SEGMENTATION OF ULTRASOUND IMAGES USING A SPATIALLY COHERENT GENERALIZED RAYLEIGH MIXTURE MODEL

Marcelo Pereyra, Nicolas Dobigeon, Hadj Batatia and Jean-Yves Tourneret

University of Toulouse, IRIT/INP-ENSEEIHT, 31071 Toulouse Cedex 7, France

{marcelo.pereyra, nicolas.dobigeon, hadj.batatia, jean-yves.tourneret}@enseeiht.fr

ABSTRACT

This paper addresses the problem of jointly estimating the statistical distribution and segmenting multiple-tissue high-frequency ultrasound images. The distribution of multiple-tissue images is modeled as a spatially coherent finite mixture of heavy-tailed Rayleigh distributions. Spatial coherence inherent to biological tissues is introduced into the model by enforcing local dependance between pixels. An original Bayesian algorithm combined with a Markov chain Monte Carlo method is then derived to jointly estimate the mixture parameters and a label vector associating each voxel to a tissue. Precisely, a hybrid Metropolis-within-Gibbs sampler is proposed to draw samples that are asymptotically distributed according to the posterior distribution of the Bayesian model. These samples are then used to compute the Bayesian estimators of the model parameters. Simulation results are conducted on synthetic data to illustrate the performance of the proposed estimation strategy. The method is then successfully applied to the segmentation of an in-vivo lesion in a high frequency 3D ultrasound image.

Index Terms— Heavy-tailed Rayleigh distribution, mixture model, Potts-Markov field, Bayesian estimation, Gibbs sampler.

1. INTRODUCTION

Ultrasound imaging is a longstanding medical imaging modality with important applications in diagnostics, preventive examinations, therapy and image-guided surgery. Recent advances in high frequency transducers and 3D probes have opened new opportunities for this modality in dermatological oncology, where ultrasound images are increasingly used for lesion assessment. This evaluation relies on two manual segmentation measures: the depth of the lesion and the number of skin layers that has been invaded.

Despite the extensive literature on the subject, automatic segmentation of ultrasound images is still a challenging task and a focus of considerable research efforts. A detailed survey of the state of the art up to 2006 is provided in [1]. Recent works have proposed the application of Bayesian mixture models to ultrasound image segmentation. Cardinal *et al.* proposed an intravascular segmentation method based on a Rayleigh mixture model [2]. This work was extended to non-Rayleigh images by Destrempes *et al.* [3], who proposed a carotid artery segmentation method based on a Nakagami mixture. In both cases the estimation of the mixture parameters and the segmentation of voxels were performed separately. First, the parameters of the mixture model were estimated using the expectation-maximization (EM) algorithm under the assumption that voxels are independent. Then, voxel classification was performed conditionally to the mixture parameters using region-based segmentation algorithms. In both cases the segmentation was initialized using a manual solution.

This paper addresses the problem of jointly estimating the parameters of the mixture model and the classification of voxels in 2D and 3D skin ultrasound images. Ultrasound images are modeled as a heavy-tailed Rayleigh mixture, as in [4]. This model, which originally described voxel classes independently, is now equipped with a Markov random field (MRF) that takes into account the spatial correlation inherent to biological tissues. MRFs have shown interesting properties to model spatial correlation in images [5]. The main contribution of this paper is to show that MRFs are appropriate to ultrasound images whose statistical properties are descried by mixtures of α-Rayleigh distributions.

2. PROBLEM STATEMENT

The mixture model used in this paper to describe ultrasound image voxels has been presented in detail in [4]. Some elements of this model are recalled in this section for consistency. Recent works on scattering in biological tissues have shown that the α-Rayleigh distribution is appropriate to model the distribution of skin ultrasound image voxels [6]. This paper considers the case where the ultrasound image can be divided into multiple regions (or classes) associated with different biological tissues, each with its own echogenicity and therefore its own speckle statistics. As a consequence, the observed ultrasound image envelope $\boldsymbol{r}$ is assumed to be by an α-Rayleigh stationary process with piecewise constant parameters.

