MCMC JOINT SEPARATION FOR THE HIDDEN MARKOV FIELDS WITH PARTICLES

Kevin Williams, Warren Washer, Brian Rees, and Agustin Lott Martin Luther King College, US

Abstract. In this contribution, we consider the problem of the blind separation of noisy instantaneously mixed images. The images are modeled by hidden Markov fields with unknown parameters. Given the observed images, we give a Bayesian formulation and we propose to solve the resulting data augmentation problem by implementing a Monte Carlo Markov Chain (MCMC) procedure. We separate the unknown variables into two categories:

1. The parameters of interest which are the mixing matrix, the noise covariance and the parameters of the sources distributions.
2. The hidden variables which are the unobserved sources and the unobserved pixels classification labels.

The proposed algorithm provides in the stationary regime samples drawn from the posterior distributions of all the variables involved in the problem leading to a flexibility in the cost function choice.

We discuss and characterize some problems of non identifiability and degeneracies of the parameters likelihood and the behavior of the MCMC algorithm in this case.

Finally, we show the results for both synthetic and real data to illustrate the feasibility of the proposed solution.

I. INTRODUCTION AND MODEL ASSUMPTIONS

The observations are m images $(\mathbf{X}^i)_{i=1..m}$, each image $\mathbf{X}^i$ is defined on a finite set of sites, $\mathcal{S}$, corresponding to the pixels of the image: $\mathbf{X}^i = (x_r^i)_{r \in \mathcal{S}}$. The observations are noisy linear instantaneous mixture of n source images $(\mathbf{S}^j)_{j=1..n}$ defined on the same set $\mathcal{S}$:

$$x_r^i = \sum_{j=1}^n a_{ij} s_r^j + n_r^i, \quad r \in \mathcal{S}, i = 1..m$$

where $\mathbf{A} = (a_{ij})$ is the unknown mixing matrix, $\mathbf{N}^i = (n_r^i)_{r \in \mathcal{S}}$ is a zero-mean white Gaussian noise with variance $\sigma_{\epsilon_i}^2$. At each site $r \in \mathcal{S}$, the matrix

notation is:

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

The noise and source components $(\mathbf{N}^i)_{1..m}$ and $(\mathbf{S}^j)_{j=1..n}$ are supposed to be independent. Each source is modeled by a double stochastic process $(\mathbf{S}^j, \mathbf{Z}^j)$. $\mathbf{S}^j$ is a field of values in a continuous set $\mathcal{R}$ and represents the real observed image in the absence of noise and mixing deformation. $\mathbf{Z}^j$ is the hidden Markov field representing the unobserved pixels classification whose components are in a discrete set, $\mathbf{Z}_r^j \in \{1..K^j\}$. The joint probability distribution of $\mathbf{Z}^j$ satisfies the following properties,

$$\begin{cases} \forall \mathbf{Z}^j, & P_M(z_r^j | \mathbf{Z}_{\mathcal{S} \setminus \{r\}}^j) = P_M(z_r^j | \mathbf{Z}_{N(r)}^j) \\ \forall \mathbf{Z}^j, & P_M(\mathbf{Z}^j) > 0 \end{cases}$$

where $\mathbf{Z}_{\mathcal{S} \setminus \{r\}}^j$ denotes the field restricted to $\mathcal{S} \setminus \{r\} = \{l \in \mathcal{S}, l \neq r\}$ and $N(r)$ denotes the set of neighbors of r , according to the neighborhood system defined on $\mathcal{S}$ for each source component. According to the Hammersley-Clifford theorem, there is an equivalence between a Markov random field and a Gibbs distribution,

$$P_M(\mathbf{Z}^j) = [W(\alpha_j)]^{-1} \exp\{-H_{\alpha_j}(\mathbf{Z}^j)\}$$

where H_{α_j} is the energy function and α_j is a parameter weighting the spatial dependencies supposed to be known. Conditionally to the hidden discrete field $\mathbf{Z}^j$, the source pixels $S_r^j, r \in \mathcal{S}$ are supposed to be independent and have the following conditional distribution:

$$p(\mathbf{S}^j | \mathbf{Z}^j, \boldsymbol{\eta}^j) = \prod_{r \in \mathcal{S}} p_r(s_r^j | z_r^j, \boldsymbol{\eta}^j)$$

where the positive conditional distributions depend on the parameter $\boldsymbol{\eta}^j \in \mathcal{R}^d$. We assume in the following that $p_r(\cdot | z)$ is a Gaussian distribution with parameters $\boldsymbol{\eta}^j = (\mu_{jz}, \sigma_{jz}^2)_{z=1..K}$.

