

MULTIPLE RESOLUTION BAYESIAN SEGMENTATION OF ULTRASOUND IMAGES

Edward A. Ashton and Kevin J. Parker

Department of Electrical Engineering
University of Rochester
Rochester, NY 14627

We propose a novel method for obtaining the maximum *a posteriori* (MAP) probabilistic segmentation of speckle-laden ultrasound images. Our technique is multiple-resolution based, and relies on the conversion of speckle images with Rayleigh statistics to subsampled images with Gaussian statistics. This conversion reduces computation time, as well as allowing accurate parameter estimation for a probabilistic segmentation algorithm. Results appear to provide improvements over previous techniques in terms of low-contrast detail and accuracy. © 1995 Academic Press, Inc.

Key words: Bayesian classification; multiple resolution; Rayleigh statistics; speckle; ultrasound.

1. INTRODUCTION

The first step toward automatic analysis or evaluation of a given image type is generally considered to be a successful segmentation. In the field of biomedical imaging, this step has been accomplished using a variety of approaches including adaptive or non-adaptive k-means type clustering algorithms for imaging modalities such as MRI [1-3]. These techniques, however, are not effective for the segmentation of ultrasound images. The primary reason for this has to do with the statistical nature of speckle-laden images. Images produced by modalities such as radiography and MRI may be characterized by Gaussian statistics, i.e. these images may be modeled as connected regions of slowly varying intensity corrupted by additive white Gaussian noise. Ultrasound images, however, are generally characterized by Rayleigh statistics [4, 5]. A Rayleigh probability density function is described by only one parameter, α , as follows:

$$p(x) = (x/\alpha^2)exp(-x^2/2\alpha^2) \tag{1}$$

The mean of a given region and its standard deviation are linearly dependent, with a constant ratio of approximately 1.91 [4].

This signal-dependent noise property makes parameter estimation for ultrasound images extremely difficult. Simple clustering algorithms like k-means fail, because the true

means are not, in fact, the minimum mean squared error (MMSE) solution. Much of the literature on segmentation of images demonstrating Rayleigh statistics has focused on nonlinear filtering and morphological operations. Sperry [6] has proposed an algorithm that makes use of successive iterations of thresholding followed by morphological filtering for the segmentation of ultrasound B-scans. Taxt et al. [7] have presented an interactive segmentation technique involving nonlinear noise reduction filtering followed by manual thresholding. More recently, Herlin and Ayache [8] have proposed a multi-step algorithm that operates on video data in polar coordinates. This technique requires median filtering, followed by edge detection and segmentation using 2-D deformable models. All these techniques share a common difficulty: they require that nonlinear operations be performed on the data prior to or during segmentation. For this reason, it is difficult to prove that they reach any sort of statistically optimal solution.

This problem is avoided by Rignot and Chellappa [9], who have proposed a segmentation algorithm that makes use of a Markov random field (MRF) model for region distribution and a two-step approach to the energy minimization problem. They make use of the iterated conditional modes (ICM) algorithm of Besag [10] to maximize the posterior probability density of the region labels, followed by a maximum posterior marginal (MPM) algorithm which maximizes the posterior marginal density of the region labels at each pixel location. This technique has been shown to be applicable to complex-amplitude SAR segmentation. However, like the Herlin algorithm, it is not designed for operation on the intensity data that is most commonly stored in ultrasound B-scan images.

A more applicable algorithm has been proposed by Muzzolini et al. [11]. Their approach calls for a quadtree-structured multiresolution texture segmentation, making use of the Metropolis algorithm [12] to obtain an optimum labeling of image blocks at each level and then splitting and merging blocks among levels to reach the lowest energy configuration. Difficulties with this algorithm include the assumption that region boundaries should fall along arbitrarily assigned block divisions, the *ad hoc* nature of the image modeling, and the use of the computationally-intensive Metropolis algorithm. Use of Besag's ICM for energy minimization would improve the applicability of this algorithm. However, ICM requires accurate initial parameter estimation, which is very difficult under Rayleigh statistics.

