Proceedings of the International Conference on Statistics: Theory and Applications (ICSTA'22)
Prague, Czech Republic July 28 - 30, 2022
Paper No. 146
DOI: 10.11159/icsta22.146

Myocardial Perfusion Classification Using A Markov Random Field Constrained Gaussian Mixture Model

Yalei Yang[1], Hao Gao[1], Colin Berry[2], Aleksandra Radjenovic[2], Dirk Husmeier[1]

[1]School of Mathematics & Statistics, University of Glasgow
University Place Glasgow, G12 8QQ, United Kingdom
y.yang.2@research.gla.ac.uk; hao.gao@glasgow.ac.uk; dirk.husmeier@glasgow.ac.uk
[2]Institute of Cardiovascular & Medical Sciences, University of Glasgow
BHF Glasgow Cardiovascular Research Centre (GCRC), 126 University Place, Glasgow, G12 8TA, United Kingdom
colin.berry@glasgow.ac.uk; aleksandra.radjenovic@glasgow.ac.uk

Abstract - Dynamic Contract Enhanced Magnetic Resonance (MR) Imaging (DCE-MRI) has been widely used as a non-invasive assessment approach to estimate the myocardial blood flow (MBF). The delineation of a hypo-perfused region (low MBF region) is important for understanding a patient's heart condition in clinical diagnosis. In this paper, a Markov random field constrained Gaussian mixture model (GMM-MRF) classification method is introduced to classify MBF maps using myocardial perfusion DCE-MRI data. The GMM-MRF method, trained with an ICM algorithm, makes use of spatial neighbourhood information to improve classification accuracy. The proposed method is applied to and assessed on both synthetic and clinical data, and compared with established classification methods.

Keywords: DCE-MRI, myocardial perfusion, classification, lesion delineation, Gaussian Mixture Model, Markov random field constrained Gaussian mixture model.

1. Introduction

Myocardial ischaemia (inadequate perfusion) is highly related to many forms of coronary heart disease (CHD). Usually, this disease occurs when the blood supply to the myocardium (the heart muscle) is limited. The limitation of the blood supply is generally caused by a narrowing of the coronary arteries. Dynamic contrast-enhanced Magnetic Resonance Imaging (DCE-MRI) is a non-invasive assessment method to measure myocardial perfusion, and further to measure the degree of the hypo-perfusion (inadequate perfusion) of the myocardium [1]. The Fermi model [2] is widely used to quantify the myocardial blood flow (MBF) and detect hypo-perfused regions in the myocardium, and it has been used in this paper to generate the MBF map of the myocardial image.

In this paper, a novel Markov random field constrained Gaussian mixture model (GMM-MRF) is introduced to classify the MBF map into two categories: lesions versus healthy tissues. The Gaussian mixture model (GMM) classification method (see details in [3], Chapter 9) is widely applied to achieve this classification task. However, a main shortcoming of the GMM classification method for myocardial perfusion DCE-MRI is that its classification map contains physiologically unrealistic small and even singular clusters [4]. Spatial prior information has been incorporated to address this issue in [4, 5] using a spatial variant finite mixture model (SVFMM) and hierarchical Bayesian model (HBM). However, the method proposed in [4] only reduces the number of small and singular clusters (rather than eliminate them altogether), and the method proposed in [5] is computationally expensive. Therefore, the main aim of the proposed method is to develop a fast and effective classification method based on the estimated MBF map. In this study, the Fermi model [2] is applied to estimate the MBF map. However, the proposed method is generic and can be applied to other types of images. Specifically, it can be applied to the MBF maps estimated by different methods, or different types of parametric maps. Moreover, it can also be directly applied to the DCE-MRI data.

In this work, the proposed method has been applied to and assessed on both synthetic and clinical data. We have used GMM and HBM, proposed in [5], as benchmark methods for comparison. Section 2 describes the statistical model. Section 3 describes the inference procedures. Section 4 provides an overview of the synthetic and clinical data, on which subsequently a comparative performance analysis is conducted. Section 5 completes this paper with a discussion and concluding remarks.

