

Bayesian Orthogonal Component Analysis for Sparse Representation

Nicolas Dobigeon, *Member, IEEE*, and Jean-Yves Tournet, *Senior Member, IEEE*

Abstract—This paper addresses the problem of identifying a lower dimensional space where observed data can be sparsely represented. This undercomplete dictionary learning task can be formulated as a blind separation problem of sparse sources linearly mixed with an unknown orthogonal mixing matrix. This issue is formulated in a Bayesian framework. First, the unknown sparse sources are modeled as Bernoulli–Gaussian processes. To promote sparsity, a weighted mixture of an atom at zero and a Gaussian distribution is proposed as prior distribution for the unobserved sources. A noninformative prior distribution defined on an appropriate Stiefel manifold is elected for the mixing matrix. The Bayesian inference on the unknown parameters is conducted using a Markov chain Monte Carlo (MCMC) method. A partially collapsed Gibbs sampler is designed to generate samples asymptotically distributed according to the joint posterior distribution of the unknown model parameters and hyperparameters. These samples are then used to approximate the joint maximum *a posteriori* estimator of the sources and mixing matrix. Simulations conducted on synthetic data are reported to illustrate the performance of the method for recovering sparse representations. An application to sparse coding on undercomplete dictionary is finally investigated.

Index Terms—Bayesian inference, dictionary learning, Markov chain Monte Carlo (MCMC) methods, sparse representation.

I. INTRODUCTION

IN recent years, sparse representations have motivated much research in the signal processing community. This issue consists of identifying a sparse decomposition of a signal on a given dictionary. Among the main motivations, such representations have been demonstrated to be an efficient alternative for regularizing ill-posed inverse problems [1]. More recently, compressive sensing has extensively benefited from sparsity to reconstruct a signal from a few projections [2], [3]. Signal reconstruction under hard sparse constraints can be mainly formulated as an optimization problem of a ℓ_0 -penalized quadratic criterion, whose numerical resolution is unfortunately an NP-complete problem. Several greedy algorithms have been proposed to approximate the signal reconstruction solutions, such as the well-known matching pursuit (MP) [4] and orthogonal matching pursuit (OMP) [5] algorithms. However, under appropriate sufficient conditions,

replacing the ℓ_0 -norm by the ℓ_1 -norm in the penalization term can lead to the same unique solution [6]. Therefore, exploiting these interesting sparseness properties, extensive works have been devoted on ℓ_1 -constrained estimation problems for sparse representation (see, for example, [7] and [8]).

In all the above works, the (generally overcomplete) dictionary on which the signal is sparsely decomposed is assumed to be *a priori* known. The joint estimation of the atoms of the dictionary and the corresponding sparse representation is a much more challenging task. In [9], Aharon *et al.* introduced an MP-based iterative method for designing overcomplete dictionaries. Of course, the overcompleteness allows redundancy in the atom decomposition. We address here the problem of recovering a sparse data representation in a lower-dimensional space defined by an undercomplete orthogonal dictionary. Some up-to-date research activities conducted in the signal processing and machine learning communities have been focusing on this still open problem. Specifically, Mishali and Eldar have introduced in [10] an alternating minimization procedure to solve a complete sparse representation problem when the sparsity level is assumed to be known. In [11], [12], and more recently in [13], the decomposition of a covariance matrix into sparse factors has been formulated as a regression problem with sparsity constraints. More generally, these matrix factorization strategies under some particular constraints, e.g., non-negativity, orthogonality and sparsity have demonstrated great interest for many different applications. These applications include representation of natural images [14] and gene expression data analysis [15].

In this paper, the undercomplete dictionary learning task is formulated as a blind source separation problem with sparsity constraints. Many applications have encouraged research on sparse signal and image deconvolution. These applications include astronomy [16], geophysics [17], audio signal decomposition [18] and, more recently, molecular imaging [19]. In the present work, we propose a hierarchical Bayesian model for blind separation of sparse sources linearly mixed by an orthogonal matrix.¹ This model is based on the choice of pertinent prior distributions for unknown parameters and hyperparameters. Following the works of Kormylo and Mendel [20], the unknown sources are assumed to be Bernoulli–Gaussian (BG) processes. Therefore, the source prior is composed of a weighted mixture of a standard Gaussian distribution and a mass at zero. Note that this distribution has been widely advocated to solve reconstruction problems in a Bayesian framework (see [21]–[24], among others). However, estimating

¹In the following, the mixing matrix is said orthogonal although it is not a square matrix. This abuse of language will mean that its columns, i.e., the dictionary atoms, are orthogonal.

Manuscript received August 31, 2009; accepted January 02, 2010. Date of publication January 26, 2010; date of current version April 14, 2010. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Z. Jane Wang.

The authors are with the University of Toulouse, IRIT/INP-ENSEEIH/TéSA, BP 7122, 31071 Toulouse Cedex 7, France (e-mail: (Nicolas.Dobigeon@enseiht.fr; Jean-Yves.Tournet@enseiht.fr).

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

Digital Object Identifier 10.1109/TSP.2010.2041594

hyperparameters involved in such prior mixture is a critical issue that drastically impacts the estimation performance. As an example, the empirical Bayes (EB) and Stein unbiased risk (SURE) approach proposed in [25] experienced instability especially at high signal-to-noise ratios (SNRs). In the adopted Bayesian estimation framework, several strategies are available to efficiently estimate these hyperparameters in an unsupervised manner. Lavielle *et al.* proposed to couple Markov chain Monte Carlo (MCMC) methods to a (stochastic) expectation-maximization (EM) algorithm [26], [27]. A popular alternative to this hybrid strategy consists of introducing a second level of hierarchy in the Bayesian model by assigning noninformative prior distributions to the unknown hyperparameters [28, p. 383]. The joint posterior distribution of the unknown model parameters and hyperparameter is then approximated from samples generated by MCMC methods. This fully Bayesian estimation technique, followed in this paper, has been recently applied to signal segmentation [29] and hyperspectral imaging [30], [31].

Besides, standard MCMC methods have shown some limitations for deconvolving BG processes. More precisely, as noticed in [16] and [32], a standard Gibbs sampler can be stuck in a particular configuration of the BG process to be recovered, leading to poor mixing properties. Ge and Idier recently demonstrated that this BG deconvolution can be easily improved by marginalizing over the amplitudes of the nonzero components [32]. The resulting MCMC scheme is a partially collapsed Gibbs sampler deeply studied by van Dyk *et al.* in [33] and [34]. Following this approach, we propose in this paper to take advantage of this MCMC strategy to estimate the sparse sources efficiently.

To avoid redundant atoms in the dictionary and, more generally, ensure a full rank mixing matrix, we address the blind source separation problem under orthogonality constraint on the mixing matrix. The main motivation for imposing this orthogonality constraint on the mixing matrix is to capture more diversity among the recovered atoms belonging to the undercomplete dictionary to be estimated. Only a few works in the signal processing literature have considered this additional property. Preliminary results on this issue have been reported in [10] that has addressed the problem of sparse source separation from orthogonal mixtures. However, the strategy was based on a strong hypothesis for the source and mixing matrix, i.e., the prior knowledge of the sparsity level shared by all the sources. Hoff recently proposed in [35] a Bayesian formulation of the dimension reduction operators. More precisely, to estimate the rank of an unobserved matrix \mathbf{X} involved in a noisy model, he has derived a Bayesian description of the singular value decomposition (SVD). The idea is to decompose the unobserved noise-free data \mathbf{X} as $\mathbf{X} = \mathbf{U}\mathbf{D}\mathbf{V}^T$, where \mathbf{U} and \mathbf{V} are matrices with orthogonal columns. The Bayesian inference on \mathbf{U} and \mathbf{V} has been finally conducted after assigning uniform prior distributions for \mathbf{U} and \mathbf{V} on their definition space, called the Stiefel manifold. This choice, coupled with the Gaussian properties of the noise, leads to von Mises-Fisher conditional posterior distributions for the columns of the matrices \mathbf{U} and \mathbf{V} . In the Bayesian orthogonal component analysis (BOCA) studied in this paper, a similar strategy is adopted by assigning a uniform distribution on the Stiefel manifold to the mixing matrix. The resulting MCMC algorithm generates mixing matrix samples

distributed according to the posterior distribution following the efficient scheme developed in [35].

This paper is organized as follows. The BOCA is formulated as a blind source separation problem under constraints in Section II. Section III derives the statistical quantities required to define the Bayesian model. The BOCA MCMC algorithm is described step-by-step in Section IV. This algorithm allows one to generate samples distributed according to the joint posterior distribution of the source and mixing matrices. Simulation results conducted on synthetic data, as well as a performance comparison with the K-SVD algorithm, are reported in Section V. Section VI illustrates the interest of the proposed algorithm by solving a sparse coding problem with an application to natural image processing. Conclusions and potential future works are considered in Section VII.