We propose avoiding these difficulties by taking advantage of the Central Limit Theorem, which tells us that the average of a large number of random variables must tend toward a Gaussian distribution around their collective mean. This leads us to a multiple-resolution segmentation algorithm, making use of local-average filtering and decimation, with each coarse-resolution segmentation used to estimate the parameters of the next finer resolution. This accurate parameter estimation allows us to make use of the computationally efficient ICM energy minimization algorithm. It should be noted that, while the filter-and-decimate process would appear to obscure small details, the target resolution of this algorithm is ultimately limited by the pulse characteristics of the imaging system, which determine the "speckle spot size," assumed to be known in our analysis. Note here that although speckle characteristics are subject to variation within an image due to beam diffraction and attenuation, our algorithm is effective as long as the "worst case" is considered.

2. MODELING

Ideally, we would like a segmentation algorithm to be able to separate an ultrasound image into large contiguous regions, each of which would represent a particular tissue type or structure. We would therefore like to incorporate a spatial constraint into our

segmentation process. This may be accomplished by using a random field model. Bouman [13] and others use the MRF. However, this model is defined in terms of local properties, making the derivation of a global joint distribution difficult. We therefore make use of the similar but more mathematically tractable GRF.

Assume that a random field x is defined over an $N \times N$ lattice (square array) of sites denoted by S , with the individual sites (pixels) denoted by $s_{1,1}, s_{1,2}, s_{1,3}, \dots, s_{N,N}$. For segmentation purposes, each site on x may take on any integer value from 1 to k , where k is the number of classes (tissue types) into which the image is being segmented. Let w be a realization of x . In order to define the random field, it is necessary first to define a neighborhood system G_s on the lattice S . The neighborhood system determines which pixels on the lattice will directly affect the classification of site s . In this work, a first order, or four-point, neighborhood system is used, so that only pixels which share a side with site s will be considered its neighbors.

A *clique* is defined as a set of sites on S such that all points are mutual neighbors. In a four-point neighborhood system, there are four two-point cliques, given by $(s_{n,m}, s_{n-1,m})$, $(s_{n,m}, s_{n+1,m})$, $(s_{n,m}, s_{n,m-1})$ and $(s_{n,m}, s_{n,m+1})$. There is also a single one-point clique, $(s_{n,m})$. However, as this adds no information, it is disregarded.

The random field x is an MRF with respect to G_s if

$$P(x = w) > 0 \text{ for all } w \tag{2}$$

and

$$P(X_s = x_s \mid X_r = x_r, r \neq s) = P(X_s = x_s \mid X_r = x_r, r \in G_s) \tag{3}$$

i.e. all configurations must have a non-zero probability, and the conditional probability of a particular site must depend only on its immediate neighbors as defined by G_s . Note that in the above equations, x is a particular realization of the random field X . This model has been used in a number of applications, including [9, 13]. However, it has certain difficulties which make the Gibbs distribution more attractive. Primary among these is that the MRF is defined only in terms of local characteristics, making the calculation of a global joint distribution very difficult.

In contrast, the GRF provides a global model for an image, specifying a probability density function of the form:

$$P(X = x) = e^{-U(x)} / Z \tag{4}$$

where $U(x)$ is the Gibbs potential, or Gibbs energy, given by the summation over all cliques C :

$$U(x) = \sum_C V_C(x), \tag{5}$$

$V_C(x)$ is the clique potential, given by

$$V_C(x_i, x_j) = \begin{cases} -\beta & \text{if } x_i = x_j \\ +\beta & \text{otherwise} \end{cases} \tag{6}$$

and Z is a normalizing constant, called the partition function, given by

$$Z = \sum_w e^{-U(X=w)}. \quad (7)$$

Since Z is the sum of the numerator exponents over all possible vectors x , it is extremely difficult to compute. This intractability prevents realization of a GRF using direct calculation, thus requiring the use of stochastic relaxation methods.

