$ ). We considered the algorithm in [20], using software from Prof. Candès' web-site. For most of the examples previously considered, even after very careful tuning of the parameters, the algorithm diverged, suggesting that the low-rank assumptions were violated. For examples for which the algorithm did work, the PSNR values were typically 4 to 5 dB worse than those reported here for our model.

## E. Interpolation of Hyperspectral Imagery

The basic BPFA, DP-BPFA, and PSBP-BPFA technology may be also applied to hyperspectral imagery, and it is here where these methods may have significant practical utility. Specifically, the amount of data that need be measured and read off a hyperspectral camera is often enormous. By selecting a small fraction of voxels for measurement and readout, uniformly selected at random, the quantity of data that need be handled is substantially reduced. Furthermore, one may simply modify existing hyperspectral cameras. We consider hyperspectral data with 106 spectral bands, from HyMAP scene A.P. Hill, VA, with permission from U.S. Army Topographic Engineering Center. Because of the significant statistical correlation across the multiple spectral bands, the fraction of data that need be read is further reduced, relative to gray-scale or RGB imagery. In this example, we considered $2 \%$ of the voxels, uniformly selected at random, and used image patches of size $4 \times 4 \times 106$. Other than the increased data dimensionality, nothing in the model was changed.

In Fig. 6, we show example (mean) inferred images, at two (arbitrarily selected) spectral bands, as computed via BPFA. All 106 spectral bands are simultaneously analyzed. The average PSNR for the data cube (size $845 \times 512 \times 106$ ) is 30.96 dB . While the PSNR value is of interest, for data of this type, the more important question concerns the ability to classify different materials based upon the hyperspectral data. In a separate forthcoming paper, we consider classification based on the full datacube, and based upon the BPFA-inferred datacube using $2 \%$ of the voxels, with encouraging results reported. We also tried the lowrank matrix completion algorithm from [20] for the hyperspectral data, and even after extensive parameter tuning, the algorithm diverged for all hyperspectral data considered.

In Table VI, we summarize algorithm performance on another hyperspectral data set, composed of 210 spectral bands. We show the PSNR values as a function of percentage of observed data and as a function of the size of the image patch. Note that the $1 \times 1$ patches only exploit spectral information, whereas the other patch sizes exploit both spatial and spectral information.

## F. CS

We consider a CS example in which the image is divided into $8 \times 8$ patches, with these constituting the underlying data $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$ to be inferred. For each of the $N$ blocks, a vector of CS measurements $y_{F}=\Sigma x_{i}$ is measured, where the number of projections per patch is $n$, and the total number of CS projections is $n \cdot N$. In our first examples, the elements of $\Sigma$ are randomly constructed, as draws from $\mathcal{N}(0,1)$; many other random projection classes may be considered [3] (and below, we also consider optimized projections $\Sigma$, matched to dictionary D). Each $\boldsymbol{x}_{i}$ is assumed represented in terms of dictionary $\boldsymbol{x}_{i}=\mathbf{D} \boldsymbol{w}_{i}+\boldsymbol{\varepsilon}_{i}$, and three constructions for $\mathbf{D}$ were considered: 1) a DCT expansion; 2) learning of $\mathbf{D}$ using BPFA, using training images; and 3) using the BPFA to perform joint CS inversion and learning of D. For 2, the training data consisted of $40008 \times 8$ patches chosen at random from 100 images selected from the Microsoft database (http://research.microsoft.com/enus/projects/ objectclassrecognition). The dictionary was set to $K=256$, and the offline BP inferred a dictionary of size $M=237$.

Representative CS reconstruction results are shown in Fig. 7 (left) based upon a DCT dictionary, for a gray-scale version of the "castle" image. The results in Fig. 7 (right) are based on a learned dictionary; except for the "online BP" results (where $\mathbf{D}$ and $\left\{\boldsymbol{w}_{i}\right\}_{i=1, N}$ are jointly learned), all of these results employ the same dictionary $\mathbf{D}$ learned offline, as previously mentioned, and the algorithms are distinguished by different ways of estimating $\left\{\boldsymbol{w}_{i}\right\}_{i=1, N}$. A range of CS-inversion algorithms are considered from the literature, and several BPFA-based constructions are considered as well for CS inversion. The online BPFA results (with no training data) are quite competitive with those based on a dictionary learned offline.

Note that results based on a learned dictionary are markedly better than those based on the DCT; similar results were achieved when the DCT was replaced by a wavelet representation. For the DCT-based results, note that the DP-BPFA and PSBP-BPFA CS inversion results are significantly better than those of all other CS inversion algorithms. The results reported here are consistent with tests we performed using over 100 images from the aforementioned Microsoft database, not reported here in detail for brevity.

In all previous results, projection matrix $\Sigma$ was randomly constituted. We now consider a simple means of matching $\Sigma$ to a $\mathbf{D}$ learned offline, based upon representative training images. Assume a learned $\mathbf{D} \in \mathbb{R}^{P \times K}$, with $K>P$, which may be represented via SVD as $\mathbf{D}$ $=\mathbf{U} \Lambda \mathbf{V}^{T} ; \mathbf{U} \in \mathbb{R}^{P \times P}$ and $\mathbf{V} \in \mathbb{R}^{K \times P}$ are each composed of orthonormal columns, and $\Lambda$ is a $P \times P$ diagonal matrix. The columns of $\mathbf{U}$ span the linear subspace of $\mathbb{R}^{P}$ in which the columns of $\mathbf{D}$ reside. Furthermore, since the columns of $\mathbf{D}$ are generally not orthonormal, each column of $\mathbf{D}$ is "spread out" when expanded in the columns of $\mathbf{U}$. Therefore, one expects that $\mathbf{U}$ and $\mathbf{D}$ are incoherent. Hence, a simple means of matching CS projections to the data is to define the rows of $\Sigma$ in terms of randomly selected columns of $\mathbf{U}$. This was done in Fig. 8 for the gray-scale "castle" image, using the same learned dictionary as considered in Fig. 7. It is observed that this procedure yields a marked improvement in CS recovery accuracy, for all CS inversion algorithms considered.

Concerning computational costs, all CS inversions were efficiently run on personal computers, with the specifics computational times dictated by the detailed MATLAB implementation and the machine run on. A rough ranking of the computational speeds, from fastest to slowest, is as follows: StOMP-CFAR, Fast BCS, OMP, BPFA, LARS/Lasso, Online BPFA, DP-BPFA, PSBP-BPFA, VB BCS, Basis Pursuit; in this list, algorithms BPFA through Basis Pursuits have approximately the same computational costs.