II. PROBLEM FORMULATION

Let $\mathbf{x}(t) = [x_1(t), \dots, x_M(t)]^T$ denote measurement vectors of \mathbb{R}^M observed at time instants $t = 1, \dots, T$ by M sensors. These observations are assumed to be related to $N < M$ unobserved sources denoted $\mathbf{s}(t) = [s_1(t), \dots, s_N(t)]^T$ via the matrix Ψ in the following noisy linear model:

$$\mathbf{x}(t) = \Psi \mathbf{s}(t) + \mathbf{n}(t) \quad (1)$$

where $\mathbf{n}(t)$ stands for an additive measurement noise. Standard matrix notations yield

$$\mathbf{X} = \Psi \mathbf{S} + \mathbf{N} \quad (2)$$

with $\mathbf{X} = [\mathbf{x}(1), \dots, \mathbf{x}(T)]$, $\mathbf{S} = [\mathbf{s}(1), \dots, \mathbf{s}(T)]$ and $\mathbf{N} = [\mathbf{n}(1), \dots, \mathbf{n}(T)]$. The $M \times 1$ noise vectors $\mathbf{n}(t)$ ($t = 1, \dots, T$) are assumed to be independent and distributed according to a centered multivariate Gaussian distribution $\mathcal{N}(\mathbf{0}_M, \sigma^2 \mathbf{I}_M)$.

In this work, the $M \times N$ matrix Ψ is assumed to be an unknown orthogonal matrix

$$\psi_i^T \psi_j = \begin{cases} 1, & \text{if } i = j \\ 0, & \text{if } i \neq j \end{cases} \quad (3)$$

which defines an orthogonal basis associated with a space where the sources to be recovered can be sparsely represented. Consequently, since only a few sources are assumed to be active at time index t , the unobserved vector of N sources $\mathbf{s}(t)$ is sparse and contains only a few components that are nonzero.

This paper proposes a Bayesian model as well as an MCMC sampling strategy to estimate the unknown sources \mathbf{S} , the orthogonal matrix Ψ and the noise variance σ^2 .

III. BAYESIAN MODEL

The unknown parameter vector associated with the mixing model defined in (1) is $\theta = \{\mathbf{S}, \Psi, \sigma^2\}$. This section gives the likelihood function of the observations and introduces prior distributions for the unknown model parameters (assumed to be *a priori* independent).

A. Likelihood Function

The Gaussian property of the additive noise yields for each observed vector $\mathbf{x}(t)$

$$f(\mathbf{x}(t) | \Psi, \mathbf{s}(t), \sigma^2) = \left(\frac{1}{2\pi\sigma^2} \right)^{\frac{M}{2}} \times \exp \left[-\frac{1}{2\sigma^2} \|\mathbf{x}(t) - \Psi \mathbf{s}(t)\|^2 \right]$$

where $t = 1, \dots, T$ and $\|\cdot\|$ stands for the standard ℓ_2 -norm. By assuming the noise vectors $\mathbf{n}(t)$ to be *a priori* independent, the full likelihood function is

$$f(\mathbf{X} | \Psi, \mathbf{S}, \sigma^2) = \left(\frac{1}{2\pi\sigma^2} \right)^{\frac{TM}{2}} \times \exp \left[-\frac{1}{2\sigma^2} \sum_{t=1}^T \|\mathbf{x}(t) - \Psi \mathbf{s}(t)\|^2 \right]. \quad (4)$$

B. Noise Variance Prior

As in numerous works, including [29], [30], [36], a conjugate inverse-Gamma distribution is chosen as prior distribution for the noise variance σ^2

$$\sigma^2 | \gamma \sim \mathcal{IG} \left(\frac{\nu}{2}, \frac{\gamma}{2} \right) \quad (5)$$

where $\nu = 2$ and γ is an unknown hyperparameter that will be estimated from the data. The main motivation for choosing conjugate prior distribution for σ^2 is to simplify the computation of the posterior distribution of interest.

C. Prior for the Mixing Matrix

The mixing matrix Ψ to be estimated is an $M \times N$ matrix with orthogonal columns whose rank is N . The set of such matrices, denoted $\mathcal{S}_{N,M}$, is called the Stiefel manifold² (see [37, p. 8]) for a general introduction of this space). To reflect the absence of any additional prior knowledge regarding the mixing matrix, a uniform distribution on this set is chosen as prior distribution for Ψ [38, p. 279]

$$f(\Psi) = \frac{1}{\text{vol}(\mathcal{S}_{N,M})} \mathbf{1}_{\mathcal{S}_{N,M}}(\Psi) \quad (6)$$

where $\mathbf{1}(\cdot)$ stands for the indicator function

$$\mathbf{1}_{\mathcal{S}_{N,M}}(\Psi) = \begin{cases} 1, & \text{if } \Psi \in \mathcal{S}_{N,M} \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

In (6), $\text{vol}(\mathcal{S}_{N,M})$ is the volume of the Stiefel manifold $\mathcal{S}_{N,M}$ given by [39, p. 70]

$$\text{vol}(\mathcal{S}_{N,M}) = \frac{2^M \pi^{\frac{NM}{2}}}{\Gamma_M \left(\frac{N}{2} \right)} \quad (8)$$

where $\Gamma_M(\cdot)$ is the M -variate Gamma function

$$\Gamma_M(u) = \pi^{\frac{M(M-1)}{4}} \prod_{m=1}^M \Gamma \left(u + \frac{1-m}{2} \right) \quad (9)$$

and $\Gamma(\cdot)$ is the Gamma function

$$\Gamma(u) = \int_0^{+\infty} t^{u-1} e^{-t} dt \quad (10)$$

²Note that for the special case $M = N$, the Stiefel manifold $\mathcal{S}_{M,M}$ is the orthogonal group $\mathcal{O}(M)$ of orthogonal $M \times M$ matrices.

with $u > 0$.

Generating samples according to (6) can be easily achieved by first sampling an $M \times N$ matrix \mathbf{V} of independent standard normal random variables and then by setting $\Psi = \mathbf{V}(\mathbf{V}^T \mathbf{V})^{-(1/2)}$ [37]. However, as highlighted by Hoff in [35], sampling Ψ via its conditional distributions is frequently required, especially within an MCMC estimation framework. Therefore, we recall below the procedure proposed in [35] to sample orthogonal matrices Ψ according to the uniform distribution (6) using the conditional distributions of its columns.

First, let $\Psi_{\mathcal{A}} = [\psi_i]_{i \in \mathcal{A}}$ denote the matrix formed by the columns of Ψ indexed by the label vector $\mathcal{A} \subset \{1, \dots, N\}$, where ψ_i stands for the i th column of Ψ . Let $\mathbf{N}_{\mathcal{A}}$ denote an orthogonal basis associated with the null space of the orthogonal matrix $\Psi_{\mathcal{A}}$. Then, as demonstrated in [35], an orthogonal $M \times N$ matrix Ψ can be uniformly drawn on the Stiefel manifold $\mathcal{S}_{N,M}$ via the following steps:

- 1) sample \mathbf{v}_1 uniformly on the unit M -sphere and set $\psi_1 = \mathbf{v}_1$;
- 2) sample \mathbf{v}_2 uniformly on the unit $(M-1)$ -sphere and set $\psi_2 = \mathbf{N}_1 \mathbf{v}_2$;
- 3) sample \mathbf{v}_3 uniformly on the unit $(M-2)$ -sphere and set $\psi_3 = \mathbf{N}_{\{1,2\}} \mathbf{v}_3$;
- \vdots
- N) sample \mathbf{v}_N uniformly on the unit $(M-N+1)$ -sphere and set $\psi_N = \mathbf{N}_{\{1,\dots,N-1\}} \mathbf{v}_N$.

Uniform sampling on a sphere required in the scheme detailed above can be easily achieved following the normal-deviate method described in [40]. Finally, we have to mention that a similar strategy will be used in Section IV to sample mixing matrices Ψ according to their conditional posterior distributions.

D. Source Prior

Since the source vectors $\mathbf{s}(t)$ are sparse, most of the elements $s_n(t)$ ($n = 1, \dots, N, t = 1, \dots, T$) of the matrix \mathbf{S} are expected to be equal to zero. Therefore, choosing a ‘‘sparse’’ prior for $s_n(t)$ is recommended. Coupling a standard probability density function (pdf) with an atom at zero is a classical strategy to ensure sparsity. This strategy has been widely used for located event detection [20] such as spike train deconvolution [17], [41], astrophysical frequency detection [16], sparse approximations of times-series [42] and reconstruction of molecular images [19]. We propose here to take advantage of this approach by choosing a BG distribution as prior for $s_n(t)$. The distribution of this BG prior is defined as the following mixture:

$$f(s_n(t) | \lambda_n, a_n^2) = (1 - \lambda_n) \delta(s_n(t)) + \lambda_n g_{a_n^2}(s_n(t)) \quad (11)$$

where $\delta(\cdot)$ is the Dirac delta function and $g_{a_n^2}(s_n(t))$ is the pdf of the centered Gaussian distribution with variance a_n^2 . In (11), the unknown hyperparameter λ_n is the prior probability of having an active source. Consequently, this hyperparameter tunes the degree of sparseness of the source vector $\mathbf{s}(t)$. Note that this probability λ_n of having an active source, as well as the nonzero component variance a_n^2 , have been assumed to be different from one source to another to provide a flexible model. This strategy have been previously adopted in [24], [42], and [43]. Another strategy would be to assume that the