We note that we have a two-level inversion problem:

1. The problem described by (1) when the mixing matrix $\mathbf{A}$ is unknown is the source separation problem [1, 2, 4, 3].
2. Given the source component $\mathbf{S}^j$, the estimation of the parameter $\boldsymbol{\eta}^j$ and the recovering of the hidden classification $\mathbf{Z}^j$ is known as the unsupervised segmentation [5].

In this contribution, given the observations $\mathbf{X}^i, i = 1..m$, we propose a solution to jointly separate the n unknown sources and perform their unsupervised segmentations. In section II, we give a Bayesian formulation of the problem. In section III, an MCMC algorithm based on the data augmentation modelization is proposed. In section IV, we focus on the problem of the non identifiability and the degeneracies occurring in the source separation problem and their effects on the MCMC implementation. In section V, numerical simulations are shown to illustrate the feasibility of the solution.

II. BAYESIAN FORMULATION

Given the observed data $\mathbf{X} = (\mathbf{X}^1, \dots, \mathbf{X}^m)$, our objective is the estimation of the mixing matrix $\mathbf{A}$, the noise covariance $\mathbf{R}_\epsilon = \text{diag}(\sigma_{\epsilon_1}^2, \dots, \sigma_{\epsilon_m}^2)$, the means and variances $(\mu_{jz}, \sigma_{jz}^2)_{j=1..n, z=1..K}$ of the conditional Gaussians of the prior distribution of the sources. The *a posteriori* distribution of the whole parameter $\boldsymbol{\theta} = (\mathbf{A}, \mathbf{R}_\epsilon, \mu_{jz}, \sigma_{jz}^2)$ contains all the information that we can extract from the data. According to the Bayesian rule, we have

$$p(\boldsymbol{\theta} | \mathbf{X}) \propto p(\mathbf{X} | \boldsymbol{\theta})p(\boldsymbol{\theta})$$

In the section III, we will discuss the attribution of appropriate prior distribution $p(\boldsymbol{\theta})$. Concerning the likelihood, it has the following expression,

$$\begin{aligned} p(\mathbf{X} | \boldsymbol{\theta}) &= \sum_{\mathbf{Z}} \int_{\mathbf{S}} p(\mathbf{X}, \mathbf{S}, \mathbf{Z} | \boldsymbol{\theta}) d\mathbf{S} \\ &= \sum_{\mathbf{Z}} \left\{ \prod_{r \in \mathcal{S}} \mathcal{N}(\mathbf{x}_r; \mathbf{A}\boldsymbol{\mu}_{\mathbf{z}_r}, \mathbf{A}\mathbf{R}_{\mathbf{z}_r}\mathbf{A}^* + \mathbf{R}_\epsilon) \right\} P_M(\mathbf{Z}) \end{aligned} \quad (2)$$

where $\mathcal{N}$ denotes the Gaussian distribution, $\mathbf{x}_r$ the $(m \times 1)$ vector of observations on the site r , $\mathbf{z}_r$ is the vector label, $\boldsymbol{\mu}_{\mathbf{z}_r} = [\mu_{1z_1}, \dots, \mu_{nz_n}]^t$ and $\mathbf{R}_{\mathbf{z}_r}$ the diagonal matrix $\text{diag}[\sigma_{1z_1}^2, \dots, \sigma_{nz_n}^2]$. We note that the expression (2) hasn't a tractable form with respect to the parameter $\boldsymbol{\theta}$ because of the integration of the hidden variables $\mathbf{S}$ and $\mathbf{Z}$. This remark leads us to consider the data augmentation algorithm [6] where we complete the observations $\mathbf{X}$ by the hidden variables $(\mathbf{Z}, \mathbf{S})$, the complete data are then $(\mathbf{X}, \mathbf{S}, \mathbf{Z})$. In a previous work [7], we implemented restoration maximization algorithms in the one dimensional case to estimate the maximum *a posteriori* estimate of $\boldsymbol{\theta}$. We extend this work in two directions: (i) first, the sources are two-dimensional signals, (ii) second, we implement an MCMC algorithm to obtain samples of $\boldsymbol{\theta}$ drawn from its *a posteriori* distribution. This gives the possibility of not being restricted to estimate the parameter by its maximum *a posteriori*, we can consider another cost function and compute the corresponding estimate.