Using the Gibbs formulation, it is possible to calculate the global probability of a given image configuration based on the clique potentials, which depend solely on the local properties of the pixels.

Additionally, it should be noted that a unique GRF exists for every MRF as long as the Gibbs field is defined in terms of cliques on a neighborhood system [14]. In other words, it is possible to uniquely specify an MRF by specifying Gibbs potentials. This means that nothing is lost through making use of the more mathematically tractable Gibbs model.

Dubes and Jain have shown that it is possible to achieve dramatic image reconstruction results using multiparameter GRF modeling and algorithms such as the Gibbs sampler [15] if all parameters of the Gibbs distribution and degrading noise are known. However, there is currently no algorithmic method for deriving these parameters from an arbitrary data set. We therefore assume a simple *pairwise interaction model*. This model stipulates only that there should be a dependence in classification between neighboring pixels, and requires only a weighting constant, β , which determines the relative importance of the two terms in the Bayesian formulation. Unfortunately, β must be optimized experimentally for a class of images. In this work, $\beta = 1.0$ was used for all segmentations.

One possible drawback to the choice of the pairwise interaction model for this application is the model's symmetry in the vertical and horizontal directions. It is known that ultrasonic images have better spatial resolution parallel to the ultrasound propagation direction than perpendicular to it. It is reasonable to suppose, therefore, that a model which takes this effect into account might provide improved results. The formulation and testing of such a model is one possible direction of future research.

A more detailed examination of random field models may be found in [10, 15, 16].

3. CLUSTERING ALGORITHM

The core of our algorithm is a modification of the adaptive clustering algorithm of Pappas [17]. The Pappas algorithm as implemented in [1] is useful because it considers class average to be a local rather than a global property, which allows for variations in the mean of a given class from region to region, compensating for the fact that the intensity of ultrasound echo signals tends to vary with depth due to non-uniform attenuation. However, this algorithm assumes that pixel intensities are given by a slowly varying class mean corrupted by globally-uniform additive white Gaussian noise. Clearly, this is not an accurate model for an image which is characterized by Rayleigh statistics.

We would like to be able to find a more accurate image model while still retaining Gaussian statistics, which are far more mathematically tractable and separable than the Rayleigh statistics which actually characterize ultrasound images. Therefore, we assume

the pixel values within a given region to be Gaussian distributed about the slowly-varying mean μ , with variance $\sigma_{x_s}^2$, where x is the classification of the pixel at site s . $\sigma_{x_s}^2$ is estimated independently for each class, and is proportional to the local class mean. This assumption is reasonably accurate for low-pass filtered (locally averaged) and decimated speckle images that were originally governed by Rayleigh statistics. We rely heavily on this model for parameter estimation. As Rayleigh statistics begin to dominate at finer resolutions, we increase our reliance on the previously mentioned Gibbs Random Field model with a strong neighborhood constraint.

3.1 Modified adaptive clustering

We would now like to find a global maximization of the probability of segmentation field x , given the observed image data field y . Using Bayes' rule, it may be shown that,

$$P(x | y) \propto P(y | x)P(x) \tag{8}$$

In this equation, $P(x)$ is the *a priori* probability of the segmentation. This is given by the GRF model, as described in Section 2.1.

We also require a functional form for $P(y | x)$. This will be dependent upon our model for the image data. Under the previously stated assumptions, the conditional probability may be written as

$$P(y | x) \propto \prod_s \left(\frac{1}{\sigma_{x_s}} \right) \exp - \sum_s \left[\frac{1}{2\sigma_{x_s}^2} (y_s - \mu_{x_s})^2 \right] \tag{9}$$

where σ_{x_s} is the overall noise variance for class x_s .