sources have a nonhomogeneous sparsity level over times as explained in [44]. By assuming that the source amplitudes $s_n(t)$ are *a priori* independent, and introducing the index subsets $\mathcal{I}_n(\epsilon) = \{t; s_n(t) = \epsilon\}$ ($\epsilon \in \{0, 1\}$), the full prior distribution for the source matrix \mathbf{S} is

$$\begin{aligned} f(\mathbf{S} | \boldsymbol{\lambda}, \mathbf{a}^2) &= \prod_{n=1}^N \left[(1 - \lambda_n)^{m_n(0)} \prod_{t \in \mathcal{I}_n(0)} \delta(s_n(t)) \right] \\ &\times \prod_{n=1}^N \left[\lambda_n^{m_n(1)} \prod_{t \in \mathcal{I}_n(1)} \left(\frac{1}{2\pi a_n^2} \right)^{\frac{1}{2}} \exp \left(-\frac{s_n(t)^2}{a_n^2} \right) \right] \end{aligned}$$

where $\boldsymbol{\lambda} = [\lambda_1, \dots, \lambda_N]^T$, $\mathbf{a}^2 = [a_1^2, \dots, a_N^2]^T$ and $m_n(\epsilon) = \text{card}\{\mathcal{I}_n(\epsilon)\}$. Note that $m_n(1) = \|\mathbf{s}_n^T\|_0$, where $\mathbf{s}_n = [s_n(1), \dots, s_n(T)]$ and $\|\cdot\|_0$ denotes the ℓ_0 -norm, is the number of active components in the source n , whereas $m_n(0) = T - m_n(1)$ is the number of components that are equal to 0 in the source n .

E. Hyperparameter Priors

A noninformative Jeffreys' prior is elected as prior distribution for the hyperparameter γ

$$f(\gamma) \propto \frac{1}{\gamma} \mathbf{1}_{\mathbb{R}^+}(\gamma). \quad (12)$$

An inverse-gamma distribution with fixed hyper-hyperparameters α_0 and α_1 (in order to obtain vague prior with large variance) is chosen as prior for the variance a_n^2 of the nonzero components in each source

$$a_n^2 \sim \mathcal{IG}(\alpha_0, \alpha_1). \quad (13)$$

A uniform distribution is chosen as prior distribution for the λ_n in each source

$$\lambda_n \sim \mathcal{U}([0, 1]). \quad (14)$$

Assuming that the individual hyperparameters are independent the full prior distribution for the hyperparameter vector $\boldsymbol{\phi} = \{\gamma, \boldsymbol{\lambda}, \mathbf{a}^2\}$ can be expressed as

$$\begin{aligned} f(\boldsymbol{\phi}) &\propto \frac{1}{\gamma} \mathbf{1}_{\mathbb{R}^+}(\gamma) \\ &\times \prod_{n=1}^N \left[\left(\frac{1}{a_n^2} \right)^{\alpha_0+1} \exp \left(-\frac{\alpha_1}{a_n^2} \right) \mathbf{1}_{[0,1]}(\lambda_n) \right]. \quad (15) \end{aligned}$$

F. Posterior Distribution

The posterior distribution of $\{\boldsymbol{\theta}, \boldsymbol{\phi}\}$ can be computed from the following hierarchical structure

$$f(\boldsymbol{\theta}, \boldsymbol{\phi} | \mathbf{X}) \propto f(\mathbf{X} | \boldsymbol{\theta}) f(\boldsymbol{\theta} | \boldsymbol{\phi}) f(\boldsymbol{\phi}) \quad (16)$$

where \propto means "proportional to"

$$f(\boldsymbol{\theta} | \boldsymbol{\phi}) = f(\mathbf{S} | \boldsymbol{\lambda}, \mathbf{a}^2) f(\boldsymbol{\Psi}) f(\sigma^2 | \gamma) \quad (17)$$

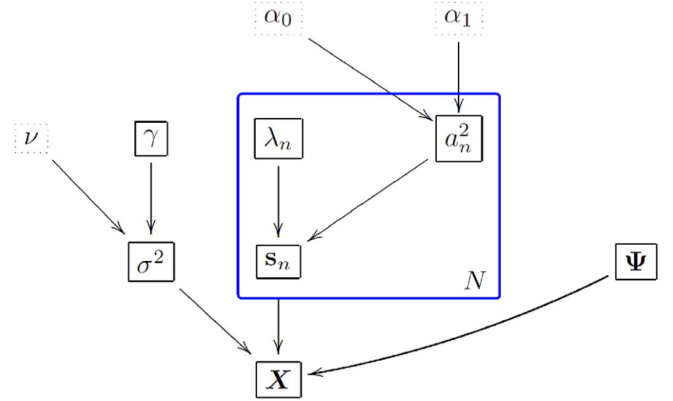


Fig. 1. DAG for the parameter priors and hyperpriors (the fixed hyperparameters appear in dashed boxes).

and where $f(\mathbf{X} | \boldsymbol{\theta})$ and $f(\boldsymbol{\phi})$ have been defined in (4) and (15). This hierarchical structure is represented as a graphical model on the directed acyclic graph (DAG) of Fig. 1. In the joint distribution (16), the nuisance parameter γ can be easily integrated out, leading to

$$\begin{aligned} f(\mathbf{S}, \boldsymbol{\Psi}, \sigma^2, \boldsymbol{\lambda}, \mathbf{a}^2 | \mathbf{X}) &\propto \mathbf{1}_{\mathcal{S}_{N,M}}(\boldsymbol{\Psi}) \left(\frac{1}{\sigma^2} \right)^{\frac{TM+\nu}{2}} \\ &\times \exp \left[-\frac{1}{2\sigma^2} \sum_{t=1}^T \|\mathbf{x}(t) - \boldsymbol{\Psi} \mathbf{s}(t)\|^2 \right] \\ &\times \prod_{n=1}^N \left[(1 - \lambda_n)^{m_n(0)} \prod_{t \in \mathcal{I}_n(0)} \delta(s_n(t)) \right] \\ &\times \prod_{n=1}^N \left[\lambda_n^{m_n(1)} \prod_{t \in \mathcal{I}_n(1)} \left(\frac{1}{2\pi a_n^2} \right)^{\frac{1}{2}} \exp \left(-\frac{s_n(t)^2}{a_n^2} \right) \right] \\ &\times \prod_{n=1}^N \left[\left(\frac{1}{a_n^2} \right)^{\alpha_0+1} \exp \left(-\frac{\alpha_1}{a_n^2} \right) \mathbf{1}_{[0,1]}(\lambda_n) \right]. \quad (18) \end{aligned}$$

Inferring the source matrix \mathbf{S} and the orthogonal matrix $\boldsymbol{\Psi}$ from (18) is not straightforward, mainly due to the combinatory problem induced by the quantities $n_1(t)$ and $n_0(t)$. In particular, closed-form expressions of the Bayesian estimators of \mathbf{S} and $\boldsymbol{\Psi}$ are difficult to obtain. We propose to use MCMC methods to generate samples that are asymptotically distributed according to the target distribution (18). These generated samples are then used to approximate the Bayesian estimators of \mathbf{S} and $\boldsymbol{\Psi}$.

IV. PARTIALLY COLLAPSED GIBBS SAMPLER FOR ORTHOGONAL COMPONENT ANALYSIS OF SPARSE SOURCES

We describe in this section an MCMC method that allows one to generate a sample collection

$$\mathcal{Y} = \left\{ \left(\tilde{\mathbf{S}}^{(h)}, \tilde{\boldsymbol{\Psi}}^{(h)}, \tilde{\sigma}^{2(h)}, \tilde{\boldsymbol{\lambda}}^{(h)}, \tilde{\mathbf{a}}^{2(h)} \right) \right\}_{h=1, \dots, N_{\text{MC}}}$$

asymptotically distributed according to the posterior distribution (18). The interested reader is invited to consult [28] for more details about MCMC methods.

The easiest way to sample according to this posterior would consist of using a standard Gibbs sampler whose main steps are as follows:

- 1) sample \mathbf{S} from $f(\mathbf{S} | \boldsymbol{\Psi}, \sigma^2, \boldsymbol{\lambda}, \mathbf{a}^2, \mathbf{X})$;

- 2) sample Ψ from $f(\Psi | \mathcal{S}, \sigma^2, \lambda, \mathbf{a}^2, \mathbf{X}) = f(\Psi | \mathcal{S}, \sigma^2, \mathbf{X})$;
- 3) sample σ^2 from $f(\sigma^2 | \mathcal{S}, \Psi, \lambda, \mathbf{a}^2, \mathbf{X}) = f(\sigma^2 | \mathcal{S}, \Psi, \mathbf{X})$;
- 4) sample λ from $f(\lambda | \mathcal{S}, \Psi, \sigma^2, \mathbf{a}^2, \mathbf{X}) = f(\lambda | \mathcal{S})$;
- 5) sample \mathbf{a}^2 from $f(\mathbf{a}^2 | \mathcal{S}, \Psi, \sigma^2, \lambda, \mathbf{X}) = f(\mathbf{a}^2 | \mathcal{S})$.