The improved performance of the CS inversion based upon the learned dictionaries is manifested as a consequence of the structure that is imposed on the underlying image while
performing inversion. The early CS inversion algorithms, of the type considered in the above comparisons, imposed that the underlying image is sparsely represented in an appropriate basis (we showed results here based upon a DCT expansion of the $8 \times 8$ blocks over which CS inversion was performed, and similar results were manifested using a wavelet expansion). The imposition of such sparseness does not take into account the additional structure between the wavelet coefficients associated with natural images. Recently, researchers have utilized such structure to move beyond sparseness and achieve even better CS-inversion quality [2], [19]; in this paper, structural relationships and correlations between basis-function coefficients are accounted for. Additionally, there has been recent statistical research that has moved beyond sparsity and that are of interest in the context of CS inversion [38]. In the tests, we omit, for brevity, the algorithms in [2] and [19] yield performance that is comparable to the best results to the right in Fig. 7. Consequently, the imposition of structure in the form of learned dictionaries can be achieved using conventional basis expansions but with more sophisticated inversion techniques, that account for structure in imagery. Therefore, the main advantage of the learned dictionary in the context of CS is that it provides a very convenient means of defining projection matrices that are matched to natural imagery, as done in Fig. 8. Once those projection matrices are specified, the CS inversion may be employed based upon dictionaries as discussed here or based upon more traditional expansions (DCT or wavelets) and newer CS inversion methods [2], [19].

## VI. Conclusion

The truncated beta-Bernoulli process has been employed to learn dictionaries matched to image patches $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$. The basic nonparametric Bayesian model is termed a BPFA framework, and extensions have been also considered. Specifically, the DP has been employed to cluster the $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$, encouraging similar dictionary-element usage within respective clusters. Furthermore, the PSBP has been used to impose that proximate patches are more likely to be similarly clustered (imposing that they are more probable to employ similar dictionary elements). All inference has been performed by a Gibbs sampler, with analytic update equations. The PBFA, DP-BPFA, and PSBP-BPFA have been applied to three problems in image processing: 1) denoising; 2) image interpolation based upon a subset of pixels selected uniformly at random; and 3) learning dictionaries for CS and also CS inversion. We have also considered jointly performing 1 and 2 . Important advantages of the proposed methods are as follows. 1) A full posterior on model parameters are inferred, and therefore, "error bars" may be placed on the inverted images. 2) The noise variance need not be known; it is inferred within the analysis and may be non-stationary, and it may be inferred in the presence of significant missing pixels. 3) While training data may be used to initialize the dictionary learning, this is not needed, and the BPFA results are highly competitive even based upon random initializations. In the context of CS, the DP-BPFA and PSBP-BPFA results are state of the art, which is significantly better than existing published methods. Finally, based upon the learned dictionary, a simple method has been constituted for optimizing the CS projections.

The interpolation problem is related to CS, in that we exploit the fact that $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$ reside on a low-dimensional subspace of $\mathbb{R}^{P}$, such that the total number of measurements is small relative to $N \cdot P\left(\right.$ recall $\left.\boldsymbol{x}_{i} \in \mathbb{R}^{P}\right)$. However, in CS, one employs projection measurements $\Sigma \boldsymbol{x}_{i}$, where $\Sigma \in \mathbb{R}^{n \times P}$, ideally with $n \ll P$. The interpolation problem corresponds to the special case in which the rows of $\Sigma$ are randomly selected rows of the $P \times P$ identity matrix. This problem is closely related to the problem of matrix completion [6], [23], [33], where the incomplete matrix $\mathbf{X} \in \mathbb{R}^{P \times N}$ has columns defined by $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$.

While the PSBP-BPFA successfully segmented the image while performing denoising and interpolation of missing pixels, we found that the PSNR performance of direct BPFA analysis performed very close to that of PSBP-BPFA in those applications. The use of PSBP-BPFA utilizes the spatial location of the image patches employed in the analysis, and therefore, it removes the exchangeability assumption associated with the simple BPFA (the location of the patches may be interchanged within the BPFA, without affecting the inference). However, since in the denoising and interpolation problems we have many overlapping patches, the extra information provided by PSBP-BPFA does not appear to be significant. By contrast, in the CS inversion problem, we do not have overlapping patches, and PSBP-BPFA provided significant performance gains relative to BPFA alone.

## Acknowledgments

This work was supported in part by DOE, by NSF, by ONR, by NGA, and by ARO. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Maya Gupta.