2. Statistical model

Let x_i represent the signal intensity (or a parameter such as MBF) of pixel i of an image (or a parametric map) where $i = 1, 2, ..., N$ and N is the number of pixels. Any pixel i in the image is either labelled as"healthy" or as " lesion". The pixel with the healthy label has the value $k_i = 0$, and on the other hand, the pixel with the lesion label has the value $k_i = 1$. An assumption of this model is that all pixels that have the label "healthy" are from one Gaussian distribution with mean μ_0 and variance σ_0^2. Similarly, all pixels that have the label "lesion" are from another Gaussian distribution with mean μ_1 and variance σ_1^2. Therefore, the statistical model can be expressed as:

$$P(x_i|\boldsymbol{\mu}, \boldsymbol{\sigma}^2, k_i) = \frac{k_i}{\sqrt{2\pi}\sigma_1} \exp\left(\frac{-(x_i - \mu_1)^2}{2\sigma_1^2}\right) + \frac{1 - k_i}{\sqrt{2\pi}\sigma_0} \exp\left(\frac{-(x_i - \mu_0)^2}{2\sigma_0^2}\right) \tag{1}$$

where $\boldsymbol{\mu} = \{\mu_0, \mu_1\}$, $\boldsymbol{\sigma}^2 = \{\sigma_0^2, \sigma_1^2\}$. Specifically, when a pixel i is in the healthy group and $k_i = 0$, its probability distribution is

$$P(x_i|\boldsymbol{\mu}, \boldsymbol{\sigma}^2, k_i = 0) = \frac{1}{\sqrt{2\pi}\sigma_0} \exp\left(\frac{-(x_i - \mu_0)^2}{2\sigma_0^2}\right). \tag{2}$$

On the other hand, if a pixel is the in the lesion group, its probability distribution has a similar form to equation (2) with replacements of the parameters from μ_0 and σ_0^2 to μ_1 and σ_1^2. The advantage of this modelling approach is that given the label k_i of pixel i, the conditional distribution of x_i is a standard Gaussian distribution. Let $X = \{x_1, x_2, ..., x_N\}$, the joint probability distribution of X is

$$P(X|\boldsymbol{\mu}, \boldsymbol{\sigma}^2, K) = \prod_{i=1}^{N} P(x_i|\boldsymbol{\mu}, \boldsymbol{\sigma}^2, k_i) \tag{3}$$

where $K = \{k_1, k_2, ..., k_N\}$. Because of the sum in equation (1), computing the log likelihood function from equation (3), which requires a marginalisation over the unknown labels, is analytically intractable. However, if the labels K are assumed to be known, the conditional likelihood function of equation (3) has a closed form. Specifically, making use of equation (2), the conditional negative log likelihood function can be expressed as

$$-\log P(X|\boldsymbol{\mu}, \boldsymbol{\sigma}^2, K) \propto N_0 \log \sigma_0 + N_1 \log \sigma_1 + \sum_{p=1}^{N_0} \frac{(x_p - \mu_0)^2}{2\sigma_0^2} + \sum_{q=1}^{N_1} \frac{(x_q - \mu_1)^2}{2\sigma_1^2} \tag{4}$$

where N_0 is the number of pixels that are labelled as healthy and N_1 is the number of pixels that are labelled as lesions. p and q are the indices for healthy and lesion pixels.

The myocardial lesion pixels are physiologically connected spatially because the blood supply to the myocardium is from three coronary arteries (which can potentially get blocked), and each artery controls a connected region (rather than isolated muscle cells). A Markov random field prior is used to model these spatial connections in the present work. Specifically, the Markov random field prior for label k_i can be expressed as

$$P(k_i|\boldsymbol{k}_{-i}) = \frac{1}{Q} \exp\left(-U(k_i|\boldsymbol{k}_{-i})\right) \tag{5}$$

where $\boldsymbol{k}_{-i}$ denotes the set of all other labels except k_i. Specifically, $\boldsymbol{k}_{-i} = \{k_j\}_{j \neq i}$. Q is a normalization constant of the prior distribution $P(k_i|\boldsymbol{k}_{-i})$. The function U is usually called the potential function. Since only the neighbours of pixel i is considered to affect the value of k_i, the potential function U is defined as

$$U(k_i|\boldsymbol{k}_{-i}) = \frac{1}{T} \sum_{j \sim i} u(k_i|k_j) \tag{6}$$

where $j \sim i$ represents all pixels j that are the neighbours of pixel i and T is the weight parameter. In this study, the neighbours