However, as highlighted in previous works, sampling BG processes following the crude Gibbs sampler detailed above [step 1)] often leads to poor mixing properties and weak estimation performance [16]. As an alternative a new MCMC algorithm for BG deconvolution was recently studied [32]. The approach relies on explicitly introducing binary variables \mathbf{Q} that indicate the presence of nonzero BG components. Then, these indicators are sampled after marginalizing over the BG variable amplitudes. This strategy casts the resulting MCMC algorithm as a partially collapsed Gibbs (PCG) sampler. Van Dyk and Park have described in [34] and [33] how PCG samplers can be efficient tools to overcome drawbacks inherent to standard Gibbs sampler, e.g., slow convergence. As detailed in the works cited above, PCG samplers consists of replacing some of the conditional distributions with marginalized conditional distribution. The resulting PCG sampling scheme can be summarized by the following steps:

- 1) sample \mathbf{Q} from $f(\mathbf{Q} | \Psi, \sigma^2, \lambda, \mathbf{a}^2, \mathbf{X})$;
- 2) sample \mathcal{S} from $f(\mathcal{S} | \mathbf{Q}, \Psi, \sigma^2, \lambda, \mathbf{a}^2, \mathbf{X})$;
- 3) sample Ψ from $f(\Psi | \mathcal{S}, \sigma^2, \mathbf{X})$;
- 4) sample σ^2 from $f(\sigma^2 | \mathcal{S}, \Psi, \mathbf{X})$;
- 5) sample λ from $f(\lambda | \mathcal{S})$;
- 6) sample \mathbf{a}^2 from $f(\mathbf{a}^2 | \mathcal{S})$.

Note that the source amplitudes have been marginalized to provide the discrete distribution appearing in step 1). The main steps of the PCG sampler are detailed in Subsections IV-A–IV-E (see also the step-by-step Algorithm 1).

Algorithm 1: Gibbs Sampling for Orthogonal Component Analysis of Sparse Sources

```

1: % sampling the sources
2: for t = 1 to T do
3:   % sampling the indicators recursively following
   Algorithm 2
4:   for n = 1 to N do
5:     sample the indicator  $q_n(t)$  following the
     probability (21),
6:   end for
7:   sample the source vector  $\mathbf{s}(t)$  from the pdf's in (22)
   and (23),
8: end for
9: % sampling the orthogonal mixing matrix
10: for n = 1 to N do
11:   compute the basis  $\mathbf{N}_n$  of the null space of  $\Psi_{-n}$ ,
12:   sample  $\mathbf{v}_n$  from the von Mises-Fisher distribution
   in (26),
13:   set  $\boldsymbol{\psi}_n = \mathbf{N}_n \mathbf{v}_n$ ,
14: end for
15: % sampling the noise variance
16: sample parameter  $\sigma^2$  from the pdf in (27),
17: % sampling the probability of having active sources
18: for n = 1 to N do
19:   sample the hyperparameter  $\lambda_n$  from the pdf in (29),
20:   % sampling the active source variances
21:   sample the hyperparameter  $a_n^2$  from the pdf in (30),
22: end for
    
```

A. Sampling the Indicator and Source Matrices

The prior independence assumption of the source vectors allows one to rewrite the joint posterior distribution of the source matrix \mathcal{S} as

$$f(\mathcal{S} | \lambda, \mathbf{a}^2, \sigma^2, \Psi, \mathbf{X}) = \prod_{t=1}^T f(\mathbf{s}(t) | \lambda, \mathbf{a}^2, \sigma^2, \Psi, \mathbf{x}(t)). \quad (19)$$

Consequently, sampling according to $f(\mathcal{S} | \lambda, \mathbf{a}^2, \sigma^2, \Psi, \mathbf{X})$ can be achieved by successively sampling the source vectors $\mathbf{s}(t)$ according to $f(\mathbf{s}(t) | \lambda, \mathbf{a}^2, \sigma^2, \Psi, \mathbf{x}(t))$ for $t = 1, \dots, T$.

It is important to note here that standard derivations similar to those in [17] allow one to state that the conditional posterior distribution of the t th component $s_n(t)$ of the n th source is a BG distribution. However, as pointed out in [16], sampling according to this distribution needs to explore the state space efficiently, which can be difficult mainly due to the difficulty of the Gibbs sampler to escape from local maxima. Recently, Ge *et al.* have introduced a performing MCMC algorithm to overcome this issue by explicitly introducing an auxiliary binary variable $q_n(t)$ that indicates the active sources [32]

$$q_n(t) = \begin{cases} 1, & \text{if } s_n(t) \neq 0, \\ 0, & \text{otherwise.} \end{cases} \quad (20)$$

Conditionally upon this indicator variable $q_n(t)$, the prior of the source component $s_n(t)$ in (11) can be easily rewritten

$$\begin{aligned} f(s_n(t) | q_n(t) = 0) &= \delta(s_n(t)), \\ f(s_n(t) | q_n(t) = 1, a_n^2) &= g_{a_n^2}(s_n(t)). \end{aligned}$$

The probability of having an active component in source n is governed by an unknown hyperparameter λ_n such that

$$\begin{aligned} \mathbb{P}[q_n(t) = 1] &= \lambda_n, \\ \mathbb{P}[q_n(t) = 0] &= 1 - \lambda_n. \end{aligned}$$

In [32], Ge *et al.* have proposed to sample the source vectors $\mathbf{s}(t)$ ($t = 1, \dots, T$) and the indicators $\mathbf{q}(t) = [q_1(t), \dots, q_N(t)]^T$ using the following two steps:

- 1) sampling according to $f(\mathbf{q}(t) | \lambda, \mathbf{a}^2, \sigma^2, \Psi, \mathbf{x}(t))$;
- 2) sampling according to $f(\mathbf{s}(t) | \mathbf{q}(t), \lambda, \mathbf{a}^2, \sigma^2, \Psi, \mathbf{x}(t))$.

As mentioned above, these two steps make the resulting Gibbs sampler a PCG sampler and are detailed below.

1) *Sampling According to $f(\mathbf{q}(t) | \lambda, \mathbf{a}^2, \sigma^2, \Psi, \mathbf{x}(t))$* : Sampling according to $f(\mathbf{q}(t) | \lambda, \mathbf{a}^2, \sigma^2, \Psi, \mathbf{x}(t))$ can be performed by updating the N components $q_n(t)$ ($n = 1, \dots, N$) successively. As noticed in [32], the posterior probability of having the source component $q_n(t)$ to be active given the other components denoted $\mathbf{q}_{-n}(t) = [q_1(t), \dots, q_{n-1}(t), q_{n+1}(t), \dots, q_N(t)]^T$ is

$$\begin{aligned} \mathbb{P}[q_n(t) = 1 | \mathbf{q}_{-n}(t), \lambda_n, a_n^2, \Psi, \sigma^2, \mathbf{x}(t)] \\ = \left[1 + \exp\left(-\frac{u_0 - u_1}{2}\right) \right]^{-1} \end{aligned} \quad (21)$$

with

$$u_\epsilon = \mathbf{x}(t)^T \mathbf{B}_\epsilon^{-1} \mathbf{x}(t) + \log |\mathbf{B}_\epsilon| + 2\epsilon \log \left(\frac{1}{\lambda_n} - 1 \right)$$

$$\mathbf{B}_\epsilon = a_n^2 \Psi \text{diag} \left\{ \mathbf{q}_{[\epsilon]}(t) \right\} \Psi^T + \sigma^2 \mathbf{I}_M$$

$$\mathbf{q}_{[\epsilon]}(t) = [q_1(t), \dots, q_{n-1}(t), \epsilon, q_{n+1}(t), \dots, q_N(t)]^T.$$

The probability in (21) can be efficiently computed following the recursive scheme initially introduced in [41] (and used in [32]), and adapted here to take into account the orthogonality property of Ψ . This numerical implementation relies on the Cholesky decomposition of B_ϵ and the matrix inversion lemma. This avoids to calculate the compute-intensive inversion of B_ϵ and the determinant $|B_\epsilon|$ at each step of the Gibbs sampler. We describe in Algorithm 2 how the component $q_n(t)$ is updated.