## References

1. Aharon M, Elad M, Bruckstein AM. K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. IEEE Trans Signal Process. Nov; 2006 54(11):4311-4322.
2. Baraniuk RG, Cevher V, Duarte MF, Hegde C. Model-based compressive sensing. IEEE Trans Inf Theory. Apr; 2010 56(4):1982-2001.
3. Baraniuk RG. Compressive sensing. IEEE Signal Process Mag. Jul; 2007 24(4):118-121.
4. Bruckstein AM, Donoho DL, Elad M. From sparse solutions of systems of equations to sparse modeling of signals and images. SIAM Rev. Jan; 2009 51(1):34-81.
5. Candès E, Tao T. Near-optimal signal recovery from random projections: Universal encoding strategies? IEEE Trans Inf Theory. Dec; 2006 52(12):5406-5425.
6. Candès EJ, Tao T. The power of convex relaxation: Near-optimal matrix completion. IEEE Trans Inf Theory. May; 2010 56(5):2053-2080.
7. Chung Y, Dunson DB. Nonparametric Bayes conditional distribution modeling with variable selection. J Amer Stat Assoc. Dec; 2009 104(488):1646-1660.
8. Dabov K, Foi A, Katkovnik V, Egiazarian K. Image denoising by sparse 3D transform-domain collaborative filtering. IEEE Trans Image Process. Aug; 2007 16(8):2080-2095. [PubMed: 17688213]
9. Donoho DL. Compressed sensing. IEEE Trans Inf Theory. Apr; 2006 52(4):1289-1306.
10. Donoho DL, Johnstone IM, Kerkyacharian G, Picard D. Wavelet shrinkage: Asymptopia. J Roy Stat Soc B. 1995; 57(2):301-369.
11. Duarte MF, Davenport MA, Takhar D, Laska JN, Sun T, Kelly KF, Baraniuk RG. Single-pixel imaging via compressive sampling. IEEE Signal Process Mag. Mar; 2008 25(2):83-91.
12. Duarte-Carvajalino JM, Sapiro G. Learning to sense sparse signals: Simultaneous sensing matrix and sparsifying dictionary optimization. IEEE Trans Image Process. Jul; 2009 18(7):1395-1408. [PubMed: 19497818]
13. Elad M, Aharon M. Image denoising via sparse and redundant representations over learned dictionaries. IEEE Trans Image Process. Dec; 2006 15(2):3736-3745. [PubMed: 17153947]
14. Elad, M.; Yavneh, I. A Weighted Average of Sparse Representations is Better Than the Sparsest One Alone. 2010. Preprint
15. Ferguson T. A Bayesian analysis of some nonparametric problems. Ann Stat. Mar; 1973 1(2):209230.
16. Geman S, Geman D. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. IEEE Trans Pattern Anal Mach Intell. Nov; 1984 PAMI-6(6):721-741. [PubMed: 22499653]
17. Gleichman S, Eldar YC. Blind Compressed Sensing Technion-Israel Inst. Technol., Haifa, Israel. CCIT Rep. 2010; 759 Preprint (on Arxiv.org).
18. Griffiths TL, Ghahramani Z. Infinite latent feature models and the Indian buffet process. Proc Adv Neural Inf Process Syst. 2005:475-482.
19. He L, Carin L. Exploiting structure in wavelet-based Bayesian compressive sensing. IEEE Trans Signal Process. Sep; 2009 57(9):3488-3497.
20. Candès EJ, Cai JF, Shen Z. A singular value thresholding algorithm for matrix completion. SIAM J Optim. Jan; 2008 20(4):1956-1982.
21. Ji S, Xue Y, Carin L. Bayesian compressive sensing. IEEE Trans Signal Process. Jun; 2008 56(6): 2346-2356.
22. Knowles D, Ghahramani Z. Infinite sparse factor analysis and infinite independent components analysis. Proc Int Conf Ind Compon Anal Signal Separation. 2007:381-388.
23. Lawrence ND, Urtasun R. Non-linear matrix factorization with Gaussian processes. Proc Int Conf Mach Learn. 2009:601-608.
24. Mairal J, Bach F, Ponce J, Sapiro G. Online dictionary learning for sparse coding. Proc Int Conf Mach Learn. 2009:689-696.
25. Mairal J, Bach F, Ponce J, Sapiro G, Zisserman A. Supervised dictionary learning. Proc Neural Inf Process Syst. 2008:1033-1040.
26. Mairal J, Bach F, Ponce J, Sapiro G, Zisserman A. Non-local sparse models for image restoration. Proc Int Conf Comput Vis. 2009:2272-2279.
27. Mairal J, Elad M, Sapiro G. Sparse representation for color image restoration. IEEE Trans Image Process. Jan; 2008 17(1):53-69. [PubMed: 18229804]
28. Mairal J, Sapiro G, Elad M. Learning multiscale sparse representations for image and video restoration. SIAM Multisc Model Simul. 2008; 7(1):214-241.
29. Paisley J, Carin L. Nonparametric factor analysis with beta process priors. Proc Int Conf Mach Learn. 2009:777-784.
30. Raina R, Battle A, Lee H, Packer B, Ng AY. Self-taught learning: Transfer learning from unlabeled data. Proc Int Conf Mach Learn. 2007:759-766.
31. Ranzato M, Poultney C, Chopra S, Lecun Y. Efficient learning of sparse representations with an energy-based model. Proc Neural Inf Process Syst. 2006:1137-1144.
32. Ren L, Du L, Dunson D, Carin L. The logistic stick breaking process. J Mach Learn Res. preprint.
33. Salakhutdinov R, Mnih A. Bayesian probabilistic matrix factorization using Markov chain Monte Carlo. Proc Int Conf Mach Learn. 2008:880-887.
34. Sethuraman J. A constructive definition of Dirichlet priors. Stat Sin. 1994; 4(2):639-650.
35. Shankar M, Pitsianis NP, Brady DJ. Compressive video sensors using multichannel imagers. Appl Opt. Feb; 2010 49(10):B9-B17. [PubMed: 20357845]
36. Teh YW, Jordan MI, Beal MJ, Blei DM. Hierarchical Dirichlet processes. J Amer Stat Assoc. Dec; 2006 101(476):1566-1581.
37. Thibaux R, Jordan MI. Hierarchical beta processes and the Indian buffet process. Proc Int Conf Artif Intell Stat. 2007:564-571.
38. Tibshirani R, Saunders M, Rosset S, Zhu J, Knight K. Sparsity and smoothness via the fused Lasso. J Roy Stat Soc B. Feb; 2005 67(1):91-108.
39. Tipping M. Sparse Bayesian learning and the relevance vector machine. J Mach Learn Res. Jun. 2001 1:211-244.
40. Wright J, Yang AY, Ganesh A, Sastry SS, Ma Y. Robust face recognition via sparse representation. IEEE Trans Pattern Anal Mach Intell. Feb; 2009 31(2):210-227. [PubMed: 19110489]
41. Yang J, Wright J, Huang T, Ma Y. Image super-resolution via sparse representation. IEEE Trans Image Process. Nov; 2009 19(11):2861-2873.
42. Zhou M, Chen H, Paisley J, Ren L, Sapiro G, Carin L. Non-parametric Bayesian dictionary learning for sparse image representations. Proc Neural Inf Process Syst. 2009:1-9.
43. Zhou M, Yang H, Sapiro G, Dunson D, Carin L. Dependent hierarchical beta process for image interpolation and denoising. Proc Int Conf Artificial Intelligence Statistics (AISTATS) JMLR W\&CP. 2011; 15:883-891.

## Biographies



Mingyuan Zhou received the B.Sc. degree in acoustics from Nanjing University, Nanjing, China, in 2005 and the M.Eng. degree in signal and information processing from the Chinese Academy of Sciences, Beijing, China, in 2008. He is currently working toward the Ph.D. degree with the Department of Electrical and Computer Engineering, Duke University, Durham, NC.

His current research interests include statistical machine learning and signal processing, with emphasis on dictionary learning, sparse coding, and image and video processing.


Haojun Chen received the M.S. and Ph.D. degrees in electrical and computer engineering from Duke University, Durham, NC, in 2009 and 2011, respectively.

His research interests include machine learning, data mining, signal processing, and image processing.

John Paisley received the B.S.E., M.S., and Ph.D. degrees from Duke University, Durham, NC, in 2004, 2007, and 2010.

He is currently a Postdoctoral Researcher with the Department of Computer Science, Princeton University, Princeton, NJ. His research interests include Bayesian nonparametrics and machine learning.


Lu Ren received the B.S. degree in electrical engineering from Xidian University, Xi'an, China, in 2002 and the Ph.D. degree in electrical and computer engineering from Duke University, Durham, NC, in 2010.

She is currently working with Yahoo! Lab as a Scientist. Her research interests include machine learning, data mining, and statistical modeling.


Lingbo Li received the B.Sc. degree in electronic information engineering with Xidian University, Xi'an, China, in 2008 and the M.Sc. degree in sensing and signals with Duke University, Durham, NC, in 2010. She is currently working toward the Ph.D. degree with the Department of Electrical and Computer Engineering, Duke University.

Her current research interests include statistical machine learning and signal processing, with emphasis on dictionary learning and topic modeling.


Zhengming Xing received the B.S. degree in telecommunications from the Civil Aviation University of China, Tianjin, China, in 2008 and the M.S. degree in electrical engineering from Duke University, Durham, NC, in 2010. He is currently working toward the Ph.D. degree in electrical and computer engineering with Duke University.

His research interests include statistical machine learning, hyperspectral image analysis, compressive sensing, and dictionary learning.


David Dunson is currently a Professor of statistical science with Duke University, Durham, NC. His recent projects have developed sparse latent factor models that scale to massive dimensions and improve performance in predicting disease and other phenotypes based on high-dimensional and longitudinal biomarkers. Related methods can be used for combining high-dimensional data from different sources and for massively dimensional variable selection. His research interests include the development and application of novel Bayesian statistical methods motivated by high-dimensional and complex data sets, with particular emphasis on nonparametric Bayesian methods that avoid assumptions, such as normality
and linearity, and on latent factor models that allow dimensionality reduction in massively dimensional settings.

Dr. Dunson is a Fellow of the American Statistical Association and the Institute of Mathematical Statistics. He was a recipient of the 2007 Mortimer Spiegelman Award for the top public health statistician, the 2010 Myrto Lefkopoulou Distinguished Lectureship at Harvard University, and the 2010 COPSS Presidents’ Award for the top statistician under 41.


