Hidden Markov Random Field Model and Segmentation of Brain MR Images

FMRIB Technical Report TR00YZ1
(A related paper has been submitted to IEEE Trans. Medical Imaging)

Yongyue Zhang, Stephen Smith, and Michael Brady

Oxford Centre for Functional Magnetic Resonance Imaging of the Brain (FMRIB),
Department of Clinical Neurology, University of Oxford, John Radcliffe Hospital,
Headley Way, Headington, Oxford, UK

Abstract

The finite mixture (FM) model is the most commonly used model for statistical segmentation of brain MR images because of its simple mathematical form and the piecewise constant nature of ideal brain MR images. However, being a histogram-based model, the FM has an intrinsic limitation – no spatial information is taken into account. This causes the FM model to work only on well-defined images with low noise level; unfortunately, this is often not the the case due to artifacts such as partial volume effect and bias field distortion. Under such conditions, FM model-based methods produce unreliable results. In this paper, we propose a novel hidden Markov random field (HMRF) model, which is a stochastic process generated by a Markov random field whose state sequence cannot be observed directly but which can be observed through observations. Mathematically, it can be shown that the FM model is a degenerate version of the HMRF model. The advantage of the HMRF model derives from the way in which the spatial information is encoded through the mutual influences of neighbouring sites. Although MRF modelling has been employed in MR image segmentation by other researchers, most reported methods are limited to using MRF as a general prior in an FM model-based approach. Moreover, they either lack a proper parameter estimation step to fit the FM model [14, 16] or the parameter estimation procedure they use, such as ML or EM [33, 27, 20], suffers greatly from the limitation of the FM model mentioned above. To fit the HMRF model, an expectation-maximization (EM) algorithm is used. We show that by incorporating both the HMRF model and the EM algorithm into a HMRF-EM framework, an accurate and robust segmentation can be achieved. Moreover, the HMRF-EM framework can easily be combined with other techniques. As an example, we show how the bias field correction algorithm of Guillemaud and Brady [13] can be incorporated into this framework. A three-dimensional fully automated approach for brain MR image segmentation is achieved and significant improvement is obtained compared to the Guillemaud-Brady algorithm.

Keywords: MRI, Segmentation, Hidden Markov Random Field, HMRF, Expectation-Maximization, Bias Field Correction.

1 Introduction

MRI is an advanced medical imaging technique providing rich information about the human soft tissue anatomy. It has several advantages over other imaging techniques enabling it to provide 3-dimensional data with high contrast between soft tissues. However, the amount of data is far too much for manual analysis/interpretation, and this has been one of the biggest obstacles in the effective use of MRI. For this reason, automatic or semi-automatic techniques of computer-aided image analysis are necessary. Segmentation of MR images into different tissue classes, especially gray matter (GM), white matter (WM) and cerebrospinal fluid (CSF), is an important task.

Brain MR images have a number of features, especially the following: Firstly, they are statistically simple: MR Images are theoretically piecewise constant with a small number of classes. Secondly, they have relatively high contrast between different tissues. Unlike many other medical imaging modalities, the contrast in an MR image depends strongly upon the way the image is acquired. By adding RF or gradient pulses, and by carefully choosing

relaxation timings, it is possible to highlight different components in the object being imaged and produce high-contrast images. These two features facilitate segmentation. On the other hand, ideal imaging conditions never realised in practice. The piecewise-constant property is degraded considerably by electronic noise, the bias field (intensity inhomogeneities in the RF field) and the partial-volume effect (multiple tissue class occupation within a voxel), all of which cause classes to overlap in the image intensity histogram. Moreover, MR images are not always high-contrast. Many T_2-weighted and proton density images have low contrast between gray matter and white matter. Therefore, it is important to take advantage of useful data while at the same time overcoming potential difficulties.

A wide variety of approaches have been proposed for brain MR image segmentation. These can be roughly divided into two categories: *structural* and *statistical*. Structural methods are based on the spatial properties of the image, such as edges and regions. Various edge detection algorithms have been applied to extract boundaries between different brain tissues [5, 10, 8]. However such algorithms are vulnerable to artifacts and noise. Region growing [6, 34] is another popular structural approach. In this approach, one begins by dividing an image into small regions, which can be considered as "seeds". Then, all boundaries between adjacent regions are examined. Strong boundaries (in terms of certain specific properties) are kept, while weak boundaries are rejected and the adjacent regions merged. The process is carried out iteratively until no boundaries are weak enough to be rejected. This method is employed in [9] to extract brain surfaces. The segmentation tool in the commercial software package ANALYZE [28] is also based on this idea. However, as concluded by Clarke *et al.* in [8] critically, the performance of the method depends on seed selection and whether the regions are well defined, and therefore is also not considered robust.

Starting from a totally different viewpoint, statistical methods label pixels according to probability values, which are determined based on the intensity distribution of the image. In their simplest form, thresholding-based methods are always chosen for scenes containing solid objects resting on a background with intensities well separated from the objects. However, as noted above, this generally is not effective for brain MR images. Therefore, thresholding-based methods are unlikely to produce reliable results [23, 7, 18].

Most statistical approaches rely on certain assumptions or models of the probability distribution function of the image intensities and its associated class labels, which can both be considered random variables. Let X and Y be two random variables for the class label and the pixel intensity, respectively, and x and y be typical instances. The class-conditional density function is $p(y|x)$. Statistical approaches attempt to solve the problem of estimating the associated class label x, given only the intensity y for each pixel. Such an estimation problem is necessarily formulated from an established criterion. Maximum *a posteriori* (MAP) or maximum likelihood (ML) principles are two such examples. But before those criteria can be assessed, the formula for the density function $p(y)$ has to be chosen carefully [4]. Many statistical segmentation methods differ in terms of models of $p(y)$. Depending on whether a specific functional form for the density model is assumed, a statistical approach can either be *parametric* or *non-parametric*. Both have been widely used in segmentation of brain MR images.

In non-parametric methods, the density model $p(y)$ relies entirely on the data itself, i.e. no prior assumption is made about the functional form of the distribution but a large number of correctly labelled training points are required in advance. One of the most widely used non-parametric methods is K-Nearest-Neighbours (K-NN). One starts by choosing a fixed K, which is the number of nearest neighbours to find in the neighbourhood of any unlabelled pixel in the y space. Then a certain distance measure between pairs of points is applied to determine their relationship [3, 15, 24]. The Parzen-window method is another example of non-parametric methods, in which the intensity density function is modelled using a Parzen-window distribution. Such a distribution can be obtained by centering a small Gaussian around each training point [17, 14].

Non-parametric methods are adaptive, but suffer from the difficulty of obtaining a large number of training points, which can be tedious and a heavy burden even for experienced people. Clearly, such methods are not fully automatic.

Unlike non-parametric approaches, parametric approaches rely on an explicit functional form of the intensity density function. For brain MR images, the only method developed to date is based on the *finite mixture* (FM) model, in particular the *finite Gaussian mixture* (FGM) model when the Gaussian likelihood distribution is assumed [31, 13, 14]. FM models have a number of elegant features and are mathematically simple. However, being a histogram-based model, the FM has an intrinsic limitation – spatial information is not taken into account because all the data points are considered to be independent samples drawn from a population. Such a limitation causes the FM model to work only on well-defined images with low level of noise; unfortunately, this is often not the case due to artifacts such as the partial volume effect and bias field distortion. Under such conditions, FM model-based

methods produce unreliable results.

In order to address this problem, we develop a hidden Markov random field (HMRF) model, which is a stochastic process generated by a Markov random field whose state sequence cannot be observed directly but which can be observed through a field of observations. The importance of the HMRF model derives from Markov random field (MRF) theory, in which the spatial information of an image is encoded through contextual constraints of neighbouring pixels. By imposing such constraints, we expect neighbouring pixels to have the same class labels (in the case of piecewise constant images) or similar intensities (in the case of piecewise continuous images). This is achieved through characterizing mutual influences among pixels using conditional MRF distributions.

To apply the HMRF model, an expectation-maximization (EM) algorithm is also derived. We show that by incorporating both the HMRF model and the EM algorithm into a mathematically sound HMRF-EM framework, an accurate and robust segmentation approach can be achieved, which is demonstrated through experiments on both simulated images and real data, and comparison made with the FM-EM framework. Being a flexible approach, the HMRF-EM can be easily combined with other techniques to improve the segmentation performance. As an example, we show how the bias field correction algorithm of Guillemaud and Brady [13] is incorporated into it.

Although MRF modelling and its application in image segmentation have been investigated by many other researchers [12, 2, 22], only in recently years has MRF theory become popular in MR image segmentation. But most reported methods use MRF only as a general prior in an FM model-based parametric approach to build the MAP estimation. They either lack a proper parameter estimation step to fit the FM model [14, 16] or the parameter estimation procedure they use, such as ML or EM [33, 27, 20], suffers from the limitation of the FM model mentioned above. In general, although an MRF prior can improve the performance, the FM assumption is still a big limitation.