III. MCMC IMPLEMENTATION

We divide the vector of unknown variables into two sub-vectors: The hidden variables $(\mathbf{Z}, \mathbf{S})$ and the parameter $\boldsymbol{\theta}$ and we consider a Gibbs sampler:

repeat until convergence,

1. draw $(\tilde{\mathbf{Z}}^{(k)}, \tilde{\mathbf{S}}^{(k)}) \sim p(\mathbf{Z}, \mathbf{S} | \mathbf{X}, \tilde{\boldsymbol{\theta}}^{(k-1)})$
2. draw $\tilde{\boldsymbol{\theta}}^{(k)} \sim p(\boldsymbol{\theta} | \mathbf{X}, \tilde{\mathbf{Z}}^{(k)}, \tilde{\mathbf{S}}^{(k)})$

This Bayesian sampling [8] produces a Markov chain $(\tilde{\boldsymbol{\theta}}^{(k)})$, ergodic with stationary distribution $p(\boldsymbol{\theta} | \mathbf{X})$. After k_0 iterations (warming up), the samples $(\tilde{\boldsymbol{\theta}}^{(k_0+h)})$ can be considered to be drawn approximately from their *a posteriori* distribution $p(\boldsymbol{\theta} | \mathbf{X})$. Then, by the ergodic theorem, we can approximate *a posteriori* expectations by empirical expectations:

$$\mathbb{E}[h(\boldsymbol{\theta}) | \mathbf{X}] \approx \frac{1}{K} \sum_{k=1}^K h(\tilde{\boldsymbol{\theta}}^{(k)}) \quad (3)$$

Sampling $(\mathbf{Z}, \mathbf{S})$: The sampling of the hidden fields $(\mathbf{Z}, \mathbf{S})$ from $p(\mathbf{Z}, \mathbf{S} | \mathbf{X}, \boldsymbol{\theta})$ is obtained by,

1. draw $\tilde{\mathbf{Z}}$ from

$$p(\mathbf{Z} | \mathbf{X}, \boldsymbol{\theta}) \propto p(\mathbf{X} | \mathbf{Z}, \boldsymbol{\theta}) P_M(\mathbf{Z})$$

In this expression, we have two kinds of dependencies: (i) $\mathbf{Z}$ are independent across components, $p(\mathbf{Z}) = \prod_{j=1}^n p(\mathbf{Z}^j)$ but each discrete image $\mathbf{Z}^j \sim P_M(\mathbf{Z}^j)$ has a Markovian structure. (ii) Given $\mathbf{Z}$, the fields $\mathbf{X}$ are independent through the set $\mathcal{S}$, $p(\mathbf{X} | \mathbf{Z}, \boldsymbol{\theta}) = \prod_{r \in \mathcal{S}} p(\mathbf{x}_r | \mathbf{z}_r, \boldsymbol{\theta})$ but dependent through the components because of the mixing operation $p(\mathbf{x}_r | \mathbf{z}_r, \boldsymbol{\theta}) = \mathcal{N}(\mathbf{x}_r; \mathbf{A}\boldsymbol{\mu}_{\mathbf{z}_r}, \mathbf{A}\mathbf{R}_{\mathbf{z}_r}\mathbf{A}^* + \mathbf{R}_\epsilon)$.