Guillermo Sapiro was born in Montevideo, Uruguay, on April 3, 1966. He received his B.Sc. (summa cum laude), M.Sc., and Ph.D. degrees from Technion Israel Institute of Technology, Hafia, Israel, in 1989, 1991, and 1993, respectively.

After his postdoctoral research with Massachusetts Institute of Technology, he became a member of the Technical Staff with the research facilities of HP Labs, Palo Alto, CA. He is currently with the Department of Electrical and Computer Engineering, University of Minnesota, where he holds the position of a Distinguished McKnight University Professor and the Vincentine Hermes-Luh Chair in electrical and computer engineering.

Lawrence Carin (F'01-SM'96) was born on March 25, 1963 in Washington, DC. He received the B.S., M.S., and Ph.D. degrees in electrical engineering from the University of Maryland, College Park, in 1985, 1986, and 1989, respectively.

In 1989, he joined the Department of Electrical Engineering, Polytechnic University, Brooklyn, Brooklyn, NY, as an Assistant Professor and, in 1994, became an Associate Professor. In September 1995, he joined the Department of Electrical Engineering, Duke University, Durham, NC. He is now the William H. Younger Professor of engineering with Duke University. He is a Cofounder with Signal Innovations Group, Inc., which is a small business where he serves as the Director of Technology. He is the author of over 200 peerreviewed papers. His current research interests include signal processing, sensing, and machine learning.

Dr. Carin is a member of the Tau Beta Pi and Eta Kappa Nu Honor Societies.

## Appendix. Gibbs Sampling Inference

The Gibbs sampling update equations are given below; we provide the update equations for the BPFA, and the DP and PSBP versions are relatively simple extensions. Below, $\Sigma_{i}$ represents the projection matrix on the data, for image patch $\boldsymbol{x}_{i}$. For the CS problem, $\Sigma_{i}$ is typically fully populated, whereas for the interpolation problem, each row of $\Sigma_{i}$ is all zeros, except for a single one corresponding to the specific pixel that is measured. The update equations are the conditional probability of each parameter, conditioned on all other parameters in the model.

```
Sample d}\mp@subsup{\boldsymbol{d}}{\boldsymbol{k}}{}\mathrm{ :
```

$$
p\left(\boldsymbol{d}_{k} \mid-\right) \propto \prod_{i=1}^{N} \mathscr{N}\left(\boldsymbol{y}_{i} ; \sum_{i} \mathbf{D}\left(\boldsymbol{s}_{i} \odot \boldsymbol{z}_{i}\right), \gamma_{\varepsilon}^{-1} \mathbf{I}_{\left\|\Sigma_{i}\right\|_{0}}\right) \times \mathscr{N}\left(\boldsymbol{d}_{k} ; 0, P^{-1} \mathbf{I}_{P}\right) .
$$

It can be shown that $\boldsymbol{d}_{k}$ can be drawn from a normal distribution as

$$
p\left(\boldsymbol{d}_{k} \mid-\right) \sim \mathscr{N}\left(\mu_{\boldsymbol{d}_{k}}, \sum_{\boldsymbol{d}_{k}}\right)
$$

with the covariance $\Sigma_{d_{k}}$ and mean $\boldsymbol{\mu}_{d_{k}}$ expressed as

$$
\begin{gathered}
\sum_{\boldsymbol{d}_{k}}=\left(P \mathbf{I}+\gamma_{\varepsilon} \sum_{i=1}^{N} z_{i k}^{2} s_{i k}^{2} \sum_{i}^{T} \sum_{i}\right)^{-1} \\
\mu_{\boldsymbol{d}_{k}}=\gamma_{\varepsilon} \sum_{\boldsymbol{d}_{k}} \sum_{i=1}^{N} z_{i k} s_{i k} \tilde{\boldsymbol{x}}_{i}^{-k}
\end{gathered}
$$

where

$$
\tilde{\boldsymbol{x}}_{i}^{-k}=\sum_{i}^{T} \boldsymbol{y}_{i}-\sum_{i}^{T} \sum_{i} \mathbf{D}\left(s_{i} \odot \boldsymbol{z}_{i}\right)+\sum_{i}^{T} \sum_{i} \boldsymbol{d}_{k}\left(s_{i k} \odot z_{i k}\right)
$$

Sample $z_{k}:=\left[z_{1 k}, z_{2 k}, \cdots z_{N k}\right]$ :

$$
p\left(z_{i k} \mid-\right) \propto \mathscr{N}\left(\boldsymbol{y}_{i} ; \sum_{i} \mathbf{D}\left(\boldsymbol{s}_{i} \odot \boldsymbol{z}_{i}\right), \gamma_{\varepsilon}^{-1} \mathbf{I}_{\left\|_{i}\right\|_{k}}\right) \text { Bernoulli }\left(z_{i k} ; \pi_{k}\right) .
$$

The posterior probability that $z_{i k}=1$ is proportional to

$$
p_{1}=\pi_{k} \exp \left[-\frac{\gamma_{\varepsilon}}{2}\left(s_{i k}^{2} \boldsymbol{d}_{k}^{T} \sum_{i}^{T} \sum_{i} \boldsymbol{d}_{k}-2 s_{i k} \boldsymbol{d}_{k}^{T} \tilde{\boldsymbol{x}}_{i}^{-k}\right)\right]
$$

and the posterior probability that $z_{i k}=0$ is proportional to

$$
p_{0}=1-\pi_{k} .
$$

Thus, $z_{i k}$ can be drawn from a Bernoulli distribution as

$$
\begin{equation*}
z_{i k} \sim \text { Bernoulli }\left(\frac{p_{1}}{p_{0}+p_{1}}\right) \tag{11}
\end{equation*}
$$

Sample $\boldsymbol{s}_{k}:=\left[s_{1 k}, s_{2 k}, \cdots s_{N k}\right]$ :

$$
p\left(s_{i k} \mid-\right) \propto \mathscr{N}\left(\boldsymbol{y}_{i} ; \sum_{i} \mathbf{D}\left(\boldsymbol{s}_{i} \odot \boldsymbol{z}_{i}\right), \gamma_{\varepsilon}^{-1} \mathbf{I}_{\left\|\Sigma_{i}\right\|_{0}}\right) \mathscr{N}\left(\boldsymbol{s}_{i} ; 0, \gamma_{s}^{-1} \mathbf{I}_{K}\right)
$$

It can be shown that $s_{i k}$ can be drawn from a normal distribution, i.e.,

$$
\begin{equation*}
p\left(s_{i k} \mid-\right) \sim \mathscr{N}\left(\mu_{s_{i k}}, \sum_{s_{i k}}\right) \tag{12}
\end{equation*}
$$

variance $\Sigma_{s_{i k}}$ and mean $\mu_{s_{i k}}$ expressed as

$$
\begin{gathered}
\sum_{s_{i k}}=\left(\gamma_{s}+\gamma_{\varepsilon} z_{i k}^{2} \boldsymbol{d}_{k}^{T} \sum_{i}^{T} \sum_{i} \boldsymbol{d}_{k}\right)^{-1} \\
\mu_{s_{i k}}=\gamma_{\varepsilon} \sum_{s_{i, k}} z_{i k} \boldsymbol{d}_{k}^{T} \sum_{i}^{T} \sum_{i} \tilde{\boldsymbol{x}}_{i}^{-k} .
\end{gathered}
$$