In this paper, segmentation is treated as a statistical model-based problem with three steps: (i)model selection; (ii)model fitting; (iii)classification. The HMRF model is presented in Section 2; it has all the advantages of the FM model while at the same time being more robust because of the MRF neighbourhood relationships. The HMRF-EM framework is presented in Section 5, which contains both the model fitting and classification steps. It enables adaptive and reliable automatic segmentation. The framework is easily extensible by combining other techniques, such as bias field correction, as shown in Section 7.

2 Hidden Markov Random Field Model

Let $\mathcal{L}$ and $\mathcal{D}$ be two alphabets:

$$\mathcal{L} = \{1, 2, \cdots, l\}, \quad \mathcal{D} = \{1, 2, \cdots, d\}.$$

Let $\mathcal{S} = \{1, 2, \cdots, N\}$ be the set of indices and $R = \{r_i, i \in \mathcal{S}\}$ denote any family of random variables indexed by $\mathcal{S}$, in which each random variable R_i takes a value z_i in its state space. Such a family R is called a random field. The joint event $(R_i = r_i, \cdots, R_N = r_N)$ is simplified to $R = r$ where $r = \{r_1, \cdots, r_N\}$ is a *configuration* of R, corresponding to a realization of this random field. Let X and Y be two such random fields whose state spaces are $\mathcal{L}$ and $\mathcal{D}$ respectively so that for $\forall i \in \mathcal{S}$ we have $X_i \in \mathcal{L}$ and $Y_i \in \mathcal{D}$. Let $\mathbf{x}$ denote a configuration of X and $\mathcal{X}$ be the set of all possible configurations so that

$$\mathcal{X} = \{\mathbf{x} = (x_1, \cdots, x_N) | x_i \in \mathcal{L}, i \in \mathcal{S}\}.$$

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

$$\mathcal{Y} = \{\mathbf{y} = (y_1, \cdots, y_N) | y_i \in \mathcal{D}, i \in \mathcal{S}\}.$$

Given $X_i = \ell$, Y_i follows a conditional probability distribution

$$p(y_i|\ell) = f(y_i; \theta_\ell), \quad \forall \ell \in \mathcal{L} \tag{1}$$

where θ_ℓ is the set of parameters. For all ℓ, the function family $f(\cdot; \theta_\ell)$ has the same known analytic form. We also assume that (X, Y) is pairwise independent, meaning

$$P(\mathbf{y}, x) = \prod_{i \in \mathcal{S}} P(y_i, x_i) \tag{2}$$

In order to develop the HMRF model we firstly take the standard FM model as a comparison.

2.1 Finite Mixture Model

For every $\ell \in \mathcal{L}$ and $i \in S$,

$$P(X_i = \ell) = \omega_\ell$$

is independent of the individual sites $i \in \mathcal{S}$ and called a *mixing parameter*. We take ϕ as the model parameter set with

$$\phi = \{\omega_\ell; \theta_\ell | \ell \in \mathcal{L}\}.$$

Consider two configurations $\mathbf{y} \in \mathcal{Y}$ and $\mathbf{x} \in \mathcal{X}$. From (1) and (2), we can compute the joint probability distribution of $\mathbf{x}$ and $\mathbf{y}$ dependent on the model parameters (ϕ is treated as a set of random variables), namely

$$p(\mathbf{x}, y|\phi) = \prod_{i \in \mathcal{S}} p(y_i, x_i|\phi) = \prod_{i \in \mathcal{S}} \{\omega_{x_i} \cdot f(y_i; \theta_{x_i})\}. \tag{3}$$

We can compute the marginal distribution of $Y_i = y$, dependent on the parameter set ϕ:

$$\begin{aligned} p(y|\phi) &= \sum_{\ell \in \mathcal{L}} p(\ell, y|\phi) \\ &= \sum_{\ell \in \mathcal{L}} \omega_\ell \cdot f(y; \theta_\ell) \end{aligned} \tag{4}$$

This is the so-called *finite mixture* (FM) model.

The FM model is the most frequently employed statistical model, due to its simple mathematical form [21, 29]. But it is clear that the FM model only describes the data statistically. No spatial information about the data is utilized. In other words, the FM model is spatially independent and can therefore be specified fully by the histogram of the data. However, images with the same intensity distribution may have totally different structural properties. For this reason, FM model is not complete.

To overcome this drawback, certain spatial properties, or constraints, have to be incorporated into the model. Here, spatial properties means quantifiable structural characteristics of an object, such as size and shape. Under certain intensity distributions, we want the model to be "adaptive" to structural information or spatially dependent in order to fit the actual image better. This leads to the consideration of MRF theory and our HMRF model.

2.2 Markov Random Field Theory

The spatial property can be modelled through different aspects, among which, the *contextual constraint* is a general and powerful one. Markov random field (MRF) theory provides a convenient and consistent way to model context-dependent entities such as image pixels and correlated features. This is achieved by characterizing mutual influences among such entities using conditional MRF distributions.

In an MRF, the sites in $\mathcal{S}$ are related to one another via a *neighbourhood system*, which is defined as $\mathcal{N} = \{\mathcal{N}_i, i \in \mathcal{S}\}$, where $\mathcal{N}_i$ is the set of sites neighbouring i, $i \notin \mathcal{N}_i$ and $i \in \mathcal{N}_j \Longleftrightarrow j \in \mathcal{N}_i$. A random field X said to be an MRF on $\mathcal{S}$ with respect to a neighbourhood system $\mathcal{N}$ if and only if

$$\begin{aligned} &P(\mathbf{x}) > 0, \forall \mathbf{x} \in \mathcal{X} \\ &P(x_i|x_{\mathcal{S}-\{i\}}) = P(x_i|x_{\mathcal{N}_i}) \end{aligned}$$

Note, the neighbourhood system can be multi-dimensional. According to the Hammersley-Clifford theorem [1], an MRF can equivalently be characterized by a Gibbs distribution. Thus,

$$P(\mathbf{x}) = Z^{-1} \exp(-U(\mathbf{x})), \tag{5}$$

where

$$Z = \sum_{\mathbf{x} \in \mathcal{X}} \exp(-U(\mathbf{x})) \tag{6}$$

is a normalizing constant called the *partition function*, and $U(\mathbf{x})$ is an *energy function* of the form

$$U(\mathbf{x}) = \sum_{c \in \mathcal{C}} V_c(\mathbf{x}), \tag{7}$$

which is a sum of *clique potentials* $V_c(\mathbf{x})$ over all possible cliques $\mathcal{C}$. A *clique* c is defined as a subset of sites in $\mathcal{S}$ in which every pair of distinct sites are neighbours, except for single-site cliques. The value of $V_c(\mathbf{x})$ depends on the local configuration on clique c. For more detail on MRF and Gibbs distribution see [12].

2.3 Hidden Markov Random Field Model

The concept of a *hidden Markov random field* model is derived from *hidden Markov models* (HMM), which are defined as stochastic processes generated by a Markov chain whose state sequence cannot be observed directly, only through a sequence of observations. Each observation is assumed to be a stochastic function of the state sequence. The underlying Markov chain changes its state according to a $l \times l$ transition probability matrix, where l is the number of states. HMMs have been applied successfully to speech recognition [26, 30] and handwritten script recognition [19].

Since original HMMs were designed as 1D Markov chains with first order neighbourhood systems, it can not directly be used in 2D/3D problems such as image segmentation. Here, we consider a special case of a HMM, in which the underlying stochastic process is a Markov random field (MRF), instead of a Markov chain, therefore not restricted to 1D. We refer to this special case as a *hidden Markov random field* (HMRF) model. Mathematically, an HMRF model is characterized by the following:

- **Hidden Random Field (MRF)**

  The Random field $X = \{X_i, i \in \mathcal{S}\}$ is an underlying MRF assuming values in a finite state space $\mathcal{L}$ with probability distribution (5). The state of X is unobservable.

- **Observable Random Field**

  $Y = \{Y_i, i \in \mathcal{S}\}$ is a random field with a finite state space $\mathcal{D}$. Given any particular configuration $\mathbf{x} \in \mathcal{X}$, every Y_i follows known conditional probability distribution $p(y_i|x_i)$ of the same functional form $f(y_i; \theta_{x_i})$, where θ_{x_i} is the involved parameters. This distribution is called the *emission probability function* and Y is also referred to as the *emitted random field*.

- **Conditional Independence**

  For any $\mathbf{x} \in \mathcal{X}$, the random variables Y_i are conditional independent,

$$P(\mathbf{y}|x) = \prod_{i \in \mathcal{S}} P(y_i|x_i). \tag{8}$$

Based on the above, we can write the joint probability of (X, Y) as

$$\begin{aligned} P(\mathbf{y}, x) &= P(\mathbf{y}|x)P(\mathbf{x}) \\ &= P(\mathbf{x}) \prod_{i \in \mathcal{S}} P(y_i|x_i). \end{aligned}$$

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

$$P(y_i, x_i|x_{\mathcal{N}_i}) = P(y_i|x_i)P(x_i|x_{\mathcal{N}_i}) \tag{9}$$

Thus, we can compute the marginal probability distribution of Y_i dependent on the parameter set θ (in this case, we treat θ as a random variable) and $X_{\mathcal{N}_i}$,

$$\begin{aligned} p(y_i|x_{\mathcal{N}_i}, \theta) &= \sum_{\ell \in \mathcal{L}} p(y_i, \ell|x_{\mathcal{N}_i}, \theta) \\ &= \sum_{\ell \in \mathcal{L}} f(y_i; \theta_\ell) p(\ell|x_{\mathcal{N}_i}), \end{aligned} \tag{10}$$

where $\theta = \{\theta_\ell, \ell \in \mathcal{L}\}$. We call this the *hidden Markov random field* (HMRF) model. Note, the concept of an HMRF is different from that of an MRF in the sense that the former is defined with respect to a pair of random variable families (X, Y) while the latter is only defined with respect to X.