Algorithm 2: Recursive Sampling of Indicator Vector $\mathbf{q}(t)$

```

1: for  $n = 1$  to  $N$  do
2:   set  $\delta_n = (-1)^{q_n(t)}$ ,
3:   set  $\mathbf{G} = [\boldsymbol{\psi}_i]_{q_i(t)=1}$ ,
4:   set  $\mu_n = (a_n^2)/(\sigma^2)$ ,
5:   set  $\tau_n = \delta_n + \mu_n \|\boldsymbol{\psi}_n\|^2 - (\mu_n^2)/(1 + \mu_n) \boldsymbol{\psi}_n^T \mathbf{G} \mathbf{G}^T \boldsymbol{\psi}_n$ ,
6:   set  $\eta_n = \mathbf{x}(t)^T \boldsymbol{\psi}_n - (\mu_n)/(1 + \mu_n) \mathbf{x}(t)^T \mathbf{G} \mathbf{G}^T \boldsymbol{\psi}_n$ ,
7:   set  $\Delta u = \log(\delta_n \tau_n) - (\mu_n)/(\sigma^2 \tau_n) \eta_n^2 + 2\delta_n \log((1/\lambda_n) - 1)$ ,
8:   sample  $w \sim \mathcal{U}_{[0,1]}$ ,
9:   if  $u > [1 + \exp(-(\Delta u)/(2))]^{-1}$  then
10:    set  $q_n(t) = q_n(t) + \delta_n$ ,
11:   end if
12: end for

```

2) *Sampling According to $f(\mathbf{s}(t) | \mathbf{q}(t), \mathbf{a}^2, \sigma^2, \Psi, \mathbf{x}(t))$:* Conditionally upon the indicator variable $q_n(t)$, the distribution of $s_n(t)$ is defined by

$$f(s_n(t) | q_n(t) = 0, \sigma^2, \Psi, \mathbf{x}(t)) = \delta(s_n(t)) \quad (22)$$

and

$$[s_n(t)]_{q_n(t)=1} | \mathbf{q}(t), \mathbf{a}^2, \sigma^2, \Psi, \mathbf{x}(t) \sim \mathcal{N}(\Lambda_1 \mathbf{G} \mathbf{x}(t), \Lambda_2) \quad (23)$$

where $[s_n(t)]_{q_n(t)=1}$ stands for $L \times 1$ vector composed of the active components in the source vector $\mathbf{s}(t)$, $L = \|\mathbf{q}(t)\|_0$, $\mathbf{G} = [\boldsymbol{\psi}_n]_{q_n(t)=1}$ is the $M \times L$ matrix composed of the columns of Ψ corresponding to the active source components and $\Lambda_1 = \text{diag}\{\mu_n/(\mu_n + 1)\}_{q_n(t)=1}$ and $\Lambda_2 = \text{diag}\{1/(\mu_n + 1)\}_{q_n(t)=1}$ are $L \times L$ diagonal matrices with $\mu_n = (a_n^2)/(\sigma^2)$. Note that this block sampling strategy, also adopted in [43], avoids to sample the nonzero source components one-by-one.

B. Sampling the Mixing Matrix

The conditional distribution $f(\Psi | \mathbf{S}, \sigma^2, \mathbf{X})$ being intractable, this section describes how Gibbs moves can be used to generate samples $\tilde{\boldsymbol{\psi}}_n^{(h)}$ according to the posterior distribution of each column of Ψ conditionally upon the others. Let $\Psi_{-n} = [\boldsymbol{\psi}_i]_{i \neq n}$ (respectively, $\mathbf{s}_{-n}(t)$) denote the matrix Ψ (respectively, vector $\mathbf{s}(t)$) whose n th column (respectively, component) has been removed. By denoting

$$\boldsymbol{\mu}_n(t) = s_n(t) [\mathbf{x}(t) - \Psi_{-n} \mathbf{s}_{-n}(t)] \quad (24)$$

straightforward computations yield

$$f(\boldsymbol{\psi}_n | \Psi_{-n}, \mathbf{S}, \sigma^2, \mathbf{X}) \propto \exp \left[\frac{1}{\sigma^2} \sum_{t=1}^T \boldsymbol{\mu}_n(t)^T \boldsymbol{\psi}_n \right] \mathbf{1}_{S_{M,N}}(\Psi). \quad (25)$$

As detailed in [35], conditionally upon Ψ_{-n} , $\boldsymbol{\psi}_n$ can be written $\boldsymbol{\psi}_n = \mathbf{N}_n \mathbf{v}_n$ where \mathbf{v}_n is uniform on the sphere and \mathbf{N}_n is

a basis for the null space of Ψ_{-n} . Therefore, from (25), the conditional distribution of \mathbf{v}_n is

$$f(\mathbf{v}_n | \Psi_{-n}, \mathbf{S}, \sigma^2, \mathbf{X}) \propto \exp \left[\frac{1}{\sigma^2} \sum_t \boldsymbol{\mu}_{n,t}^T \mathbf{N}_n \mathbf{v}_n \right] \quad (26)$$

which is a von Mises-Fisher distribution with parameter $(\sum_t \boldsymbol{\mu}_{n,t}^T \mathbf{N}_n) / \sigma^2$. A standard method to sample according to this distribution is given in [45]. To summarize, the columns of Ψ can be iteratively sampled conditionally upon the others by drawing samples \mathbf{v}_n from a von Mises-Fisher distribution and setting $\boldsymbol{\psi}_n = \mathbf{N}_n \mathbf{v}_n$.

C. Sampling the Noise Variance

Looking carefully at (18), the conditional distribution of the noise variance is an inverse-gamma distribution such that

$$\sigma^2 | \mathbf{S}, \Psi, \mathbf{X} \sim \mathcal{IG} \left(\frac{TM}{2}, \frac{1}{2} \sum_{t=1}^T \|\mathbf{x}(t) - \Psi \mathbf{s}(t)\|^2 \right). \quad (27)$$

D. Sampling the Probability of Having an Active Source

The conditional distribution of the hyperparameter λ_n (i.e., the probability of the source n to be active) can be computed from (18)

$$f(\lambda_n | \mathbf{s}_n) \propto (1 - \lambda_n)^{m_n(0)} \lambda_n^{m_n(1)} \quad (28)$$

where $m_n(\epsilon)$, $\epsilon \in \{0, 1\}$, has been defined in Section III-D. Therefore, sampling according to $f(\lambda_n | \mathbf{s}_n)$ is achieved as follows:

$$\lambda_n | \mathbf{s}_n \sim \text{Be}(m_n(0) + 1, m_n(1) + 1). \quad (29)$$

where $\text{Be}(a, b)$ is a Beta distribution with shape parameters a and b .

E. Sampling the Variance of the Active Sources

Straightforward computations leads to the following IG distribution as conditional posterior the variance of nonzero BG components in the source n

$$a_n^2 | \mathbf{s}_n \sim \mathcal{IG} \left(\frac{1}{2} m_n(1) + \alpha_0, \frac{1}{2} \|\mathbf{s}_n^T\|^2 + \alpha_1 \right). \quad (30)$$

F. Inferring the Sources and the Mixing Matrix

The main objective of the proposed Bayesian algorithm is to estimate the source matrix \mathbf{S} and the mixing matrix Ψ from the data, independently from the nuisance parameters σ^2 , λ_n and a_n^2 . The MCMC algorithm detailed in Sections IV-A–IV-E generates samples asymptotically distributed according to the posterior distribution (18). Consequently, the MMSE estimators of \mathbf{S} and Ψ can be approximated by empirical averages over the $N_{\text{MC}} - N_{\text{bi}}$ drawn samples as follows:

$$\hat{\mathbf{S}}_{\text{MMSE}} = \mathbb{E}[\mathbf{S} | \mathbf{X}] \approx \frac{1}{N_{\text{MC}} - N_{\text{bi}}} \sum_{h=N_{\text{bi}}+1}^{N_{\text{MC}}} \tilde{\mathbf{S}}^{(h)} \quad (31)$$

$$\hat{\Psi}_{\text{MMSE}} = \mathbb{E}[\Psi | \mathbf{X}] \approx \frac{1}{N_{\text{MC}} - N_{\text{bi}}} \sum_{h=N_{\text{bi}}+1}^{N_{\text{MC}}} \tilde{\Psi}^{(h)} \quad (32)$$

where N_{bi} denotes the number of burn-in iterations of the sampler and N_{MC} is the total number of Monte Carlo iterations.

Another important property of the proposed MCMC algorithm is that the generated pairs

$$\left\{ \left(\tilde{\mathbf{S}}^{(1)}, \tilde{\Psi}^{(1)} \right), \dots, \left(\tilde{\mathbf{S}}^{(N_{\text{MC}})}, \tilde{\Psi}^{(N_{\text{MC}})} \right) \right\}$$

form also a Markov chain whose stationary distribution is the joint marginal distribution $f(\mathbf{S}, \Psi | \mathbf{X})$. Therefore, the joint MAP estimator of (\mathbf{S}, Ψ) can be computed by retaining among the collection

$$\mathcal{X} = \left\{ \left(\tilde{\mathbf{S}}^{(h)}, \tilde{\Psi}^{(h)} \right) \right\}_{h=N_{\text{bi}}, \dots, N_{\text{MC}}}$$

the sample that maximizes the marginalized distribution $f(\mathbf{S}, \Psi | \mathbf{X})$ [46, p. 165]

$$\begin{aligned} (\hat{\mathbf{S}}_{\text{MAP}}, \hat{\Psi}_{\text{MAP}}) &= \underset{(\mathbf{S}, \Psi) \in \mathbb{R}^{N \times T} \times \mathcal{S}_{M, N}}{\operatorname{argmax}} f(\mathbf{S}, \Psi | \mathbf{X}) \\ &\approx \underset{(\mathbf{S}, \Psi) \in \mathcal{X}}{\operatorname{argmax}} f(\mathbf{S}, \Psi | \mathbf{X}). \end{aligned} \quad (33)$$

Note that this joint marginal distribution $f(\mathbf{S}, \Psi | \mathbf{X})$ can be easily computed from the hierarchical structure (16) that allows one to integrate out the hyperparameter vector ϕ and the noise variance σ^2 in the full posterior distribution (18), yielding

$$\begin{aligned} f(\mathbf{S}, \Psi | \mathbf{X}) &\propto \frac{\prod_n B(1 + m_n(1), 1 + m_n(0))}{\left[\sum_{t=1}^T \|\mathbf{x}(t) - \Psi \mathbf{s}(t)\|^2 \right]^{\frac{T M}{2}}} \\ &\times \prod_n \frac{\Gamma\left(\alpha_0 + \frac{m_n(1)}{2}\right)}{\left[\alpha_1 + \frac{\|\mathbf{s}_n^T\|^2}{2} \right]^{\alpha_0 + \frac{m_n(1)}{2}}} \end{aligned} \quad (34)$$

where $B(a, b) = (\Gamma(a)\Gamma(b))/(\Gamma(a + b))$.