Note that $z_{i k}$ is equal to either 1 or 0 , and $\Sigma_{s_{i k}}$ and $\mu_{s_{i k}}$ can be further expressed as

$$
\begin{aligned}
& \sum_{s_{i k}}= \begin{cases}\left(\gamma_{s}+\gamma_{\varepsilon} \boldsymbol{d}_{k}^{T} \sum_{i}^{T} \sum_{i} \boldsymbol{d}_{k}\right)^{-1} & \text { if } z_{i k}=1 \\
\gamma_{s}^{-1} & \text { if } z_{i k}=0\end{cases} \\
& \mu_{s_{i k}}= \begin{cases}\gamma_{\varepsilon} \sum_{s_{i k}} \boldsymbol{d}_{k}^{T} \sum_{i}^{T} \sum_{i} \tilde{x}_{i}^{-k} & \text { if } z_{i k}=1 \\
0 & \text { if } z_{i k}=0 .\end{cases}
\end{aligned}
$$

Sample $\pi_{k}$ :

$$
p\left(\pi_{k} \mid-\right) \propto \operatorname{Beta}\left(\pi_{k} ; \mathrm{a} / \mathrm{K}, \mathrm{~b}(\mathrm{~K}-1) / \mathrm{K}\right) \prod_{i=1}^{N} \operatorname{Bernoulli}\left(Z_{i k} ; \pi_{k}\right) .
$$

It can be shown that $\pi_{k}$ can be drawn from a Beta distribution as

$$
p\left(\pi_{k} \mid-\right) \sim \operatorname{Beta}\left(\frac{a}{K}+\sum_{i=1}^{N} z_{i k} \frac{b(K-1)}{K}+N-\sum_{i=1}^{N} z_{i k}\right)
$$

Sample $\gamma_{s}$ :

$$
p\left(\gamma_{s} \mid-\right) \propto \Gamma\left(\gamma_{s} ; c_{0}, d_{0}\right) \prod_{i=1}^{N} \mathscr{N}\left(\mathbf{s}_{i} ; 0, \gamma_{s}^{-1} \mathbf{I}_{K}\right)
$$

It can be shown that $\gamma_{S}$ can be drawn from a Gamma distribution as

$$
p\left(\gamma_{s} \mid-\right) \sim \Gamma\left(c_{0}+\frac{1}{2} K N, d_{0}+\frac{1}{2} \sum_{i=1}^{N} \mathbf{s}_{i}^{T} \mathbf{s}_{i}\right) .
$$

Sample $\gamma_{\varepsilon}$ :

$$
\begin{equation*}
p\left(\gamma_{\varepsilon} \mid-\right) \propto \Gamma\left(\gamma_{\varepsilon} ; e_{0}, f_{0}\right) \prod_{i=1}^{N} \mathscr{N}\left(\boldsymbol{y}_{i} ; \sum_{i} \mathbf{D}\left(s_{i} \odot \boldsymbol{z}_{i}\right), \gamma_{\varepsilon}^{-1} \mathbf{I}_{\left\|\Sigma_{i}\right\|_{0}}\right) \tag{13}
\end{equation*}
$$

It can be shown that $\gamma_{\varepsilon}$ can be drawn from a Gamma distribution as

$$
\begin{equation*}
p\left(\gamma_{\varepsilon} \mid-\right) \sim \Gamma\left(e_{0}+\frac{1}{2} \sum_{i=1}^{N}\left\|\sum_{i}\right\|_{0}, f_{0}+\frac{1}{2} \sum_{i=1}^{N}\left\|\sum_{i}^{T} \boldsymbol{y}_{i}-\sum_{i}^{T} \sum_{i} \mathbf{D}\left(\boldsymbol{s}_{i} \odot \boldsymbol{z}_{i}\right)\right\|_{\ell_{2}}^{2}\right) . \tag{14}
\end{equation*}
$$

Note that $\sum_{i}^{T} \sum_{i}$ is a sparse identity matrix, $\Sigma_{d_{k}}$ is a diagonal matrix, and $\mathbf{Z}$ is a sparse matrix; it is easy to find that only basic arithmetical operations are needed, and many unnecessary calculations can be avoided, leading to fast computation and low memory requirement.


Fig. 1.
(From left to right) The original horses image, the noisy horses image with the noise standard deviation of 25, the denoised image, and the inferred dictionary (from top left) with its elements ordered in the probability to be used. The low-probability dictionary elements are never used to represent $\left\{\boldsymbol{x}_{i}\right\}_{i=1, N}$ and are drawn from the prior, showing the ability of the model to learn the number of dictionary elements needed for the data.


Fig. 2.
Images with $80 \%$ of the RGB pixels missing at random. Although only $20 \%$ of the actual pixels are observed, in these figures, the missing pixels are estimated based upon averaging all observed neighboring pixels within a $5 \times 5$ spatial extent. (Left) Castle image ( $22.58-\mathrm{dB}$ PSNR). (Right) Mushroom image ( 24.85 dB ).


Fig. 3.
PSBP-BPFA analysis with $80 \%$ of the RGB pixels uniformly missing at random (see Fig. 2). The analysis is based on $8 \times 8 \times 3$ image patches, considering all possible (overlapping) parches. For a given pixel, the results are the average based upon all patches in which it is contained. For each example: (left) recovered image based on an average of Gibbs collection samples and (right) each color representing one of the PSBP mixture components.


Fig. 4.
Expected variance of each pixel for the (Mushroom) data considered in Fig. 3.


Fig. 5.
(From left to right) The original barbara256 image, the noisy and incomplete barbara256 image with the noise standard deviation of 15 and $70 \%$ of its pixels missing at random (displayed based upon imputing missing pixels as the average of all observed pixels in a $5 \times 5$ neighborhood), the restored image, and the inferred dictionary with its elements ordered in the probability to be used (from top left).


Fig. 6.
Comparison of recovered band (average from Gibbs collection iterations) for hyperspectral imagery with 106 spectral bands. The interpolation is performed using $2 \%$ of the hyperspectral datacube, uniformly selected at random. The analysis employs $4 \times 4 \times 106$ patches. All spectral bands are analyzed at once, and here, the data (recovered and original) are shown (arbitrarily) for bands (top) 1 and (bottom) 50. Results are computed using the BPFA model.


Fig. 7.
(Left) CS results on gray-scale Castle image, based on DCT dictionary D. CS projection matrix $\Sigma$ is randomly constituted, with elements drawn i.i.d. from $\mathcal{N}(0,1)$. Results are shown using the DP-BPFA and PSBP-BPFA models in Section IV. Comparisons are also made with several CS inversion algorithms from the literature. (Right) Same as on the left but based on learned dictionary $\mathbf{D}$, instead of DCT. The online BP results employ BPFA to learn $\mathbf{D}$ and do CS inversion jointly. All other results are based upon learned $\mathbf{D}$ with learning performed offline using distinct training images.


Fig. 8.
CS results on gray-scale Castle image, based on learned dictionary $\mathbf{D}$ (learning performed offline, using distinct training data). The projection matrix $\Sigma$ is matched to $\mathbf{D}$, based upon an SVD of $\mathbf{D}$.
TABLE I
Gray－Scale Image Denoising PSNR Results，Comparing KSVD［13］and BPFA，Using Patch Size $8 \times 8$ ．The Top and Bottom Parts of Each Cell Are Results of KSVD and BPFA，Respectively

| 易 | $\begin{array}{\|l\|l} \hline \underset{\sim}{\mathrm{B}} \\ \hline \end{array}$ | $\begin{aligned} & \text { H } \\ & \text { è } \end{aligned}$ | $\stackrel{\underset{m}{m}}{\substack{n}}$ | $\begin{aligned} & \bar{m} \\ & \underset{m}{n} \end{aligned}$ | $\underset{\mathrm{m}}{\mathrm{~F}}$ | $\begin{array}{\|c} \hline \mathrm{d} \\ \hline \mathbf{m} \end{array}$ | $\stackrel{\infty}{\stackrel{\infty}{8}}$ | $\begin{gathered} \text { f } \\ \dot{\sim} \end{gathered}$ | $\stackrel{\infty}{\underset{\sim}{i}}$ | $\begin{aligned} & \text { in } \\ & \vdots \end{aligned}$ | $\begin{gathered} \underset{\sim}{\underset{\sim}{c}} \\ \underset{\sim}{2} \end{gathered}$ | $\begin{aligned} & \bar{\infty} \\ & \underset{\sim}{0} \end{aligned}$ | $\begin{aligned} & \stackrel{\infty}{\underset{\sim}{c}} \\ & \underset{\sim}{2} \end{aligned}$ | $\stackrel{\text { ̃ }}{\text { d }}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{\|l\|} \hline 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{array}$ |  | N. | $\begin{gathered} \tilde{n} \\ \underset{\sim}{n} \end{gathered}$ | $\stackrel{\substack{\mathrm{N}}}{ }$ | $\frac{\mathfrak{n}}{\dot{m}}$ | $\stackrel{n}{\sim}$ | $\begin{aligned} & \stackrel{8}{\dot{~}} \\ & \hline \end{aligned}$ | $\begin{aligned} & \stackrel{\rightharpoonup}{\mathrm{m}} \\ & \stackrel{e}{2} \end{aligned}$ | $\underset{\sim}{\infty}$ | $\begin{gathered} \infty \\ \underset{\sim}{c} \end{gathered}$ | $\begin{aligned} & \underset{\sim}{N} \\ & \underset{\sim}{n} \end{aligned}$ | $\begin{gathered} \text { む̀ } \\ \text { べ } \end{gathered}$ | $\begin{aligned} & \stackrel{8}{0} \\ & \text { di } \end{aligned}$ | $\stackrel{\rightharpoonup}{\dot{\sim}}$ |
|  | $\begin{aligned} & n \\ & \stackrel{n}{6} \\ & \underset{\sim}{2} \end{aligned}$ | స్లి | $\begin{gathered} \underset{\sim}{2} \\ \text { in } \end{gathered}$ | $\underset{\sim}{\underset{\sim}{\mathrm{I}}}$ | $\begin{gathered} \stackrel{\circ}{\circ} \\ \stackrel{e}{\circ} \end{gathered}$ | $\begin{gathered} \underset{\sim}{\mathrm{I}} \\ \stackrel{\ominus}{\circ} \end{gathered}$ | $\underset{\sim}{\dot{\infty}}$ | $\underset{\underset{\sim}{\infty}}{\underset{\sim}{N}}$ | $\underset{\text { ה̀ }}{\substack{n}}$ | $\stackrel{\infty}{n} \stackrel{n}{\wedge}$ | $\underset{\underset{\sim}{\mathrm{N}}}{\substack{\mathrm{~N}}}$ |  | $\stackrel{\underset{\sim}{\infty}}{\stackrel{\infty}{\infty}}$ | $\stackrel{\square}{\text { ¢ }}$ |
|  | $\underset{\sim}{\underset{\sim}{n}}$ | 然 | $\begin{gathered} \text { d } \\ \underset{\sim}{j} \end{gathered}$ | $\stackrel{\text { ®}}{\infty}$ | $\underset{\sim}{\underset{m}{n}}$ | $\frac{\stackrel{\rightharpoonup}{\mathrm{a}}}{\bar{m}}$ | $\begin{gathered} \infty \\ \underset{\sim}{c} \end{gathered}$ | $\begin{aligned} & \stackrel{\circ}{\dot{\sim}} \\ & \text { in } \end{aligned}$ | $\stackrel{\sim}{c}$ | $\stackrel{\stackrel{\rightharpoonup}{\dot{~}}}{\stackrel{1}{2}}$ | $\begin{gathered} \text { n} \\ \underset{\sim}{n} \end{gathered}$ | $\begin{aligned} & \text { in } \\ & \text { in } \end{aligned}$ | $\begin{aligned} & \stackrel{\rightharpoonup}{\infty} \\ & \text { ત̀ } \end{aligned}$ | $\underset{\sim}{\underset{\sim}{N}}$ |
|  | $\begin{gathered} \infty \\ \infty \\ \infty \\ \infty \\ \hline \end{gathered}$ | $\underset{\underset{\sim}{c}}{\stackrel{\rightharpoonup}{2}}$ | $\begin{aligned} & \text { I } \\ & \underset{\sim}{f} \end{aligned}$ | $\begin{aligned} & \stackrel{o}{6} \\ & \underset{\sim}{2} \end{aligned}$ | $\begin{gathered} \stackrel{\rightharpoonup}{i} \\ \underset{\sim}{2} \end{gathered}$ | $\begin{array}{\|c} \stackrel{\rightharpoonup}{i} \\ \text { in } \end{array}$ | $\begin{aligned} & \infty \\ & \stackrel{\infty}{\infty} \\ & \hline \end{aligned}$ | $\stackrel{0}{\dot{m}}$ | $\begin{aligned} & \text { B. } \\ & \stackrel{\text { N }}{ } \end{aligned}$ | $$ | $\begin{aligned} & \underset{\sim}{\text { H. }} \\ & \underset{\sim}{n} \end{aligned}$ | $\stackrel{\substack{\mathrm{N}}}{\substack{0}}$ | $\begin{aligned} & \stackrel{\rightharpoonup}{\infty} \\ & \stackrel{\rightharpoonup}{\lambda} \end{aligned}$ | $\begin{aligned} & \underset{\sim}{\mathrm{i}} \end{aligned}$ |
| $\begin{array}{\|l\|l} \underline{y y} \\ \hline \end{array}$ | $\begin{aligned} & \underset{\sim}{8} \\ & \infty \\ & \infty \\ & \hline \end{aligned}$ | $\begin{gathered} \stackrel{\rightharpoonup}{\circ} \\ \infty \\ \hline \end{gathered}$ | $\underset{\sim}{f}$ | $\underset{\substack{\text { ®i } \\ \hline}}{ }$ | $\underset{\sim}{\underset{m}{2}}$ | $\underset{\sim}{\infty}$ | $\underset{\underset{\sim}{\infty}}{\substack{\infty \\ \hline}}$ | $\begin{aligned} & \stackrel{\rightharpoonup}{\mathrm{i}} \\ & \text { R } \end{aligned}$ | $\stackrel{\sim}{c}$ | $\begin{gathered} \substack{o \\ \hdashline \\ \hline} \end{gathered}$ | $\stackrel{\imath}{\grave{\wedge}}$ | $\underset{\underset{\sim}{\infty}}{\stackrel{\rightharpoonup}{c}}$ | $\begin{aligned} & \text { + } \\ & \underset{\sim}{c} \end{aligned}$ | $\underset{\sim}{\text { ® }}$ |
| $$ | $\stackrel{\infty}{\stackrel{\infty}{\infty}}$ | $\underset{\underset{\sim}{\mathrm{N}}}{\substack{\mathrm{a}}}$ | $\begin{gathered} \infty \\ \underset{\sim}{\infty} \end{gathered}$ | $\begin{aligned} & \bar{m} \\ & \underset{m}{2} \end{aligned}$ | $\begin{gathered} \tilde{y} \\ \underset{\sim}{2} \end{gathered}$ | $\begin{gathered} \substack{o \\ \text { din } \\ \text { n }} \end{gathered}$ | $\begin{aligned} & \infty \\ & \infty \\ & \dot{\sim} \end{aligned}$ | $\stackrel{9}{\dot{m}}$ | $\stackrel{\text { N}}{\underset{\sim}{c}}$ | $\begin{aligned} & \stackrel{8}{6} \\ & \stackrel{e}{2} \end{aligned}$ | $\stackrel{m}{\stackrel{m}{\dot{c}}}$ | $\begin{aligned} & \circ \\ & \stackrel{\leftrightarrow}{\circ} \\ & \text { in } \end{aligned}$ | $\stackrel{n}{\underset{\sim}{i}}$ | $\stackrel{\text { N }}{\text { ® }}$ |
|  | $\underset{\sim}{\stackrel{\rightharpoonup}{c}}$ | $\stackrel{\infty}{\infty}$ | $$ | $\begin{aligned} & \text { సे } \\ & \text { む̀ } \end{aligned}$ | $\begin{gathered} \underset{\sim}{c} \\ \underset{\sim}{2} \end{gathered}$ | $\begin{aligned} & \underset{n}{n} \\ & \dot{\sim} \end{aligned}$ | $\begin{gathered} \underset{\sim}{c} \end{gathered}$ | $\begin{aligned} & \text { N్ల } \\ & \underset{\sim}{2} \end{aligned}$ | $\underset{\sim}{i}$ | $\underset{\substack{\mathrm{j}}}{\underset{\sim}{2}}$ | $\stackrel{\curvearrowleft}{\grave{へ}}$ | $\begin{aligned} & \stackrel{\rightharpoonup}{\dot{\alpha}} \\ & \stackrel{\text { ® }}{2} \end{aligned}$ | $\underset{\underset{\sim}{N}}{\underset{\sim}{2}}$ | $\stackrel{\text { ch }}{\substack{\text { c }}}$ |
| $\begin{array}{\|l\|l\|} \hline \end{array}$ | $\stackrel{\stackrel{\infty}{\infty}}{\stackrel{\infty}{\infty}}$ | $\stackrel{\underset{\sim}{\sim}}{\underset{\sim}{n}}$ | $\underset{\sim}{\infty}$ | $\begin{gathered} \underset{\sim}{i} \\ \text { in } \end{gathered}$ | $\stackrel{\underset{\sim}{\mathcal{F}}}{\underset{\sim}{2}}$ | $\stackrel{\stackrel{+}{\underset{\sim}{m}}}{\substack{2}}$ | बे | $\begin{aligned} & \stackrel{o}{0} \\ & \stackrel{e}{0} \end{aligned}$ | $\stackrel{\infty}{\infty} \underset{\sim}{\infty}$ | $\underset{\underset{\sim}{\dot{\alpha}}}{\stackrel{\rightharpoonup}{2}}$ | $\begin{gathered} \underset{\sim}{n} \\ \underset{\sim}{c} \end{gathered}$ | $\begin{aligned} & \stackrel{\rightharpoonup}{6} \\ & \stackrel{\rightharpoonup}{n} \end{aligned}$ | $\stackrel{\stackrel{\rightharpoonup}{i}}{\stackrel{1}{2}}$ | $\stackrel{\aleph}{\stackrel{\sim}{2}}$ |
| $\bigcirc$ |  | n |  | $\bigcirc$ |  |  |  |  |  |  |  |  |  |  |