More precisely, an HMRF model can be described by the following:

- $X = \{X_i, i \in \mathcal{S}\}$ – hidden MRF, with prior distribution $p(\mathbf{x})$;
- $Y = \{Y_i, i \in \mathcal{S}\}$ – observable random field, with emission probability distribution $p(y_i|x_i)$ for each y_i;
- $\theta = \{\theta_\ell, \ell \in \mathcal{L}\}$ – the set of parameters involved in the above distributions.

If we assume the random variables X_i are independent of each other, which means that for $\forall \ell \in \mathcal{L}$ and $i \in \mathcal{S}$, we have

$$P(\ell|x_{\mathcal{N}_i}) = P(\ell) = \omega_\ell,$$

then equation(10) reduces to

$$p(y|\theta) = \sum_{\ell \in \mathcal{L}} \omega_\ell \cdot f(y; \theta_\ell),$$

which is the definition of the finite mixture model. Therefore a FM model is a degenerate special case of an HMRF model.

It is obvious from the above that the fundamental difference between the FM model and the HMRF model lies in their different spatial properties. The FM model is spatially independent whereas the HMRF model may be spatially dependent. Therefore, the HMRF model is more flexible for image modelling in the sense that it has the ability to encode both the statistical and spatial properties of an image.

With a Gaussian emission distribution, the FM model is usually known as the *finite Gaussian Mixture* (FGM) or *finite normal mixture* (FNM) model. More specifically, the observable random variables have the following density function:

$$p(y|\phi) = \sum_{\ell \in \mathcal{L}} \omega_\ell \cdot g(y; \theta_\ell) \tag{11}$$

where

$$g(y; \theta_\ell) = \frac{1}{\sqrt{2\pi\sigma_\ell^2}} \exp\Big(-\frac{(y-\mu_\ell)^2}{2\sigma_\ell^2}\Big) \quad \text{and} \quad \theta_\ell = (\mu_\ell, \sigma_\ell)^T. \tag{12}$$

Similarly, an HMRF model with a Gaussian emission distribution can be specified as:

$$p(y_i|x_{\mathcal{N}_i}, \theta) = \sum_{\ell \in \mathcal{L}} g(y_i; \theta_\ell) p(\ell|X_{\mathcal{N}_i}), \tag{13}$$

where g and θ_ℓ are defined as in (12). We refer to this type of HMRF as the *Gaussian hidden Markov random field* (GHMRF) model.

3 Model Simulation and Image Synthesis

Simulation is often used to verify statistical models. In this case, simulation is used to generate synthetic images by drawing random samples from the model distribution using stochastic sampling methods. Here, the Gibbs sampler proposed by Geman and Geman [12] is employed. Many different experiments have been carried out to compare the FGM model and the GHMRF model. Figure 1 shows two examples, in which the number of intensity levels was set equal to the number of classes and the Gaussian emission distributions have the same standard deviation for all classes. For the GHMRF model, a homogeneous and isotropic MRF model is employed to generate the prior distribution with clique potential $V_c(\mathbf{x}) = -\delta(x_i - x_j)$. The two rows in Figure 1 correspond, respectively, to simulations with 3 and 5 classes. The first column is the sample drawn from the FGM model while the other three columns show samples drawn from the GHMRF model with different standard deviations. Apparently, the FGM model generates meaningless noise images whereas the GHMRF model generates images with controllable spatial structures – the smaller the standard deviation, the clearer the spatial structures.

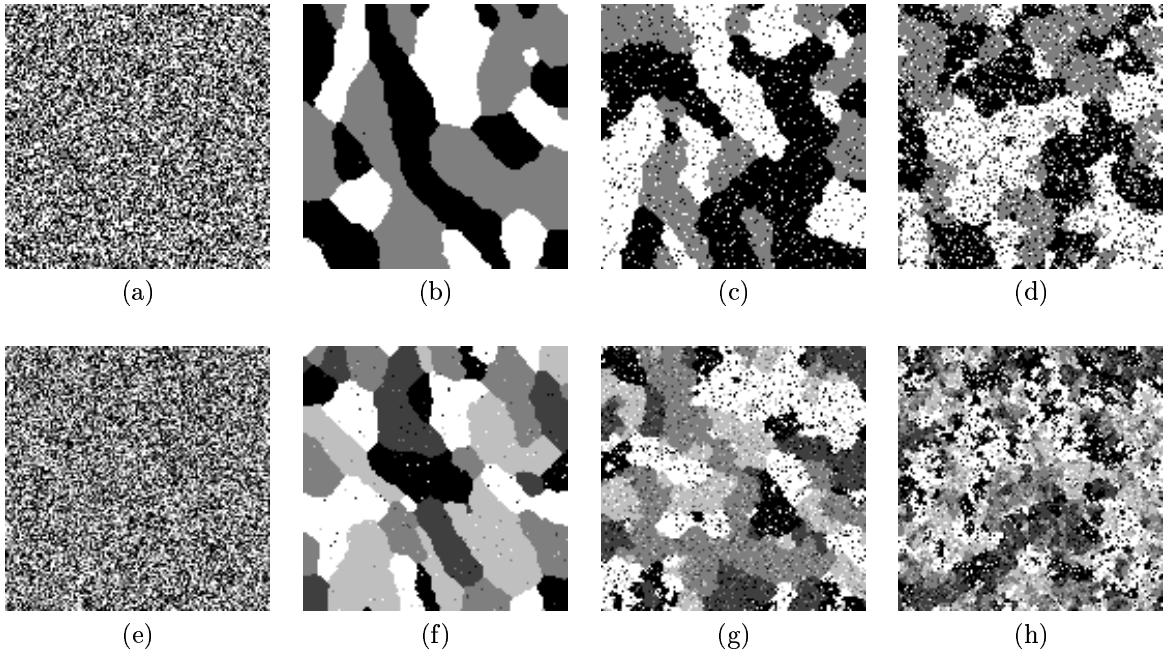

Figure 1: Image simulation by the FGM model and the GHMRF model. The first row shows the 3-class case. (a) FGM model; (b)-(d) GHMRF mode with standard deviation 0.23, 0.4, 0.5, respectively. The second row shows the 5-class case. (e) FGM model; (f)-(h) GHMRF model with standard deviation 0.3, 0.47, 0.55, respectively.

4 MRF-MAP Classification

The image classification problem we consider involves assigning to each pixel a class label taking a value from the set $\mathcal{L}$. The pixels are indexed by a two-dimensional rectangular lattice $\mathcal{S}$ and each pixel is characterized by an intensity value y_i from the set $\mathcal{D}$. A labelling of $\mathcal{S}$ will be denoted by $\mathbf{x}$, where $x_i, i \in \mathcal{S}$ is the corresponding class label of pixel i. We write $\mathbf{x}^*$ for the true but unknown labelling configuration and $\hat{\mathbf{x}}$ for an estimate of $\mathbf{x}^*$, both of which are interpreted as particular realizations of a random field X, which is an MRF with a specified distribution $P(\mathbf{x})$. The observable image itself is denoted by $\mathbf{y}$, which is a realization of a GHMRF as described in Section 2. The problem of classification is to recover $\mathbf{x}^*$, given the observed image $\mathbf{y}$.