Equivalently, each ultrasound image voxel is a mixture of K α-Rayleigh distributions, where K is the number of classes in the image and w_k represents the relative weight (or proportion) of the kth class. More precisely, we assume there is a set of K stationary classes $\{C_1, ..., C_K\}$ with proportions $\{w_1, ..., w_K\}$ such that

$$r_n \sim \sum_{k=1}^{K} w_k \alpha\mathcal{R}(\alpha_k, \gamma_k) \tag{1}$$

where $r_n \in \mathbb{R}^+$ is the nth image voxel and $\alpha\mathcal{R}(\alpha_k, \gamma_k)$ denotes the α-Rayleigh distribution with parameters α_k and γ_k (associated with the class $\mathcal{C}_k$). Note that the number of classes K is assumed to be known in this study. This assumption is based on the fact that sonography technicians can quickly assess the number of tissues within a region of interest by visual inspection. A Bayesian model was introduced in [4] to estimate the α-Rayleigh mixture model parameters $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_K)^T$ and $\boldsymbol{\gamma} = (\gamma_1, \ldots, \gamma_K)^T$ and the class of each voxel of an ultrasound image. The next section generalizes this Bayesian model to exploit spatial correlations between adjacent voxels of the image.

3. BAYESIAN MODEL

In order to characterize the class of the nth image voxel, we introduce a label vector $\boldsymbol{z} = \{z_1, \ldots, z_N\}$ such that $z_n = k$ if and only if $r_n \in C_k$. This label vector will allow one to characterize each image observation and to discriminate different kinds of tissues. This section defines a Bayesian model associated with the unknown parameter vector $\boldsymbol{\theta} = (\boldsymbol{\alpha}^T, \boldsymbol{\gamma}^T, \boldsymbol{z}^T)^T$ for a mixture of α-Rayleigh distributions. This model requires to define the likelihood of the observation vector $\boldsymbol{r}$ and the priors for the unknown parameters.

3.1. Likelihood

Assuming that the observations r_n are conditionally independent and using the mixture model (1), the likelihood of the proposed Bayesian model can be written as in [4]

$$p(\boldsymbol{r}|\boldsymbol{\theta}) = \prod_{k=1}^{K} \prod_{\{n|z_n=k\}} p_{\alpha\mathcal{R}}(r_n|\alpha_k, \gamma_k) \tag{2}$$

where

$$p_{\alpha\mathcal{R}}(r_n|\alpha_k, \gamma_k) = r_n \int_0^\infty \lambda \exp\left[-(\gamma_k\lambda)^{\alpha_k}\right] J_0(r_n\lambda)\, d\lambda \tag{3}$$

and where J_0 is the 0th order Bessel function of the first kind.

3.2. Parameter priors

3.2.1. Labels

We propose in this paper to exploit the spatial coherence inherent to biological tissues. More precisely, it is natural to consider that there is some correlation between the characteristics of a given voxel and those of its neighbors. Since the seminal work of Geman [7], MRFs have become very popular to introduce spatial correlation in images. MRFs assume that the distribution of a pixel conditionally to all other pixels of the image equals the distribution of this pixel conditionally to its neighbors. Consequently, it is important to properly define the neighborhood structure. The neighborhood relation between two pixels (or voxels) i and j has to be symmetric: if n is a neighbor of m then m is also a neighbor of n. There are several neighborhood structures that have been used in the literature. In the bidimensional case, neighborhoods defined by the four or eight nearest voxels represented in Fig. 1 are the most commonly used. Similarly, in the tridimensional case the most frequently used neighborhoods are defined by the six or fourteen nearest voxels represented in Fig 2. In the rest of this paper 4-pixel neighborhoods will be considered for 2D images and 6-voxel neighborhoods for 3D images. Therefore, the associated set of neighbors, or cliques, have only vertical, horizontal and depth possible configurations (see [7,8] for more details).

Once the neighborhood structure has been established, the MRF can be defined. Let z_n denote the random variable indicating the class of the nth image voxel. In the case of K classes, the random variables $z_1, z_2, \ldots, z_N$ take their values in the finite set $\{1, \ldots, K\}$. The whole set of random variables $z_1, z_2, \ldots, z_N$ forms a random field. An MRF is then defined when the conditional distribution of z_n given the other pixels $\boldsymbol{z}_{-n} = (z_1, ..., z_{n-1}, z_{n+1}, ..., z_N)$ only depend on its neighbors $\boldsymbol{z}_{\mathcal{V}(n)}$, i.e.,

$$\mathrm{P}[z_n|\boldsymbol{z}_{-n}] = \mathrm{P}[z_n|\boldsymbol{z}_{\mathcal{V}(n)}] \tag{4}$$

where $\mathcal{V}(n)$ is the considered neighborhood structure.

