Evaluation Of Cardiac Ultrasound Data by Bayesian Probability Maps

Mattias Hansson[1], Sami Brandt[1,2], Petri Gudmundsson[3], and Finn Lindgren[4]

[1] School of Technology, Malmö University, Sweden, `mattias.hansson@mah.se`.
[2] Information Processing Laboratory, Oulu University, Finland.
[3] Faculty of Health and Society, Malmö University, Sweden.
[4] Centre for Mathematical Sciences, Lund University, Sweden.

Abstract. In this paper we present improvements to our Bayesian approach for describing the position distribution of the endocardium in cardiac ultrasound image sequences. The problem is represented as a latent variable model, which represents the inside and outside of the endocardium, for which the posterior density is estimated. We start our construction by assuming a three-component Rayleigh mixture model: for blood, echocardiographic artifacts, and tissue. The Rayleigh distribution has been previously shown to be a suitable model for blood and tissue in cardiac ultrasound images. From the mixture model parameters we build a latent variable model, with two realizations: tissue and endocardium. The model is refined by incorporating priors for spatial and temporal smoothness, in the form of total variation, connectivity, preferred shapes and position, by using the principal components and location distribution of manually segmented training shapes. The posterior density is sampled by a Gibbs method to estimate the expected latent variable image which we call the Bayesian Probability Map, since it describes the probability of pixels being classified as either heart tissue or within the endocardium. By sampling the translation distribution of the latent variables, we improve the convergence rate of the algorithm. Our experiments show promising results indicating the usefulness of the Bayesian Probability Maps for the clinician since, instead of producing a single segmenting curve, it highlights the uncertain areas and suggests possible segmentations.

1 Introduction

Echocardiography is more accessible, mobile and inexpensive compared to other imaging techniques and has become a widely used diagnostic method in cardiology in recent years. Unfortunately ultrasound images struggle with inherent problems which in large part stem from noise, and is often referred to as *speckle contamination*. Speckle is the result of interference between echoes, which are produced when the ultrasound beam is reflected from tissue, and has the properties of a random field, see [1, 2]. The use of the Rayleigh distribution in modeling speckle in ultrasonic B-scan images is well-established through early works, such as [3, 1], and more recently in [4].

There is much previous work done in the field of segmentation of cardiac ultrasound images, of which an excellent overview is given in [5]. Here we will only mention those works which, like our algorithm, treat segmentation of blood and tissue as a pixel-classification or region-based problem. Our model makes a dependency assumption

of neighboring pixels via total variation. A similar approach is employed in [6–10], where Markov random field (MRF) regularization is used. Like our model, in [7, 9–11] a Bayesian framework is used, although the construction of the posterior density function is different. Our approach uses priors on location and shape; of the forementioned, only one work [9] uses a shape prior. Also in [9] probabilistic pixel class prediction is used, which is reminiscent of the proposed Bayesian Probability Maps (BPM).

In this paper, we present improvements to our method [12] for determining the position of the endocardium in ultrasound sequences. Information about position may be used for determining ejection fraction (by comparing systolic and diastolic volume) and assessment of regional wall abnormalities of the heart; measures used in diagnosis of ischaemic heart disease. The problem is represented as a latent variable model, which represents the inside and outside of the endocardium. The method uses priors for spatial and temporal smoothness, in the form of total variation, connectivity, preferred shapes and location, by using the principal components and location distribution of manually segmented training shapes. The main steps of the method are: 1) We assume a three-component Rayleigh mixture model for the pixel intensities (of blood, echocardiographic artifacts, and tissue) and estimate the parameters by expectation maximization. 2) A latent variable model with two realizations, tissue and endocardium, is built using the estimated mixture model parameters. The posterior distribution of the latent variables is then sampled. 3) The mean of the posterior gives us the Bayesian probability map, which describes the position distribution of the endocardium. Instead of giving a single segmenting curve, the certainty of which may vary along the curve, our method provides a more versatile measure.

Our method shares some analogy with other region-based methods, but our approach of describing the position of the endocardium as the expected latent variable image and incorporating priors on location, connectivity, shape and smoothness in space and time, is in its construction novel to our knowledge.

The improvement of our previous model consists of: 1) the use of a three-component mixture model, improving the sensitivity of the algorithm in distinguishing between tissue and blood. 2) A connectivity prior ensuring that samples are spatially simply connected. 3) An atrium prior, which prevents blood in the atrium being classified as within the endocardium. 4) Sampling translation distribution, which improves the convergence of the algorithm.