4.1 MRF-MAP Estimation

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

$$\hat{\mathbf{x}} = \arg\max_{\mathbf{x}\in\mathcal{X}} \{P(\mathbf{y}|x)P(\mathbf{x})\}. \tag{14}$$

From (14), we need to compute the prior probability of the class and the likelihood probability of the observation. Since $\mathbf{x}$ is considered as a realization of an MRF, its prior probability can be derived from

$$P(\mathbf{x}) = \frac{1}{Z}\exp(-U(\mathbf{x})). \tag{15}$$

It is also assumed that the pixel intensity y_i follows a Gaussian distribution with parameters $\theta_i = \{\mu_\ell, \sigma_\ell\}$, given the class label $x_i = \ell$,

$$p(y_i|x_i) = g(y_i;\theta_\ell) = \frac{1}{\sqrt{2\pi\sigma_\ell^2}}\exp\Big(-\frac{(y_i-\mu_\ell)^2}{2\sigma_\ell^2}\Big). \tag{16}$$

Based on the conditional independence assumption of $\mathbf{y}$ (8), we have the joint likelihood probability

$$\begin{aligned} P(\mathbf{y}|x) &= \prod_{i\in\mathcal{S}} p(y_i|x_i) \\ &= \prod_{i\in\mathcal{S}} \Big[\frac{1}{\sqrt{2\pi}} \exp\Big(-\frac{(y_i-\mu_{x_i})^2}{2\sigma_{x_i}^2} - \log(\sigma_{x_i})\Big)\Big] \\ &= \frac{1}{(2\pi)^{\frac{N}{2}}} \exp\Big[\sum_{i\in\mathcal{S}}\Big(-\frac{(y_i-\mu_{x_i})^2}{2\sigma_{x_i}^2} - \log(\sigma_{x_i})\Big)\Big] \end{aligned}$$

which can be written as

$$P(\mathbf{y}|x) = \frac{1}{Z'}\exp(-U(\mathbf{y}|x)), \tag{17}$$

with the *likelihood energy*

$$U(\mathbf{y}|x) = \sum_{i\in\mathcal{S}} U(y_i|x_i) = \sum_{i\in\mathcal{S}}\Big[\frac{(y_i-\mu_{x_i})^2}{2\sigma_{x_i}^2} + \log(\sigma_{x_i})\Big], \tag{18}$$

and the constant normalization term $Z' = (2\pi)^{\frac{N}{2}}$. It is easy to show that

$$\log P(\mathbf{x}|y) \propto -U(\mathbf{x}|\mathbf{y}), \tag{19}$$

where

$$U(\mathbf{x}|y) = U(\mathbf{y}|x) + U(\mathbf{x}) + const \tag{20}$$

is the *posterior energy*. The MAP estimation is equivalent to minimizing the posterior energy function

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}\in\mathcal{X}}\{U(\mathbf{y}|x) + U(\mathbf{x})\}. \tag{21}$$

Although mathematically simple, this type of MAP estimation clearly presents a computationally infeasible problem. Therefore, optimal solutions are usually computed using some iterative optimization (minimization) techniques. In this paper, we adopt a *Iterated conditional modes* (ICM) algorithm proposed by Besag [2], which uses the "greedy" strategy in the iterative local minimization and convergence is guaranteed after only a few iterations. Given the data $\mathbf{y}$ and the other labels $x_{\mathcal{S}-\{i\}}^{(k)}$, the algorithm sequentially updates each $x_i^{(k)}$ into $x_i^{(k+1)}$ by minimizing $U(x_i|\mathbf{y}, x_{\mathcal{S}-\{i\}})$, the conditional posterior probability, with respect to x_i.

5 Model Fitting Using the EM Algorithm

A statistical model is complete only if both its functional form and its parameters are determined. The procedure for estimating the unknown parameters is known as *model fitting*. For an HMRF model, the parameter set $\theta = \{\theta_\ell, \ell \in \mathcal{L}\}$ is what should be estimated. If the Gaussian emission function is assumed for the observable random variable y, the mean and standard deviation of each Gaussian class are the parameters, so that $\theta_\ell = (\mu_\ell, \sigma_\ell)$.

Since both the class label and the parameters are unknown and they are strongly inter-dependent, the data set is said to be "incomplete" and the problem of parameter estimation is regarded as an "incomplete-data" problem. Many techniques have been proposed to solve this problem, among which the *expectation-maximization* (EM) algorithm [11] is the one most widely used.

The strategy underlying the EM algorithm consists of the following: estimate the missing part as $\hat{\mathbf{x}}$ given the current θ estimate and then use it to form the complete data set $\{\hat{\mathbf{x}}, \mathbf{y}\}$; new θ can be estimated by maximizing the expectation of the complete-data log likelihood, $\mathcal{E}\Big[\log P(\mathbf{x}, \mathbf{y}|\theta)\Big]$. Mathematically, the EM algorithm can be described by:

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

The E-step Calculate the conditional expectation

$$\begin{aligned} Q(\theta|\theta^{(t)}) &= \mathcal{E}\Big[\log P(\mathbf{x},\mathbf{y}|\theta)|\mathbf{y},\theta^{(t)}\Big] \\ &= \sum_{\mathbf{x}\in\mathcal{X}} p(\mathbf{x}|y,\theta^{(t)})\cdot\log p(\mathbf{x},y|\theta), \end{aligned} \tag{22}$$

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

$$\theta^{(t+1)} = \arg\max_{\theta} Q(\theta|\theta^{(t)}). \tag{23}$$

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

Under certain reasonable conditions, EM estimates converge locally to the ML estimates [32].

For the GHMRF field model, the intensity distribution function, dependent on the parameter set θ, is

$$p(y_i|\theta) = \sum_{\ell\in\mathcal{L}} \frac{1}{\sqrt{2\pi\sigma_\ell^2}} \exp\Big(-\frac{(y-\mu_\ell)^2}{2\sigma_\ell^2}\Big)\cdot p(\ell|x_{\mathcal{N}_i}), \tag{24}$$

where $p(\ell|x_{\mathcal{N}_i})$ is the locally dependent probability of $x_i = \ell$ and the parameter set $\theta = \{\mu_\ell, \sigma_\ell|\ell\in\mathcal{L}\}$.

The Q-function is then formulated as

$$Q = \sum_{i\in\mathcal{S}}\sum_{\ell\in\mathcal{L}} \left\{ P^{(t)}(\ell|y_i)\left[\log p^{(t)}(\ell|\mathbf{x}_{\mathcal{N}_i}) - \log\sigma_\ell - \frac{(y_i-\mu_\ell)^2}{2\sigma_\ell^2}\right] + C\right\}, \tag{25}$$

where $C = -0.5\log(2\pi)$.

Applying the EM algorithm, we obtain

$$\mu_\ell^{(t+1)} = \frac{\sum_{i\in\mathcal{S}} P^{(t)}(\ell|y_i)y_i}{\sum_{i\in\mathcal{S}} P^{(t)}(\ell|y_i)} \tag{26}$$

$$\big(\sigma_\ell^{(t+1)}\big)^2 = \frac{\sum_{i\in\mathcal{S}} P^{(t)}(\ell|y_i)(y_i-\mu_\ell)^2}{\sum_{i\in\mathcal{S}} P^{(t)}(\ell|y_i)}, \tag{27}$$

which are the same update equations for the FGM model [4], except that

$$P^{(t)}(\ell|y_i) = \frac{g^{(t)}(y_i;\theta_\ell)\cdot P^{(t)}(\ell|x_{\mathcal{N}_i})}{p(y_i)}. \tag{28}$$

The calculation of the conditional probability $P^{(t)}(\ell|x_{\mathcal{N}_i})$ involves estimation of the class labels, which are obtained through MRF-MAP estimation – Equation (21). We refer to this HMRF model-based EM algorithm as a HMRF-EM algorithm and the standard FM model-based EM algorithm as a FM-EM algorithm.

6 Segmentation Using the HMRF-EM Framework

The EM algorithm presented in Section 5 not only provides an effective method for parameter estimation, but also a complete framework for unsupervised classification using iterative updating.

6.1 Initial Parameter Estimation

Since both the EM model fitting algorithm and the ICM labelling algorithm converge locally, the choice of initial conditions, including the initial parameter set and the classification, is important.

Without prior information, histogram analysis is widely used to estimate statistics such as means and variances of a distribution. Histogram fitting is the most commonly adopted method, especially for the Gaussian mixture model [29]. But in the case when the image has high noise level, the Gaussian mixture assumption may not hold

and fitting may fail. From the standpoint of classification, we want the classes to be widely separated from each other while at the same time having relatively low intra-class variances. Whether the intensity PDF follows a Gaussian mixture or not is relatively unimportant. According to this, we adopt an initial estimation in this paper a discriminant measure-based thresholding method proposed by Otsu [25].

The basic idea is to find thresholds maximizing the inter-class variances thus also minimizing the intra-class variances. According to theories of discriminant analysis, such thresholds are optimal solutions. Once the optimal thresholds have been determined, the mean μ and the standard deviation σ for each class type can then be used as the initial parameters for further estimation. The initial classification can also be obtained either directly through the thresholding, or through an ML estimation with those known parameters.

6.2 Experiments

Various experiments have been carried out to test the performance of the HMRF-EM framework. An example is shown in the following figures. Figure 2(a) shows a simulated 3-class image sampled from an MRF model using the Gibbs sampler. The intensities for the three classes are 30, 125 and 220 respectively. Figure 2(b)-(e) show the same images with added Gaussian noise with standard deviation of 28, 47, 66, and 95. Because image contrast is what we are most interested in for examining qualities of an image, a measurement of the noise is more meaningful with image contrast being taken into account. Thus we define a measure, the noise-to-contrast ratio (NCR) as the following:

$$NCR = \frac{\text{standard deviation of the noise}}{\text{mean inter-class contrast}}.$$

Thus, the NCRs of the four test images are 0.3, 0.5, 0.7 and 1.0, respectively. Figure 2(f)-(k) show their intensity histograms. Except for the first, each histogram exhibits severe overlap. The true parameters for the test images are listed in Table 1.