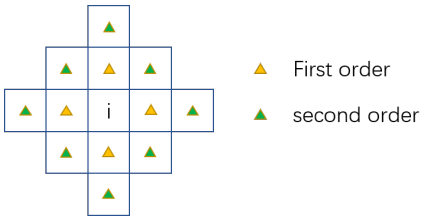

Fig. 1: This figure shows pixel i and its neighbours. The four pixels marked by yellow triangles are the first order neighbours. The eight pixels marked by green triangles are the second order neighbours.

of pixel i is defined in Fig. 1, and furthermore u is defined as

$$u(k_i|k_j) = \begin{cases} -\left(\frac{1}{2}\right)^{o-1} & k_i = k_j, \\ \left(\frac{1}{2}\right)^{o-1} & k_i \neq k_j \end{cases} \tag{7}$$

where o is the degree of order for the neighbours. For instance, if pixel j is the first order neighbour of pixel i, $o = 1$, and if j is the second order neighbour, $o = 2$. In this work, up to second order neighbours have been used: $o \in \{1, 2\}$.

According to Bayes theorem, the posterior distribution of k_i is

$$P(k_i|x_i, \boldsymbol{\mu}, \boldsymbol{\sigma}^2, \boldsymbol{k}_{-i}) \propto P(x_i|k_i, \boldsymbol{\mu}, \boldsymbol{\sigma}^2)P(k_i|\boldsymbol{k}_{-i}) \tag{8}$$

and the posterior distribution of K is

$$P(K|X, \boldsymbol{\mu}, \boldsymbol{\sigma}^2, K^-) \propto \prod_{i=1}^{N} P(x_i|k_i, \boldsymbol{\mu}, \boldsymbol{\sigma}^2)P(k_i|\boldsymbol{k}_{-i}) \tag{9}$$

where K^- is the set of the neighbouring labels, $K^- = \{\boldsymbol{k}_{-i}\}_{i=1}^N$. This posterior distribution is intractable because of the summation in equation (1). However, an algorithm called iterated conditional modes (ICM) [6] can be applied to find the maximum a posteriori (MAP) estimate of equation (9).

3. ICM algorithm

In order to apply the ICM algorithm, two conditions must be satisfied. Firstly, the observations x_i have to be conditionally independent given their corresponding labels k_i. Secondly, the conditional dependence structure of the labels k_i can be described by a Markov random field. Both conditions are satisfied in this study, and the ICM algorithm can thus be applied.

The negative logarithm of the posterior distribution $P(k_i|x_i, \boldsymbol{\mu}, \boldsymbol{\sigma}^2, k_{-i})$ in equation (8) can be expressed as

$$L_i = -\log P(k_i|x_i, \boldsymbol{\mu}, \boldsymbol{\sigma}^2, k_{-i}) \propto \begin{cases} \log\sigma_0 + \dfrac{(x_i-\mu_0)^2}{2\sigma_0^2} + \dfrac{1}{T}\displaystyle\sum_{j\sim i} u(k_i|k_j) & k_i = 0, \\ \log\sigma_1 + \dfrac{(x_i-\mu_1)^2}{2\sigma_1^2} + \dfrac{1}{T}\displaystyle\sum_{j\sim i} u(k_i|k_j) & k_i = 1. \end{cases} \tag{10}$$

The ICM algorithm aims to decrease the negative logarithm of the conditional posterior probability by updating the value of label k_i. Specifically, given the old value k_i^{old}, the old negative logarithm of the conditional posterior probability L_i^{old} can be calculated. Then, by proposing a new value k_i^{new}, L_i^{new} can be calculated. If $L_i^{\text{new}} < L_i^{\text{old}}$, then the old value is updated from k_i^{old} to k_i^{new}. By iteratively processing these steps for all pixels, the updated labels K^{new} are obtained. Given $K = K^{\text{new}}$, the

new global negative logarithm of the conditional posterior probability can be calculated as

$$G^{\text{new}} = -\log P(K = K^{\text{new}}|X, \boldsymbol{\mu}, \boldsymbol{\sigma}^2, K^-) \propto N_0 \log \sigma_0 + N_1 \log \sigma_1 + \sum_{p=1}^{N_0} \frac{(x_p - \mu_0)^2}{2\sigma_0^2} + \sum_{q=1}^{N_1} \frac{(x_q - \mu_1)^2}{2\sigma_1^2} + \frac{1}{T} \sum_{i=1}^{N} \sum_{j \sim i} u(k_i|k_j). \tag{11}$$

This equation can also be used to calculate the old negative logarithm of the global conditional posterior probability G^{old} by replacing K^{new} by K^{old}.