V. SIMULATION RESULTS

A. Performance Analysis

This section first considers a toy example to provide comprehensive and extensive results. We generate $N = 2$ sources of length $T = 100$ according to the prior distribution in (11) with the probabilities of having active sources $\lambda_1 = 0.05$, $\lambda_2 = 0.1$ and the active source variances $a_1^2 = 100$ and $a_2^2 = 10$. These two sources, represented in Fig. 2 (red), are mixed to obtain $T = 100$ observation vectors $\mathbf{x}(t) \in \mathbb{R}^{50}$ of dimension $M = 50$.

The generated orthogonal mixing matrix Ψ is composed of $N = 2$ basis vectors proportional to sinusoidal functions, $\psi_{mn} \propto \cos(2\pi f_n m + \pi)$ ($m = 1, \dots, M$) with two different frequencies $f_1 = 0.02$ and $f_2 = 0.04$. These two vectors are represented in Fig. 2 (red).

The $T = 100$ observation vectors are corrupted by an additive white Gaussian noise with variance $\sigma^2 = 1.3 \times 10^{-3}$, corresponding to a signal-to-noise ratio (SNR) $\text{SNR}_{\text{dB}} = 15$ dB, where

$$\text{SNR}_{\text{dB}} = 10 \log_{10} \left(\frac{\sum_{t=1}^T \|\Psi \mathbf{s}(t)\|^2}{M T \sigma^2} \right). \quad (35)$$

The proposed Gibbs algorithm is applied on these noisy observations with $N_{\text{MC}} = 1000$ iterations including $N_{\text{bi}} = 100$ burn-in iterations. Note that these numbers of iterations have been chosen to ensure convergence of the Markov chains. More




Fig. 2. Actual sources (circles, red) and corresponding MAP estimates (stars, blue).

precisely, as a first convergence assessment, the outputs of the Markov chains have been monitored for different parameters of interest. As examples, the source parameters λ_n and a_n^2 ($n = 1, 2$) generated by the proposed BOCA algorithm have been depicted as functions of the iteration number in Fig. 4 (top and middle, respectively). Note that the generated values converge towards the actual values of the corresponding parameters after very few iterations. As an additional convergence criterion, the reconstruction error

$$e(h) = \sqrt{\sum_{t=1}^T \|\mathbf{x}(t) - \hat{\Psi}^{(h)} \hat{\mathbf{s}}^{(h)}(t)\|^2} \quad (36)$$

has been evaluated as a function of the iteration number h ($h = N_{\text{bi}} + 1, \dots, N_{\text{MC}}$) where $\hat{\Psi}^{(h)}$ and $\hat{\mathbf{s}}^{(h)}(t)$ are the MMSE estimates of Ψ and $\mathbf{s}(t)$ computed following (31) after h iterations, respectively. The results depicted in Fig. 4 (bottom) show that $N_{\text{MC}} = 1000$ iterations are sufficient to ensure a small reconstruction error. The computation time required by 1000 MCMC iterations is 17 s for an unoptimized MATLAB 2007b 32-bit implementation on a 2.2-GHz Intel Core 2. Of course, for more challenging problems, more MCMC iterations may be required.

The obtained joint MAP estimates of the sources and mixing matrices are represented in Fig. 2 (stars, blue) and Fig. 3 (blue), respectively. These results show that the proposed BOCA algorithm allows one to estimate the sources and basis vectors for this simple example. Note that the first active component in the second source signal has not been detected due to its very low amplitude.

The proposed algorithm implicitly generates binary variables $q_n(t)$ that indicate the presence/absence of nonzero source components (see Section IV-A). Thus, these indicators can be used to compute interesting statistics regarding the probability of having active components. As an example, the number K_n ($n = 1, \dots, N$) of active components in a given source $\mathbf{s}(n)$ can be estimated by

$$K_n = \sum_{t=1}^T q_n(t). \quad (37)$$




Fig. 3. Actual basis vectors (red) and corresponding MAP estimates (blue).




Fig. 4. Top (respectively, middle): Source parameters λ_n (respectively, a_n^2) generated by the proposed Gibbs sampler for the first (blue) and second (green) sources. Bottom: Error reconstruction as a function of iteration number.




Fig. 5. Estimated histograms of numbers of active components in the sources. The actual numbers appear in red dotted line.

The posterior probabilities of \hat{K}_1 and \hat{K}_2 estimated by the proposed method for the considered synthetic example are represented as histograms in Fig. 5. The actual numbers of active




Fig. 6. Posterior probabilities of having active components. The actual active nonzero components appear in red dotted line.

components $K_1 = 5$ and $K_2 = 12$ are represented by red dotted lines in these figures.

The binary variables $q_n(t)$ have Bernoulli distributions. Therefore, the MMSE estimator of $q_n(t)$ provides the posterior probability of $s_n(t)$ to be active. Following (31), the MMSE estimates of the indicators $q_n(t)$ are computed and represented in Fig. 6. These posterior probabilities are in good agreement with the actual positions of the active components, represented by red dotted lines in these figures. Note that these probabilities allow one to locate the first active component in the second source signal that has been previously omitted by the MAP estimator.

B. Performance Comparison

We propose here to compare the proposed algorithm with an up-to-date dictionary learning technique referred to as K-SVD algorithm [9]. This algorithm has been widely applied for various signal and image processing problems and has demonstrated promising results (see, for example, [47], [48], and [49]). For fixed signal dimensions $M = 128$ and $T = 256$, an $M \times T$ sparse matrix \mathbf{S} is randomly drawn according to the prior in (11) with $\lambda_1 = \dots = \lambda_N = 0.05$ and $a_1^2 = \dots = a_N^2 = 10$. Then, an $M \times N$ random orthogonal matrix Ψ is selected as the first N columns of the left orthogonal matrix provided by the singular value decomposition (SVD) of an $M \times M$ matrix whose elements have been drawn according to a $\mathcal{N}(0, 1)$ distribution. The number of dictionary atoms N is set to three different values: $N = 4, 8,$ and 16 . The observations, computed following (2), are corrupted by an additive Gaussian noise with SNR ranging from 0 to 20 dB. The proposed Bayesian algorithm is applied on the generated data with $N_{MC} = 300$ Monte Carlo iterations and $N_{bi} = 50$ burn-in iterations. The accuracy of the MAP estimates of the source and mixing matrices $\hat{\mathbf{S}}_{MAP}$ and $\hat{\Psi}_{MAP}$ are compared with the results provided by the K-SVD algorithm with a total number of 80 iterations (as in [9]). The performance of these algorithms are expressed in terms of reconstruction error and sparsity. More precisely, the root mean-square error (RMSE) between the actual noise-free




Fig. 7. RMSEs as functions of the noise level SNR for different values of the number of sources N .

data $\Psi\mathbf{S}$ and the estimated reconstructed data $\hat{\Psi}\hat{\mathbf{S}}$ is computed as

$$\text{RMSE} = \sqrt{\frac{1}{MT} \sum_{t=1}^T \|\Psi\mathbf{s}(t) - \hat{\Psi}\hat{\mathbf{s}}(t)\|^2}. \quad (38)$$

The sparsity of the estimated source vectors is measured using the following score function \tilde{K} :

$$\tilde{K} = 1 - \frac{1}{NT} \sum_{n=1}^N \|\hat{\mathbf{s}}(t)\|_0 = 1 - \frac{1}{NT} \sum_{n=1}^N \sum_{t=1}^T \hat{q}_n(t). \quad (39)$$

Note that $\tilde{K} \in [0, 1]$ where $\tilde{K} = 1$ means that $\hat{\mathbf{S}}$ is the $N \times T$ matrix of zeros whereas $\tilde{K} = 0$ means that $\hat{\mathbf{S}}$ is a matrix containing NT nonzero elements. The RMSEs and score functions \tilde{K} , computed over 100 Monte Carlo trials, are depicted in Figs. 7 and 8 as functions of the noise level SNR and for different values of the number of sources N .

Fig. 7 indicates that the Bayesian orthogonal component analysis (BOCA) generally outperforms the K-SVD algorithm in term of reconstruction RMSE. In other words, the proposed strategy provides a combination of source and mixing matrices $\hat{\mathbf{S}}$ and $\hat{\Psi}$ that better fit the observed data than the solution provided by K-SVD, especially at low SNR. Moreover, Fig. 8 shows that when the number of sources and the SNR are low, the sources estimated by BOCA are much sparser than the sources identified by K-SVD. Note that the sparsity level of the solutions provided by K-SVD is implicitly fixed by the number of nonzero entries in the orthogonal-matching-pursuit (OMP) subprocedure, whereas the introduced BOCA estimates this sparsity degree via an unsupervised framework.