Table 1: True model parameters of Figure 2(b)-(e).

class	Class 1			Class 2			Class 3		
parameter	μ_1	σ_1	ω_1	μ_2	σ_2	ω_2	μ_3	σ_3	ω_3
$NCR = 0.3$	30	28	0.372	125	28	0.299	220	28	0.329
$NCR = 0.5$	30	47	0.372	125	47	0.299	220	47	0.329
$NCR = 0.7$	30	66	0.372	125	66	0.299	220	66	0.329
$NCR = 1.0$	30	95	0.372	125	95	0.299	220	95	0.329

Table 2: Initial parameter estimation using discriminant measure-based thresholding.

class	Class 1			Class 2			Class 3		
parameter	μ_1	σ_1	ω_1	μ_2	σ_2	ω_2	μ_3	σ_3	ω_3
$NCR = 0.3$	31.5	23.6	0.389	124.8	25.7	0.27	219.4	24.2	0.341
$NCR = 0.5$	29.3	27.0	0.436	123.6	29.2	0.20	220.6	28.6	0.364
$NCR = 0.7$	25.8	26.6	0.478	123.8	29.8	0.145	225.2	28.7	0.387
$NCR = 1.0$	22.1	26.0	0.477	127.6	30.2	0.111	232.6	26.1	0.412

The discriminant measure-based thresholding method is then applied to each of the four test images to estimate the initial parameters. Table 2 shows the results. Comparing it with Table 1, we can see that the estimates are acceptable, especially when the noise level is low.

The standard FM-EM algorithm and the HMRF-EM algorithm are then applied to the four test images until there is no significant change in the value of the Q-function. To measure the segmentation accuracy, we also define

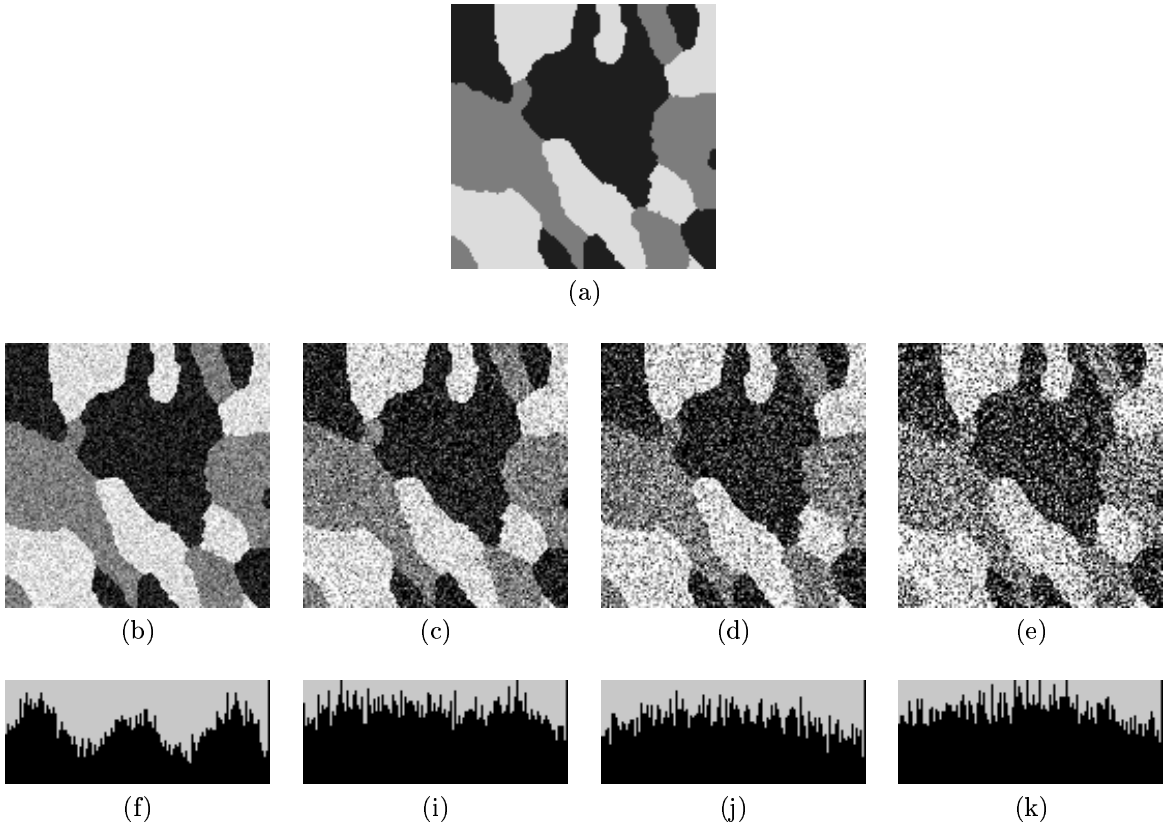

Figure 2: Test images for parameter estimation. (a) the original image; (b)-(e) noisy images with NCR 0.3, 0.5, 0.7, and 1.0; (f)-(k) histogram of (b)-(e).

the misclassification ratio (MCR), which is

$$MCR = \frac{\text{number of mis-classified pixels}}{\text{total number of pixels}}.$$

The standard FM-EM algorithm only converges for the first image, which has the lowest noise level (NCR=0.3). In this case, the estimation results and the number of iterations K are shown in Table 3. With those estimated parameters, we reconstruct the histogram and obtain the segmentation, as shown in Figure 3. Note that, the parameter estimation is not accurate when compared with their true values listed in Table 1.

Table 3: Parameter estimation using the FM-EM algorithm

class	Class 1			Class 2			Class 3			
parameter	μ_1	σ_1	ω_1	μ_2	σ_2	ω_2	μ_3	σ_3	ω_3	K
$NCR = 0.3$	25.6	19.8	0.309	122.6	41.5	0.409	224.5	21.2	0.282	48

The HMRF-EM algorithm rapidly converges for all the four test images. Table 4 and Figure 4 show the results. Taking the true parameters shown in Table 1 as the references and comparing the results from the two methods, it can be seen that: (1) the HMRF-EM algorithm gives more accurate estimates for parameters; (2) the HMRF-EM algorithm provides automatic segmentation with much lower MCR.

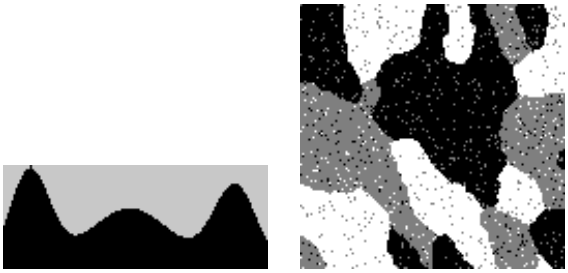

(a) (b)

Figure 3: Parameter estimation for Figure 2(b) using the standard FM-EM algorithm. (a) the reconstructed histogram; (b) the segmentation with MCR 5.82%.

Table 4: Parameter estimation using the HMRF-EM algorithm

class	Class 1			Class 2			Class 3			
parameter	μ_1	σ_1	ω_1	μ_2	σ_2	ω_2	μ_3	σ_3	ω_3	K
$NCR = 0.3$	32.0	24.6	0.378	124.9	27.8	0.300	219.2	24.8	0.322	7
$NCR = 0.5$	36.1	35.6	0.377	124.6	46.1	0.304	213.5	38.1	0.320	9
$NCR = 0.7$	42.2	46.3	0.372	123.0	62.7	0.307	207.9	49.3	0.321	10
$NCR = 1.0$	52.1	61.7	0.355	127.7	80.5	0.363	203.6	62.1	0.281	31

7 Segmentation of Brain MR Images With Bias Field Correction

Since it is a complete approach to segmenting piecewise-constant images, the HMRF-EM framework can be applied to brain MR images. However, MR images are often corrupted by a low-frequency spatially varying artifact known as the *bias field*. This is mainly caused by the variation of the interaction between the human body and the magnetic field. The bias field can also be due to the differential attenuation of signals and nonlinearities of the sensitivities of the receiver coils, etc. Although such an artifact has little impact on visual diagnosis, the performance of most automatic image analysis techniques, especially intensity-based segmentation, can degrade dramatically. Therefore, a robust, automatic, and inexpensive way of correcting for this artifact is required.