\footnotetext{
॥эา9ャ1
RGB Image Denoising PSNR Results Comparing KSVD [27] and BPFA, Both Using a Patch Size of $7 \times 7$. The Top and Bottom Parts of Each Cell Show the Results of KSVD and BPFA, Respectively

|  | $\begin{aligned} & \stackrel{\circ}{\circ} \\ & \stackrel{\rightharpoonup}{2} \end{aligned}$ | $\begin{aligned} & \stackrel{\circ}{\circ} \\ & \stackrel{\rightharpoonup}{2} \end{aligned}$ | $\left\lvert\, \begin{aligned} & \stackrel{\rightharpoonup}{2} \\ & \underset{\sim}{2} \end{aligned}\right.$ | $\begin{aligned} & \text { స̇ } \\ & \text { ju } \end{aligned}$ | $\underset{\sim}{\underset{m}{\infty}}$ | $\underset{\sim}{\circ}$ | $\underset{\sim}{\infty}$ | $\stackrel{\substack{\infty \\ \infty \\ \infty}}{\text { ¢ }}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{\|l\|l} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{array}$ | $\begin{array}{\|l\|l} \mathrm{g} \\ \mathrm{G} \end{array}$ | $$ | $\begin{array}{\|c} \substack{f \\ \text { w } \\ \hline} \end{array}$ | $\begin{aligned} & \infty \\ & \substack{\infty \\ \dot{N} \\ \hline} \end{aligned}$ | $\begin{array}{\|c} \stackrel{\circ}{c} \\ \text { ㄹ } \end{array}$ | $\begin{gathered} \infty \\ \underset{\sim}{i} \\ \underset{\sim}{2} \end{gathered}$ | $\begin{aligned} & \bar{\infty} \\ & \dot{\alpha} \end{aligned}$ | $\stackrel{7}{\text { F }}$ |
| $\begin{array}{\|l\|l} \hline E \\ \hline \end{array}$ | $\begin{aligned} & \stackrel{\circ}{j} \\ & \stackrel{y}{2} \end{aligned}$ | $\begin{aligned} & \infty \\ & \stackrel{\infty}{2} \end{aligned}$ | $$ | $\begin{aligned} & \hline \infty \\ & \stackrel{\leftrightarrow}{4} \\ & \underset{\sim}{2} \end{aligned}$ | $\stackrel{\stackrel{\overbrace{}}{\mathrm{N}}}{\underset{\mathrm{~N}}{2}}$ | $\begin{aligned} & \underset{\sim}{n} \\ & \underset{\sim}{2} \end{aligned}$ | $\underset{\sim}{\infty}$ | $\stackrel{\infty}{\infty}$ |
|  | $\begin{array}{\|c} \text { N } \\ \text { ® } \end{array}$ | $\underset{\sim}{\underset{\sim}{2}}$ | $\begin{aligned} & \stackrel{\rightharpoonup}{6} \\ & \stackrel{\sim}{m} \end{aligned}$ | $\begin{aligned} & \underset{\sim}{2} \\ & \underset{\sim}{2} \end{aligned}$ | $\underset{\sim}{\infty} \underset{\sim}{\infty}$ | $\begin{gathered} \underset{\sim}{j} \\ \underset{\sim}{2} \end{gathered}$ | $\begin{array}{\|l\|l} \stackrel{\rightharpoonup}{1} \\ \text { ej } \end{array}$ | $\begin{gathered} \text { ƠO } \\ \text { Ben } \end{gathered}$ |
|  | $\begin{array}{\|l\|l} \hat{m} \\ \dot{q} \end{array}$ |  |  |  | $\begin{aligned} & \stackrel{\infty}{\infty} \\ & \underset{\sim}{2} \end{aligned}$ | $\begin{aligned} & \dot{d} \\ & \dot{j} \\ & \hline \end{aligned}$ | $\stackrel{\stackrel{\rightharpoonup}{\dot{m}}}{\stackrel{1}{2}}$ | $\underset{\underset{\sim}{\top}}{\underset{\sim}{+}}$ |
| $\bigcirc$ |  | $n$ |  | $\bigcirc$ |  |  |  | ภ |