Substituting into Eq. (8), and summing the exponentials, we obtain

$$P(x | y) \propto \prod_s \left(\frac{1}{\sigma_{x_s}} \right) \exp - \left[\sum_s \left[\frac{1}{2\sigma_{x_s}^2} (y_s - \mu_{x_s})^2 \right] + \sum_C V_C(x) \right] \tag{10}$$

We wish to maximize this function, or, equivalently, to maximize its natural logarithm, given by

$$\ln(P(x | y)) \propto - \sum_s \left[\ln(\sigma_{x_s}) + \frac{1}{2\sigma_{x_s}^2} (y_s - \mu_{x_s})^2 \right] - \sum_C V_C(x). \tag{11}$$

This is achieved in an iterative fashion, using ICM. We assume an initial segmentation and an initial estimation of the parameters μ_{x_s} and σ_{x_s} . Because the ICM algorithm is a local minimization only, mathematically equivalent to simulated annealing with instantaneous freezing [10], it is important that these estimations be reasonably accurate.

An $I \times J \times M$ matrix is formed, where I is the width of the image, J is the height of the image, and M is the number of classes into which the image is to be segmented. Each point (i, j, m) in that matrix then stores the average of all pixels of class m within a rectangular window of width W centered at point (i, j) . W begins as the size of the image, and is iteratively reduced to some final value W_{min} .

Direct calculation of the local class averages at each point in the image is computationally prohibitive, particularly in the case of a large window size. For this reason, it

is assumed that the class averages vary smoothly and slowly. Under this assumption, it is an acceptable approximation to directly calculate the class averages at a regularly-spaced grid of points in the image and fill in the remainder of the points through bilinear interpolation. The maximum spacing of the grid points will depend upon the expected smoothness of the class averages. In this work, as in the Pappas paper, a spacing of $W/2$ is used. An illustration of the sample grid with sliding estimation window is given in figure 1.

At each window size, Eq. 10 is maximized on a point-by-point basis, with μ_{x_s} re-estimated after each iteration. When a local maximum is reached, σ_{x_s} is re-estimated, and the estimation window is reduced. Eq. (10) is then remaximized. This process is repeated until the smallest desired estimation window is reached. A more detailed examination of the ICM energy minimization process may be found in [10].

3.2 Multiple resolution segmentation

The most obvious problem with this algorithm is the assumed initial segmentation. As has been mentioned, standard clustering algorithms such as k-means or the closely related LBG algorithm [18] fail under Rayleigh statistics. However, we have claimed that, through low-pass filtering and decimation, we are able to force our image statistics toward Gaussian. This leads us to a multiple resolution segmentation scheme, which is summarized below.

1. Lowpass filter the original image using an FIR filter of support $2n \times 2n$. Optimization of n is discussed in the following section. Decimate by a factor of n .
2. Obtain initial cluster estimate using k-means algorithm.
3. Set the GRF weighting $\beta = \beta_0$




Fig. 1. Sampling grid for local class average estimation. All pixels in the window defined by W are used to estimate mean and variance at site s . Values for sites between the grid points are filled in using bilinear interpolation.

4. Apply clustering algorithm.
5. Expand cluster matrix (the matrix containing class assignments) to next larger size by replicating by n , then decimating by $n-1$.
6. Let $n = n-1$.
7. Lowpass filter the original image using an FIR filter of support $2n \times 2n$. Decimate by a factor of n .
8. If $n > 1$, let $\beta = \beta + \delta\beta$. Go to 4.
9. Apply clustering algorithm to original image.
10. STOP.

The parameters of interest in this algorithm are n , which determines the maximum coarseness of resolution, β , which determines the low-resolution weighting of the GRF model, and $\delta\beta$, which determines the rate of increase in the reliance on the GRF model at finer resolutions. As has been mentioned, there is at present no algorithmic means for optimizing these parameters. They must therefore be optimized experimentally. Fortunately, our experiments, as well as the many applications of the GRF model in the segmentation literature, indicate that the β parameters may be optimized a single time for application to a broad class of images. All experiments described below were carried out with $\beta \approx 1.0$ and $\delta\beta = 0.55$. Details of the optimization of n , which is not well covered in the literature, will be given in the following section.