7.1 Bias Field Correction through Modified EM Algorithm

One of the most successful methods for dealing with the bias field problem was developed by Wells *et al.* [31], in which the bias field $B = (b_1, \ldots, b_N)$ is modelled as a multiplicative N-dimensional random vector with zero mean Gaussian prior probability density $p(B) = G_{\psi_B}(B)$, where ψ_B is the $N \times N$ covariance matrix. Let $I = (I_1, \cdots, I_N)$ and $I^* = (I_1^*, \cdots, I_N^*)$ be the observed and the ideal intensities of a given image respectively. The degradation effect of the bias field at pixel $i, 1 \leq i \leq N$ can be expressed as follows:

$$I_i = I_i^* \times b_i, \tag{29}$$

After logarithmic transformation of the intensities, the bias field effect can be treated as an additive artifact. Let $\mathbf{y}$ and $\mathbf{y}^*$ denote respectively the observed and the ideal log-transformed intensities: then $\mathbf{y} = \mathbf{y}^* + B$. Given the class labels $\mathbf{x}$, it is further assumed that the ideal intensity value at pixel i follows a Gaussian distribution with parameter $\theta(x_i) = (\mu_{x_i} \sigma_{x_i})$:

$$p(y_i^* | x_i) = g(y_i^*; \theta(x_i)). \tag{30}$$

With the bias field b_i taken into account, the above distribution can be written in terms of the observed intensity y_i as

$$p(y_i | x_i, B) = g(y_i - b_i; \theta(x_i)). \tag{31}$$

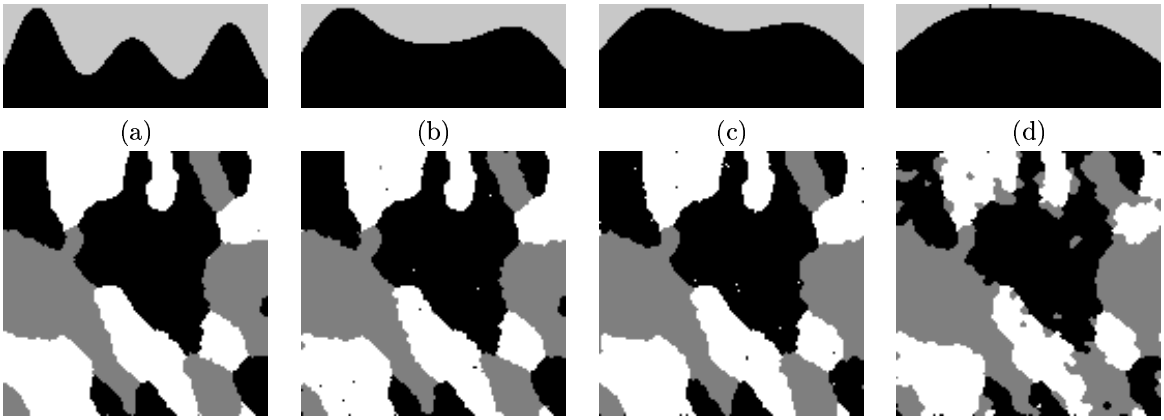

Figure 4: Parameter estimation for Figure 2(b)-(e) using the HMRF-EM algorithm. top row: the reconstructed histograms; bottom row: the segmentations.

Thus, the intensity distribution is modelled as a Gaussian mixture, given the bias field. It follows that

$$p(y_i|B) = \sum_{j\in\mathcal{L}} \{g(y_i - b_i; \theta(j))P(j)\}. \tag{32}$$

The MAP principle is then employed to obtain the optimal estimate of the bias field, given the observed intensity values:

$$\hat{B} = \arg\max_B p(\mathbf{y}|B)p(B), \tag{33}$$

A zero-gradient condition is then used to assess this maximum, which leads to (see [31] for detail):

$$W_{ij} = \frac{p(y_i|x_i,\beta)p(x_i)}{p(y_i|\beta)} \tag{34}$$

$$b_i = \frac{[FR]_i}{[F\psi^{-1}1]_i}, \text{ with } 1 = (1,1,\cdots,1)^T, \tag{35}$$

where R is the *mean residual* for pixel i

$$R_i = \sum_{j\in\mathcal{L}} \frac{W_{ij}(y_i - \mu_j)}{\sigma_j^2}, \tag{36}$$

ψ is the *mean inverse covariance*

$$\psi_{ik}^{-1} = \begin{cases} \sum_{j\in\mathcal{L}} W_{ij}\sigma_j^{-2}, & \text{if } i = k \\ 0 & \text{otherwise} \end{cases} \tag{37}$$

and F is a lowpass filter. W_{ij} is the posterior probability that pixel i belongs to class j given the bias field estimate.

The EM algorithm is applied to Equations (34) and (35). The E step assumes that the bias field is known and calculates the posterior tissue class probability W_{ij}. In the M step, the bias field B is estimated given the estimated W_{ij} in the E step. Once the bias field is obtained, the original intensity I^* is restored by dividing I by the inverse log of B. Initially, the bias field is assumed to be zero everywhere.

Wells *et al.*'s algorithm is found to be problematic when there are classes in an image that do not follow a Gaussian distribution. The variance of such a class tends to be very large and consequently the mean can not be considered representative[13]. Such situations are commonly seen in the regions of CSF, pathologies and other

non-brain classes. Bias field estimation can be significantly affected by this type of problem. To overcome this problem, Guillemaud and Brady [13] unify all such classes into an outlier class, which is called "other", with uniform distribution. Let $\mathcal{L}_G$ denote the set of labels for Gaussian classes and l_o the class label for the "other" class. The intensity distribution of the image is still a finite mixture except with an additional non-Gaussian class,

$$p(y_i|b_i) = \sum_{j \in \mathcal{L}_G} \{g(y_i - b_i; \theta(j))P(j)\} + \lambda P(l_o), \tag{38}$$

where λ is the density of the uniform distribution. Due to the large variance of the uniform distribution, the bias field is only estimated with respect to the Gaussian classes. The same iterative EM method can be applied, except for a slight modification to the formulation of mean residual R_i (36)

$$R_i = \sum_{j \in \mathcal{L}_G} \frac{W_{ij}(y_i - \mu_j)}{\sigma_j^2}. \tag{39}$$

With such a modification, the performance of the EM algorithm can be significantly improved in certain situations. This approach is referred to as the modified EM (MEM) algorithm.

7.2 HMRF-EM Framework for Brain MR Image Segmentation

As has been stated in the previous section, FM model-based segmentation methods do not utilize any spatial information and therefore are not robust and appropriate in certain cases. The MEM algorithm for brain MR image segmentation suffers from the same problem. But as an effective way to remove bias field, the MEM algorithm is worth improving by overcoming this drawback. We present in this section how our HMRF-EM framework can be easily extended to incorporate an additional bias field correction step. More specifically, we seek an EM solution for three dependent unknowns: the bias field, the image classification and the involved model parameters. In the E step, we calculate the MAP estimate of the bias field and the class labels to form the Q-function. In the M step, we calculate the ML estimate of the parameters using the estimated bias field and the class labels in the E step.

- **E step**:

$$\text{MAP estimation} \quad B^{(t)} = \arg\max_B p(B|\mathbf{y}, \mathbf{x}^{(t-1)}, \theta^{(t)}) \tag{40}$$

$$\text{MAP estimation} \quad \mathbf{x}^{(t)} = \arg\min_{\mathbf{x} \in \mathcal{X}} U(\mathbf{x}|y, B^{(t)}, \theta^{(t)}) \tag{41}$$

- **M step**:

$$\text{ML estimation} \quad \theta^{(t+1)} = \arg\max_\theta P(\mathbf{y}|\theta, \mathbf{x}^{(t)}, B^{(t)}) \tag{42}$$

The complete algorithm is described in Figure 5.

8 Experiments

Various experiments have been carried out both on real and simulated data, in both 2D and 3D. For the MEM algorithm, the parameters used in the experiments take their actual values for the simulated images and are manually estimated for real data, since the algorithm does not deal with parameter estimation itself. For the HMRF-EM algorithm, parameters are estimated automatically.

The first experiment shown here tests the noise sensitivity of the two algorithms. Two images consisting of two constant regions with the same simulated bias field but with different white noise were generated (Figure 6(a),(b)). Two Gaussian classes, corresponding to the two regions, are used. For Figure 6(a), both algorithms give perfect estimation results, as shown in Figures 6(c) and (d). However, for Figure 6(b), the HMRF-EM algorithm gives much better results than the MEM algorithm.

The second experiment tests the performance of the two algorithms on real 2D MR images but with a simulated bias field. Two Gaussian distributions are used for the two tissue classes (white matter and grey matter) and a uniform distribution (density = 0.3) is used for the rest. Figure 7(a) is the original T_2-weighted image and Figure

1. Perform image thresholding for initial parameter estimation and segmentation.
2. Estimate the bias field

$$b_i^{(t)} = \frac{[FR]_i}{[F\psi^{-1}1]_i}, \text{ with } 1 = (1, 1, \cdots, 1)^T$$

3. Calculate the likelihood distribution

$$p^{(t)}(y_i|x_i, B) = g^{(t)}(y_i - b_i; \theta(x_i))$$

4. Estimate the class labels by MRF-MAP estimation

$$\mathbf{x}^{(t)} = \arg\max_{\mathbf{x}\in\mathcal{X}}\{P(\mathbf{y}|x, \theta^{(t)}) + P(\mathbf{x})\}$$

5. Calculate the posterior distribution

$$P^{(t)}(\ell|y_i) = \frac{g^{(t)}(y_i; \theta_\ell) \cdot p^{(t)}(\ell|x_{\mathcal{N}_i})}{p(y_i)}$$

6. Update parameters by

$$\begin{aligned} \mu_\ell^{(t+1)} &= \frac{\sum_{i\in\mathcal{S}} P^{(t)}(\ell|y_i)y_i}{\sum_{i\in\mathcal{S}} P^{(t)}(\ell|y_i)} \\ \left(\sigma_\ell^{(t+1)}\right)^2 &= \frac{\sum_{i\in\mathcal{S}} P^{(t)}(\ell|y_i)(y_i - \mu_\ell)^2}{\sum_{i\in\mathcal{S}} P^{(t)}(\ell|y_i)} \end{aligned}$$

7. $t \leftarrow t + 1$ and repeat from 2 until enough iterations have been performed.