The conditional posterior distributions for $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}^2$ have closed forms. Since the prior distribution of k_i is independent on $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}^2$, the estimations of $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}^2$ can be obtained using the log conditional likelihood function. Specifically, setting the derivatives with respect to $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}^2$ of equation (4) to zero gives the following maximum likelihood estimates (MLE):

$$\hat{\mu}_0 = \frac{1}{N_0} \sum_{m=1}^{N_0} x_m \tag{12}$$

$$\hat{\sigma}_0^2 = \frac{1}{N_0} \sum_{m=1}^{N_0} \left(x_m - \frac{1}{N_0} \sum_{m=1}^{N_0} x_m \right)^2. \tag{13}$$

$$\hat{\mu}_1 = \frac{1}{N_1} \sum_{n=1}^{N_1} x_n \tag{14}$$

$$\hat{\sigma}_1^2 = \frac{1}{N_1} \sum_{n=1}^{N_1} \left(x_n - \frac{1}{N_1} \sum_{n=1}^{N_1} x_n \right)^2. \tag{15}$$

Given the expression of the conditional posterior distribution and the MLE of the parameters, the details of the MAP estimation using the ICM algorithm can be found in algorithm 1.

4. Data and results

4.1. Data

The cardiac MRI exams were performed with a Siemens MAGNETOM Avanto (Erlangen, Germany) 1.5-Tesla scanner with a 12-element phased array cardiac surface coil at Golden Jubilee National Hospital, Glasgow, UK. The assessment of resting myocardial perfusion was performed during intravenous administration of 0.075 mmol/kg of contrast agent (gadoterate meglumine, Dotarem, Guebert S.A.). The mid slice of the myocardial perfusion DCE-MRI data is used to estimate the MBF by the Fermi method [2]. The choice of the clinical data is consistent with the clinical data used in [5] for easier comparisons.

The synthetic data is designed to mimic the clinical data used in this work. Similarly, for easier comparison, the synthetic data is also consistent with the data generated in [5]. A summary of the design is illustrated here. A double exponential curve (16) is used to generate the true signal intensity values given time point t for healthy tissues and lesions:

$$s(t) = \frac{p_2 p_3}{(p_2 - p_1)} \times (e^{-p_1 t} - e^{-p_2 t}). \tag{16}$$

In other words, for all pixels with the same label, “healthy" versus “lesion", the corresponding true signal intensities are the same. Gaussian noise (general noise) and Rician noise (commonly assumed for MRI data [8]) are added to the signals to generate the synthetic data. In this work, the parameters for healthy tissue are set to $p_1 = 0.01$, $p_2 = 0.4$, $p_3 = 25$, and the parameters for lesions are set to $p_1 = 0.02$, $p_2 = 0.3$, $p_3 = 20$. The details of the choices of these values can be found in [5].

4.2. Results

Algorithm 1: ICM algorithm for the Markov random field constrained Gaussian mixture model

Data: $X = \{x_1, x_2, ..., x_N\}$
Output: $K = \{k_1, k_2, ..., k_N\}$
Generate initial values of K and set it as K^{old}. The initial values can be obtained using standard established clustering algorithms, e.g. k-means [7] or GMM. To avoid the local minima issue, multiple initial values can be generated ;
Given K^{old}, estimate $\boldsymbol{\mu}^{\text{old}}$ and $(\boldsymbol{\sigma}^2)^{\text{old}}$ using equations (12)-(15);
Given K^{old}, $\boldsymbol{\mu}^{\text{old}}$ and $(\boldsymbol{\sigma}^2)^{\text{old}}$, calculate global negative logarithm posterior G^{old} using equation (11).;
for $i \leftarrow 1$ **to** N **do**
    Given K^{old}, calculate $\boldsymbol{\mu}^{\text{old}}$ and $(\boldsymbol{\sigma}^2)^{\text{old}}$ using equations (12)-(15);
    Given K^{old}, calculate old local negative logarithm posterior L_i^{old} for pixel i using equation (10);
    Give a new $k_i^{\text{new}} = |k_i^{\text{old}} - 1|$;
    Given K^{new}, calculate new $\boldsymbol{\mu}^{\text{new}}$ and $(\boldsymbol{\sigma}^2)^{\text{new}}$ using equations (12)-(15);
    Given new k_i^{new}, $\boldsymbol{\mu}^{\text{new}}$ and $(\boldsymbol{\sigma}^2)^{\text{new}}$, calculate new local negative logarithm posterior L_i^{new} for pixel i using equation (10);
    If $L_i^{\text{new}} < L_i^{\text{old}}$, update old k_i^{old} using new k_i^{new}.
end
Given K, calculate $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}^2$ using equations (12)-(15);
Given K, $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}^2$, calculate global negative logarithm posterior G^{new} using equation (11);
Given a small threshold $\epsilon = 0.01$;
if $G^{old} - G^{new} > \epsilon$ **then**
    Update G^{old} using G^{new};
    Go back to the for loop above;
else
    Stop the algorithm and the latest updated values of K are the estimated MAP;
end

Figure 2 shows the synthetic data and results generated given Rician noise with scale $\sigma^2 = 1.8^2$. Panel (c) shows the GMM classification map, which contains many small and singular clusters, as mentioned in Section 1. Panels (d) (f) and (g) present the classification maps using the proposed GMM-MRF method, for different weight parameters T. For $T = 10$, the spatial prior is not strong enough, and there are still a few small and singular clusters left. However, when the strength of the spatial prior is increased, $T = 1$ and $T = 0.1$, all singular clusters are removed, and only a single connected lesion is predicted. However, even the best performance (the one that is closest to the ground truth in panel (b)), $T = 1$, is still not as good as the classification map generated by the HBM method proposed in [5]. However, the HBM model is more complex and incurs much higher computational costs, as we discuss in more detail further below.

To quantitatively compare these methods, 10 sets of synthetic data given different types of noise and different values of scales are generated. The average number of misclassified pixels is used as the assessment criterion to compare different methods, which is shown in Table 1. It can be seen that the proposed GMM-MRF method with different weight parameters performs much better than the GMM method, but is not as good as the HBM method proposed in [5]. In particular, when $T = 1$, the GMM-MRF method performs better than with $T = 0.1$ and $T = 10$. When the weight parameter is too large ($T = 10$), the spatial information is too weak to eliminate all singular and small clusters. On the other hand, when the weight parameter is too small ($T = 0.1$), the spatial information is too strong and more borderline pixels are misclassified.

As shown in Table 1 and Fig. 2, the proposed GMM-MRF method is not as accurate as the HBM method in [5]. The reason is that the proposed method is designed as a computationally cheaper alternative. In fact, the choice of different models can be considered as a trade-off between computational efficiency and accuracy. While the proposed GMM-MRF method is not as accurate as the HBM method in [5], it is much faster to calibrate and apply. In our studies, it took less than 1 minute to apply the GMM-MRF method, while the HBM method required run times of several hours using the same hardware.

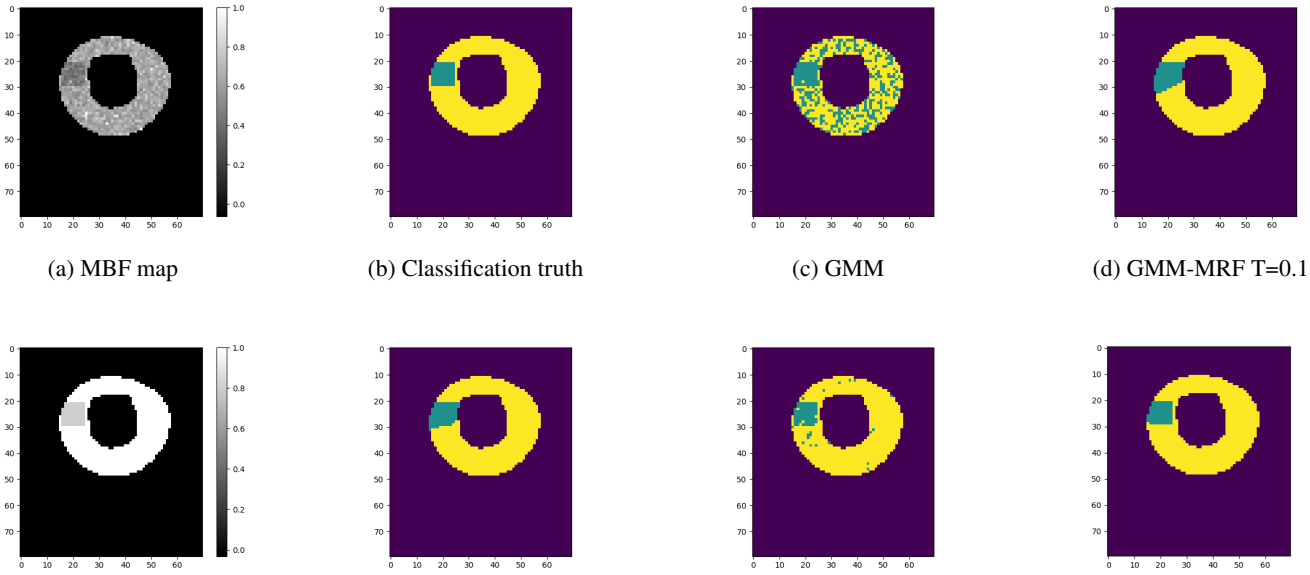

(a) MBF map (b) Classification truth (c) GMM (d) GMM-MRF T=0.1

(e) MBF truth (f) GMM-MRF T=1 (g) GMM-MRF T=10 (h) HBM

Fig. 2: Synthetic data and results generated given Rician noise with scale $\sigma^2 = 1.8^2$. Panel (a) shows the estimated MBF map (rescaled to [0,1]). Panel (b) shows the ground truth of the classification map. Panel (c) shows the classification map generated with the GMM method. Panel (e) shows the ground truth of the MBF map (rescaled to [0,1]). Panels (d) (f) and (g) show the classification maps generated with the proposed GMM-MRF method, given $T = 0.1, 1, 10$ (smaller T indicates stronger spatial prior). Panel (h) shows the classification map generated with the HBM method proposed in [5]. Yellow pixels indicate healthy tissues. Dark green pixels indicate lesions. Black pixels represent the background, which can be ignored.

The proposed GMM-MRF method has also been tested on the clinical data. Figure 3 shows the clinical data and classification maps generated by different methods. In panel (b), the classification map generated by the GMM method shows a few small and singular clusters at the right bottom of the myocardial image. The GMM-MRF method eliminates all singular clusters when $T = 0.1$ and $T = 1$. However, it can be seen that a connected group of misclassified pixels is located in the right bottom of the myocardial image in panels (d) - (f), which may be a consequence of spatial correlated noise. There is no such issue in the synthetic data, where the noise is assumed to be independent. However, for the clinical data, because of some artifacts, e.g. intensity inhomogeneity, the image may have spatial correlated noise, which affects the parametric map (panel (a) in Figure 3). Recall that the ICM algorithm is a greedy optimization algorithm, which may get stuck in local optima, and some of these local optima appear to be related to artifacts caused by the spatial correlation structure of the noise. The HBM method, on the other hand, is a sampling-based method, using an MCMC algorithm, which avoids any entrapment in suboptimal local optima, albeit at substantially increased computational costs. Interestingly, our results suggest that if the spatial correlation of the noise is not too strong, the novel proposed GMM-MRF method can also accurately delineate the position of the hypo-perfusion region (see panels (g) - (i) in Figure 3), without the erroneous prediction of any spurious lesions.

5. Discussion and Conclusion

The main aim of the present paper has been to develop a fast spatially constrained classification method to automatically detect lesions in CMR images. The proposed GMM-MRF method is a GMM imbued with spatial information using a Markov random field prior. A fast greedy optimization algorithm, the ICM algorithm, has been used to train the proposed GMM-MRF model and find the MAP. The difference between this algorithm and the more commonly used EM algorithm [9] has been described in [10]. The EM algorithm aims to optimize the parameters $\boldsymbol{\theta}$ ($\boldsymbol{\mu}$ and $\boldsymbol{\sigma}^2$ in our case) by iteratively proceeding:

Table 1: The average number of misclassified pixels

	SNR/sd*(Gaussian)		**SNR/sd(Rician)**
	4.1/2	3.27/2.5	2.91/1.8
GMM	64.0	234.4	183.4
GMM-MRF $T = 0.1$	27.2	27.4	43.6
GMM-MRF $T = 1$	3.8	27.1	20.6
GMM-MRF $T = 10$	40.3	60.7	31.6
HBM Proposed in [5]	2.9	4	3.5

*The abbreviation "sd" in the table stands for standard deviation of the Gaussian or Rician noise, and "SNR" means signal-to-noise ratio. The results are averages over 10 independent synthetic datasets.

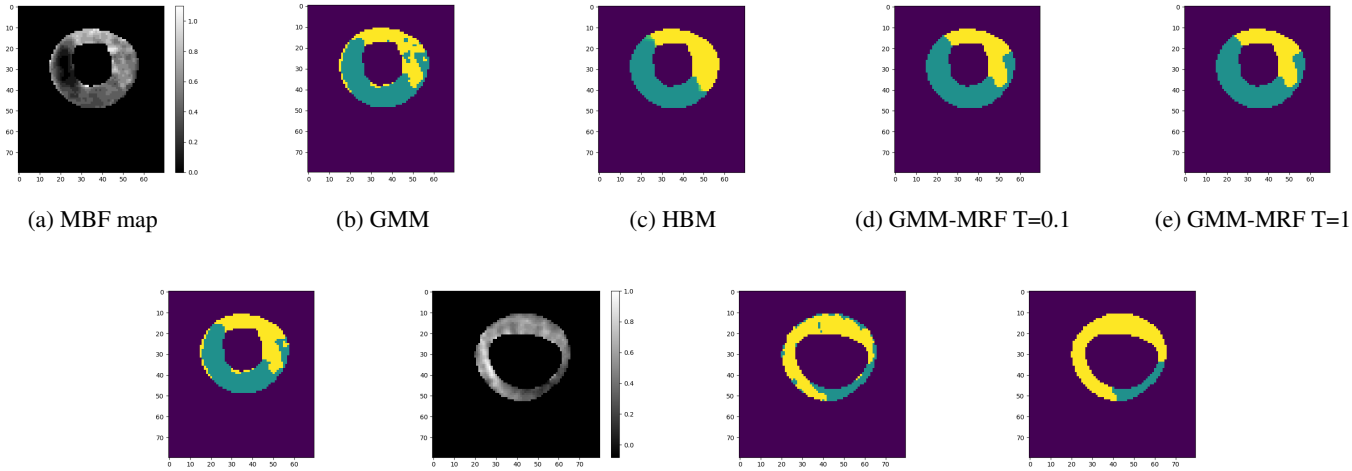

(a) MBF map (b) GMM (c) HBM (d) GMM-MRF T=0.1 (e) GMM-MRF T=1

(f) GMM-MRF T=10 (g) MBF extra (h) GMM extra (i) GMM-MRF extra

Fig. 3: Clinical data and model predictions. Panel (a) shows the estimated MBF map (rescaled to [0,1]). Panel (b) shows the classification map generated by the GMM method. Panel (c) shows the classification map generated by the HBM method proposed in [5]. Panels (d) - (f) show the classification maps generated by the GMM-MRF method, given $T = 0.1, 1, 10$ (a smaller T indicates a stronger spatial prior). Panels (g) - (i) show the MBF map, GMM classification map and GMM-MRF classification map, respectively, for an extra data set. Yellow pixels indicate healthy tissues. Dark green pixels indicate lesions. Black pixels represent the background, which can be ignored.

1. $\boldsymbol{\theta}^* = \arg\max_\theta E\left[\log P(\boldsymbol{\theta}, K, \mathbf{x})\right]_{\gamma(K)}$
2. $\gamma(K) = P(K|\boldsymbol{\theta}^*, \mathbf{x})$.

The ICM algorithm, on the other hand, aims to optimize the parameter vector $\boldsymbol{\theta}$ and the hidden labels K with the following iterative procedure:

1. $\boldsymbol{\theta}^* = \arg\max_\theta P(K^*, \boldsymbol{\theta}, \mathbf{x})$
2. $K^* = \arg\max_K P(K, \boldsymbol{\theta}^*, \mathbf{x})$.

In this way, the ICM algorithm can be seen as a (faster) modal approximation of the EM algorithm, and it has been shown to be equivalent to the coordinate descent algorithm [11]. The adoption of the ICM algorithm in our work reflects a trade-off between computational efficiency and accuracy [10].

We have assessed the performance of the proposed GMM-MRF model on synthetic and real clinical data, and compared it with two alternative methods: a standard GMM without the incorporation of spatial neighbourhood information, and the HBM of [5]. The new GMM-MRF model substantially outperforms the standard GMM, in that it avoids the singular and physiologically unrealistic small clusters that the latter is susceptible to. The delineation of the lesions is not as accurate as achieved with the HBM, though (see Figure 2), and it can occasionally lead to spurious lesion predictions (see Figure 3 c-f). This is a consequence of the greedy nature of the training scheme, which renders it susceptible to entrapment in local optima. The MCMC-based sampling scheme of the HBM in [5] avoids this problem and is more robust. However, this performance improvement comes with substantially higher computational costs, typically in the order of hours, whereas training the new GMM-MRF model with the ICM algorithm achieves results in the order of minutes. For clinical applications, where decisions need to be made in real time, this improvement in computational efficiency is important. We therefore conclude that the two methods complement each other and should be used in combination: First run the proposed GMM-MRF model with the ICM algorithm to get a first-pass prediction in real time, for preliminary clinical assessment, while running the HBM model in parallel for refined follow-up diagnosis available at some later time (in the order of an hour). Our future work will focus on including spatial correlation in the noise as a potential remedy to avoid some of the current artifacts [12].

Acknowledgements

This work was funded by EPSRC, grant reference numbers EP/T017899/1 and EP/S020950/1. Yalei Yang is funded by a grant from GlaxoSmithKline plc.

References

[1] M. Jerosch-Herold, "Quantification of myocardial perfusion by cardiovascular magnetic resonance," *Journal of Cardiovascular Magnetic Resonance*, vol. 12, no. 1, p. 57, 2010.

[2] M. Jerosch-Herold, N. Wilke, A. E. Stillman, and R. F. Wilson, "Magnetic resonance quantification of the myocardial perfusion reserve with a Fermi function model for constrained deconvolution," *Medical physics*, vol. 25, no. 1, pp. 73–84, 1998.

[3] C. M. Bishop, *Pattern recognition and machine learning*. Springer, 2006.

[4] Y. Yang, H. Gao, C. Berry, A. Radjenovic, and D. Husmeier, "Quantification of myocardial perfusion lesions using spatially variant finite mixture modelling of DCE-MRI," *Proceedings of the International Conference on Statistics: Theory and Applications (ICSTA)*, 2019.

[5] Y. Yang, H. Gao, C. Berry, D. Carrick, A. Radjenovic, and D. Husmeier, "Classification of myocardial blood flow based on dynamic contrast enhanced magnetic resonance imaging using hierarchical Bayesian models," *Journal of the Royal Statistical Society: Series C (Applied Statistics), in press*, 2022.

[6] J. Besag, "On the statistical analysis of dirty pictures," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 48, no. 3, pp. 259–279, 1986.

[7] A. Likas, N. Vlassis, and J. J. Verbeek, "The global k-means clustering algorithm," *Pattern recognition*, vol. 36, no. 2, pp. 451–461, 2003.

[8] H. Gudbjartsson and S. Patz, "The Rician distribution of noisy MRI data," *Magnetic resonance in medicine*, vol. 34, no. 6, pp. 910–914, 1995.

[9] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 39, no. 1, pp. 1–22, 1977.

[10] M. Welling and K. Kurihara, "Bayesian K-means as a maximization-expectation algorithm," in *Proceedings of the 2006 SIAM international conference on data mining*, pp. 474–478, SIAM, 2006.

[11] T. Hebert and R. Leahy, "A generalized EM algorithm for 3-D Bayesian reconstruction from poisson data using Gibbs priors," *IEEE transactions on medical imaging*, vol. 8, no. 2, pp. 194–202, 1989.

[12] A. Banerjee and P. Maji, "Rough sets and stomped normal distribution for simultaneous segmentation and bias field correction in brain MR images," *IEEE Transactions on Image Processing*, vol. 24, no. 12, pp. 5764–5776, 2015.