Fig. 2. Normalized histogram of a 512x512 Rayleigh speckle field, together with a Gaussian curve with identical mean and variance. Clearly, Gaussian assumptions are not well suited to Rayleigh statistics.

4. EXPERIMENTAL PROCEDURES AND RESULTS

We first wished to confirm that our filter-and-decimate procedure did, in fact, force our Rayleigh-type images toward Gaussian statistics. To this end, a histogram was taken for a 512×512 artificially produced single-level speckle field. This image, like all other artificial images described below, was produced by convolving a white noise field with a pulse of support 7×7 pixels. The pulse had a Gaussian profile in the x direction with $\sigma^2 = 3.0$, and a sinusoidal profile in the y direction with frequency $\pi/3$ radians/pixel. See [18] for a more detailed discussion of this method of speckle generation. The envelope of the resultant image was then taken using the Hilbert transform. A plot of this histogram, together with a normalized Gaussian curve created using the mean and variance found in the image, is given in figure 2. Clearly, the Gaussian distribution is a poor model for a Rayleigh-distributed image. This speckle image was then blurred and decimated with filter support ranging from $n = 1$ to $n = 8$. Histograms of the resulting images were taken, and these were compared to normalized Gaussian curves. A plot of the mean squared error between the two curves at each decimation level is given in figure 3. A plot of the histogram of the image created with $n = 6$, together with the appropriate Gaussian curve, is given in figure 4.

The curve given in figure 3 was used in the optimization of n . Higher values of n produce more nearly Gaussian statistics. However, they also result in coarser resolution and loss of fine detail. We therefore chose the point at which the error curve begins to flatten, and thus the lowest value of n that would not incur serious estimation error. All experiments described below were carried out with an initial value of $n = 6$.

In order to test the efficacy of this algorithm against known results, four artificial, multilevel speckle fields were created. figures 5(a) and 5(b) are three-level fields. figure 5(c) is a five-level field, and figure 5(d) is a five-level decreasing contrast and target-




Fig. 3. Mean squared error between histograms of the filtered and decimated images and normalized Gaussian curves for filter support of $n = 1$ to $n = 8$.




Fig. 4. Histogram of a 512x512 Rayleigh speckle field which has been blurred and decimated with filter support $n = 6$. Note that the histogram is accurately modeled by an ideal Gaussian curve with identical mean and variance.




Fig. 5. Artificial test images.




Fig. 6. Segmentations of test images in figure 5.




Fig. 7. (a) One slice of our three dimensional ultrasound phantom. (b) Ideal volume-rendered reconstruction of the full phantom. (c) Reconstruction using noise-reduction filtering followed by ideal thresholding. (d) Reconstruction using our multiple resolution segmentation algorithm.

size test pattern. We tested the segmentation algorithm on all four test images, and compared the results to the known region distributions. Note that in these cases the correct number of classes was known *a priori*. Segmentations are given in figures 6(a) - 6(d). Percentages of misclassified pixels were 5.94 percent, 2.98 percent, 3.13 percent and 6.31 percent, respectively.

We next wished to compare our algorithm to previous techniques, particularly with respect to low-contrast accuracy. A multiple slice two level artificial speckle field was created, each slice of which resembled a circular lesion on a uniform background. The slices were segmented using our technique, and by noise-reduction filtering followed by ideal thresholding [7]. The results were then volume-rendered for comparison to an ideal reconstruction created from the original phantom. Results of this experiment are given in figure 7.