2   Model

Our goal is to determine the position of the endocardium in an ultrasound sequence. To accomplish this we represent the endocardium by the latent variable model with values one and zero for the inside and outside, respectively and estimate the posterior distribution of the latent variable model

$$P(\mathbf{u}|\mathbf{z}, \theta) \propto p(\mathbf{z}|\mathbf{u}, \theta)P(\mathbf{u}|\theta) , \tag{1}$$

where $\mathbf{u}$ is the vector of latent variables, $\mathbf{z}$ represent image intensities stacked into a single vector and θ are parameters. The Rayleigh distribution has been reported to be an appropriate for modeling blood and tissue in cardiac ultrasound images, see [3, 1, 4].

Therefore to construct the likelihood $p(\mathbf{z}|\mathbf{u}, \theta)$, we assume a Rayleigh mixture model for pixels intensities in the ultrasound images, as described in Section 2.1. In Section 2.2, we construct the prior distribution $P(\mathbf{u}|\theta)$ by using the prior knowledge such as temporal and spatial smoothness, connectivity, shape and location.

2.1 Likelihood

By empirical observation we model the ultrasound data as a three-component mixture model: one for the blood intensities, one for echocardiographic artifacts, and finally one for the tissue. By artifacts we refer to areas with tissue-like intensity caused by cordae, papillary muscles, ribs or local increase in signal strength. Denoting the intensity value of pixel k in an ultrasound image by z_k, we assume that

$$p(z_k|\theta) = \alpha_1 p_{\mathrm{rayl}}(z_k|\sigma_1) + \alpha_2 p_{\mathrm{rayl}}(z_k|\sigma_2) + \alpha_3 p_{\mathrm{rayl}}(z_k|\sigma_3) , \qquad (2)$$

where $\theta = \{\alpha_i, \sigma_i; i = 1, 2, 3\}$ are the mixture model parameters and $p_{\mathrm{rayl}}(z|\sigma) = \frac{z}{\sigma}\exp(-\frac{z^2}{2\sigma})$, $\sigma > 0$ is the Rayleigh probability density function. In our previous work we employed two-component mixture model, but we have found that a three-component model better discriminates between tissue and blood, by adding a category for echocardiographic artifacts. Pixels are assumed to be independent in the mixture model. The likelihood is then defined as

$$p(\mathbf{z}|\mathbf{w}, \theta) = \prod_j \prod_k P(W_j = k|z_j, \sigma_k)^{\delta(w_j - k)}, k = 1, 2, 3 \qquad (3)$$

where W_j and w_j is the random latent variable and its realization, respectively, corresponding to z_j, and δ denotes the Kronecker delta function. $w_j = 1$ if x_j is in the blood pool, $w_j = 2$ if x_j is in an echocardiographic artefact, $w_j = 3$ if x_j is in the cardiac tissue, and $P(W_j \in \mathrm{i}|z_j, \theta) = \alpha_i p_{\mathrm{rayl}}(z_j|\sigma_i)/\sum_{k=1}^{3} \alpha_k p_{\mathrm{rayl}}(z_j|\sigma_k), i = 1, 2, 3$.

We use the parameters θ to build a latent variable model, with only two realizations: tissue (0) and endocardium (1). The likelihood of this model is defined as

$$p(\mathbf{z}|\mathbf{u}, \theta) = \prod_j P(U_j \in \mathrm{endocardium}|z_j, \sigma)^{u_j} P(U_j \in \mathrm{tissue}|z_j, \sigma)^{1-u_j}, \quad (4)$$

$$\sigma = \{\sigma_1, \sigma_2, \sigma_3\}$$

where U_j and u_j are the random latent variable and its realization, respectively, corresponding to z_j and $P(U_j \in \mathrm{tissue}|z_j, \theta) = \alpha_1 p_{\mathrm{rayl}}(z_j|\sigma_1)/\sum_{i=1}^{3} \alpha_i p_{\mathrm{rayl}}(z_j|\sigma_i)$ and $P(U_j \in \mathrm{endocardium}|\theta) = 1 - P(U_j \in \mathrm{backgr}|z_j, \theta)$.

2.2 Prior

Our prior model

$$P(\mathbf{u}|\theta) = P_{\mathrm{B}}(\mathbf{u}|\theta)P_{\mathrm{TV|B}}(\mathbf{u}|\theta)P_{\mathrm{shape|B,TV}}(\mathbf{u}|\theta)P_{\mathrm{atrium|B,TV,shape}}(\mathbf{u}|\theta)\times \quad (5)$$
$$P_{\mathrm{con|B,TV,shape,atrium}}(\mathbf{u}|\theta)P_{\mathrm{location|B,TV,shape,atrium,con}}(\mathbf{u}|\theta)$$
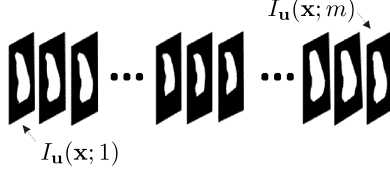
Fig. 1. The symbol $\mathbf{u}$ represents the m latent variable images $I_{\mathbf{u}}$ stacked into a single vector. Each $I_{\mathbf{u}}$ corresponds to an image in the ultrasound sequence of length m.

consists of six components, where each characterizes different kinds of properties preferred. The Bernoulli component P_{B} is the discrete latent variable distribution following from the Rayleigh mixture model. The total variation $P_{\mathrm{TV|B}}$ enforces spatial and temporal smoothness for latent variable images, and possible shape variations around the mean shape are characterized by trained eigenshapes of manually segmented images through $P_{\mathrm{shape|B,TV}}$. The sequence of ultrasound images is divided into subsequences, to take the temporal variations of the endocardium into account, and so for each part of the ultrasound sequence a corresponding set of eigenshapes and mean is used.