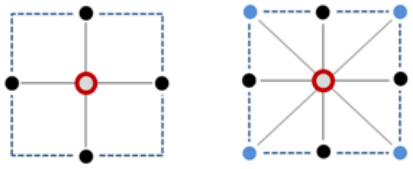

Fig. 1. 4-pixel (left) and 8-pixel (right) neighborhood structures. The considered pixel appears as a void red circle whereas its neighbors are depicted in full black and blue.

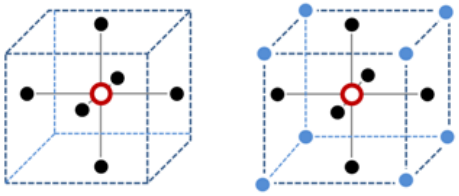

Fig. 2. 6-voxel (left) and 14-voxel (right) neighborhood structures. The considered voxel appears as a void red circle whereas its neighbors are depicted in full black and blue.

In this study we will first consider 2D and 3D Potts Markov fields as prior distributions for $\boldsymbol{z}$. More precisely, 2D MRFs are considered for single-slice (2D) ultrasound images whereas 3D MRFs are investigated for multiple-slice (3D) images. Note that Potts Markov fields are particularly well suited for label-based segmentation as explained in [9]. By the Hammersley-Clifford theorem the corresponding prior for $\boldsymbol{z}$ can be expressed as follows:

$$p(\boldsymbol{z}) = \frac{1}{C(\beta)} \exp\left[\sum_{n=1}^{N} \sum_{n' \in \mathcal{V}(n)} \beta\delta(z_n - z_{n'})\right] \tag{5}$$

where β is the granularity coefficient, $C(\beta)$ is the normalizing constant or partition function [10] and $\delta(\cdot)$ is the Kronecker function. The hyperparameter β tunes the degree of homogeneity of each region in the image. Some simulations have been conducted to show the influence of this parameter on image homogeneity. Synthetic 3D images have been generated using a Potts-Markov model with $K = 3$ classes (corresponding to three gray levels in the image) and a first order neighborhood structure. Fig. 3 shows different realizations corresponding to different values of β. This figure shows that a small value of β induces a noisy image with a large number of regions, contrary to a large value of β that leads to few and large homogeneous regions. In this work, the granularity coefficient β will be fixed a priori. However, it is interesting to mention that the estimation of β has recently received a great attention in the literature [11–14]. Estimating the granularity coefficient using one of these methods is clearly an interesting problem that will be investigated in future works. Finally, it is interesting to note that despite not knowing $C(\beta)$, drawing labels $\boldsymbol{z} = \{z_1, \ldots, z_N\}$ from the distribution (5) can be easily achieved by using a Gibbs sampler [15].

3.2.2. α-Rayleigh parameters

The priors used for the parameters α_k and γ_k $(k = 1, \ldots, K)$ have been introduced and motivated in [4]: α_k has a uniform distribution on $(0, 2)$

$$\alpha_k \sim \mathcal{U}(0, 2) \tag{6}$$

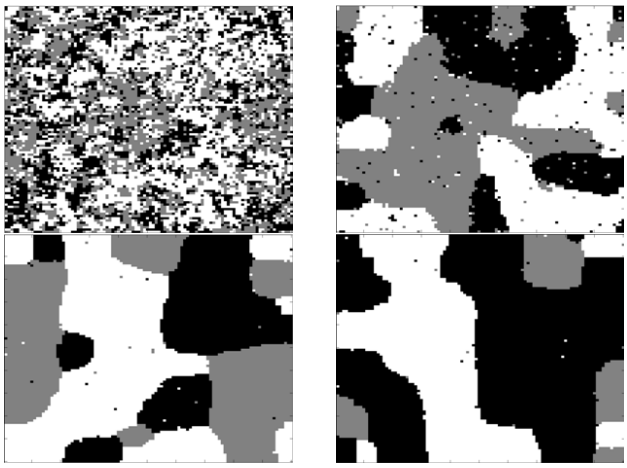

Fig. 3. Synthetic images generated from a 3D Potts-Markov model with (from upper-left to bottom-right) $\beta = 0.6, 1, 1.2, 1.4$

since the interval $(0, 2)$ covers all possible values of this parameter. The prior for the spread γ_k is a conjugate inverse gamma distribution with hyperparameters a_0 and b_0

$$\gamma_k \sim \mathcal{IG}(a_0, b_0), \quad k = 1, \ldots, K \tag{7}$$

where the hyperparameters are fixed in order to obtain a vague prior ($a_0 = 1$ and $b_0 = 1$ will be used in our simulations). Assuming a priori independence between the parameters α_k, γ_k and $\boldsymbol{z}$ the following result is obtained

$$p(\boldsymbol{\theta}) = p(\boldsymbol{\alpha})p(\boldsymbol{\gamma})p(\boldsymbol{z}) = \prod_{k=1}^{K} p(\alpha_k)p(\gamma_k)p(\boldsymbol{z}). \tag{8}$$

3.3. Posterior Distribution of $(\boldsymbol{\theta}, \boldsymbol{z})$

Using Bayes theorem the posterior distribution of $\boldsymbol{\theta}$ can be expressed as follows

$$p(\boldsymbol{\theta}|\boldsymbol{r}) \propto p(\boldsymbol{r}|\boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{z})p(\boldsymbol{\alpha})p(\boldsymbol{\gamma})p(\boldsymbol{z}) \tag{9}$$

where the likelihood $p(\boldsymbol{r}|\boldsymbol{\theta})$ and the prior distribution $p(\boldsymbol{\theta})$ have been defined in (2) and (8), respectively. Unfortunately the posterior distribution (9) is too complex to derive the minimum mean square error (MMSE) or maximum a posteriori (MAP) estimators of the unknown parameters $\boldsymbol{\alpha}$, $\boldsymbol{\gamma}$ and $\boldsymbol{z}$. One can think of using the EM algorithm that has received much attention for mixture problems (see [2,3] for applications to ultrasound images). However, the EM algorithm cannot be used easily for mixtures of α-Rayleigh distributions since the expectation of (9) is not defined if there is at least one characteristic index $\alpha_k < 1$. Also, as explained in [16, p. 259], the EM algorithm suffers from the initialization issue and can converge to local maxima or saddle points of the log-likelihood function. An interesting alternative is to use an MCMC method generating samples that are asymptotically distributed according to the target distribution (9) [15]. The generated samples are then used to approximate the Bayesian estimators. This strategy has been used successfully in many image processing applications [17–21]. One sampling technique allowing the parameters of ultrasound images to be estimated is studied in the next section.

4. HYBRID GIBBS SAMPLER

This section studies a hybrid Metropolis-within-Gibbs sampler that generates samples that are asymptotically distributed according to (9). The conventional Gibbs sampler draws samples according to the conditional distributions associated with the distribution of interest (here the posterior (9)). When a conditional distribution cannot be sampled easily, one can resort to a Metropolis-Hastings move, which generates samples according to an appropriate proposal and accept or reject these generated samples with a given probability. The resulting sampler is referred to as Metropolis-within-Gibbs sampler (see [15] for more details about MCMC methods). The sampler investigated in this section is based on the conditional distributions $\mathrm{P}[\boldsymbol{z}|\boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{r}]$, $p(\boldsymbol{\alpha}|\boldsymbol{z}, \boldsymbol{\gamma}, \boldsymbol{r})$ and $p(\boldsymbol{\gamma}|\boldsymbol{z}, \boldsymbol{\alpha}, \boldsymbol{r})$ that are provided in the next paragraphs.

4.1. Approximation of the Likelihood

Evaluating the likelihood function defined in (2) involves the computation of the following indefinite integral

$$\int_0^\infty \lambda \exp\left[-(\gamma_k\lambda)^{\alpha_k}\right] J_0(r_n\lambda)\, d\lambda. \tag{10}$$

This computation is very time-consuming and is required at every step of the sampler. An efficient way to compute the likelihood (2) is to use the following asymptotic expansions [22, 23]

$$p_{\alpha\mathcal{R}}(r_n|\alpha_k, \gamma_k) = \sum_{p=0}^{P} a_p r_n^{2p+1} + o\left(r_n^{2(P+1)+1}\right) \tag{11}$$

as $r_n \to 0$ and

$$p(r_n|\alpha_k, \gamma_k) = \sum_{p=1}^{P} b_p r_n^{-\alpha_k p-1} + o\left(r_n^{-\alpha_k(P+1)-1}\right) \tag{12}$$

as $r_n \to \infty$ where the coefficients a_p and b_p are

$$a_p = \frac{1}{\alpha_k\gamma_k} \frac{(-1)^p}{(p!)^2 2^{2p}} \Gamma\left(\frac{2p+2}{\alpha_k}\right) \gamma_k^{-2p-1}$$

$$b_p = \frac{(-1)^{p-1}\, 2^{p\alpha_k+1}}{p!\, \pi\gamma_k} \Gamma^2\left(\frac{p\alpha_k+2}{2}\right) \sin\left(\frac{p\pi\alpha_k}{2}\right) \gamma_k^{p\alpha_k+1}$$

The choice of P and other considerations regarding the implementation of (11) and (12) have been studied in [22].

4.2. Conditional probability $\mathrm{P}[\boldsymbol{z}|\boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{r}]$

For each pixel or voxel $n = \{1, 2, \ldots, N\}$, the class label z_n is a discrete random variable whose conditional distribution is fully characterized by the probabilities

$$\mathrm{P}[z_n = k|z_{-n}, r_n, \alpha_k, \gamma_k] \propto p(r_n|z_n = k, \boldsymbol{\alpha}, \boldsymbol{\gamma})p(z_n|\boldsymbol{z}_{-n}) \tag{13}$$

where $k = 1, \ldots, K$ and where it is recalled that K is the number of classes and $\boldsymbol{z}_{-n}$ is the vector $\boldsymbol{z}$ whose nth element has been removed. These posterior probabilities can be expressed as

$$\begin{aligned} &\mathrm{P}[z_n = k|z_{-n}, r_n, \alpha_k, \gamma_k] \\ &\qquad \propto \exp\left[\sum_{n' \in \mathcal{V}(n)} \beta\delta(z_n - z_{n'})\right] \\ &\qquad \times r_n \int_0^\infty \lambda \exp\left[-(\gamma_k\lambda)^{\alpha_k}\right] J_0(r_n\lambda)\, d\lambda. \end{aligned} \tag{14}$$

Note that the posterior probabilities of the label vector $\boldsymbol{z}$ in (14) define an MRF. Sampling from this conditional distribution can be achieved by drawing a discrete value in the finite set $1, \ldots, K$ with probabilities (14). The integral $r_n \int_0^\infty \lambda \exp\left[-(\gamma_k \lambda)^{\alpha_k}\right] J_0(r_n \lambda)\, d\lambda$. is computed using the approximations presented in paragraph 4.1.

4.3. Conditional probability density function $p(\boldsymbol{\alpha}|\boldsymbol{\gamma}, \boldsymbol{z}, \boldsymbol{r})$

The conditional density $p(\boldsymbol{\alpha}|\boldsymbol{\gamma}, \boldsymbol{z}, \boldsymbol{r})$ can be expressed as follows

$$p(\boldsymbol{\alpha}|\boldsymbol{\gamma}, \boldsymbol{z}, \boldsymbol{r}) \propto p(\boldsymbol{r}|\boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{z}) p(\boldsymbol{\alpha})$$

where $p(\boldsymbol{r}|\boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{z})$ is defined in (2) and $p(\boldsymbol{\alpha}) = \prod_{k=1}^{K} p(\alpha_k)$. The generation of samples according to $p(\boldsymbol{\alpha}|\boldsymbol{\gamma}, \boldsymbol{z}, \boldsymbol{r})$ is not easy to perform. We propose in this paper to sample $\boldsymbol{\alpha}$ coordinate-by-coordinate using Metropolis-Hastings (MH) moves, leading to a Metropolis-within-Gibbs algorithm as in [4]. Note that the likelihoods $p_{\alpha\mathcal{R}}(r_n|\alpha_k^*, \gamma_k)$ and $p_{\alpha\mathcal{R}}(r_n|\alpha_k^{(t-1)})$ have been computed using the approximations described in paragraph 4.1.

4.4. Conditional probability density function $p(\boldsymbol{\gamma}|\boldsymbol{\alpha}, \boldsymbol{z}, \boldsymbol{r})$

The conditional density $p(\boldsymbol{\gamma}|\boldsymbol{\alpha}, \boldsymbol{z}, \boldsymbol{r})$ can be expressed as follows

$$p(\boldsymbol{\gamma}|\boldsymbol{\alpha}, \boldsymbol{z}, \boldsymbol{r}) \propto p(\boldsymbol{r}|\boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{z}) p(\boldsymbol{\gamma})$$

where $p(\boldsymbol{r}|\boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{z})$ is defined in (2) and $p(\boldsymbol{\gamma}) = \prod_{k=1}^{K} p(\gamma_k)$. Again, the generation of samples according to this distribution is not easy. Consequently, we propose to update each component of $\boldsymbol{\gamma}$ by using an MH move. The proposal distribution associated with this move is a normal distribution centered on the previous value of the chain with variance σ^2

$$\gamma_k^* \sim \mathcal{N}\left(\gamma_k^{(t-1)}, \sigma^2\right). \tag{15}$$

This strategy is referred to as random walk MH algorithm, as in [4]. Again, the likelihoods $p_{\alpha\mathcal{R}}(r_n|\alpha_k, \gamma_k^*)$ and $p_{\alpha\mathcal{R}}(r_n|\alpha_k, \gamma_k^{(t-1)})$ have been computed using the approximations described in paragraph 4.1.

5. SIMULATIONS RESULTS

This section presents experimental results conducted on synthetic and real data to assess the performance of the proposed α-Rayleigh mixture model and the associated Bayesian estimation algorithm.

5.1. Synthetic Data

The proposed algorithm was first applied to a synthetic 3-component α-Rayleigh mixture displayed in Fig. 4(a). The parameters associated with the mixture components of the 3 different 2D regions are $\boldsymbol{\alpha} = [1.99; 1.99; 1.8]$ and $\boldsymbol{\gamma} = [1; 5; 10]$. Figure 4(b) shows the resulting observation vector $\boldsymbol{r}$, which is the only input provided to the algorithm. Note that the different observations are clearly spatially correlated. The proposed Gibbs sampler has been run for this example using a two dimensional random field with a 4-pixel neighborhood structure. The averaged MMSE estimates and the corresponding standard deviations for the different parameters are reported in Table 1. These estimates have been computed from a single Markov chain of $25,000$ iterations whose first 100 iterations (burn-in period) have been removed. The estimated parameters are clearly in good agreement with the actual values of the α-Rayleigh mixture components. Figure 4(c) shows the class labels estimated by the MAP rule applied to the last samples of the Markov chain. The three classes are recovered even if there are a few misclassifications due to the complexity of the problem.

Table 1. Parameter estimation

	true value	MMSE estimates	standard deviation
α_1	1.99	1.99	0.002
α_2	1.99	1.99	0.003
α_3	1.80	1.79	0.006
γ_1	1.00	1.00	0.003
γ_2	5.00	5.01	0.025
γ_3	10.00	9.96	0.036

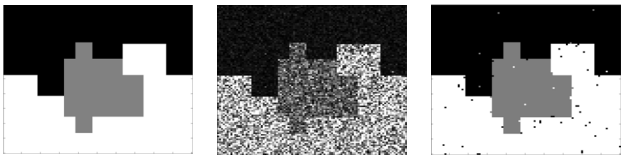

(a) (b) (c)

Fig. 4. (a) True labels, (b) observations $\boldsymbol{r}$, (c) MAP estimates for a 3-class mixture.

5.2. Application to real data

After validating the proposed Gibbs sampler on synthetic data, this section applies the proposed algorithm to the segmentation of a skin lesion. This experiment was conducted using a 3D high frequency ultrasound image of in-vivo skin tissues, acquired at 100MHz with a focalized 25MHz 3D probe. Fig. 5(a) shows a skin lesion outlined by the red rectangle. This region is displayed with coarse expert annotations in figure 5(b). It should be noted that annotations are used to locate the lesion and do not represent ground truth. The proposed Bayesian algorithm was used to label each voxel of an ultrasound image as *healthy* or *lesion* tissue. The estimated labels obtained using a bidimensional random field with granularity coefficient $\beta = 1$ are displayed in figure 5(c). For comparison purposes figure 5(d) shows the estimation results when labels are considered a priori independent, as in [4]. Due to the proposed MRF prior for the labels, the spatial correlations between image voxels are clearly recovered.

Furthermore, the proposed algorithm was also applied in three dimensions using a tridimensional random field with granularity coefficient $\beta = 1$. Three slices of the 3D B-mode image associated with the region of interest depicted in figure 5(a) are shown in figures 6(a), 6(b) and 6(c). Figures 6(d), 6(e) and 6(f) show the results obtained when labels are considered a priori independent, as in [4]. The labels estimated with the proposed 3D method are displayed in figures 6(g), 6(h) and 6(i) where healthy voxels are represented in white and lesion voxels in red. We observe that most of the MAP labels are in very good agreement with the expert annotations. The improvement obtained from considering correlations in the third dimension can be assessed by comparing figures 5(c) and 6(h), which have been computed from the same data slice. We observe that using a 3D MRF reduces significantly the number of misclassifications and improves the agreement with the expert annotations.

6. CONCLUDING REMARKS

A spatially coherent finite mixture of α-Rayleigh distributions was proposed to represent the distribution of envelope ultrasound images backscattered from multiple tissues. This work extends the mixture

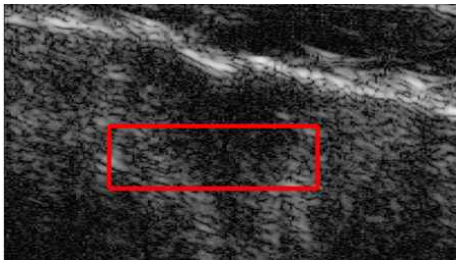

(a) Dermis view with skin lesion (ROI $= 100 \times 100 \times 3$).

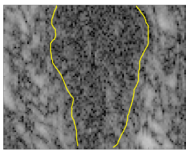

(b) ROI (slice 2)

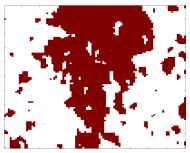

(c) MRF Labels $\boldsymbol{z}$

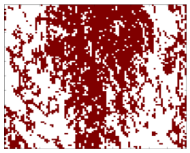

(d) Independent Labels $\boldsymbol{z}$

Fig. 5. Log-compressed ultrasound images of skin lesion and the corresponding estimated labels (*healthy* = white, *lesion* = red). Figure (c) proposed 2D algorithm, (d) algorithm of [4]

.

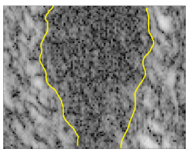

(a) ROI (slice 1)

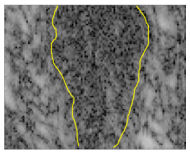

(b) ROI (slice 2)

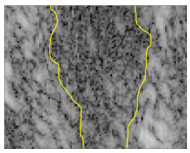

(c) ROI (slice 3)

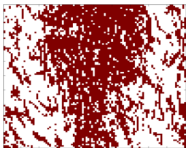

(d) Ind. Labels $\boldsymbol{z}$ (1)

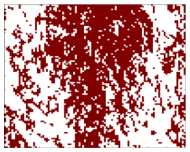

(e) Ind. Labels $\boldsymbol{z}$ (2)

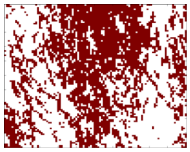

(f) Ind. Labels $\boldsymbol{z}$ (3)

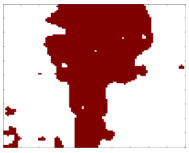

(g) MRF Labels $\boldsymbol{z}$ (1)

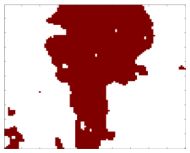

(h) MRF Labels $\boldsymbol{z}$ (2)

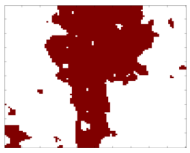

(i) MRF Labels $\boldsymbol{z}$ (3)

Fig. 6. Log-compressed ultrasound images of skin lesion and the corresponding estimated labels (*healthy* = white, *lesion* = red). Figures (d)-(f) show the results obtained by considering that voxel labels are independent, as in [4]. Figure (g)-(i) shows the results obtained with the proposed 3D Markov random field (MRF) method.

model proposed in [4] which did not consider spatial information. Spatial correlation was introduced into the model by a Markov random field that promotes dependance between neighbor pixels. Based on the proposed model, a Bayesian segmentation method was derived. Bidimensional and tridimensional implementations of this segmentation method have been presented using a Markov chain Monte Carlo algorithm that jointly estimates the unknown parameters of the mixture model and classifies voxels into different tissues. The method was successfully applied to the segmentation of an in-vivo lesion in a high frequency 3D ultrasound image. Future works include the study of estimation algorithms for the granularity coefficient defining the Markov random field prior.

7. REFERENCES

[1] J. Noble and D. Boukerroui, "Ultrasound image segmentation: a survey," *IEEE Trans. Med. Imag.*, vol. 25, no. 8, pp. 987 –1010, Aug. 2006.

[2] M. H. R. Cardinal, J. Meunier, G. Soulez, R. L. Maurice, E. Therasse, and G. Cloutier, "Intravascular ultrasound image segmentation: a three-dimensional fast-marching method based on gray level distributions," *IEEE Trans. Med. Imag.*, vol. 25, no. 5, pp. 590–601, May 2006.

[3] F. Destrempes, J. M. J., M. F. Giroux, G. Soulez, and G. Cloutier, "Segmentation in ultrasonic B-mode images of carotid arteries using mixture of Nakagami distributions and stochastic optimization," *IEEE Trans. Med. Imag.*, vol. 28, no. 2, pp. 215–229, Feb. 2009.

[4] M. A. Pereyra, N. Dobigeon, H. Batatia, and J.-Y. Tourneret, "Labeling skin tissues in ultrasound images using a generalized Rayleigh mixture model," in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Proc. (ICASSP)*, Prague, Czech Republic, May 2011.

[5] S. Z. Li, *Markov Random Field Modeling in Image Analysis, 3rd ed.* New York, London: Springer, 2009.

[6] M. A. Pereyra and H. Batatia, "A Levy flight model for ultrasound in skin tissues," in *Proc. IEEE Ultrason. Symp.*, San Diego (CA), USA, Oct. 2010.

[7] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 6, no. 6, pp. 721–741, Nov. 1984.

[8] J. Besag, "Spatial interaction and the statistical analysis of lattice systems," *J. Roy. Stat. Soc. Ser. B*, vol. 36, no. 2, pp. 192–236, 1974.

[9] F. Y. Wu, "The Potts model," *Rev. Mod. Phys.*, vol. 54, no. 1, pp. 235–268, Jan. 1982.

[10] R. Kindermann and J. L. Snell, *Markov random fields and their applications*. Providence: RI: Amer. Math. Soc., 1980.

[11] Z. Zhou, R. Leahy, and J. Qi, "Approximate maximum likelihood hyperparameter estimation for Gibbs prior," *IEEE Trans. Image Process.*, vol. 6, no. 6, pp. 844–861, June 1997.

[12] X. Descombes, R. Morris, J. Zerubia, and M. Berthod, "Estimation of Markov random field prior parameters using Markov chain Monte Carlo maximum likelihood," *IEEE Trans. Image Process.*, vol. 8, no. 7, pp. 945–963, June 1999.

[13] J. Moller, A. N. Pettitt, R. Reeves, and K. K. Berthelsen, "An efficient Markov chain Monte Carlo method for distributions with intractable normalising constants," *Biometrika*, vol. 93, no. 2, pp. 451–458, June 2006.

[14] C. McGrory, D. Titterington, R. Reeves, and A. Pettitt, "Variational Bayes for estimating the parameters of a hidden Potts model," *Statistics and Computing*, vol. 19, pp. 329–340, 2009.

[15] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*. New York: Springer-Verlag, 1999.

[16] J. Diebolt and E. H. S. Ip., "Stochastic EM: method and application," in *Markov Chain Monte Carlo in Practice*, W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, Eds. London: Chapman & Hall, 1996, pp. 259–273.

[17] N. Dobigeon and J.-Y. Tourneret, "Bayesian orthogonal component analysis for sparse representation," *IEEE Trans. Signal Process.*, vol. 58, no. 5, pp. 2675–2685, May 2010.

[18] N. Dobigeon, A. Hero, and J.-Y. Tourneret, "Hierarchical Bayesian sparse image reconstruction with application to MRFM," *IEEE Trans. Image Process.*, vol. 18, no. 9, pp. 2059 –2070, Sept. 2009.

[19] T. Vincent, L. Risser, and P. Ciuciu, "Spatially adaptive mixture modeling for analysis of fMRI time series," *IEEE Trans. Med. Imag.*, vol. 29, no. 4, pp. 1059 –1074, April 2010.

[20] K. Kayabol, E. Kuruoglu, and B. Sankur, "Bayesian separation of images modeled with MRFs using MCMC," *IEEE Trans. Image Process.*, vol. 18, no. 5, pp. 982 –994, May 2009.

[21] M. Mignotte, "Image denoising by averaging of piecewise constant simulations of image partitions," *IEEE Trans. Image Process.*, vol. 16, no. 2, pp. 523 –533, Feb. 2007.

[22] Z. Sun and C. Han, "Heavy-tailed Rayleigh distribution: A new tool for the modeling of sar amplitude images," in *Proc. (IGARSS 08). IEEE Int. Geosc. and Remote Sensing Symp.*, vol. 4, July 2008, pp. 1253–1256.

[23] C. L. Nikias and M. Shao, *Signal Processing with Alpha-Stable Distribution and Applications*. New York (NJ): Wiley, 1995.