The algorithm was then tested against an ensemble of 110 two-level artificial speckle fields, each of which resembled a lesion on a uniform background, with varying contrast levels and lesion sizes. Contrast ratio, which is defined here as the ratio of the mean value of all pixels in class 2 to the mean value of all pixels in class 1, was varied from 1.2 to 2.0. The lesions were circular and centered in the images, and lesion radius varied from 6 pixels to 36 pixels. Image size was 128 by 128 pixels. A representative sampling of these images and their segmentations is given in figure 8. A plot of contrast ratio vs. percent pixel error for each lesion size from 6 to 21 pixels is given in figure 9. The results at larger lesion sizes were nearly identical to those at $r = 21$. The results of this analysis are important in that they demonstrate that accurate segmentations can be made under the most demanding conditions - small target size and low contrast. For example, with a pulse length of 7 pixels and contrast ratio of 1.5, lesions of radius 26 to 36 pixels are resolved with percent pixel error less than 5 percent. In contrast, Sperry [6] has reported low levels of accuracy in segmentation at contrast ratios exceeding 2.0.

Finally, we wished to confirm that this technique was applicable to actual ultrasound data. We segmented a transesophageal echocardiogram (TEE) scan of a human heart (figure 10a), and a standard ultrasound scan of a human gallbladder (figure 11a). In these cases, where the correct number of classes is not necessarily known, selection of k becomes an issue. Our experiments have shown that good results are usually obtained in the range of $k = 3$ to $k = 5$, depending on the complexity of the image and the features which are of interest. Segmentations are given in figures 10(b) and 11(b). The image in figure 10 was segmented with $k = 4$. The image in figure 11 was segmented with $k = 3$. Note that the gradual loss of intensity with depth of figure 10(a) is successfully segmented by our algorithm. figure 10(b) demonstrates that the top two (brightest) classes encircle the left ventricular chamber despite the loss of intensity. In comparison, abrupt shadows such as those caused by the gallstone in figure 11 result in boundaries between segmented classes.

5. DISCUSSION

Our experiments indicate that this technique provides an accurate and reproducible method for segmentation of multilevel speckle fields. Pixel misclassification was minimal in all test images other than those with very low contrast. High errors in the small size lesion tests are attributable not to the minimum target resolution of the algorithm, but rather to the fact that when a very small proportion of all pixels are in the lesion class and contrast is low, the algorithm achieves a lower MSE between data points and cluster means by dividing the background class. Sensitivity and contrast discrimination were both better than results given in [6], and although the "correct" segmentation of the




Fig. 8. Low-contrast lesion images, with segmentations. Each image contains a centered circular lesion. Contrast ratios vary from 1.40 in (a) to 1.92 in (f). Note that previous techniques have had difficulty with contrast ratios exceeding 2.0.




Fig. 9. Percent misclassified from low-contrast lesion test for lesion radius ranging from $r = 6$ to $r = 21$. Note that the pulse length used in forming these images was 7 pixels.




Fig. 10. (a) TEE scan of a human heart. (b) Segmentation of (a).

actual data given in figures 10 and 11 is somewhat subjective, our segmentations seem to agree with visual separation of distinct regions. Additionally, our algorithm is only moderately computationally intensive, requiring approximately 6 minutes of CPU time on a Sun SPARC 1 workstation for the segmentation of a three-class, 256 x 256 image.

The primary assumptions required for the use of our algorithm are that a speckle spot size for the imaging system is known, that approximately Gaussian statistics are generated by our filter and decimation steps, that shading of intensity by attenuation or depth varying gain is smoothly varying, and that a limited number of classes of tissue exist within an image. These assumptions appear to be valid for a wide range of imaging situations, and the results show that the new algorithm is tolerant of deviations from these assumptions. We believe the results should be useful in assisting with 2-D measurements of anatomical structures, and in 3-D reconstructions of complex geometries from ultrasound B-scan images.




Fig. 11. (a) Ultrasound scan of a human gallbladder, with gallstone. (b) Segmentation of (a).

6. ACKNOWLEDGMENTS

This work was carried out with the support of the National Science Foundation, NIH, and the Center for Electronic Imaging Systems.