## TABLE III

Comparison of Interpolation of the Castle and Mushroom Images, Based Upon Observing 20\% of the Pixels, Uniformly Selected at Random. Results Are Shown Using BPFA and PSBP-BPFA, and the Analysis Is Separately Performed Using $8 \times 8 \times 3$ and $5 \times 5 \times 3$ Image Patches

|  | Castle $\mathbf{8} \times \mathbf{8} \times \mathbf{3}$ | Castle $\mathbf{5} \times \mathbf{5} \times \mathbf{3}$ | Mushroom $\mathbf{8} \times \mathbf{8} \times \mathbf{3}$ | Mushroom $\mathbf{5} \times \mathbf{5} \times \mathbf{3}$ |
| :---: | :---: | :---: | :---: | :---: |
| BPFA | 29.32 | 28.48 | 31.63 | 31.17 |
| PSBP-BPFA | 29.54 | 28.46 | 32.03 | 31.27 |

TABLE IV
Top：BPFA Gray－Scale Image Interpolation PSNR Results，Using Patch Size $8 \times 8$ ．Bottom：BPFA RGB Image Interpolation PSNR Results，Using Patch Size $7 \times 7$

| data ratio | C．man | House | Peppers | Lena | Barbara | Boats | F．print | Man | Couple | Hill |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $20 \%$ | 24.11 | 30.12 | 25.92 | 31.00 | 24.80 | 27.81 | 26.03 | 28.24 | 27.72 | 29.33 |
| $30 \%$ | 25.71 | 33.14 | 28.19 | 33.31 | 27.52 | 30.00 | 09.01 | 30.06 | 30.00 | 31.21 |
| $50 \%$ | 28.90 | 38.02 | 32.58 | 36.94 | 33.17 | 33.78 | 33.53 | 33.29 | 35.56 | 34.23 |
| $80 \%$ | 34.70 | 43.03 | 37.73 | 41.27 | 40.76 | 39.50 | 40.17 | 39.11 | 38.71 | 38.75 |


|  | $\begin{aligned} & \hat{n} \\ & \underset{\sim}{2} \end{aligned}$ | $\begin{gathered} \vec{j} \\ \underset{\sim}{j} \end{gathered}$ | $\stackrel{\underset{\sim}{\mathrm{m}}}{\underset{\sim}{c}}$ |  |
| :---: | :---: | :---: | :---: | :---: |
|  | $\begin{aligned} & \stackrel{\rightharpoonup}{\mathrm{A}} \end{aligned}$ | $\begin{aligned} & \tilde{N} \\ & \text { N } \end{aligned}$ | $\underset{\substack{\mathrm{N}}}{\stackrel{\rightharpoonup}{4}}$ | $\stackrel{\stackrel{3}{7}}{7}$ |
| $\begin{aligned} & \text { 要 } \\ & \hline \end{aligned}$ | $\begin{aligned} & \text { in } \\ & \underset{\sim}{n} \end{aligned}$ | $\begin{aligned} & \stackrel{8}{+} \\ & \stackrel{y}{n} \end{aligned}$ | $\begin{array}{\|c} \underset{\text { ci }}{ } \\ \text { in } \end{array}$ | ¢ |
|  | $\frac{\stackrel{n}{n}}{\stackrel{m}{m}}$ | $\begin{aligned} & \mathfrak{o} \\ & \underset{\sim}{f} \end{aligned}$ | $\begin{aligned} & \infty \\ & \infty \\ & \infty \\ & \infty \end{aligned}$ | $\xrightarrow{\sim}$ |
|  | $\begin{aligned} & \underset{\sim}{A} \\ & \hline \end{aligned}$ | $\begin{gathered} \underset{\sim}{\mathrm{I}} \\ \text { din } \end{gathered}$ | $\begin{aligned} & \text { f } \\ & \text { én } \end{aligned}$ | $\stackrel{\square}{\square}$ |
| $\begin{aligned} & \text { 品 } \\ & \frac{5}{5} \\ & \frac{\pi}{5} \end{aligned}$ | $\stackrel{\circ}{\mathrm{c}}$ | 术 | $\begin{gathered} \circ 8 \\ \text { 鬲 } \end{gathered}$ | － |



## TABLE VI

BPFA Hyperspectral Image Interpolation PSNR Results. For This Example, the Test Image Is a $150 \times 150$ Urban Image With 210 Spectral Bands. Results Are Shown as a Function of the Percentage of Observed Voxels, for Different Sized Patches (e.g., the $4 \times 4$ Case Corresponds to $4 \times 4 \times 210$ "Patches")

| Observed data (\%) | $\mathbf{1 \times 1}$ | $\mathbf{2 \times 2}$ | $\mathbf{3} \times \mathbf{3}$ | $\mathbf{4} \times \mathbf{4}$ |
| :---: | :---: | :---: | :---: | :---: |
| 2 | 15.34 | 21.09 | 22.72 | 23.46 |
| 5 | 17.98 | 23.58 | 25.30 | 25.88 |
| 10 | 20.41 | 25.27 | 26.36 | 26.68 |
| 20 | 22.22 | 26.50 | 27.02 | 27.16 |


[^0]:    © 2011 IEEE
    Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

[^1]:    ${ }^{1}$ Proximity can be modeled as in "spatial proximity," as here developed in detail, or "feature proximity" as in nonlocal means and related approaches; see [26] and references therein.