VI. APPLICATION TO SPARSE CODING WITH UNDERCOMPLETE ORTHOGONAL DICTIONARIES

In this section, we present the ability of the proposed procedure to perform sparse coding with undercomplete orthogonal dictionaries. A fraction of the well-known Barbara natural image is analyzed by BOCA. This 256×256 -pixel image, depicted in Fig. 9 (column 1), is decomposed into $T = 16^2$ block




Fig. 8. Sparsity as function of the noise level SNR for different values of the number of sources N .

patches of size $M = 16 \times 16$ pixels. The proposed Bayesian strategy and the K-SVD algorithm are applied on these observations for different values of the number of sources N (i.e., different numbers of dictionary atoms). The images reconstructed by the algorithms after estimating the source and mixing matrices are depicted in Fig. 9 (columns 2–5) for different values of N ranging from $N = 4$ to $N = 32$. As in Section V-B, the estimation performances of both algorithms are evaluated in terms of reconstruction error (RMSE) and sparsity level (\tilde{K} expressed as a percentage). The RMSEs are reported for each reconstructed image and the sparsity measures appear between brackets. These results clearly indicate the reliability of BOCA for fitting the observed data and its ability of identifying a sparse representation.

The Bayesian algorithm generates a collection of mixing matrices $\{\hat{\Psi}^{(h)}\}_{h=1, \dots, N_{\text{MC}}}$ that can be used to approximate the MAP estimator of Ψ following (33). The MAP estimate of the dictionary atoms (i.e., the mixing matrix), formatted as N block patches of size 16×16 , are represented in Fig. 10 (left) for the corresponding values of the number of sources N . As an illustration, the dictionary atoms estimated by K-SVD have been also depicted in Fig. 10 (right).

VII. CONCLUSIONS AND FUTURE WORK

We introduced in this paper a new Bayesian algorithm for sparse representation with undercomplete orthogonal dictionary. This problem was formulated as a blind separation problem of sparse sources mixed by an orthogonal matrix. The proposed approach relied on appropriate prior distributions for the unknown model parameters. The sparse sources to be estimated were modeled as Bernoulli–Gaussian processes. A uniform distribution on the Stiefel manifold was elected as prior distribution for the mixing matrix. The hyperparameters associated with this prior model were estimated from the data in an unsupervised fully Bayesian framework. A partially collapsed Gibbs sampler was studied to generate samples distributed according to the joint posterior distribution of the mixing matrix, source matrix, the noise variance and the model hyperparameters. The Bayesian estimators of the unknown model parameters were then approximated by using the generated

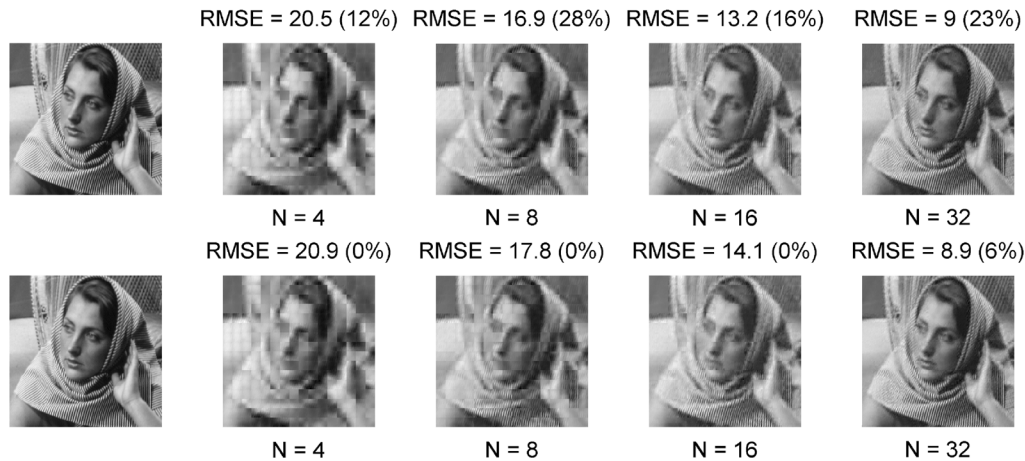


Fig. 9. Sparse coding of the Barbara image obtained by BOCA (first row) and K-SVD (second row) with different values of N . Corresponding RMSEs and sparsity levels (expressed as percentage) are indicated above each image.

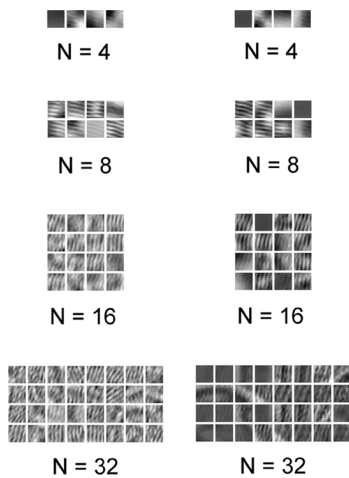


Fig. 10. Estimated dictionary atoms by BOCA (left) and K-SVD (right) for different values of the number of dictionary atoms N .

samples. The estimation performance of the proposed algorithm was evaluated from simulations conducted on synthetic data. A comparison with the K-SVD algorithm showed very promising results in favor of the proposed Bayesian method. An application of the proposed sparse coding technique for natural image processing was also investigated.

Future works include the unsupervised estimation of the number of sources N (i.e., dimension of the subspace) using a reversible-jump MCMC algorithm as in [35]. Extension of the proposed linear decomposition model to union of orthogonal dictionaries as in [8] is currently under investigation. Finally, it would be interesting to apply BOCA to sparse coding in transformed domains for compression problems.

ACKNOWLEDGMENT

The authors would like to thank Prof. A. O. Hero (University of Michigan) for interesting suggestions related to this work. They are also very grateful to G. Morin and M. Chabert (University of Toulouse) for their valuable feedback regarding the application of BOCA on natural images considered in this work.

REFERENCES

- [1] I. Gorodnitsky and B. Rao, "Sparse signal reconstruction from limited data using FOCUSS: A re-weighted minimum norm algorithm," *IEEE Trans. Signal Process.*, vol. 45, no. 3, pp. 600–616, Mar. 1997.
- [2] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [3] R. G. Baraniuk, "Compressive sensing," *IEEE Signal Process. Mag.*, vol. 24, no. 4, pp. 118–121, Jul. 2007.
- [4] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3397–3415, Dec. 1993.
- [5] J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Trans. Inf. Theory*, vol. 50, no. 10, pp. 2231–2242, Oct. 2004.
- [6] D. Donoho and X. Huo, "Uncertainty principles and ideal atomic decompositions," *IEEE Trans. Inf. Theory*, vol. 47, no. 10, pp. 2845–2862, Nov. 2001.
- [7] S. Alliney and S. A. Ruzinsky, "An algorithm for the minimization of mixed l_1 and l_2 norms with application to Bayesian estimation," *IEEE Trans. Signal Process.*, vol. 42, no. 3, pp. 618–627, Mar. 1994.
- [8] R. Gribonval and M. Nielsen, "Sparse representations in unions of bases," *IEEE Trans. Inf. Theory*, vol. 49, no. 12, pp. 3320–3325, Dec. 2003.
- [9] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4311–4319, Nov. 2006.
- [10] M. Mishali and Y. C. Eldar, "Sparse source separation from orthogonal mixtures," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, Apr. 2009, pp. 3145–3148.
- [11] H. Zou, T. Hastie, and R. Tibshirani, "Sparse principal component analysis," *J. Comput. Graph. Stat.*, vol. 15, no. 2, pp. 262–286, 2006.
- [12] A. d'Aspremont, L. El Ghaoui, M. I. Jordan, and G. R. G. Lanckriet, "A direct formulation for sparse PCA using semidefinite programming," *SIAM Rev.*, vol. 49, no. 3, pp. 434–448, 2007.
- [13] A. d'Aspremont, F. Bach, and L. El Ghaoui, "Optimal solutions for sparse principal component analysis," *J. Mach. Learn. Res.*, vol. 9, pp. 1269–1294, 2008.
- [14] P. O. Hoyer, "Non-negative matrix factorization with sparseness constraints," *J. Mach. Learn. Res.*, vol. 5, no. 2, pp. 1457–1469, 2004.
- [15] J.-P. Brunet, P. Tamayo, T. R. Golub, and J. P. Mesirov, "Metagenes and molecular pattern discovery using matrix factorization," *Proc. Nat. Academy Sci.*, vol. 101, no. 12, pp. 4164–4169, Mar. 2004.
- [16] S. Bourguignon and H. Carfantan, "Bernoulli-Gaussian spectral analysis of unevenly spaced astrophysical data," in *Proc. IEEE Workshop Stat. Signal Process. (SSP)*, Bordeaux, France, Jul. 2005, pp. 811–816.
- [17] Q. Cheng, R. Chen, and T.-H. Li, "Simultaneous wavelet estimation and deconvolution of reflection seismic signals," *IEEE Trans. Geosci. Remote Sens.*, vol. 34, no. 2, pp. 377–384, Mar. 1996.
- [18] C. Févotte and S. J. Godsill, "A Bayesian approach for blind separation of sparse sources," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 14, no. 6, pp. 2174–2188, Nov. 2006.
- [19] N. Dobigeon, A. O. Hero, and J.-Y. Tournier, "Hierarchical Bayesian sparse image reconstruction with application to MRFM," *IEEE Trans. Image Process.*, vol. 18, no. 9, pp. 2059–2070, Sep. 2009.