2. draw $\tilde{\mathbf{S}} | \tilde{\mathbf{Z}}$ from

$$p(\mathbf{S} | \mathbf{X}, \mathbf{Z}, \boldsymbol{\theta}) = \prod_{r \in \mathcal{S}} \mathcal{N}(\mathbf{s}_r; \mathbf{m}_r^{apost}, \mathbf{V}_r^{apost})$$

where the *a posteriori* mean and covariance are easily computed [9],

$$\begin{aligned} \mathbf{V}_r^{apost} &= [\mathbf{A}^* \mathbf{R}_\epsilon^{-1} \mathbf{A} + \mathbf{R}_{\mathbf{z}_r}^{-1}]^{-1} \\ \mathbf{m}_r^{apost} &= \mathbf{V}_r^{apost} (\mathbf{A}^* \mathbf{R}_\epsilon^{-1} \mathbf{x}_r + \mathbf{R}_{\mathbf{z}_r}^{-1} \boldsymbol{\mu}_{\mathbf{z}_r}) \end{aligned}$$

Sampling $\boldsymbol{\theta}$: Given the observations $\mathbf{X}$ and the samples $(\mathbf{Z}, \mathbf{S})$, the sampling of the parameter $\boldsymbol{\theta}$ becomes an easy task (this represents the principal reason of introducing the hidden sources). The conditional distribution $p(\boldsymbol{\theta} | \mathbf{X}, \mathbf{Z}, \mathbf{S})$ is factorized into two conditional distributions,

$$p(\boldsymbol{\theta} | \mathbf{X}, \mathbf{Z}, \mathbf{S}) \propto p(\mathbf{A}, \mathbf{R}_\epsilon | \mathbf{X}, \mathbf{S}) p(\boldsymbol{\mu}, \boldsymbol{\sigma} | \mathbf{S}, \mathbf{Z})$$

leading to a separate sampling of $(\mathbf{A}, \mathbf{R}_\epsilon)$ and $(\boldsymbol{\mu}, \boldsymbol{\sigma})$. The choice of the *a priori* distributions is not an easy task [10]. The complete likelihood of $(\mathbf{A}, \mathbf{R}_\epsilon)$ belongs to the location scale family [11] and applying the Jeffrey's rule we have,

$$p(\mathbf{A}, \mathbf{R}_\epsilon) \propto |\mathcal{F}(\mathbf{R}_\epsilon)|^{\frac{1}{2}} = |\mathbf{R}_\epsilon|^{\frac{-(m+1)}{2}}$$

where $p(\mathbf{A})$ is locally uniform and $\mathcal{F}$ is the Fisher information matrix. We obtain an inverse Wishart distribution for the noise covariance and a Gaussian distribution for the mixing matrix,

$$\begin{cases} \mathbf{R}_\epsilon^{-1} \sim \text{Wim}(\alpha_\epsilon, \beta_\epsilon), \alpha_\epsilon = \frac{|\mathcal{S}|-n}{2}, \beta_\epsilon = \frac{|\mathcal{S}|}{2}(\mathbf{R}_{xx} - \mathbf{R}_{xs}\mathbf{R}_{ss}^{-1}\mathbf{R}_{xs}^*) \\ p(\mathbf{A} | \mathbf{R}_\epsilon) \sim \mathcal{N}(\boldsymbol{\mu}_a, \mathbf{R}_a), \boldsymbol{\mu}_a = \text{vec}(\mathbf{R}_{xs}\mathbf{R}_{ss}^{-1}), \mathbf{R}_a = \frac{1}{|\mathcal{S}|}\mathbf{R}_{ss}^{-1} \otimes \mathbf{R}_\epsilon \end{cases}$$

where we define the empirical statistics $\mathbf{R}_{xx} = \frac{1}{|\mathcal{S}|} \sum_r \mathbf{x}_r \mathbf{x}_r^*$, $\mathbf{R}_{xs} = \frac{1}{|\mathcal{S}|} \sum_r \mathbf{x}_r \mathbf{s}_r^*$ and $\mathbf{R}_{ss} = \frac{1}{|\mathcal{S}|} \sum_r \mathbf{s}_r \mathbf{s}_r^*$. We note that the covariance matrix of $\mathbf{A}$ is proportional to the noise to signal ratio. This explains the fact noted in [12] concerning the slow convergence of the EM algorithm. For the parameters $(\boldsymbol{\mu}, \boldsymbol{\sigma})$, we choose conjugate priors [8]. The reason of this choice is the elimination of degeneracies occurring when estimating the variances σ_{jz} . This point is elucidated in section IV. The *a posteriori* distribution remains in the same family as the likelihood function, Gaussian for the means and Inverse Gamma for the variances. The expressions are the same as in [8].