REFERENCES

- [1] Chang, M.M., Sezan, M.I., and Tekalp, A.M., Bayesian Segmentation of MR Images Using 3-D Gibbsian Priors, in *SPIE Conference on Image and Video Processing*, San Jose, CA (1993).
- [2] Berg, M.J., Ashton, E.A., Ketonen, L., Parker, K.J., Chen, C.W., and Weisberg, J., Novel Method for Automatically Determining Volumes of Hippocampus and Other Anatomic Structure on MRI Scans, in *Epilepsia*, Miami, FL (1993).
- [3] Ashton, E.A., Berg, M.J., Parker, K.J., Chen, C.W., Weisberg, J., and Ketonen, L., Segmentation and Feature Extraction Techniques, with Applications to MRI Head Studies, *Magnetic Resonance in Medicine* 33, 670 - 677 (1995).
- [4] Burckhardt, C., Speckle in Ultrasound B-mode Scans, *IEEE Trans. Sonics, Ultrasonics SU-25*, 1 - 6 (1978).
- [5] Tuthill, T.A., Sperry, R.H., and Parker, K.J., Deviations from Rayleigh Statistics in Ultrasonic Speckle, *Ultrasonic Imaging* 10, 81 - 89 (1988).
- [6] Sperry, R.H., and Parker, K.J., Segmentation of Speckle Images Based on Level-crossing Statistics, *J. Opt. Soc. Am.* 8, 490 - 498 (1991).
- [7] Taxt, T., Lundervold, A., and Angelsen, B., Noise Reduction and Segmentation in Time-varying Ultrasound Images, in *10th Int. Conf. on Patt. Recogn.*, (Atlantic City, NJ (1990).
- [8] Herlin, I., and Ayache, N., Features Extraction and Analysis Methods for Sequences of Ultrasound Images, *Image and Vision Computing* 10, 673 - 682 (1992).
- [9] Rignot, E., and Chellappa, R., Algorithms for Segmentation of Complex-amplitude SAR Data, *NASA Tech Briefs* 17 28 - 30 (1993).
- [10] Besag, J., On the Statistical Analysis of Dirty Pictures, *J. Royal Stat. Soc.* 48, 259 - 302 (1986).
- [11] Muzzolini, R., Yang, Y.-H., and Pierson, R., Multiresolution Texture Segmentation with Application to Diagnostic Ultrasound Images, *IEEE Trans. Med. Imaging* 12, 108 - 123 (1993).
- [12] Metropolis, N., Rosenbluth, A., and Teller, A., Equations of State Calculations by Fast Computing Machines, *J. Chem. Phys.* 21, 1087 - 1091 (1953).
- [13] Bouman, C., and Liu, B., Multiple Resolution Segmentation of Textured Images, *IEEE Trans. Pattern Anal. Mach. Intel.* 13, 99 - 113 (1991).
- [14] Besag, J., Spatial Interaction and the Statistical Analysis of Lattice Systems, *J. Royal Stat. Soc.* 36 192 - 326 (1974).
- [15] Geman, S., and Geman, D., Stochastic Relaxation, Gibbs Distribution and the Bayesian Restoration of Images, *IEEE Trans. Pattern Anal. Mach. Intel.* 6, 721 - 741 (1984).
- [16] Dubes, R.C., and Jain, A.K., Random Field Models in Image Analysis, *J. Applied Stat.* 16 131 - 164 (1989).
- [17] Pappas, T.N., An Adaptive Clustering Algorithm for Image Segmentation, *IEEE Trans. Signal Processing SP-40*, 901 - 914 (1992).
- [18] Wagner, R.F., Smith, S.W., Sandrik, J.M., and Lopez, H., Statistics of Speckle in Ultrasound B-Scans, *IEEE Trans. Sonics, Ultrasonics* 30, 156 - 163 (1983).