

A Spatio-temporal Analysis of Contrast Ultrasound Image Sequences for Assessment of Tissue Perfusion

Quentin R. Williams and J. Alison Noble

Medical Vision Laboratory, University of Oxford, UK
{quentin,noble}@robots.ox.ac.uk

Abstract. The evaluation of tissue perfusion in various parenchymatous organs is important in the diagnosis and determination of the severity of ischemic disease. Contrast ultrasound perfusion imaging can be used for this purpose. This paper describes a method that identifies different areas of perfusion in a contrast ultrasound perfusion study. Pixels in an image sequence are automatically classified into different classes, by analysing their distinct temporal relationships. A novel method is presented that uses a Bayesian Factor Analysis Model set in a Markov Random Field framework; utilising both the temporal and spatial characteristics of the pixels for classification. Preliminary results are demonstrated for simulated data, and a myocardial perfusion in-vivo dataset.

1 Introduction

Tissue perfusion imaging is becoming an increasingly employed method to assess internal organ blood supply and flow in clinical applications. For example the assessment of tissue viability through perfusion imaging, has many clinical applications which include: detection of coronary artery disease (CAD) in asymptomatic and symptomatic patients; estimation of the severity of CAD; risk stratification for coronary events, and analysis of disease state in other organs such as the liver, kidneys and the brain. Tissue perfusion has been traditionally analysed by nuclear medicine imaging procedures, like T1-SPECT, which measure cell membrane integrity, or more recently by PET-FDG, which shows metabolism and blood flow rates. Limitations of these techniques are low spatial resolution and the use of ionising radiation.

The availability, low cost and non-invasiveness of echocardiography have fostered an increasing interest in the use of this modality to provide an accurate and quantitative diagnosis of tissue perfusion. The intravenous injection of contrast agents (microbubbles) has allowed the visualisation of blood flow information and regional perfusion. The contrast agents submerged in blood, increase the echo-backscatter of perfused tissues and blood pool in cavities, and are ideal as they remain intravascular and have a particle size similar to red blood cells [1]. However, interpretation of the results is very difficult, and doctors must often rely on visual assessment, which is subjective and results in various discrepancies [2]. Techniques have been proposed to quantify the dynamics of perfusion by evaluating the signal intensities with time in the ultrasound images. For instance, Mor-Avi et al. [3] proposed methods for quantifying both regional blood flow distribution and transit time, by frequency domain

analysis of regional time curves. Wei et al. [4] introduced the ‘negative bolus indicator dilution technique’ which is based on high power ultrasound induced destruction of microbubbles and the assessment of their replenishment during a constant venous infusion of a contrast agent. The replenishment (wash-in) curves, showing the refilling of microbubbles, were then fitted to exponential models to extract parameters involving the microbubble velocity and myocardial blood volume. These pixel-based techniques all suffer from ad hoc smoothing in space and time, and the loss of temporal information that is available by analysing the correlation between pixels. Noise present in the images can also greatly influence the shape of the intensity-time curves, while registration has mostly been ignored in these methods. Factor Analysis of Dynamic Studies (FADS) has also been suggested to analyse perfusion curves. Much of the work on FADS has been carried out on dynamic nuclear medicine studies [5], as well as MRI [6]. These methods attempt to rotate the final factor solution in such a way that it satisfies certain constraints (e.g. the positivity constraint) that endeavor to make the results more interpretable. These constraints are however subjective and often not sufficient on their own to produce a unique solution.

In this paper a method is presented to assess myocardial perfusion by automatically classifying the ultrasound images into different regions of perfusion, using a novel spatio-temporal technique. This is done in a global manner by analysing the temporal pattern of relationships between pixels, using a Bayesian Factor Analysis model, and incorporating spatial information through a Markov Random Field. In this manner pixels can be classified into particular types of perfusion (from which quantitative parameters can be obtained if necessary), and the nature and structure of any perfusion study can be examined. This probabilistic view of factor analysis allows a unique solution to be found automatically, without the use of subjective constraints, and has to our knowledge not before been applied to tissue perfusion studies. An interesting and novel way of interlinking the Bayesian Factor Analysis with a Markov Random Field is also presented in this paper. Although not limited to these applications, results are shown for a myocardial perfusion study.

2 Methods

2.1 Bayesian Classification Using a Markov Random Field Prior Model

In this paper classification is treated as a statistical problem, which involves assigning to each pixel a class label taking a value from the set $L = \{1, 2, \dots, l\}$. The pixels are indexed by a two-dimensional rectangular lattice $S = \{1, 2, \dots, n\}$ and each pixel is characterised by a p -variate vector of intensity values $\mathbf{y}_i = (y_{i1}, \dots, y_{ip})$, $i \in S$. In this case each observation vector \mathbf{y}_i represents an intensity-time curve for a single pixel location (i.e. there are p images), where each intensity value is taken from the same image location in consecutive timeframes in the image sequence.

A labelling of S will be denoted \mathbf{x} , where $x_i, i \in S$ is the corresponding class label of pixel i . The true but unknown labelling configuration, \mathbf{x}^* , and the configuration estimate, $\hat{\mathbf{x}}$, are both interpreted as particular realisations of a random field X . The p -

variate observation vectors on n pixels are then, $\mathbf{Y}' = (\mathbf{y}_1, \dots, \mathbf{y}_n)$, which is also a realisation of a random variable, Y . The problem of classification is to estimate \mathbf{x}^* , given the observed intensity time vectors \mathbf{Y} . In particular, the maximum a posteriori (MAP) estimate of \mathbf{x} is used:

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x} \in X} \{P(\mathbf{Y} | \mathbf{x})P(\mathbf{x})\}. \quad (1)$$

The right-hand side of the above equation contains two parts: $P(\mathbf{Y} | \mathbf{x})$ and $P(\mathbf{x})$, which are defined below as a Bayesian Factor Analysis likelihood distribution and a Markov Random Field prior distribution, respectively.

2.2 Bayesian Factor Analysis Model

In the Bayesian Factor Analysis model it is assumed that the correlation between individual pixel intensity-time curves can be explained in terms of a small number of underlying hidden factors [8]. A generative latent variable model is constructed,

$$(\mathbf{y}_i | \mu, \Lambda, \mathbf{f}_i) = \mu + \Lambda \mathbf{f}_i + \epsilon_i, \quad (2)$$

for each observation vector \mathbf{y}_i ($i=1, \dots, n$), where μ is the overall population mean, Λ is a matrix of constants called the factor loading matrix; $\mathbf{f}_i = (f_{i1}, \dots, f_{iL})$, $l \in L$, is the factor score vector for pixel i ; and the ϵ_i 's are assumed to be mutually uncorrelated and Normally distributed $N(0, \Psi)$ variables. The factor loading matrix, Λ , expresses how each hidden factor loads onto the observed variables, therefore giving an indication of how the hidden factors might look. In the case of a perfusion study, each column in the factor loading matrix will represent an intensity-time curve associated with each different type of perfusion present in the dataset. The factor scores give the estimated value ("weight") of the observations on the hidden factors. Therefore, if each hidden factor represents a class, the factor score vector gives an indication of how much an observation belongs to each class. Since the parameters μ , Λ , the \mathbf{f}_i 's, and Ψ are all unobservable, a Normal likelihood distribution for each \mathbf{y}_i is assumed, and written as:

$$p(\mathbf{y}_i | \mu, \Lambda, \mathbf{f}_i, \Psi) = (2\pi)^{-\frac{p}{2}} |\Psi|^{-\frac{1}{2}} e^{-\frac{1}{2}(\mathbf{y}_i - \mu - \Lambda \mathbf{f}_i)' \Psi^{-1} (\mathbf{y}_i - \mu - \Lambda \mathbf{f}_i)}. \quad (3)$$

Assuming independence between the observations, the joint likelihood becomes:

$$p(\mathbf{Y} | \mu, \Lambda, \mathbf{f}_i, \Psi) = (2\pi)^{-\frac{np}{2}} |\Psi|^{-\frac{n}{2}} e^{-\frac{1}{2} \sum_{i=1}^n (\mathbf{y}_i - \mu - \Lambda \mathbf{f}_i)' \Psi^{-1} (\mathbf{y}_i - \mu - \Lambda \mathbf{f}_i)}. \quad (4)$$

2.3 Markov Random Field

Markov Random Field (MRF) theory provides a convenient and consistent way to model the spatial relationships of context-dependent entities such as image pixels. In an MRF, only neighbouring sites have direct interactions with each other and they tend to have the same class labels. The probability of an MRF realisation, \mathbf{x} , is given by the Gibbs distribution:

$$P(\mathbf{x}) = Z^{-1} e^{(-\omega U(\mathbf{x}))}, \quad (5)$$

where

$$U(\mathbf{x}) = \sum_{c \in C} V_c(\mathbf{x}), \quad (6)$$

is the energy function which is a sum of *clique potentials* $V_c(\mathbf{x})$ over all possible cliques C . Z is a normalisation term and ω is a positive constant which controls the size of clustering(*). A *clique* c is defined as a subset of sites in S in which every pair of distinct sites are neighbours, except for single-site cliques. In this paper, only cliques of size two are counted. The clique potential is then of the form:

$$V_c(x_i) = -\delta_{x_i=x_i'}. \quad (7)$$

This is the same potential function as used by Xiao et al. [7] where $\delta_{x_i=x_i'} = 1$, if $x_i = x_i'$, and 0 otherwise. It is easy to show from Eqns. (1), (4), (6) and (7) that the MAP estimate of the classification is found by

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x} \in X} \{U(\mathbf{x} | \mathbf{Y})\}, \quad (8)$$

where the posterior energy $U(\mathbf{x} | \mathbf{Y})$ is given by

$$U(\mathbf{x} | \mathbf{Y}) = \sum_{i=1}^n \frac{1}{2} [(\mathbf{y}_i - \boldsymbol{\mu} - \boldsymbol{\Lambda} \mathbf{f}_i)' \boldsymbol{\Psi}^{-1} (\mathbf{y}_i - \boldsymbol{\mu} - \boldsymbol{\Lambda} \mathbf{f}_i) + \log |\boldsymbol{\Psi}|] - \omega \sum_{c \in C} \delta_{x_i=x_i'}. \quad (9)$$

MAP estimation is then computed using the iterated conditional modes (ICM) algorithm suggested in [7].

2.4 Connection Between Factors Scores and the Classification

It has been noted that the factor score vector indicates how much an observation belongs to a particular class. It is therefore assumed that the prior probability of the factor scores matrix, $\mathbf{F}' = (\mathbf{f}_1, \dots, \mathbf{f}_n)$, follows the same prior probability of the classification configuration, \mathbf{x} , and in fact that the factor score for each hidden factor (or class) is equivalent to the posterior probability of the class label. For every $l \in L$ and $i \in S$

$$f_{il} = P(\mathbf{y}_i | l) P(x_i = l). \quad (10)$$

Using the prior probability and the likelihood function with respect to x_i and f_{il} gives

$$f_{il} = Z^{-1} e^{(-\omega U(x_i))} \times (2\pi)^{-\frac{p}{2}} |\boldsymbol{\Psi}|^{-\frac{1}{2}} e^{-\frac{1}{2}(\mathbf{y}_i - \boldsymbol{\mu} - \boldsymbol{\Lambda}_l f_{il})' \boldsymbol{\Psi}^{-1} (\mathbf{y}_i - \boldsymbol{\mu} - \boldsymbol{\Lambda}_l f_{il})}. \quad (11)$$

Therefore the posterior probability values obtained through the MRF-MAP classification can directly be used as the factor scores.

* ω is used to avoid confusion with the replenishment curve blood velocity, β , in Eqn. 21.

2.5 μ, Λ, Ψ Estimation

The parameters μ , Λ , and Ψ are still unknown and thus require estimation. Rowe [8] proposed using an ICM MAP estimation procedure that maximises the posterior conditional distributions of the unknown parameters by cycling through their modes. The approach is summarised here, and the reader is referred to [8] for more detail.

In [8] it is assumed that the parameters μ , Λ , and Ψ are random variables with generalised natural conjugate prior distributions:

$$p(\mu) \propto |\Gamma|^{-1/2} e^{-\frac{1}{2}(\mu - \mu_0)' \Gamma^{-1} (\mu - \mu_0)}, \quad (12)$$

$$p(\Lambda) \propto |\Delta|^{-1/2} e^{-\frac{1}{2} \text{tr}(\Lambda - \Lambda_0)' \Delta^{-1} (\Lambda - \Lambda_0)}, \quad (13)$$

$$\text{and } p(\Psi) \propto |\Psi|^{-\nu/2} e^{-\frac{1}{2} \text{tr} \Psi^{-1} B}, \quad (14)$$

with $\Gamma, \Delta, B, \Psi > 0$ and B a diagonal matrix. ν is set as described in [8].

The posterior distributions are then obtained through Bayes rule giving:

$$p(\mu | F, \Lambda, \Psi, Y) \propto e^{-\frac{1}{2}(\mu - \tilde{\mu})' [(n\Gamma)^{-1} + \Psi^{-1}] (\mu - \tilde{\mu})}, \quad (15)$$

$$p(\Lambda | \mu, F, \Psi, Y) \propto e^{-\frac{1}{2} \text{tr}(\Lambda - \tilde{\Lambda}) [\Delta^{-1} + \Psi^{-1} \otimes F' F] (\Lambda - \tilde{\Lambda})'}, \quad (16)$$

$$p(\Psi | \mu, \mathbf{y}, \Lambda, Y) \propto |\Psi|^{-\frac{(n+\nu)}{2}} e^{-\frac{1}{2} \text{tr} \Psi^{-1} [(Y - e_n \otimes \mu' - F\Lambda')' (Y - e_n \otimes \mu' - F\Lambda') + B]}, \quad (17)$$

where

$$\tilde{\mu} = [(n\Gamma)^{-1} + \Psi^{-1}]^{-1} [(n\Gamma)^{-1} \mu_0 + \Psi^{-1} (\bar{\mathbf{y}} - \Lambda \bar{\mathbf{f}})], \quad (18)$$

$$\tilde{\Lambda} = [\Delta^{-1} + \Psi^{-1} \otimes F' F]^{-1} [\Delta^{-1} \Lambda_0 + (\Psi^{-1} \otimes F' F) ((F' F)^{-1} F' (Y - e_n \otimes \tilde{\mu}'))']. \quad (19)$$

The ICM procedure is then used to estimate the parameters by cycling through their modes $\tilde{\mu}$, $\tilde{\Lambda}$ (as defined above), and

$$\tilde{\Psi} = \frac{(Y - e_n \otimes \mu' - F\Lambda')' (Y - e_n \otimes \mu' - F\Lambda') + B}{n + \nu}, \quad (20)$$

respectively. The Kroneker product is denoted by \otimes and \mathbf{e}_n is a vector of ones with length n .

Thus, the strategy underlying this algorithm can be summarised as follows: (1) Estimate the labelling configuration, $\hat{\mathbf{x}}$, using the current estimate of the parameters; (2) use it to specify the factor scores matrix, \mathbf{F} ; (3) and estimate new values for the parameters, μ , Λ , and Ψ , using the ICM approach outlined above. These steps are iteratively followed until suitable convergence is reached.

3 Experimental Analysis

Examples of applying the method to both simulated and in-vivo data are given in this section. In-vivo results are shown for a myocardial perfusion study.

3.1 Results in Simulated Data Sets

In order to illustrate the method, a simple experiment was constructed. Figure 1(a) shows 3 intensity time-curves generated using the exponential curve model:

$$\mathbf{y}_i = A(1 - e^{-\beta t}). \quad (21)$$

According to myocardial perfusion literature [1], [4]; the replenishment curves can be successfully modelled using the above equation. In this experiment, A was kept constant at a maximum intensity value of $A = 204$, while β was selected as 0.2, 0.6 and 1, respectively, to represent 3 different classes. A classification image template (Fig. 1(b)), showing various random regions classified into 3 different classes, was used to give the original classification of each pixel based on its location in the image template. Image sequences, each consisting of 14 different images ($t = \{0, 1, \dots, 13\}$) could then be generated, where each image was created by selecting an intensity value for every pixel using Eqn. 21 and adding Gaussian noise with mean = 0 and varying standard deviation. The pixel class determined which β value to use in the equation. Therefore each pixel had an intensity time profile (with added Gaussian noise), similar to the replenishment curve of the class that the pixel belonged to. Figure 2(a)-(c), illustrates a perfect image sequence (noise = 0). It can be seen that pixel intensities in each class region, goes from black (intensity values close to zero), to almost white (intensity values close to 204), in different ways depending on the replenishment curve of their class. Since classes 2 and 3, reach the same level of intensity after the 8th timeframe, it can be seen that they overlap in later images in the sequence. In this experiment 3 such image sequences were simulated, where the standard deviation of the Gaussian noise, was 10% (20.4), 15% (30.6) and 45% (91.8) of the maximum intensity value of 204, respectively. Figure 3(a), (b) and (d) show the classification obtained from applying the BFA-MRF method on these image sequences, while (c) shows the BFA model applied without the MRF (noise = 15%). The misclassification ratio (*MCR*) is also shown. From the figures it can be seen that the method performed reasonably well and correct classification was obtained despite the random shapes of the regions. Overlaps and high misclassification ratios, were only seen at very high levels of noise (> 25%), which is as expected. For the BFA model alone, overlaps were already seen at a noise level of 15%, and without the MRF, a much higher misclassification ratio was obtained. Figure 3(d), (noise = 45%), shows that most of the misclassification occurs between regions of classes 2 and 3, while regions of class 1 was correctly classified. The reason for this is the similarity between the 2 classes, as can be seen from their intensity-time profiles in Fig. 1(a). This was deliberately done to test the sensitivity of the method. These results show that the BFA-MRF method can be successfully used to classify perfusion data, even in the presence of noise, and despite certain similarity of classes.




Fig. 1. The three simulated classes represented as (a) 3 replenishment curves; and (b) the classification template where the 3 classes are black ($\beta_1 = 0.2$), grey ($\beta_2 = 0.6$) and white ($\beta_3 = 1$).




Fig. 2. Changing of intensities in the image sequence from very dark (a) to very bright (c).




Fig. 3. Classification results for the BFA-MRF method, along with the misclassification ratio for 3 different noise levels where noise = (a) 10%, (b) 15%, (c) 15% (with the BFA model applied without MRF) and (d) 45% (BFA-MRF again), respectively.

3.2 In Vivo Dataset

A myocardial perfusion study was obtained for a healthy patient, using contrast echocardiography and the Power Pulse Inversion (PPI) technique described in [1]. The dataset consisted of 14 images, acquired in 2-chamber apical view and ECG-triggered. Only the left ventricle was of interest in the classification task. Figure 4(a)-(d) shows 4 images taken from the study, while Fig. 4(e) shows the classification results obtained. Visually, the classification appears good, showing that 3 different regions of perfusion were found. The 3 regions was correctly classified as the myocardium where normal replenishment occurs (grey), the left ventricular cavity where no replenishment occurs (dark grey), and the lateral segments where ultrasound attenuation occurred (black).




Fig. 4. a) Constant venous infusion of contrast agent. b) high power frame to destroy the microbubbles. c) and d) 1st and 5th frame after destruction (low power). e) Classification obtained showing 3 different types of perfusion present.

4 Discussion and Conclusion

This paper has presented a novel method for automatic classification of tissue perfusion ultrasound images. In particular it has been shown how both the spatial and temporal characteristics of the pixels can be incorporated into a single statistical model for classification, which automatically identifies different types of perfusion. Initial experimental results have shown that the model works well even in the presence of high-levels of noise. The preliminary results on in-vivo data are encouraging, showing plausible classification results. There are still problems where classes that are very similar to each other overlap and cause misclassification errors. In future work the method will be tested on patients with proven coronary stenosis to show that the method detects regions where an infarct might have occurred. Comparison with manual semi-quantitative ROI techniques as used in [9] will also be done.

References

1. H. Becher and P. Burns. *Handbook of Contrast Echocardiography*. Springer Verlag, 2000.
2. B.P. Paelinck and J.D. Kasprzak. Contrast-enhanced echocardiography: review and current role. *Acta Cardiol*, vol. 54:195-201, 1999.
3. V. Mor-Avi, S. Akselrod, D. David, L. Keselbrener, and Y. Bitton. Myocardial transit time of the echocardiographic contrast media. *Ultrasound Med Biol*, vol. 19:635-48, 1993.
4. K. Wei, et al. Quantification of myocardial blood flow with US induced destruction of microbubbles administered as a constant venous infusion. *Circulation*, vol. 97:473-83, 1998.
5. D.G. Pavel, et al. FA of dynamic renal studies in urology, *J. Nucl. Med*, 29:P816, 1988
6. A. Martel, A.R. Moody, et al. Extracting parametric images from dynamic contrast enhanced MRI studies of the brain using factor analysis. *Med Image Anal*, vol. 5:29-39, 2002.
7. G. Xiao, J.M. Brady, J.A. Noble, and Y. Zhang. Intensity Inhomogeneity Correction and Segmentation of Ultrasound Images, *IEEE Trans. Medical Imaging*, vol. 21(1):48-57, 2002.
8. Daniel Rowe. A Bayesian Factor Analysis Model with generalized Prior information. *Social science working paper 1099*, California institute of technology, August 2000.
9. A.Z. Linka, J. Sklenar, K. Wei, et al. Assessment of transmural distribution of myocardial perfusion with contrast echo. *Circulation*, vol. 18:1912-20, 1998.