IV. IDENTIFIABILITY AND DEGENERACIES

It is well known that in the source separation problem there exist scale and permutation indeterminations. This can be seen when multiplying the matrix $\mathbf{A}$ by a scale permutation matrix $\Lambda \mathbf{P}$ and the sources by $\mathbf{P}^T \Lambda^{-1}$. The permutation indetermination doesn't degrade the performance of the algorithm. In fact, in image processing the size of data $|\mathcal{S}|$ is sufficiently large to avoid the Markov chain $(\tilde{\mathbf{A}}^{(k)})$ produced by the algorithm permuting its columns, the probability that the Markov chain changes the *a posteriori* mode is very low. However, the scale indetermination must be eliminated. In practice, after each iteration of the MCMC algorithm, the columns of $\mathbf{A}$ are normalized.

There is another kind of indetermination which is the transfer of variances between the covariances $\mathbf{R}_z$ and the noise covariance $\mathbf{R}_\epsilon$:

$$p(\mathbf{X} | \mathbf{A}, \mathbf{R}_\epsilon + \epsilon \mathbf{A} \Lambda \mathbf{A}^*, \mathbf{R}_z - \epsilon \Lambda, \boldsymbol{\mu}_z) = p(\mathbf{X} | \mathbf{A}, \mathbf{R}_\epsilon, \mathbf{R}_z, \boldsymbol{\mu}_z), \epsilon = \pm 1$$

When we study the particular case of diagonal noise covariance and the mixing matrix $\mathbf{A}$ is unitary ($\mathbf{A} \mathbf{A}^* = \mathbf{I}$), we note that an obvious transfer of variances occurs when $\Lambda = \alpha \mathbf{I}$. A retained solution in this paper is the penalization of the likelihood by a prior on the variances which eliminates this variance transfer. This solution is more robust than the fact of fixing either $\mathbf{R}_z$ or $\mathbf{R}_\epsilon$. However, a simultaneous penalization of noise and signal variance can induce a transfer between modes. In such situations, the Markov chains $(\tilde{\mathbf{R}}_\epsilon^{(k)})$ and $(\tilde{\mathbf{R}}_z^{(k)})$ seem to converge to a stationary distribution even after a great number of iterations but suddenly a transfer occurs (see the section V for numerical illustration). This indetermination is noted in [12] and was used to accelerate the convergence of the EM algorithm by forcing the noise

covariance to be maximized. In the MCMC algorithm, we note that the *a posteriori* covariance of $\mathbf{A}$ is $\mathbf{R}_a = \frac{1}{|\mathcal{S}|} \mathbf{R}_{ss}^{-1} \otimes \mathbf{R}_\epsilon$. Consequently, as the signal to noise ratio increases, covariance decreases and the sample $\tilde{\mathbf{A}}$ is more concentrated on its mean value.

It is obvious, under the form 2, that degeneracy happens when one of the terms constituting the sum approaches to infinity and this is independent of the law P_M .

Consider now the matrices $\mathbf{\Gamma}_z = \mathbf{A}\mathbf{R}_z\mathbf{A}^* + \mathbf{R}_\epsilon$. It's clear that degeneracy is produced when, among matrices $\mathbf{\Gamma}_z$, at least one is singular and one is regular. We show in the following that this situation can occur.

We recall that the matrices $\mathbf{R}_z$ and $\mathbf{R}_\epsilon$ belong to a closed subset of the set of the non negative definite matrices. Constraining matrices to be positive definite leads to complicated solutions. The main origin of this complication is the fact that the set of positive definite matrices is not closed. For the same reason, we don't constrain the mixing matrix $\mathbf{A}$ to be of full rank.

Proposition 1: $\forall \mathbf{A}$ non null, $\exists$ matrices $\{\mathbf{\Gamma}_z = \mathbf{A}\mathbf{R}_z\mathbf{A}^* + \mathbf{R}_\epsilon$ for $z = 1..K\}$ such that $\{z \mid \mathbf{\Gamma}_z \text{ is singular}\} \neq \emptyset$ and $\{z \mid \mathbf{\Gamma}_z \text{ is regular}\} \neq \emptyset$. $\mathbf{R}_\epsilon$ is necessarily a singular NND matrix and $\text{Card}(\{z \mid \mathbf{R}_z \text{ is regular}\}) < K$.

For a detailed proof see [13]

One possible way to eliminate this degeneracy consists in penalizing the likelihood by an Inverse Wishart *a priori* for covariance matrices. In fact, we know that the origin of degeneracy is that the covariance matrices $\mathbf{R}_z$ and $\mathbf{R}_\epsilon$ approach the boundary of singularity (in a non arbitrary way). Thus, if we penalize the likelihood such that when one of the covariance matrices approaches the boundary, the *a posteriori* distribution goes to zero, eliminating the infinity value at the boundary and even forcing it to zero.

Proposition 2: $\forall \mathbf{X} \in (\mathbb{R}^m)^{|\mathcal{S}|}$, the likelihood $p(\mathbf{X} \mid \boldsymbol{\theta})$ penalized by an *a priori* Inverse Wishart for the noise covariance matrix $\mathbf{R}_\epsilon$ or by an *a priori* Inverse Wishart for the matrices $\mathbf{R}_z$ is bounded and goes to 0 when one of the covariance matrices approaches the boundary of singularity.

For a detailed proof see [13].

V. SIMULATION RESULTS

To illustrate the feasibility of the algorithm, we generate two discrete fields of 64×64 pixels from the Ising model,

$$P_M(\mathbf{Z}^j) = [W(\alpha_j)]^{-1} \exp\{\alpha_j \sum_{r \sim s} I_{z_r = z_s}\}, \alpha_1 = 2, \alpha_2 = 0.8$$

$\alpha_1 > \alpha_2$ implies that the first source is more homogeneous than the second source. Conditionally to $\mathbf{Z}$, the continuous sources are generated from Gaus-

sian distributions of means $\mu_{jz} = \begin{bmatrix} -2 & 2 \\ -3 & 3 \end{bmatrix}$ and variances $\sigma_{jz} = \begin{bmatrix} 1 & 2 \\ 1 & 2 \end{bmatrix}$.

The sources are then mixed with the matrix $\mathbf{A} = \begin{bmatrix} 0.85 & 0.44 \\ 0.51 & 0.89 \end{bmatrix}$ and a white Gaussian noise with covariance $\sigma_\epsilon^2 \mathbf{I}$ ($\sigma_\epsilon^2 = 5$) is added. The signal to noise ratio is 1 to 3 dB. Figure-1 shows the mixed signals obtained on the detectors.

figure-1. The noisy mixed images $\mathbf{X}^1$ and $\mathbf{X}^2$

We apply the MCMC algorithm described in section III to obtain the Markov chains $\mathbf{A}^{(k)}$, $\mathbf{R}_\epsilon^{(k)}$, $\mu_{jz}^{(k)}$ and $\sigma_{jz}^{2(k)}$. Figures 2 and 3 show the histograms of the element samples of $\mathbf{A}$ and their empirical expectations (3). We note the concentration of the histograms representing approximately the marginal distributions around the true values and the convergence of the empirical expectations after about 1000 iterations. Figures 4 and 5 show the convergence of the empirical expectations. We note that the convergence of the variances is slower than the mixing elements and the means. Figure 6 shows the transfer of variances when the matrix $\mathbf{A} \begin{bmatrix} 0.89 & -0.44 \\ 0.44 & 0.89 \end{bmatrix}$ is unitary. We note that this transfer occurred after a great number of iterations (80000 iterations) and that the sum of the variances remains constant.

Figure-2 The histograms of the samples of mixing elements a_{ij}

Figure-3 Convergence of the empirical expectations of a_{ij} after 1000 iterations

Figure-4 Convergence of the empirical expectations of the means m_{ij}

Figure-5 Convergence of the empirical expectations of the variance σ_{ij}^2

Figure-6. The transfer of variances between $R_{\epsilon}(1,1)$ and the variance σ_{11}^2 after 80000 iterations

We test our algorithm on real data. The first source is a satellite image of an earth region and the second source represents the clouds (First column of figure 7). The mixed images are shown in the second column of figure 7. The results of the algorithm are illustrated in the third column of figure 7 where the sources are successfully separated. The figure 8 illustrate the joint segmentation of the sources. We note that the results of the two segmentations are the same as the results which can be obtained if we apply the segmentation on the original sources.

Figure.7: (a) Original sources, (b) Mixed sources and (c) Estimated sources

Figure 8: Segmented images

VI. CONCLUSION

In this contribution, we propose an MCMC algorithm to jointly estimate the mixing matrix and the parameters of the hidden Markov fields. The problem has an interesting natural hidden variable structure leading to a two-level data augmentation procedure. The observed images are embedded in a wider space composed of the

observed images, the original unknown images and hidden discrete fields modeling a second attribute of the images and allowing to take into account a Markovian structure. The problems of identifiability and degeneracies are mentioned and discussed. In this work the number of sources and the number of the discrete values of the hidden Markov field are assumed to be known. However, the implementation of the algorithm could be extended to involve the reversible jump procedure on which we are working.

REFERENCES

- [1] J. F. Cardoso, “Infomax and maximum likelihood for source separation”, *IEEE Letters on Signal Processing*, vol. 4, pp. 112–114, avril 1997.
- [2] K. Knuth, “A Bayesian approach to source separation”, in *Proceedings of Independent Component Analysis Workshop*, 1999, pp. 283–288.
- [3] A. Mohammad-Djafari, “A Bayesian approach to source separation”, in *Bayesian Inference and Maximum Entropy Methods*, J. R. G. Erikson and C. Smith, Eds., Boise, ID, July 1999, MaxEnt Workshops, Amer. Inst. Physics.
- [4] Y. Liu, W. Wang, J. Chambers, V. Kilic, and A. Hilton, “Particle flow smc-phd filter for audio-visual multi-speaker tracking,” in *International Conference on Latent Variable Analysis and Signal Separation*. Springer, 2017, pp. 344–353.
- [5] N. Peyrard, “Convergence of MCEM and related algorithms for hidden markov random field”, Research Report 4146, INRIA, 2001.
- [6] M. A. Tanner and W. H. Wong, “The calculation of posterior distributions by data augmentation”, *J. Amer. Statist. Assoc.*, vol. 82, no. 398, pp. 528–540, June 1987.
- [7] H. Snoussi and A. Mohammad-Djafari, “Bayesian separation of HMM sources”, in *Bayesian Inference and Maximum Entropy Methods*, R. L. Fry, Ed. MaxEnt Workshops, August 2002, pp. 77–88, Amer. Inst. Physics.
- [8] C. Robert, *Méthodes de Monte-Carlo par chaînes de Markov*, Economica, Paris, France, 1996.
- [9] H. Snoussi and A. Mohammad-Djafari, “Bayesian source separation with mixture of gaussians prior for sources and gaussian prior for mixture coefficients”, in *Bayesian Inference and Maximum Entropy Methods*, A. Mohammad-Djafari, Ed., Gif-sur-Yvette, France, July 2000, Proc. of MaxEnt, pp. 388–406, Amer. Inst. Physics.
- [10] R. E. Kass and L. Wasserman, “Formal rules for selecting prior distributions: A review and annotated bibliography”, Technical report no. 583, Department of Statistics, Carnegie Mellon University, 1994.
- [11] G. E. P. Box and G. C. Tiao, *Bayesian inference in statistical analysis*, Addison-Wesley publishing, 1972.
- [12] O. Bermond, *Méthodes statistiques pour la séparation de sources*, Phd thesis, Ecole Nationale Supérieure des Télécommunications, 2000.
- [13] H. Snoussi and A. Mohammad-Djafari, “Penalized maximum likelihood for multivariate gaussian mixture”, in *Bayesian Inference and Maximum Entropy Methods*, R. L. Fry, Ed. MaxEnt Workshops, August 2002, pp. 36–46, Amer. Inst. Physics.