The atrium contains blood which should not be classified as being within the endocardium. The atrium prior $P_{\mathrm{atrium|B,TV,shape}}$ reduces the likelihood exponentially, of a pixel being classified as being within the endocardium, with distance from the horizontal position of the ventricle to the bottom of the ultrasound image. Although some blood may still be misclassified, this prior will prevent large misclassifications.

The connectivity prior $P_{\mathrm{con|B,TV,shape,atrium}}$ enforces that all samples u are spatially simply connected.

The location prior $P_{\mathrm{location|B,TV,shape,atrium,con}}$ is constructed from the mean of the unregistered binary training shapes. The location prior describes the experimental probability value for each pixel location being either inside or outside of the endocardium, thus allowing only similar latent variable values as observed in the training data.

The Bernoulli prior is defined as

$$P_{\mathrm{B}}(\mathbf{u}|\theta) \propto \prod_{j} \alpha^{u_j}(1-\alpha)^{1-u_j} \tag{6}$$

and is thus a prior on the proportion of zeros and ones in $\mathbf{u}$ and $j \in \{1,...,N\}$, where N is the total number of latent variables in $\mathbf{u}$.

Let $I_{\mathbf{u}}(\mathbf{x}; n)$ be a latent variable image, where $\mathbf{x}$ and n are its spatial and temporal coordinates, respectively (see Figure 1). The total variation prior is then given by

$$P_{\mathrm{TV|B}}(\mathbf{u}|\theta) \propto \exp\{-\lambda_{\mathrm{TV}}||I_u(\mathbf{x}; n) * h||_{L_1}\} \, ,$$

where h is a three dimensional Laplacian kernel and $*$ denotes convolution.

Let $I_{\mathbf{u},\mathrm{r}}(\mathbf{x}; n)$ be the transitionally registered latent variable image, corresponding to $I_{\mathbf{u}}(\mathbf{x}; n)$, where the center of mass has been shifted to the origin; $\mathbf{u}_r^n$ and $\bar{\mathbf{u}}_r^n$ are the

corresponding latent variable vectors. The shape prior is defined as

$$P_{\text{shape}|\text{B,TV}}(\mathbf{u}|\theta) \propto \prod_n \exp\{-\lambda_{\text{shape}}(\mathbf{u}_r^n - \bar{\mathbf{u}}_r^n)^T (\mathbf{C}_n + \lambda_0 \mathbf{I})^{-1}(\mathbf{u}_r^n - \bar{\mathbf{u}}_r^n)\} \ , \quad (7)$$

where $\mathbf{C}_n$ represents the truncated covariance of the training shapes, whose center of mass has been shifted to the origin, and $\lambda_0 \mathbf{I}$ is the Tikhonov regularizer [13].

The atrium prior is defined as

$$P_{\text{atrium}|\text{B,TV,shape}}(\mathbf{u}|\theta) \propto \prod_{\mathbf{x}} \prod_n a(\mathbf{x}) I_{\mathbf{u}_{\text{train}}}^r(\mathbf{x}; n) \ , \quad (8)$$

where

$$a(\mathbf{x}) = \begin{cases} 1 & \text{if } x_2 > x_a \\ 1 - \exp\{\frac{x_2 - x_a}{\max\limits_{x_2}(x_2 - x_a)}\} \ , & \text{otherwise} \end{cases} \quad (9)$$

where $x_a = \max\limits_n\{\operatorname*{argmin}\limits_{x_2}\{I_{\mathbf{u}_{\text{train}}}^r(\mathbf{x}; n) > 0\}\}$, $\mathbf{x} = (x_1, x_2)$. In every training image $I_{\mathbf{u}_{\text{train}}}^r$ there is a least x_2-coordinate x_l s.t $I_{\mathbf{u}_{\text{train}}}^r((x_1, x_l)) > 0$; x_a is the largest out of all x_l. This gives an approximate location of the ventricle, where the atrium starts.

The connectivity prior is defined as

$$P_{\text{con}|\text{B,TV,shape,atrium}}(\mathbf{u}|\theta) \propto \begin{cases} 1 & \text{if } \mathbf{u} \in \mathcal{N} \\ 0 \ , & \text{otherwise} \end{cases} \quad (10)$$

where $\mathcal{N} = \{\mathbf{u} : \mathbf{u} \text{ spatially simply connected}\}$.

The location prior is defined as

$$P_{\text{location}|\text{B,TV,shape,atrium,con}}(\mathbf{u}|\theta) \propto \quad (11)$$
$$\begin{cases} 1 \text{ if } \frac{1}{\sum_j u_j} \sum_n \sum_x h\left(g * \bar{I}_{\mathbf{u}_{\text{train}}}(\mathbf{x}; n)\right) I_{\mathbf{u}}(\mathbf{x}; n) = 1 \\ 0 \text{ otherwise} \end{cases}$$

where $\bar{I}_{\mathbf{u}_{\text{train}}} = \frac{1}{K} \sum_k I_{u_{\text{train}}}^k$ is the mean training image and K is the number of training images. g is a Gaussian kernel and h is the step function s.t. $h(t) = 1$ for $t > 0$, otherwise $h(t) = 0$. This component has the effect that when sampling individual latent variables outside of the (smoothed) mean shape, the result of sampling will be that the latent variable is set to zero. Inside the (unregistered) mean shape the sampling is unaffected.

Three parameters control the influence of the priors: λ_{TV}, λ_{shape} and λ_0. By increasing λ_{TV} we can regularize our sampling, while increasing λ_{shape} makes the influence of the shape prior larger. Finally λ_0 increases the influence of the mean shape in the formation of the shape prior; this is crucial when segmenting very noisy images, that do not respond well to the subtle control of the a shape prior with small λ_0.

3   Algorithm

Our algorithm for generating Bayesian Probability Maps can be divided into three parts. First the mixture model parameters are estimated by the EM algorithm from our ultrasound data; these parameters are used to compute the class posterior possibilities for
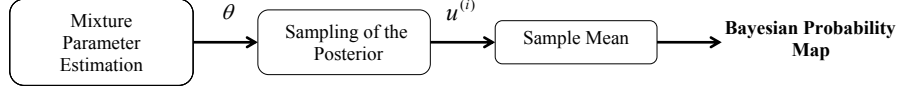
Fig. 2. Summary of the proposed algorithm to construct the Bayesian probability map.

each latent variable after seeing the corresponding image values — these probabilities are used as an input in following step in constructing the likelihood function that is further transformed to the posterior probabilities in the second step. The posterior is then sampled by Gibbs sampling and the samples are used to compute the Bayesian probability map. The algorithm is summarized in Fig. 2.

3.1   Estimation of mixture model parameters

The complete data likelihood is represented according to the latent variable model as

$$p(\mathbf{z}, \mathbf{w}|\theta) = \prod_j \prod_i p_{\mathrm{rayl}}(z_j|\sigma_i)^{\delta(w_j - i)} , \tag{12}$$

where $\mathbf{z}$ are the pixel intensity values and $\mathbf{w} = (w_1, ..., w_N)$ are interpreted as missing data. The mixture parameters $\theta = \{\alpha_i, \sigma_i; i = 1, .., 3\}$ are estimated by Expectation Maximization (EM) [14]. That is, on the E-step, we build the expected complete data loglikelihood, conditioned on the measured data and the previous parameter estimates, or

$$\chi(\theta, \hat{\theta}^{(n-1)}) = \mathrm{E}_{\mathbf{w}|\mathbf{z}, \hat{\theta}^{(n-1)}}\{\log p(\mathbf{z}, \mathbf{w}|\theta))\}$$

$$= \sum_{j=1}^{N} \sum_{k=1}^{3} P(W_j = k|z_j, \hat{\theta}^{(n-1)}) \log p_{\mathrm{rayl}}(z_j|\theta) . \tag{13}$$

On the M-step, the expected complete data loglikelihood is maximized to obtain an update for the parameters

$$\hat{\theta}^{(n)} = \underset{\theta}{\operatorname{argmax}} \, \chi(\theta, \hat{\theta}^{(n-1)}) , \tag{14}$$

and the steps are iterated until convergence.

3.2   Sampling of the Posterior

To improve convergence, the sampling of the posterior (1) was performed by alternating between conventional Gibbs sampling [15, 16] and sampling of latent variable image translations. On the Gibbs step, we draw the elements of the sample latent variable vector $\mathbf{u}$ from the conditional distribution

$$P(u_j|u_1^{(i)}, \ldots, u_{j-1}^{(i)}, u_{j+1}^{(i-1)}, \ldots, u_N^{(i-1)})$$

$$= \left\{ P(u_j = k|u_1^{(i)}, \ldots, u_{j-1}^{(i)}, u_{j+1}^{(i-1)}, \ldots, u_N^{(i-1)}) \right\}_{k=0}^{1} , \, j = 1, 2, \ldots, N . \tag{15}$$

Then, to obtain sample vector $\mathbf{u}^{(i)}$, we sample the distribution of translations which spatially move the latent variable image $I_\mathbf{u}$. The details of the translation sampling step are as follows.

We want to sample the conditional translation distribution

$$P(\mathbf{t}|\mathbf{u}, \mathbf{z}, \theta) \equiv P(\mathbf{u}'|\mathbf{u}, \mathbf{z}, \theta) \;, \tag{16}$$

where the latent variable vector $\mathbf{u}'$ is obtained from $\mathbf{u}$ by spatially translating the latent variable image $I_\mathbf{u}$ by $\mathbf{t}$. Now we may write,

$$
\begin{aligned}
P(\mathbf{t}|\mathbf{u}, \mathbf{z}, \theta) &\propto \left(\prod_{j=1}^{N} p_{\mathrm{rayl}}(z_j|\sigma_1)^{u'_j} p_{\mathrm{rayl}}(z_j|\sigma_2)^{1-u'_j} \right) p(\mathbf{u}'|\theta) \\
&\propto \prod_{j=1}^{N} \left(p_{\mathrm{rayl}}(z_j|\sigma_1) v_{1,j} \right)^{u'_j} \left(p_{\mathrm{rayl}}(z_j|\sigma_2) v_{2,j} \right)^{1-u'_j} \;,
\end{aligned}
\tag{17}
$$

where we have used the fact that, apart from the location prior, the conditional translation distribution is independent of the priors; and the location prior is encoded in the mask vectors $\mathbf{v}_1 = \lceil v \rceil$ and $\mathbf{v}_2 = 1 - \lfloor v \rfloor$, where v is the vector corresponding to the matrix $g * \bar{I}_{\mathbf{u}_{\mathrm{train}}}$, cf. (11). It follows that

$$\log P(\mathbf{t}|\mathbf{u}, \mathbf{z}, \theta) = \sum_{j=1}^{N} u'_j \log \left(p_{\mathrm{rayl}}(z_j|\sigma_1) v_{1,j} \right) + (1 - u'_j) \log \left(p_{\mathrm{rayl}}(z_j|\sigma_2) v_{2,j} \right) + C \;,$$
$$\tag{18}$$

where C is a constant that does not depend on the translation. The sums above represent correlations between the translated latent variable image and the masked log probability densities. Hence, the logarithms of the conditional translation probabilities can be computed by the correlation theorem, after which we are able to draw the translation sample and finally obtain the sample $\mathbf{u}^{(i)} = \mathbf{u}'$.

After iteration the center of mass of each latent variable image is calculated, which determines the area of influence of the shape prior.

3.3 Sample Mean

To characterize the posterior distribution, we compute estimate conditional mean of the latent variable vector over the posterior

$$\mathrm{E}\{\mathbf{u}|\mathbf{z}, \theta\} \approx \frac{1}{M} \sum_i \mathbf{u}^{(i)} = \left(\hat{P}(U_k \in \mathrm{obj}) \right)_{k=1}^{N} \equiv \hat{\mathbf{u}}_{\mathrm{CM}} \tag{19}$$

by the latent variable sample vectors $\mathbf{u}^{(i)}$. By the strong law of large numbers $\hat{\mathbf{u}}_{\mathrm{CM}} \rightarrow \mathrm{E}\{\mathbf{u}|\mathbf{z}, \theta\}$ when $n \rightarrow \infty$. The corresponding image $I_{\hat{\mathbf{u}}_{\mathrm{CM}}}$ represents the Bayesian probability map.

4   Experiments

4.1   Material

The ultrasound data used in this paper consists of cardiac cycles of two-chamber (2C) apical long-axis views of the heart. The sequences were obtained using the echocardiogram machines Philips Sonos 7500, Philips iE33 or GE Vivid 7, from consecutive adult patients referred to the echocardiography laboratory at Malmö University hospital, Sweden which has a primary catchment area of 250,000 inhabitants. Expert outlines of the endocardium in the sequences have been provided by the same hospital.

4.2   Initialization

As an initial estimate of mixture model parameters we set $\alpha^{(0)}$ to the proportion of object pixels in the training images, and σ_1 and σ_2 are set to maximum likelihood estimate $\hat{\sigma} = (\frac{1}{2Q} \sum_{i=1}^{Q} x_i^2)^{\frac{1}{2}}$ of object and background pixels in the training data, where Q is the number of pixels in the training set. The Gibbs sampling algorithm is seeded by a sample obtained by Bayesian classification of the mean of the annotated images for each category of the heart cycle. Prior parameters $\lambda_{\text{TV}}, \lambda_{\text{shape}}, \lambda_0$ are set manually.

4.3   Evaluation

We divide our data into two sets: training set and validation set. The training set consists of 20 cardiac cycles. The training set is further divided into sets, corresponding to parts of the cardiac cycle. The validation set consists of 4 different cardiac cycles.

As evaluation measure the expected misclassification E_{mc} of a pixel, w.r.t the expert outline, is used. Let $I_{\text{true}}(\mathrm{x}; n)$ be ground truth images corresponding to the data $\mathbf{z}$. Then the expected misclassification of a pixel in the examined sequence is given by

$$E_{\text{mc}} = \frac{1}{N} \sum_n \sum_x \left(1 - I_{\text{true}}(\mathbf{x}; n)\right) P\big(I_{\mathbf{u}}(\mathbf{x}; n) = 0\big) + I_{\text{true}}(\mathrm{x}; n) P\big(I_{\mathbf{u}}(\mathbf{x}; n) = 1\big) \,.$$

(20)

This measure is needed since it is impossible to present the entire sequence in images. A low E_{mc} guarantees that the Bayesian Probability Map correctly describes the position of the endocardium in the entire sequence, not just for a few selected images.

4.4   Results

Figure 4 and 5 contain Bayesian Probability Maps (BPM) formed from 150-200 samples; the approx. number samples needed to reach a stationary distribution. Running time for sampling is approx. 3 hours for the entire ultrasound sequence on a Intel Xeon 2.33 GHz, 9 Gb RAM server. The probability map spans colors from red to blue with degree of probability,of area being within the endocardium. Hence, red indicates the

Table 1. Parameters and E_{mc} for Validation Sequences

Validation Sequence	λ_{TV}	λ_{shape}	λ_0	E_{mc}
A	0.75	125	100	0.0335
B	1.5	60	100	0.0367
C	4	400	800	0.0507
D	5	500	1000	0.0569

highest probability. Table 1 contains the parameter settings and the E_{mc} for the validation sequences. From each sequence four frames are displayed, two at systole and two at diastole. Sequence A and B have quite modest parameter values, as the underlying estimate of tissue and blood, is quite satisfactory and does not need much intervention in the form of priors. Sequence C and D required large λ_0, due to artifacts in the chamber (in the case of D due to rib of patient). Overall these results are superior to those we have previously published.

4.5   Comparison with Graph Cut Method

We compare our results with a Graph Cut method as described in [17–19]. In Figure 3 we observe that the Graph Cut method fails to identify the location as clearly as the proposed method for sequence A. For validation sequences B,C and D no results were obtained since a singular covariance matrix was obtained. This may be attributed to the very noisy nature of these sequences. This comparison is limited and given to show the differences between the proposed method and pure Graph Cut algorithms, as there are some fundamental similarities such as pixel dependencies; similar to the methods described in [6, 7, 9–11], the Graph Cut method uses MRF for this. However, more complex methods share many similarities with our method, e.g. as described in [7], which we plan include in future comparative study.

5   Conclusion and future work

We have presented improvements to our approach [12] to cardiac ultrasound segmentation, which consists of modeling the endocardium by latent variables. The latent vari-
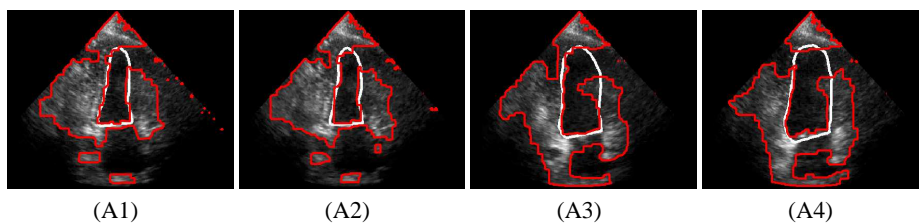

| (A1) | (A2) | (A3) | (A4) |

Fig. 3. (Color online) Graph Cut (red) applied to Validation sequence A with expert outline (white).
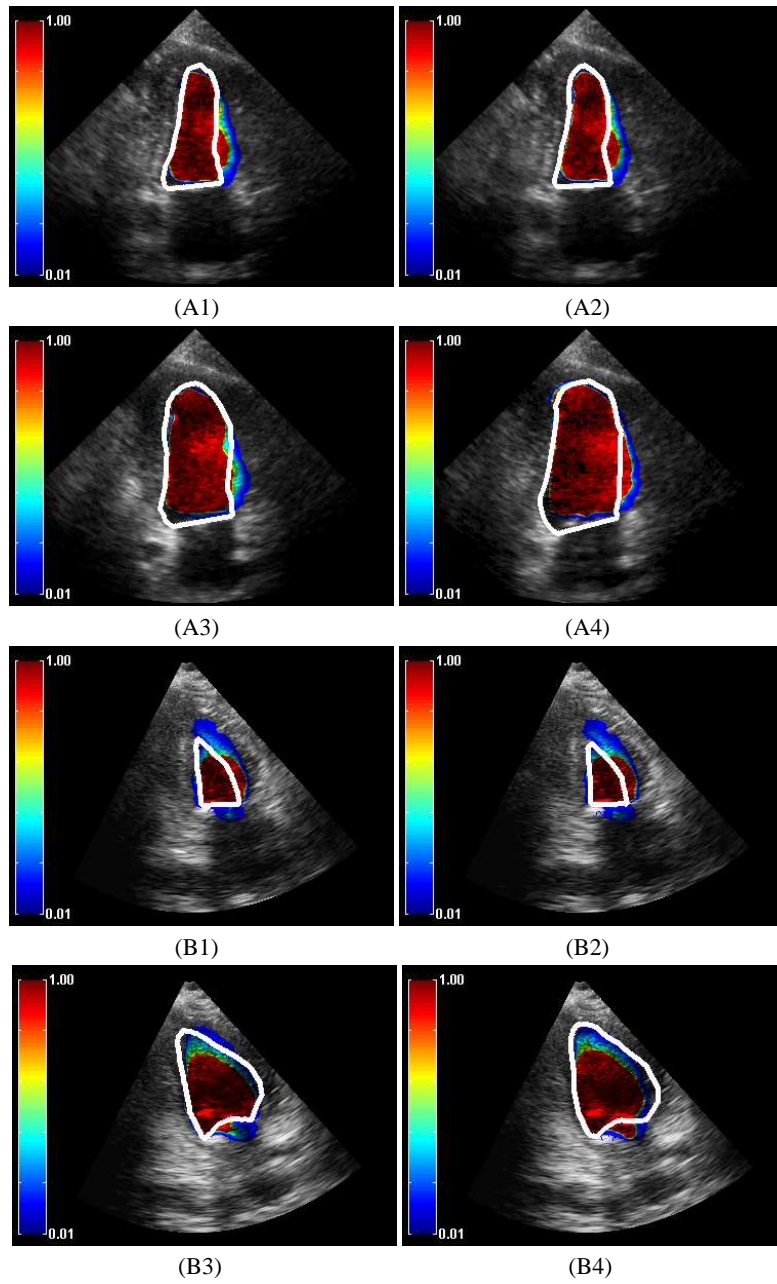
Fig. 4. (Color online) Validation sequences A (43 frames) and B (40 frames). BPM with overlaid expert outline (white). Systole (A1-2,B1-2) and Diastole (A3-4,B3-4).
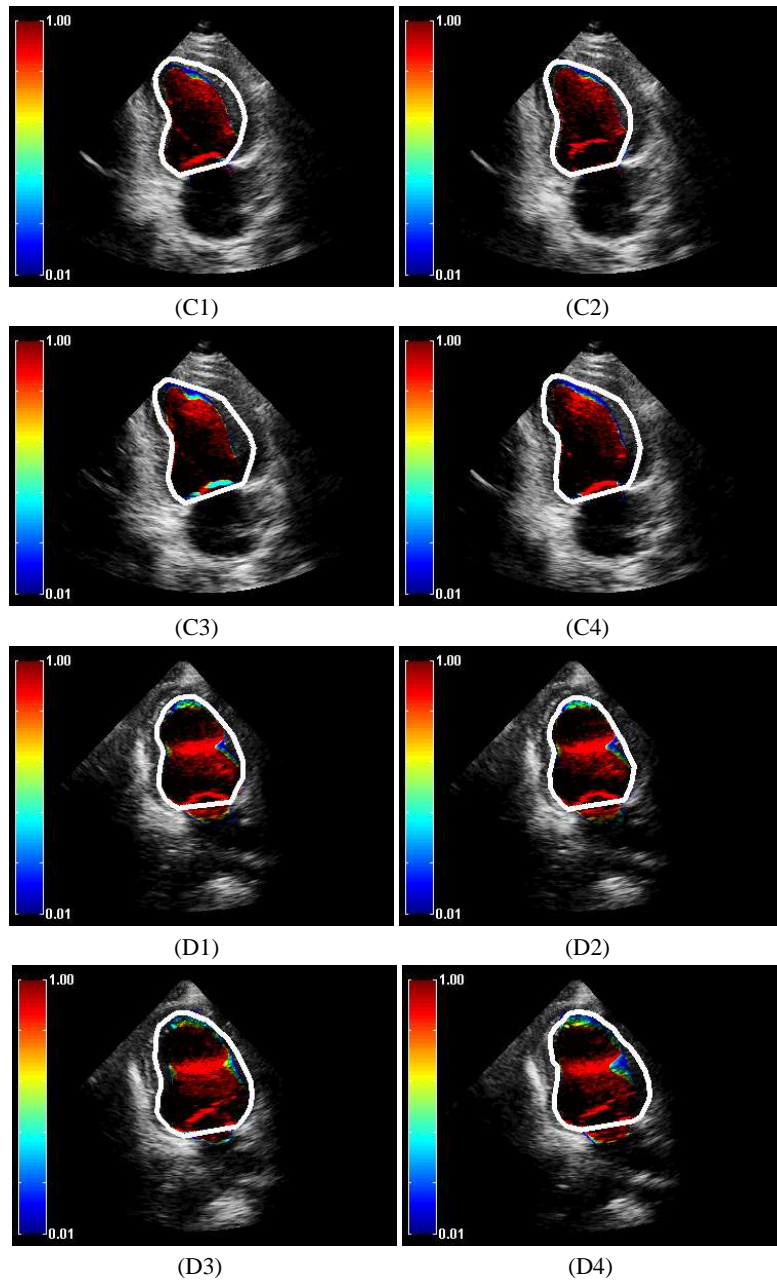
Fig. 5. (Color online) Validation sequences C (21 frames) and D (35 frames). BPM with overlaid expert outline (white). Systole (C1-2,D1-2) and Diastole (C3-4,D3-4).

able distribution is then sampled which yields Bayesian Probability Map, which describes the location of the endocardium. The improvements consist of a three-component mixture model, connectivity, an atrium prior and sampling translation distribution.

We plan to introduce a method of estimating the prior parameters, and by this refining our results. We will expand our comparative evaluation with graph-based methods and those akin to the approach proposed by [7].

References

1. Burckhardt, C.B.: Speckle in Ultrasound B-Mode Scans. IEEE Trans. Sonics and Ultrasonics **25** (1978) 1–6
2. Wagner, R.F., Smith, S.W., Sandrik, J.M., Lopez, H.: Statistics of Speckle in Ultrasound B-Scans. IEEE Trans. Sonics and Ultrasonics **30** (1983) 156–163
3. Goodman, J.: Laser Speckle and Related Phenomenon. New York: Springer-Verlag (1975)
4. Tao, Z., Tagare, H., Beaty, J.: Evaluation of four probability distribution models for speckle in clinical cardiac ultrasound images. MedImg **25** (2006) 1483–1491
5. Noble, J.A., Boukerroui, D.: Ultrasound image segmentation: A survey. Medical Imaging, IEEE Transactions on **25** (2006) 987–1010
6. Friedland, N., Adam, D.: Automatic ventricular cavity boundary detection from sequential ultrasound images using simulated anneal. IEEE Trans. Med. Imag. **8** (1989) 344–353
7. Boukerroui, D., Baskurt, A., Noble, J.A., Basset, O.: Segmentation of ultrasound images: multiresolution 2d and 3d algorithm based on global and local statistics. Pattern Recogn. Lett. **24** (2003) 779–790
8. Dias, J.M.B., Leitao, J.M.N.: Wall position and thickness estimation from sequences of echocardiographic images. IEEE Trans. Med. Imag **15** (1996) 25–38
9. Song, M., Haralick, R., Sheehan, F., Johnson, R.: Integrated surface model optimization for freehand three-dimensional echocardiography. Medical Imaging, IEEE Transactions on **21** (2002) 1077–1090
10. Xiao, G., Brady, J.M., Noble, A.J., Zhang, Y.: Contrast enhancement and segmentation of ultrasound images: a statistical method. Medical Imaging 2000: Image Processing **3979** (2000) 1116–1125
11. Figueiredo, M., Leitao, J.: Bayesian estimation of ventricular contours in angiographic images. Medical Imaging, IEEE Transactions on **11** (1992) 416–129
12. Hansson, M., Brandt, S., Gudmundsson, P.: Bayesian probability maps for the evaluation of cardiac ultrasound data. In: MICCAI Workshop: Probabilistic Models for Medical Image Analysis. (2009)
13. Tikhonov, A.: Solutions of Ill Posed Problems. Vh Winston, Scripta series in mathematics (1977)
14. Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood form incomplete data via the EM algorithm (with discussion). J. R. Statist. Soc. Ser. B-Stat. Methodol. **39** (1977) 1–38
15. Geman, S., Geman, D.: Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. IEEE Trans. Pattern Analysis and Machine Intelligence **6** (1984) 721–741
16. MacKay, D.: Information Theory, Inference and Learning Algorithms. Cambridge University Press (2003) http://www.inference.phy.cam.ac.uk/mackay/itila/book.html.
17. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. IEEE Trans. Pattern Analysis and Machine Intelligence **20** (2001) 1222–1239
18. Kolmogorov, V., Zabih, R.: What energy functions can be minimized via graph cuts? IEEE Trans. Pattern Analysis and Machine Intelligence **26** (2004) 147–159
19. Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms. IEEE Trans. Pattern Analysis and Machine Intelligence **26** (2004) 1124–1137