- [20] J. J. Kormylo and J. M. Mendel, "Maximum likelihood detection and estimation of Bernoulli–Gaussian processes," *IEEE Trans. Inf. Theory*, vol. 28, no. 3, pp. 482–488, May 1982.
- [21] J. Idier and Y. Goussard, "Stack algorithm for recursive deconvolution of Bernoulli–Gaussian processes," *IEEE Trans. Signal Process.*, vol. 28, no. 5, pp. 67–79, Sep. 1990.
- [22] M. Lavielle, "Bayesian deconvolution of Bernoulli–Gaussian processes," *Signal Process.*, vol. 33, no. 1, pp. 67–79, Jul. 1993.
- [23] A. Doucet and P. Duvaut, "Bayesian estimation of state-space models applied to deconvolution of Bernoulli–Gaussian processes," *Signal Process.*, vol. 57, no. 2, pp. 147–161, Mar. 1997.
- [24] C. Févotte, B. Torrèsani, L. Daudet, and S. J. Godsill, "Sparse linear regression with structured priors and application to denoising of musical audio," *IEEE Trans. Audio, Speech, Language Process.*, vol. 16, no. 1, pp. 174–185, Jan. 2008.
- [25] M. Ting, R. Raich, and A. O. Hero, "Sparse image reconstruction for molecular imaging," *IEEE Trans. Image Process.*, vol. 18, no. 6, pp. 1215–1227, Jun. 2009.
- [26] M. Lavielle and E. Lebarbier, "An application of MCMC methods for the multiple change-points problem," *Signal Process.*, vol. 81, no. 1, pp. 39–53, Jan. 2004.
- [27] E. Kuhn and M. Lavielle, "Coupling a stochastic approximation version of EM with an MCMC procedure," *ESAIM Probab. Statist.*, vol. 8, pp. 115–131, 2004.
- [28] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*, 2nd ed. New York: Springer, 2004.
- [29] E. Punsakaya, C. Andrieu, A. Doucet, and W. Fitzgerald, "Bayesian curve fitting using MCMC with applications to signal segmentation," *IEEE Trans. Signal Process.*, vol. 50, no. 3, pp. 747–758, Mar. 2002.
- [30] N. Dobigeon, J.-Y. Tourneret, and C.-I. Chang, "Semi-supervised linear spectral unmixing using a hierarchical Bayesian model for hyperspectral imagery," *IEEE Trans. Signal Process.*, vol. 56, no. 7, pp. 2684–2695, Jul. 2008.
- [31] N. Dobigeon, S. Moussaoui, M. Coulon, J.-Y. Tourneret, and A. O. Hero, "Joint Bayesian endmember extraction and linear unmixing for hyperspectral imagery," *IEEE Trans. Signal Process.*, vol. 57, no. 11, pp. 4355–4368, Nov. 2009.
- [32] D. Ge, J. Idier, and E. L. Carpentier, "A new MCMC algorithm for blind Bernoulli–Gaussian deconvolution," presented at the EURASIP Eur. Signal Process. Conf. (EUSIPCO), Lausanne, Switzerland, Aug. 2008.
- [33] D. A. van Dyk and T. Park, "Partially collapsed Gibbs samplers: Theory and methods," *J. Amer. Stat. Soc.*, vol. 103, no. 482, pp. 790–796, Jun. 2008.
- [34] T. Park and D. A. van Dyk, "Partially collapsed Gibbs samplers: Illustrations and applications," *J. Comput. and Graph. Stat.*, vol. 18, no. 2, pp. 283–305, Jun. 2009.
- [35] P. D. Hoff, "Model averaging and dimension selection for the singular value decomposition," *J. Amer. Stat. Soc.*, vol. 102, no. 478, pp. 674–685, Jun. 2007.
- [36] N. Dobigeon, J.-Y. Tourneret, and M. Davy, "Joint segmentation of piecewise constant autoregressive processes by using a hierarchical model and a Bayesian sampling approach," *IEEE Trans. Signal Process.*, vol. 55, no. 4, pp. 1251–1263, Apr. 2007.
- [37] Y. Chikuse, *Statistics on Special Manifolds*, ser. Lectures notes in statistics. New York: Springer-Verlag, 2003, vol. 174.
- [38] A. K. Gupta and D. K. Nagar, *Matrix Variate Distributions*, ser. Monographs and Surveys in Pure and Applied Mathematics. Boca Raton, FL: Chapman & Hall/CRC, 2000, vol. 104.
- [39] R. J. Muirhead, *Aspects of Multivariate Statistical Theory*. Hoboken, NJ: Wiley, 2005.
- [40] B. Von Hohenbalken, "Uniform sampling on simplices and spheres," *ACM SIGAPL APL Quote Quad*, vol. 11, no. 1, pp. 16–18, Sep. 1980.
- [41] F. Champagnat, Y. Goussard, and J. Idier, "Unsupervised deconvolution of sparse spike trains using stochastic approximation," *IEEE Trans. Signal Process.*, vol. 44, no. 12, pp. 2988–2998, Dec. 1996.
- [42] T. Blumensath and M. E. Davies, "Monte-Carlo methods for adaptive sparse approximations of time-series," *IEEE Trans. Signal Process.*, vol. 55, no. 9, pp. 4474–4486, Sep. 2007.
- [43] C. Févotte, "Bayesian audio source separation," in *Blind Speech Separation*, S. Makino, T.-W. Lee, and H. Sawada, Eds. Dordrecht, The Netherlands: Springer, 2007, pp. 305–335.
- [44] N. Dobigeon and J.-Y. Tourneret, "Bayesian orthogonal component analysis for sparse representation. Extension to non-homogeneous sparsity level over times," Univ. of Toulouse, IRIT/INP-ENSEEIH-TéSA, Toulouse, France, Tech. Rep., Nov. 2009 [Online]. Available: <http://dobigeon.perso.enseeiht.fr/publis.html>
- [45] A. T. A. Wood, "Simulation of the von Mises Fisher distribution," *Comm. Statist. Simulation Comput.*, vol. 23, pp. 157–164, 1994.
- [46] J.-M. Marin and C. P. Robert, *Bayesian Core: A Practical Approach to Computational Bayesian Statistics*. New York: Springer, 2007.
- [47] N. Bertin, R. Badeau, and G. Richard, "Blind signal decompositions for automatic transcription of polyphonic music: NMF and K-SVD on the benchmark," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, Apr. 2007, vol. 1, pp. I-65–I-68.
- [48] O. Bryt and M. Elad, "Compression of facial images using the K-SVD algorithm," *J. Vis. Commun. Image Represent.*, vol. 19, no. 4, pp. 270–283, May 2008.
- [49] J. Mairal, M. Elad, and G. Sapiro, "Sparse representation for color image restoration," *IEEE Trans. Image Process.*, vol. 17, no. 1, pp. 53–69, Jan. 2008.



Nicolas Dobigeon (S'05–M'08) was born in Angoulême, France, in 1981. He received the Eng. degree in electrical engineering from ENSEEIHT, Toulouse, France, and the M.Sc. degree in signal processing from the National Polytechnic Institute of Toulouse, both in 2004, and the Ph.D. degree in signal processing from the National Polytechnic Institute of Toulouse in 2007.

From 2007 to 2008, he was a Postdoctoral Research Associate at the Department of Electrical Engineering and Computer Science, University of

Michigan. Since 2008, he has been an Assistant Professor with the National Polytechnic Institute of Toulouse (ENSEEIH—University of Toulouse), within the Signal and Communication Group of the IRIT Laboratory. His research interests are centered around statistical signal and image processing with a particular interest to Bayesian inference and Markov chain Monte Carlo (MCMC) methods.



Jean-Yves Tourneret (M'94–SM'08) received the Ingénieur degree in electrical engineering from École Nationale Supérieure d'Électronique, d'Électrotechnique, d'Informatique et d'Hydraulique in Toulouse (ENSEEIH), France, in 1989 and the Ph.D. degree from the National Polytechnic Institute, Toulouse, France, in 1992.

He is currently a Professor in ENSEEIHT and a member of the IRIT laboratory (UMR 5505 of the CNRS). His research activities are centered around statistical signal processing with a particular interest

to classification and Markov chain Monte Carlo methods.

Dr. Tourneret was the Program Chair of the European Conference on Signal Processing (EUSIPCO), which was held in Toulouse, France, in 2002. He was also member of the organizing committee for the international conference ICASSP 2006, which was held in Toulouse, France, in 2006. He has been a member of different technical committees, including the Signal Processing Theory and Methods (SPTM) Committee of the IEEE Signal Processing Society (2001–2007). He is currently serving as an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING.