An HMRF-EM Algorithm for Partial Volume Segmentation of Brain MRI

FMRIB Technical Report TR01YZ1

Yongyue Zhang, J. Michael Brady[2] and Stephen M. Smith

FMRIB Centre, Oxford University
[2]Department of Engineering Science, Oxford University

Abstract

The hidden Markov random field (HMRF) model, which represents a stochastic process generated by a Markov random field whose state sequence cannot be observed directly but which can be indirectly estimated through observations, has been successfully applied on segmenting piecewise constant images, especially for human brain MR images [Zhang et al., 2001]. In this paper, we extend this model, and its associated expectation-maximization (EM) algorithm from dealing only with discrete class labels to working with continuous vectors. In particular, the problem of partial volume effect classification of brain MR images is addressed. In our method, the underlying partial volume classification, as well as its interaction with the observed image intensities, is modelled as a spatially correlated hidden Markov random field, with parameters estimated through an EM algorithm. A deterministic annealing algorithm is then used to obtain the optimal solution; this method is shown to work both for multi-spectral data (where a unique solution exists) and for data with insufficient spectral channels (where only an 'optimal' solution can be obtained). Quantitative evaluations are presented to examine the accuracy and the repeatability of this algorithm.

Keywords

MRI, Partial Volume Effect Classification, Segmentation, Hidden Markov Random Field, Expectation-Maximization.

1 Introduction

A voxel in an MR image may 'contain' a single type of tissue or a combination of different types; this is known as the *partial volume effect* (PVE). A voxel-level 'hard' segmentation does not model the tissue distribution inside a voxel; instead, a sub-voxel classification scheme is required for detailed modelling of such data.

Intuitively, because the partial volume effect is caused by the limited spatial resolution of MRI, it should only appear at boundaries of different tissue types, and have a maximum influence width of 1 or 2 voxels across the boundaries. In practice, however, the PVE affects a much wider area. The reason lies in the blurring effect of the imaging process, which has a diffuse point spread function, effectively mixing the intensity of each voxel with many of its neighbours. Together with the PVE at tissue boundaries, the blurring effect, depending on its strength and the imaging quality, makes the mixing effect of intensities from different tissue types appear in a wider area of the image.

Given the observed image, the objective of PVE classification is to estimate, for each partial volume voxel, the proportion of each tissue type, i.e., form a partial volume vector, with each element being a 'fraction' of a specific tissue type and having a sum of 1. For consistency, we can also consider pure-tissue voxels to be special cases of partial volume voxels in the sense they have a partial volume vector that has only one non-zero member.

Depending on the assumption of the possible number of tissue types seen in a partial volume voxel, currently available approaches to PVE estimation can be grouped into two categories: the two-tissue mixing model and the multi-tissue mixing model.

A two-tissue mixing model was proposed by Santago and Gage [Gage et al., 1992, Santago and Gage, 1993]. It is assumed that every voxel is either a pure-tissue voxel or a PVE voxel of a mixture of two tissue types. Accordingly, they proposed a six-class model, representing the three primary tissue (pure-tissue) types GM, WM and CSF, and the three two-tissue classes GW (GM+WM), CG (CSF+GM) and CW (CSF+WM). It is also assumed that the three primary tissue classes are all normally distributed with class-dependent means but the same standard deviation. For a given PVE voxel, its two underlying tissue types, a and b for example, have the fractions of α and $1-\alpha$ respectively, α being a uniformly distributed random variable, so that given α the intensity also follows a Gaussian distribution with mean $\alpha\mu_a + (1-\alpha)\mu_b$ and standard deviation of σ. Thus, the intensity distribution of the image can be written as

$$p(y) = \sum_t p(y|t)\omega_t,$$

where t is one of the six tissue types. The parameter set for this model includes the means of each primary tissue type, the standard deviation of the noise and the six mixing parameters ω_t. Since the model does not use any spatial information, the estimation of those parameters was carried out via least-squares fitting of the density function to the image histogram using a tree annealing scheme. Once the model is fitted, a labelling of the image with regard to each voxel's tissue type (one of the six classes), as well as the quantification of each primary tissue type can easily be derived. However, this model is not able to provide information about the proportion of each tissue type within a specific voxel, and consequently it is not able to solve the problem of PVE classification. Moreover, ambiguity exists for voxels having intensities close to the mean of the intermediate primary class, GM for T1-weighted images for example, since it is impossible to know whether they are GM voxels or mixtures of CSF and WM. The advantage of this approach is that it works for single-spectral MR images.

A multi-spectral model that enables partial volume classification was proposed by Choi *et al.* [Choi et al., 1991], and referred to as the 'mixel' model. It is assumed that all voxels are PVE ones. Given the true partial volume vector (mixel) $\vec{z}_i$ at pixel i, the observed intensity $\vec{y}_i$ follows a multi-dimensional Gaussian distribution with mean

$$\mathbf{E}_i = \sum_{\ell \in \mathcal{L}} \vec{z}_{i\ell}\boldsymbol{\mu}_\ell = \boldsymbol{\mu}^T \vec{z}_i \tag{1}$$

and covariance matrix Ψ, where $\boldsymbol{\mu}$ is a $k \times m$ matrix with each row being the mean vector of the corresponding tissue type and m being the number of spectral channels. It is assumed that the biological variations in intensity are negligible and the observed variation arises from the noise only. Thus the intensity distribution at pixel i, given the true mixel $\vec{z}_i$, can be written as follows:

$$p(\vec{y}_i|\vec{z}_i) = g(\vec{y}_i; \mathbf{E}_i, \Psi). \tag{2}$$

An MRF prior model is used to impose contextual constraints and the voxel-by-voxel estimates of $\vec{z}_i$ are found through an MAP approach:

$$\hat{\vec{z}}_i = \arg\max_{\vec{z}_i}\{P(\vec{y}_i|\vec{z}_i)P(\vec{z}_i|\vec{z}_{\mathcal{N}_i})\} \tag{3}$$

subject to $\sum_{\ell \in \mathcal{L}} \vec{z}_{i\ell} = 1$ and $0 \leq \vec{z}_{i\ell} \leq 1$, $\ell \in \mathcal{L}$. This minimization problem can be further written as a minimization of the quadratic form

$$\frac{1}{2}(\vec{z}_i^T Q \vec{z}_i + \vec{b}^T \vec{z}_i + \vec{z}_i^T \vec{c} + \vec{d})$$

where $Q = \boldsymbol{\mu}\Psi^{-1}\boldsymbol{\mu}^T$ is a $k \times k$ matrix. A sufficient condition for this minimization to have a unique solution is that the rank of matrix Q should be at least the same as the number of unknowns in $\vec{z}_i$, which is $k - 1$. Because it is only if $\boldsymbol{\mu}$ has a rank of $k - 1$ that Q has the rank of $k - 1$, it is equivalent to say that the sufficient condition for this minimization to have a unique solution is that $m \geq k - 1$. Give this sufficient condition, the estimation of the mean vector $\boldsymbol{\mu}$ is achieved iteratively using a least-squares fitting by solving the following linear equation:

$$\vec{y} \cong \hat{\boldsymbol{\mu}}\vec{z},$$

so that

$$\hat{\boldsymbol{\mu}} = (\vec{z}^T\vec{z})^{-1}\vec{z}^T\vec{y}.$$

Once a new estimate of $\hat{\boldsymbol{\mu}}$ is obtained, it will be used in Equation 3 to get the next estimate $\vec{z}_i$ for each voxel and such an iterative process continues until convergence. The covariance matrix of the noise, as well as the initial mean vector of each tissue type, is estimated from the training data obtained manually by operators.

This approach has established a solid computational framework for PVE classification for multi-tissue and multi-spectral MR data. There are only a few drawbacks. First, the least-square estimation of the mean matrix may be difficult to solve given the size of the equation and the presence of noise. Second, the assumption of an equal covariance matrix for all classes may not be valid given that some tissue types, such as CSF, normally have much larger variations than others. Finally, this method can only be used for multi-spectral data sets which are not always available.

There are other methods dealing with PVE quantification or classification problems but most of them are similar to the two approaches mentioned above [Santago and Gage, 1995, Jaggi et al., 1998, Pham and Prince, 1998, Ruan et al., 2000]. In this paper, we extend the our previously proposed HMRF-EM framework [Zhang et al., 2001] in such a way that the hidden entities are no longer discrete labels but continuous vectors that describe the sub-voxel-level proportions of each tissue class. Applying this new framework we are able to achieve classification with sub-voxel accuracy.

2 An HMRF Model for Partial Volume Effect

The PVE classification problem differs from the tissue segmentation problem in that, in the former a continuous vector is assigned to each voxel, while in the latter, a discrete label is assigned to each voxel. Apart from that, however, these two problems are similar. In this sense, the hidden Markov random field model we have derived previously is also applicable to the PVE classification problem but with the hidden random field being a set of continuous vectors instead of discrete labels. For consistency we assume the data is multi-spectral, but the model is also valid for single-spectral data with slight changes in the functional form of the Gaussian.

2.1 Notation and Assumptions

Let $\mathcal{S} = \{1, 2, \cdots, N\}$ be the set of indices, $\boldsymbol{Y}$ and $\boldsymbol{Z}$ be two multi-dimensional random fields indexed by $\mathcal{S}$. We assume that $\forall i \in \mathcal{S}$, $\boldsymbol{Y}_i$ is a m-dimensional random vector and $\boldsymbol{Z}_i$ a k-dimensional random vector satisfying

$$\sum_{\ell \in \mathcal{L}} Z_{i\ell} = 1 \text{ and } 0 \leq \vec{Z}_{i\ell} \leq 1,\ \ell \in \mathcal{L}, \tag{4}$$

where $\mathcal{L} = \{1, \cdots, k\}$ is the set of labels.

Let $\mathbf{y}$ be a configuration of $\boldsymbol{Y}$ and $\mathcal{Y}$ be the set of all possible configurations so that

$$\mathcal{Y} = \{\mathbf{y} = (\vec{y}_1, \cdots, \vec{y}_N)\}.$$

Similarly, let $\mathbf{z}$ denote a configuration of $\boldsymbol{Z}$ and $\mathcal{Z}$ be the set of all possible configurations so that

$$\mathcal{Z} = \{\mathbf{z} = (\vec{z}_1, \cdots, \vec{z}_N)\}.$$

Furthermore, it is assumed that given the value of $\boldsymbol{Z}_i = \vec{z}_i$ at i, $\boldsymbol{Y}_i = \vec{y}_i$ follows an m-dimensional Gaussian distribution with mean Σ_i and covariance matrix Ψ_i, which are defined as follows:

$$\Sigma_i = \sum_{\ell \in \mathcal{L}} z_{i\ell} \boldsymbol{\mu}_\ell = \boldsymbol{\mu}^T \vec{z}_i \tag{5}$$

where $\boldsymbol{\mu}$ is a $k \times m$ matrix and $\boldsymbol{\mu}_\ell$ is an m-dimensional vector. $\boldsymbol{\mu}$ is referred to as the mean matrix and $\boldsymbol{\mu}_\ell$ the mean vector of class ℓ. The covariance matrix at pixel i is defined as

$$\Psi_i = \sum_{\ell \in \mathcal{L}} z_{i\ell}^2 \boldsymbol{\psi}_\ell, \tag{6}$$

where $\boldsymbol{\psi}_\ell$ is the $m \times m$ covariance matrix for class ℓ.

We can write the probability distribution of $\vec{y}_i$, given $\vec{z}_i$, as

$$\begin{aligned} p(\vec{y}_i|\vec{z}_i) &= g(\vec{y}_i; \Sigma_i, \Psi_i) \\ &= (2\pi)^{-m/2} |\Psi_i|^{-1/2} \\ &\quad \exp\Big(-\frac{1}{2} (\vec{y}_i - \Sigma_i)^T \Psi_i^{-1} (\vec{y}_i - \Sigma_i) \Big). \end{aligned} \tag{7}$$

With respect to our particular problem of image PVE classification, $\mathbf{y}$ corresponds to a particular m-spectral image and $\mathbf{z}$ its PVE classification. Thus, each row of the mean matrix $\boldsymbol{\mu}$ actually represents the mean of a primary tissue type.

2.2 Gaussian Hidden Markov Random Field Model

Applying the above to the hidden Markov random field model, we obtain

- **Hidden Random Field (MRF)**

  The underlying random field $\boldsymbol{Z} = \{\vec{Z}_i, i \in \mathcal{S}\}$ is modelled as a multivariate MRF with probability distribution

  $$P(\mathbf{z}) = \frac{1}{M} \exp\{-\beta U(\mathbf{z})\}, \tag{8}$$

  where M is a normalizing factor, β is the neighbourhood parameter and

  $$U(\mathbf{z}) = \sum_{c \in \mathcal{C}} V_c(\mathbf{z}). \tag{9}$$

  The clique potential function V_c is determined by the distance measure:

  $$V_c(\vec{z}_i, \vec{z}_j) = ||\vec{z}_i - \vec{z}_j||^2 \tag{10}$$

- **Observable Random Field**

  The observed random field $\boldsymbol{Y} = \{\vec{Y}_i, i \in \mathcal{S}\}$ is also multivariate. Given any particular configuration $\mathbf{z}$, every $\vec{Y}_i$ follows a multi-dimensional Gaussian distribution $p(\vec{y}_i|\vec{z}_i) = g(\vec{y}_i; \theta_i)$, where θ_i is the parameter set.

- **Conditional Independence**

  For any $\mathbf{z}$, the random variables $\vec{Y}_i$ are conditional independent:

  $$P(\mathbf{y}|z) = \prod_{i \in \mathcal{S}} P(\vec{y}_i|\vec{z}_i).$$

According to the local characteristics of MRFs, the joint probability of any pair of $(\vec{Z}_i, \vec{Y}_i)$, given $\vec{Z}_i$'s neighbourhood configuration $\vec{Z}_{\mathcal{N}_i}$, is:

$$P(\vec{y}_i, \vec{z}_i|\vec{z}_{\mathcal{N}_i}) = P(\vec{y}_i|\vec{z}_i) P(\vec{z}_i|\vec{z}_{\mathcal{N}_i}).$$

Thus, we can compute the marginal probability distribution of $\vec{Y}_i$, depending on the parameter set θ and $\vec{Z}_{\mathcal{N}_i}$,

$$\begin{aligned} p(\vec{y}_i|\vec{z}_{\mathcal{N}_i},\theta_i) &= \int_{\vec{z}} p(\vec{y}_i,\vec{z}_i|\vec{z}_{\mathcal{N}_i},\theta_i)d\vec{z} \\ &= \int_{\vec{z}} g(\vec{y}_i;\theta_i)p(\vec{z}_i|\vec{z}_{\mathcal{N}_i})d\vec{z}. \end{aligned} \tag{11}$$

This is the GHMRF distribution for the PVE.

3 An HMRF-EM Framework for Partial Volume Effect Classification

The partial volume classification problem we deal with involves assigning to each voxel i a mixel $\vec{z}_i$, which is a continuous vector containing a proportion for each tissue type, so that $\vec{z}_i = [z_{i1}, \cdots, z_{ik}]^T$, with k being the number of classes, and $\sum_\ell \vec{z}_{i\ell} = 1$ for all $i \in \mathcal{S}$. The image is indexed by a two- or three-dimensional lattice $\mathcal{S}$ and characterized by grey-level intensities $\vec{y}_i = [y_{i1}, \cdots, y_{im}]^T$, with m being the number of spectral channels. A particular PVE classification of the image is denoted by $\mathbf{z}$.

It is assumed that each of the k classes of the image follows an m-dimensional Gaussian distribution with mean $\boldsymbol{\mu}_\ell$ and covariance matrix $\boldsymbol{\psi}_\ell$. Thus the mixture of different tissue types in one voxel, given the mixing proportion of each tissue type, makes the intensity distribution of that voxel an m-dimensional Gaussian as well. It is also desirable that the underlying mixels vary smoothly across the image. Combining all this information, we can naturally apply the HMRF model we derived in the previous section to the intensity distribution of the image, with respect to the underlying PVE classification.

3.1 MRF-MAP Objective Function

We seek a PVE classification $\hat{\mathbf{z}}$ of an image, which is an estimate of the true classification $\mathbf{z}^*$, according to the MAP criterion,

$$\hat{\mathbf{z}} = \arg\max_{\mathbf{z}}\{P(\mathbf{y}|z)P(\mathbf{z})\} = \arg\min_{\mathbf{z}}\{U(\mathbf{y}|z) + U(\mathbf{z})\},$$

where the prior energy function $U(\mathbf{z})$ is defined by Equation (9) and the likelihood energy is given by

$$U(\mathbf{y}|z) = \sum_{i\in\mathcal{S}} \left[\frac{1}{2}(\vec{y}_i - \Sigma_i)^T \Psi_i^{-1}(\vec{y}_i - \Sigma_i) + \log(\sqrt{|\Psi_i|})\right], \tag{12}$$

where Σ_i and Ψ_i are defined by Equation (5) and (6) respectively.

Thus the final objective function for estimating an optimal PVE classification is given by

$$\begin{aligned} \hat{\mathbf{z}} &= \arg\min_{\mathbf{z}} \sum_{i\in\mathcal{S}} \Big[\frac{1}{2}(\vec{y}_i - \Sigma_i)^T \Psi_i^{-1}(\vec{y}_i - \Sigma_i) \\ &\quad + \log(\sqrt{|\Psi_i|}) + \beta \sum_{j\in\mathcal{N}_i} ||\vec{z}_i - \vec{z}_j||^2\Big]. \end{aligned} \tag{13}$$

3.2 Minimization through Deterministic Sampling

Given that the state space of mixels $\vec{z}_i$ is an infinite vector space, the minimization of Equation (13) is much more difficult than the minimization of the energy function for the problem of discrete segmentation that we have discussed previously, which only involves a finite scalar state space. Therefore, finding a global minimum for a PVE classification would be more infeasible than that for a segmentation. For this reason, we apply the similar 'greedy' strategy as used by the *Iterated Conditional Modes* (ICM) algorithm [Besag, 1986]; that is, to minimize the locally dependent energy function voxel-by-voxel:

$$\begin{aligned} \vec{z}_i &= \arg\min_{\vec{z}_i}\{U(\vec{y}_i|\vec{z}_i) + U(\vec{z}_i|\vec{z}_{\mathcal{N}_i})\} \\ &= \arg\min_{\vec{z}_i} \Big[\frac{1}{2}(\vec{y}_i - \Sigma_i)^T \Psi_i^{-1}(\vec{y}_i - \Sigma_i) \\ &\quad + \log(\sqrt{|\Psi_i|}) + \beta \sum_{j\in\mathcal{N}_i} ||\vec{z}_i - \vec{z}_j||^2\Big]. \end{aligned} \tag{14}$$

Such a local minimization problem can be easily solved for discrete estimation problems, such as segmentation, where the number of possible solutions for each position is very limited so that the optimal one can be found by trying all of them and then choosing the one with the minimal energy. However, for PVE classification, where the number of possible solutions at each position is infinite, such an enumeration scheme is infeasible. Given the configuration of Equation (14), it is also very difficult or time-consuming to obtain a solution by any gradient descent method. Moreover, in order to have a unique solution for this equation, we must have a sufficient number of image channels ($>$ number of classes - 2), which is not always available in practice.

To overcome all of these difficulties we have adopted a sampling method similar to the Metropolis algorithm. Unlike Metropolis sampling, which is a stochastic method because it may accept changes toward a higher energy state during the iterations, the sampling scheme we use is a deterministic one in the sense that only changes leading to a lower energy state can be accepted. More specifically, at each voxel we generate a random PVE vector that satisfies Equation (4), and compare the energy state calculated using Equation (14) with the previous energy of this voxel. Only if the new energy is lower than the previous one will we adopt the generated PVE vector. The advantage of such a scheme is that it will not lead to a worse solution in terms of the objective function, given any initial conditions, even with an insufficient number of image channels.

3.3 Model Fitting Using the EM Algorithm

To further improve the accuracy of the classification scheme, we simultaneously estimate the involved model parameters during the iteration using the Expectation-Maximization (EM) algorithm, which is commonly used to solve the 'incomplete-data' problem [Dempster et al., 1977]. Under certain reasonable conditions, EM estimates converge locally to the ML estimates [Wu, 1983]. For the HMRF PVE model specified above, the underlying parameters are actually the mean vector and the covariance matrix of each primary tissue type ℓ, that is $\theta_\ell = \{\boldsymbol{\mu}_\ell, \boldsymbol{\psi}_\ell\}$, since the parameters of each voxel are only combinations of them. To apply the EM algorithm, we first estimate the missing information $\hat{\mathbf{z}}$ given the current θ estimate and then use it to form the complete data set $\{\mathbf{y}, \hat{\mathbf{z}}\}$; the new θ can then be estimated by maximizing the expectation of the complete-data log likelihood. Formally, the EM algorithm can be described as:

Start An initial estimate $\theta^{(0)}$.

The E-step Calculate the conditional expectation

$$Q(\theta|\theta^{(t)}) = \mathcal{E}\Big[\log P(\mathbf{y}, z|\theta)|\mathbf{y}, \theta^{(t)}\Big]. \tag{15}$$

The M-step maximize $Q(\theta|\theta^{(t)})$ to obtain the next estimate

$$\theta^{(t+1)} = \arg\max_{\theta} Q(\theta|\theta^{(t)}). \tag{16}$$

Let $\theta^{(t+1)} \rightarrow \theta^{(t)}$ and repeat from the E-step.

More specifically, when applying the EM algorithm on the PVE HMRF model we can derive the Q-function as the following:

$$\begin{aligned}
Q &= \mathcal{E}\Big\{\log\big[P(\mathbf{y}|z,\theta)P(\mathbf{z}|\theta)\big]|\mathbf{y};\theta^{(t)}\Big\} \\
&= \mathcal{E}\Big\{\big[\log P(\mathbf{y}|z,\theta) + \log P(\mathbf{z}|\theta)\big]|\mathbf{y};\theta^{(t)}\Big\} \\
&= \mathcal{E}\Big\{\big[\sum_{i\in\mathcal{S}} \log P(\vec{y}_i|\vec{z}_i,\theta) + \log P(\mathbf{z})\big]|\mathbf{y};\theta^{(t)}\Big\} \\
&= \mathcal{E}\Big\{\sum_{i\in\mathcal{S}} \log P(\vec{y}_i|\vec{z}_i,\theta)|\mathbf{y};\theta^{(t)}\Big\} + C \\
&= \sum_{i\in\mathcal{S}} \int_{\vec{z}_i} \log P(\vec{y}_i|\vec{z}_i,\theta) P(\vec{z}_i|\mathbf{y};\theta^{(t)}) + C
\end{aligned} \tag{17}$$

As Equation (17) involves an integration over all possible $\vec{z}_i$, it is difficult either to calculate or maximize with respect to θ. For this reason, we have adopted a simplified approach in which we use the partial volume proportion of tissue type ℓ at each voxel i to

approximate the posterior probability of this voxel discretely being classified as that tissue $x_i = \ell$ given its neighbouring sites $x_{\mathcal{N}_i}$ and its intensity $\vec{y}_i$; that is, we have $\forall \ell \in \mathcal{L}$

$$P(x_i = \ell | \vec{y}_i, x_{\mathcal{N}_i}) \simeq z_{i\ell}. \tag{18}$$

With this approximation, we then apply the EM algorithm for the discrete image segmentation, derived in [Zhang et al., 2001], to estimate the parameters. That is, we use an approximated Q-function:

$$\hat{Q} = \sum_{i \in \mathcal{S}} \sum_{x_i \in \mathcal{L}} \log P(\vec{y}_i | x_i, \theta) P(x_i | \mathbf{y}; \theta^{(t)}) + C, \tag{19}$$

where x_i is the class label for pixel i and $\mathcal{L}$ is the set of all possible labels.

Mathematically,

$$P(x_i | \mathbf{y}; \theta^{(t)}) = \frac{P(\mathbf{y} | x_i; \theta^{(t)}) P(x_i)}{P(\mathbf{y})},$$

which is expensive to calculate given all the conditions and assumptions. Therefore, we make the following approximation to make the calculation feasible:

$$P(x_i | \mathbf{y}; \theta^{(t)}) \approx P(x_i | \vec{y}_i, x_{\mathcal{N}_i}; \theta^{(t)}), \tag{20}$$

which is just the posterior probability of x_i given its neighbouring sites $x_{\mathcal{N}_i}$ and $\vec{y}_i$. Such an approximation is justified because the influence of the intensities of the whole image on the label of the current pixel is embodied through the influence of its own intensity and its neighbours' labelling configuration. For simplicity, we write $P(x_i = \ell | \vec{y}_i, x_{\mathcal{N}_i}; \theta^{(t)})$ as $P^{(t)}(\ell | \vec{y}_i)$

Substituting $P(x_i | \mathbf{y}; \theta^{(t)})$ in Equation (19) with the above approximation and using the approximation we have

$$\begin{aligned} \hat{Q} &= \sum_{i \in \mathcal{S}} \sum_{\ell \in \mathcal{L}} P^{(t)}(\ell | \vec{y}_i) \cdot \log P(\vec{y}_i | \ell, \theta) + C \\ &= \sum_{i \in \mathcal{S}} \sum_{\ell \in \mathcal{L}} -P^{(t)}(\ell | y_i) \frac{1}{2} \Big[\log |\boldsymbol{\psi}_\ell| + \\ &\quad (\vec{y}_i - \boldsymbol{\mu}_\ell)^T \boldsymbol{\psi}_\ell^{-1} (\vec{y}_i - \boldsymbol{\mu}_\ell) + C' \Big] + C, \end{aligned} \tag{21}$$

where C' is another expression independent of θ.

Applying the zero-gradient condition on the $\hat{Q}$-function, which is effectively the M-step of the EM algorithm, we can obtain the following updating equations:

$$\begin{aligned} \boldsymbol{\mu}_\ell^{(t+1)} &= \frac{\sum_{i \in \mathcal{S}} P^{(t)}(\ell | \vec{y}_i) \vec{y}_i}{\sum_{i \in \mathcal{S}} P^{(t)}(\ell | \vec{y}_i)} \\ \boldsymbol{\psi}_\ell^{(t+1)} &= \frac{\sum_{i \in \mathcal{S}} P^{(t)}(\ell | \vec{y}_i) (\vec{y}_i - \boldsymbol{\mu}_\ell^{(t+1)}) (\vec{y}_i - \boldsymbol{\mu}_\ell^{(t+1)})^T}{\sum_{i \in \mathcal{S}} P^{(t)}(\ell | \vec{y}_i)}, \end{aligned}$$

where $P^{(t)}(\ell | y_i)$ is given by the approximation defined in Equation (18). Thus we have a complete EM algorithm for our PVE HMRF model.

3.4 Initial Classification

Since the minimization method we employ is a deterministic sampling scheme, we must have a fairly good initial classification in order to make the sampling iterations converge to a satisfactory configuration. Given that our HMRF-EM-based hard segmentation algorithm for brain MR images generally gives both a good image segmentation and well estimated parameters for each tissue type, it is natural to start the PVE classification with a hard segmentation and then use the likelihood probability of each voxel belonging to a specific tissue type as the initial PVE constituent of this tissue type within this voxel. Another advantage of this initialization is that the bias field can be removed before the PVE classification is carried out.

With this HMRF model-based PVE classification framework, we are able to achieve a robust and reliable classification given either single- or multi-channel MR data. Examples are shown in Figures 1 and 2.

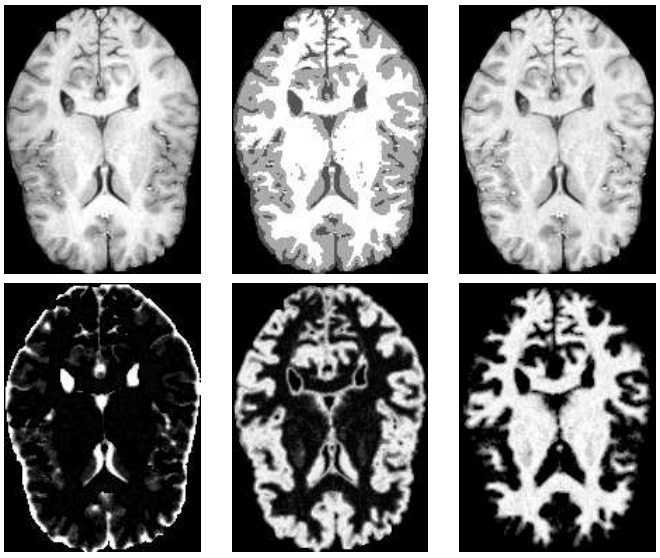

Figure 1: One-channel three-Class PVE classification. Top row: the original image, the hard segmentation, the bias-field removed image. Second row: the PVE classification of CSF, GM and WM.

4 Evaluation

For any segmentation algorithm that has the potential to be used in real clinical applications, there are two relatively important aspects: **accuracy** and **repeatability**. To evaluate the accuracy of a segmentation means to examine whether it is close to the ground truth. Given this definition, we also implicitly mean that there are two components in an accuracy evaluation, which are the measurements — the way to examine — and the ground truth, against which the classification is to be examined. We are also concerned about repeatability, which examines whether the algorithm can consistently work on similar data or the same data under slightly different conditions. For example, given images taken from the same subject, repeatability means that classifications produced by the algorithm are close to identical.

To have a more complete evaluation, we also incorporate results obtained using the segmentation tool in the popular SPM99 software package. The underlying method behind SPM99 is a simple Gaussian mixture fitting algorithm which uses *a-priori* tissue probability maps [Ashburner and Friston, 1997].

4.1 Accuracy Study on Simulated Brain MR Images

The first experiment was carried out using simulated brain MR images for which the ground truth partial volume maps are known. Four simulated brain images were obtained from the BrainWeb simulated brain database at the McConnell Brain Imaging Center of the Montreal Neurological Institute, McGill University [Cocosco et al., 1997]. All are T1-weighted images and are simulated with 1 mm cubic voxels and a size of $146 \times 184 \times 181$. The amount of noise added to the four images is 3%, 5%, 7% and 9%, and they all have the highest level of bias field provided by BrainWeb, 40%.

Figure 3 shows the underlying true hard segmentation and three partial volume maps, upon which all the simulated brains were built. Figure 4 shows one typical slice from 3 of the images used in this experiment, along with their histograms.

We ran both our HMRF-EM algorithm and the SPM99 segmentation tool on these simulated images to produce partial volume classifications. We then use the measure of mean-square errors (MSE), to test their accuracies against the ground truth:

$$MSE = \sqrt{\frac{\sum_{i \in \mathcal{S}} ||\mathbf{z}_i - \mathbf{t}_i||^2}{N}}, \tag{22}$$

where $\mathbf{z}_i$ is the estimated PVE vector of voxel i, $\mathbf{t}_i$ is the ground truth and N is the number of voxels. Note, only brain tissue voxels are used in this equation; background and other non-brain tissue voxels are removed in advance and are not counted.

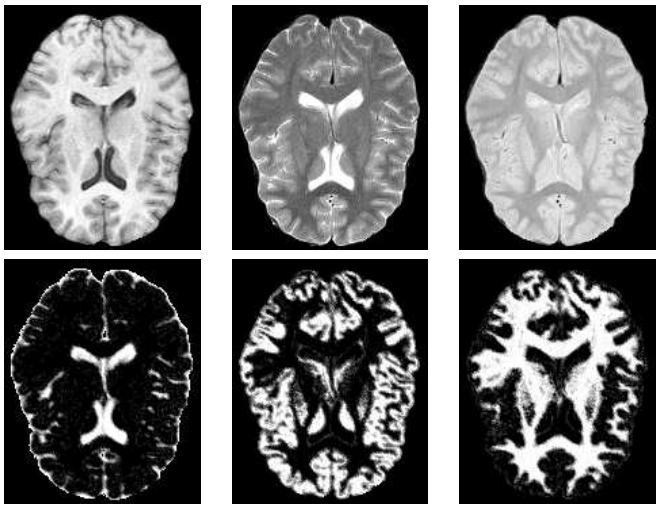

Figure 2: Three-channel three-Class PVE classification. Top row: the T1, T2 and PD channels of the original image. Second row: the PVE classification of CSF, GM and WM.

The results are shown in Figure 5. As can be seen, the advantage of the HMRF-EM method over SPM99 is significant.

We also measure the percentage of brain tissue volume over the whole brain (brain tissue + CSF) and compare the results from both algorithms against the ground truth, which is 80.98%. Figure 6 shows the their percentage errors against the ground truth. Note, with 7% noise, the estimate from the HMRF-EM method is 80.974%; hence one can barely see the error bar in the figure.

4.2 Repeatability Study Based On Brain Volume Estimation

In this experiment we test the repeatability of our algorithm. The basic idea is to process scans from the same subject but acquired at different times using the same imaging conditions, and then compare the estimated brain volumes. Because the imaging conditions may be slightly different and the subject will have slightly different positions in the scanner, we expect the segmentations to be different, unless registration is carried out. However, since the brain volume of a given person is a fairly fixed quantity, at least during a short period, the ground truth is that different scans should have the same brain volume. Given both the observations, we can conclude that a reliable segmentation algorithm should achieve very similar segmentations for images from the same subject; the less the difference is, the better the algorithm is.

The data we have acquired are from 15 subjects with two scans taken over one or two days for each subject. All images are T1-weighted MR scans of 2mm slice-thickness. Figure 7 shows the two scans from one of the subjects with their segmentations. Because all scans were tuned to have low contrast between GM and WM, we cannot expect a high quality partial volume classification of them but the classification between brain tissue (GM+WM) and non-brain tissue is important. Therefore, after a three-tissue classification of all the scans, we combine the classification of GM and WM to generate the classification for brain tissues and then calculate the volumes of both brain and non-brain tissues.

More specifically, given the partial volume vector $\vec{z}_i = \{z_{i0}, z_{i1}\}$ (assuming 1 for brain tissue and 0 for non-brain tissue) for each voxel, the volumes of brain and non-brain tissue are given by:

$$V_{brain} = \sum_{i \in \mathcal{S}} z_{i1} * v \qquad V_{non-brain} = \sum_{i \in \mathcal{S}} z_{i0} * v,$$

with v being the volume of one voxel in mm. Thus, the difference of the brain volume (error) estimates of the two scans from each subject is given by:

$$error = \frac{|V^1_{brain} - V^2_{brain}|}{avg(V^1_{brain}, V^2_{brain})}$$

Given the this error measure, we ran both the HMRF-EM algorithm and SPM99 for all images from the 15 subjects; the result is shown in Figure 8. The average error of HMRF-EM is 0.598%, compared to 0.863% of SPM99.

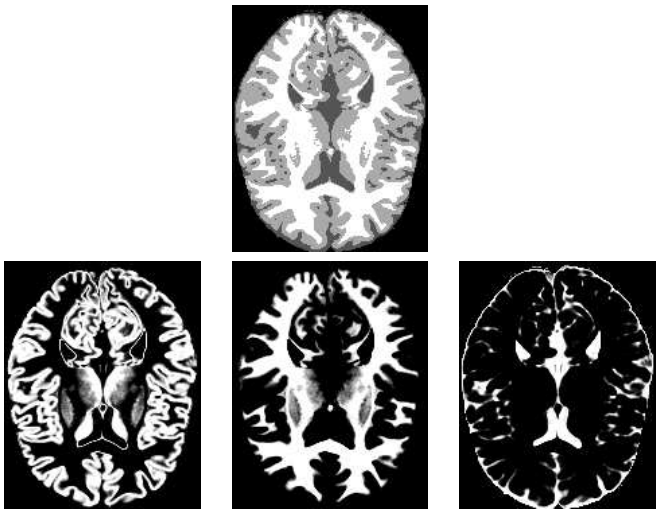

Figure 3: Ground truth of the image segmentation and PVE classifications. Top: the segmentation. Bottom: PVE classification of GM, WM and CSF.

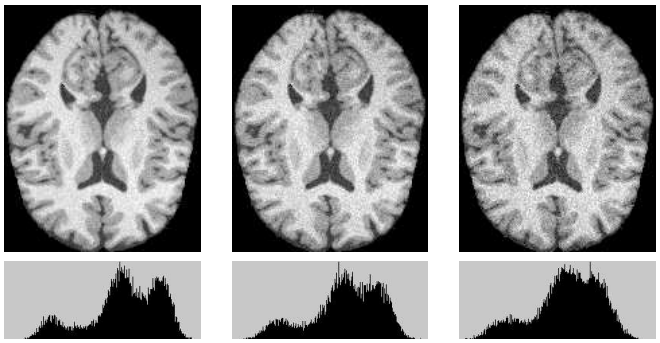

Figure 4: Simulated T1-weighted single-spectral brain MR images. Top row: from left to right, image with noise level 5%, 7% and 9%. Bottom row: histograms of the top row.

5 Discussions and Conclusions

We have presented in this paper a novel approach to the classification of human brain MR images. The method is an extension to our previously proposed HMRF-EM algorithm for discretely segmenting piecewise constant images. We have proposed a hidden Markov random field model in which the hidden random field is composed of continuous vectors with each element corresponding to the partial volume proportion of each spectral channel, and the observed random field corresponding to the actual image. These two random fields interact with each other through the HMRF distribution. An Expectation-Maximization algorithm is also derived to enable the fitting of the HMRF model to real data. Finally the PVE classification is achieved through the ICM method and a deterministic annealing approach. The advantage of this HMRF-EM algorithm is that it can work both with multi-spectral data where the unique solution exists and single-spectral data where only an 'optimal' solution can be found. Quantitative evaluation of our method with comparison to the popular SPM99 software tool has also been conducted, in which clear advantage of our method in terms of both accuracy and repeatability can be observed.

The method described in this paper has been implemented as FAST (FMRIB's Automated Segmentation Tool - a pair of standalone C++ programs, one for single channel and one for multi-channel data). FAST is freely available as part of FSL (FMRIB's Software Library - www.fmrib.ox.ac.uk/fsl).

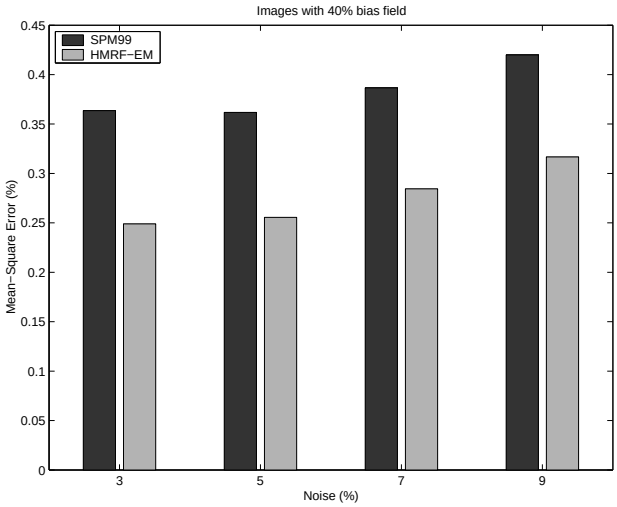

Figure 5: HMRF-EM vs. SPM99: Mean-square error of PVE classification on simulated T1 images.

Acknowledgments

The authors would like to thank Paul Matthews, Mark Jenkinson, Christian Beckmann, Mark Woolrich, Peter Bannister and Tim Behrens for helpful discussions. They acknowledge support from the UK MRC and EPSRC.

References

[Ashburner and Friston, 1997] Ashburner, J. and Friston, K. J. (1997). Multimodal image coregistration and partitioning - a unified framework. *NeuroImage*, 6:344–352.

[Besag, 1986] Besag, J. (1986). On the statistical analysis of dirty pictures (with discussion). *J. of Royal Statist. Soc., series B*, 48(3):259–302.

[Choi et al., 1991] Choi, H. S., Haynor, D. R., and Kim, Y. (1991). Partial volume tissue classification of multichannel magnetic resonance images–a mixel model. *IEEE Trans. Medical Imaging*, 10(3):395–406.

[Cocosco et al., 1997] Cocosco, C. A., Kollokian, V., Kwan, R., and Evans, A. C. (1997). Brainweb: Onlie interface to a 3D MRI simulated brain database. *Neuroimage*, 5(4):425. Available at http://www.bic.mni.mcgill.ca/brainweb.

[Dempster et al., 1977] Dempster, A. P., Laird, N. M., and Bubin, D. B. (1977). Maximum likelihood from incomplete data via EM algorithm. *J. of Royal Statist. Soc., series B*, 39(1):1–38.

[Gage et al., 1992] Gage, H. D., Snyder, W. E., and Santago, P. (1992). Quantification of brain tissue through incorporation of partial volume effects. In *Proc. Medical Imaging VI. Image Capture and Display, SPIE*.

[Jaggi et al., 1998] Jaggi, C., Ruan, S., and Bloyet, D. (1998). Mixture modeling applied to the partial volume effect in MRI data. *Proc. the 20th Annual Inter. Conf. IERR Eng. in Medical and Biology Society*, 20(2):693–695.

[Pham and Prince, 1998] Pham, D. L. and Prince, J. L. (1998). Adaptive fuzzy segmentation of magnetic resonance images. *IEEE Trans. Medical Imaging*, pages 819–822.

[Ruan et al., 2000] Ruan, S., C., J., J., X., J., F., and D., B. (2000). Brain tissue classification of magnetic resonance images using partial volume modeling. *IEEE Trans. Medical Imaging*, 19(12):1179–1187.

[Santago and Gage, 1993] Santago, P. and Gage, H. D. (1993). Quantification of MR brain images by mixture density and partial volume modeling. *IEEE Trans. Medical Imaging*, 12(3):566–574.

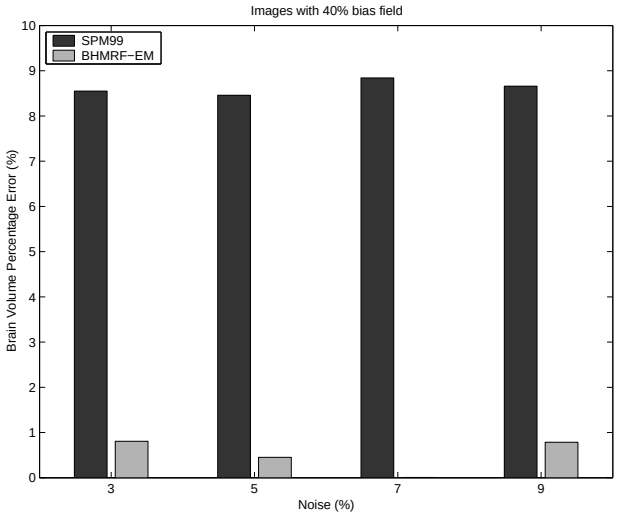

Figure 6: HMRF-EM vs. SPM99 on single-spectral brain volume estimation. images with 40% bias field but increasing amounts of noise.

[Santago and Gage, 1995] Santago, P. and Gage, H. D. (1995). Statistical models of partial volume effect. *IEEE Trans. Medical Imaging*, 11(4):1531–1540.

[Wu, 1983] Wu, C. F. J. (1983). On the convergence properties of the EM algorithm. *The Annals of Statistics*, 11:95–103.

[Zhang et al., 2001] Zhang, Y., Brady, M., and Smith, S. (2001). Segmentation of brain MR images through a hidden Markov random field model and the expectation maximization algorithm. *IEEE Trans. Medical Imaging*, 20(1):45–57.

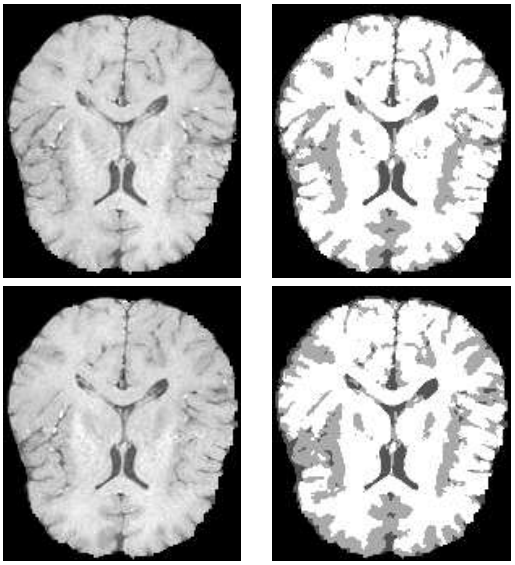

Figure 7: Two scans from the same subject with their segmentations

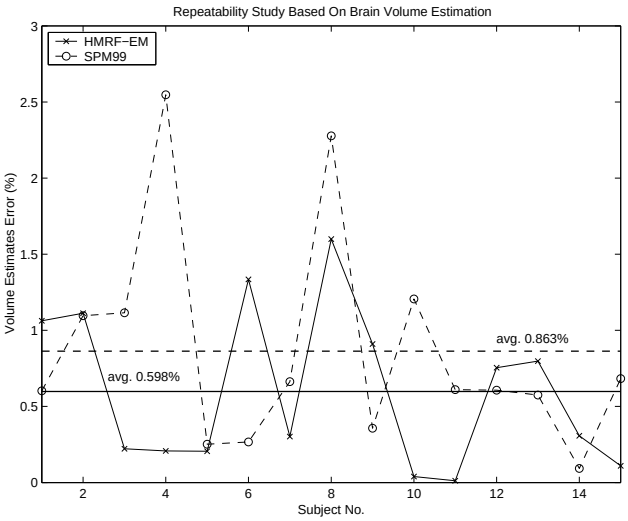

Figure 8: HMRF-EM vs. SPM99 in brain volume estimates of 15 subjects. The average error of 0.598% of HMRF-EM and 0.863% of SPM99 are shown as the solid and dash lines respectively.