Figure 5: HMRF-EM algorithm for brain MR image segmentation and bias field correction

7(b) is the image with a simulated circular bias field. Figure 7(c) is the histogram of Figure 7(b), from which a substantial intensity overlap between WM and GM can be seen. Figure 7(d) shows the best result that can be obtained from (a) using global thresholding. The second row shows the result of applying the MEM algorithm. The last row shows the results from the HMRF-EM algorithm. Comparing the segmentations, we see that without losing any significant structure, the results from the HMRF-EM algorithm are much cleaner than from the MEM method, which still looks noisy. The restored image from HMRF-EM algorithm also shows good intensity uniformity, while the histogram output in the MEM method still shows WM/GM overlap.

The last example we show is for a real 3D data set from the Montreal Neurological Institute, McGill University, courtesy of D. Arnold. The original volume has 50 256×256 slices with voxel size $0.977 \times 0.977 \times 3.0$ mm. Figure 8 shows the results of five different slices from the HMRF-EM algorithm.

9 Discussions

We note from the above experiments that using the HMRF-EM algorithm, not only does the segmentations improve significantly but the bias field estimates do too, although the HMRF-EM algorithm itself is not directly related to bias field estimation. This is due to the fact the bias field estimation relies strongly on the classification.

A practical issue has to be addressed in the 3D implementation of the HMRF-EM algorithm. Theoretically the MRF neighbourhood system should be 3-dimensionally isotropic. However, the slice thickness of a 3D volume is often larger than the intra-slice voxel dimensions. In such a situation, an isotropic neighbourhood system may cause problems. Therefore an anisotropic 3D neighbourhood system is used with a smaller weight across slices.

Although the HMRF-EM framework itself is theoretically sound, the initial estimation based on thresholding is rather heuristic. Due to the high variability of brain MR images in terms of their intensity ranges and contrasts between brain tissues, it is not guaranteed that the thresholding procedure will produce perfect results. In most cases, however, the final segmentation results are stable even with slightly different initial estimates. This is largely attributable to the robust HMRF-EM algorithm. However, as a local minimization method, the EM algorithm can

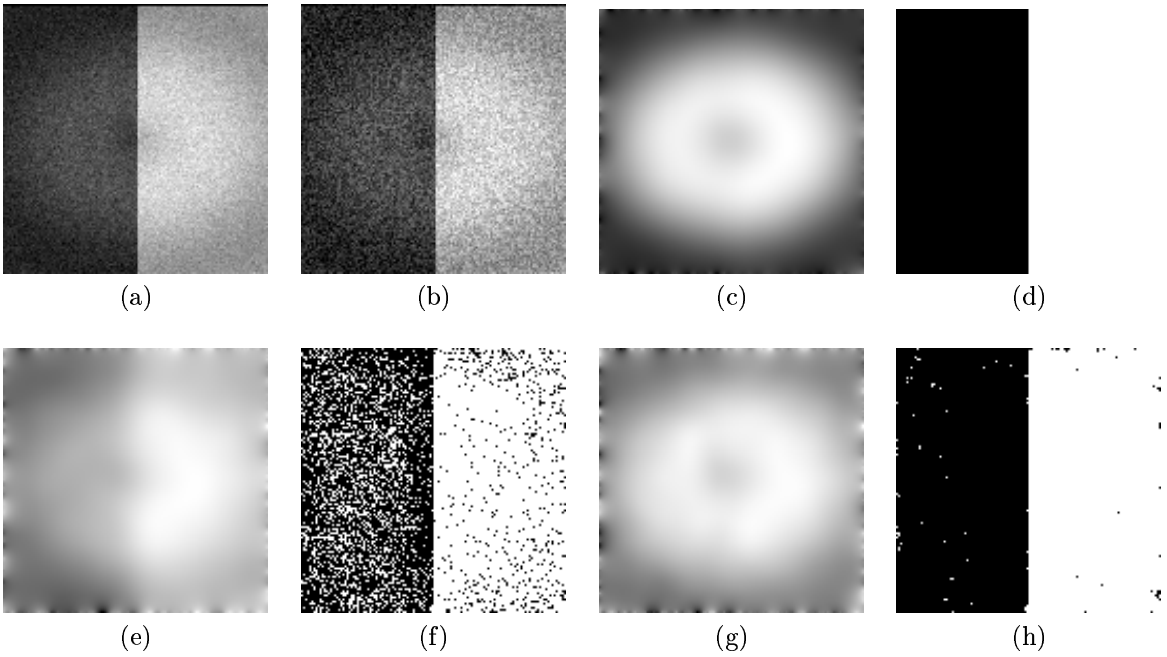

Figure 6: Comparison of the MEM and the HMRF-EM algorithm on simulated 2D images. (a) the original image with 3% noise. (b) the original image with 10% noise. (c) bias field estimation for (a) by both the algorithms. (d) segmentation for (a) by both the algorithms. (e) bias field estimation for (b) by the MEM algorithm. (f) segmentation for (b) by the MEM algorithm. (g) bias field estimation for (b) by HMRF-EM algorithm. (h) segmentation for (b) by the HMRF-EM algorithm.

be trapped in a local minimum. In some cases, where the image is poorly-defined, the thresholding procedure may fail to find the right thresholds for brain tissues, especially the threshold for GM and WM. With an initial condition far from normal, the EM procedure is likely to give a wrong final segmentation. In general, the initial estimation is a difficult problem and it will certainly be an important issue to be studied in future work.

With respect to the computational load, the whole algorithm is slightly slower than the original FM model-based MEM algorithm due to the additional MRF-MAP classification and the EM fitting procedure. However, by employing the fast deterministic ICM method and certain optimizations to the program, it runs reasonably quickly. Currently, it takes 16 seconds for a 256×256 2D image and about 10 minutes for a 3D volume data with 40 256×256 slices in an Intel PII 400MHZ-based system.

10 Conclusion

We have presented a fully automatic approach for the segmentation of brain MR images. The method is based on a HMRF-EM framework, which is a combination of the hidden Markov random field (HMRF) model and the associated MRF-MAP estimation and the EM fitting procedures. The HMRF model is proposed in this paper as a substitute for the widely used FM model, which is considered as sensitive to noise and therefore not robust. As a general method, the HMRF-EM algorithm could be applied to many other image segmentation problems.

We also show that the framework can easily be extended by incorporating other techniques in order to improve its performance on certain problems. As an example, we demonstrated how the bias field correction algorithm by Guillemaud and Brady [13] can be incorporated into this framework. As a result, a three-dimensional fully automatic approach for brain MR image segmentation is achieved and significant improvements have been observed in terms of both the bias field estimation and the tissue classification.

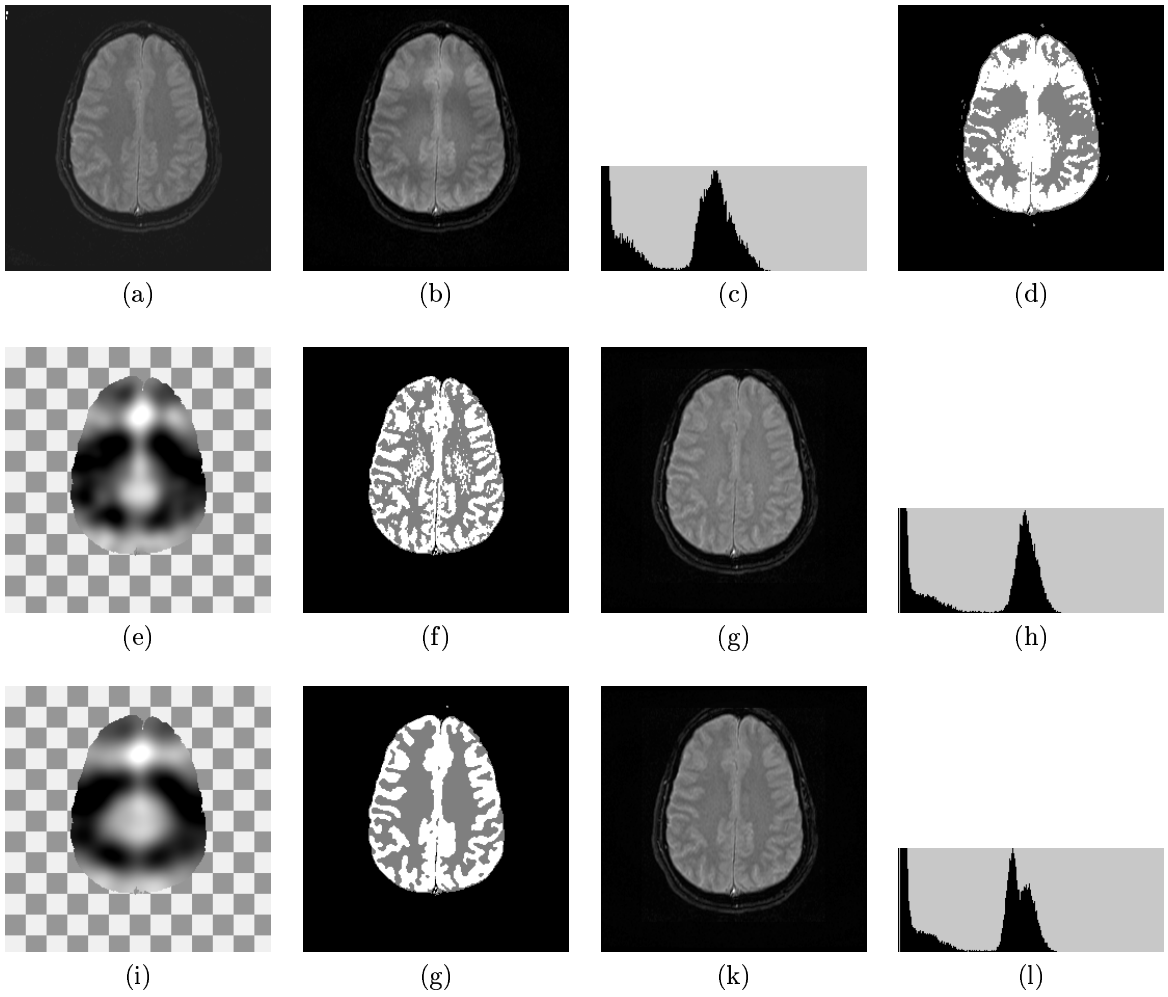

Figure 7: Comparison of the MEM and the HMRF-EM algorithm on real 2D MR images with simulated bias field. (a) the original image; (b) the image with simulated bias field; (c) histogram of (b); (d) best thresholding on (b); (e)-(h) the results from the MEM algorithm; (i)-(l) the results from HMRF-EM algorithm. For the last two rows, from left to right: the estimated bias field (the checkerboard is used to represent the background which is assumed to have no bias field), the segmentation, the restored image and the histogram of the restored image.

References

[1] J. Besag. Spatial interaction and the statistical analysis of lattice systems (with discussion). *J. of Royal Statist. Soc., series B*, 36(2):192–326, 1974.

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

[3] J.C. Bezdek, L.O. Hall, and L.P. Clarke. Review of MR image segmentation techniques using pattern recognition. *Medical Physics*, 20(4):1033–1048, 1993.

[4] C. M. Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, 1995.

[5] M. Bomans, K. H. Hohne, L. A. Kramer, and J. M. Fletcher. 3-D segmentation of MR images of the head for 3-D display. *IEEE Trans. Medical Imaging*, 18(1):25–34, 1990.

[6] C. R. Brice and C. L. Fennema. Scene analysis using regions. *Artificial Intelligence*, 1(3):205–226, 1970.

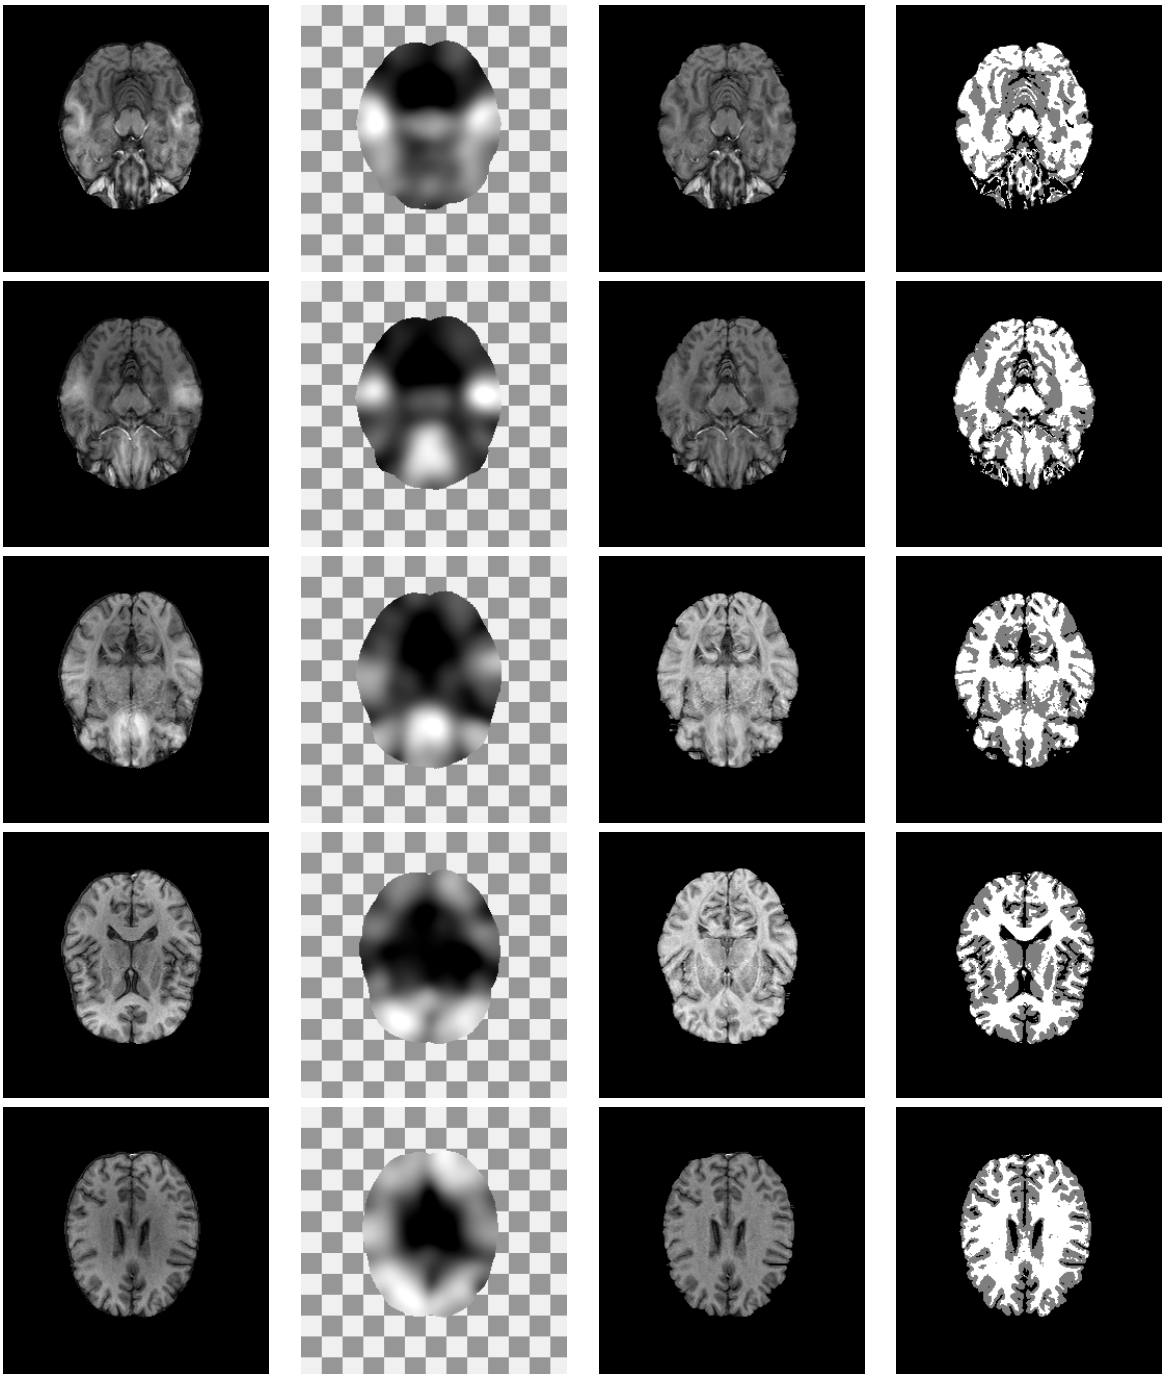

Figure 8: Five slices of a 3D MR volume image with real bias field. In each row, from left to right: the original slice, the estimated bias field, the restored slice, and the segmentation.

[7] M. E. Brummer. Optimized intensity thresholds for volumetric analysis of magnetic resonance imaging data. *Proceedings of the SPIE*, 1808:299–310, 1992.

[8] L.P. Clarke, R.P. Velthuizen, M.A. Camacho, J.J. Heine, M. Vaidyanathan, L.O. Hall, R.W. Thatcher, and M.L. Silbiger. MRI segmentation: methods and applications. *Magnetic Resonance Imaging*, 13(3):343–368, 1995.

[9] H. E. Cline, C. L. Dumoulin, H. R. Hart, W. E. Lorense, and S. Ludke. 3D reconstruction of the brain from magnetic resonance images using a connectivity algorithm. *Magnetic Resonance Imaging*, 5:345–352, 1987.

[10] S. Dellepiane. Image segmentation: Errors, sensitivity, and uncertainty. *Proc. 13th IEEE-Eng. Med. Bio. Society*, 13:253–254, 1991.

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

[12] S Geman and D Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Machine Intell.*, 6(6):721–741, 1984.

[13] R Guillemaud and J. M. Brady. Estimating the bias field of MR images. *IEEE Trans. Medical Imaging*, 16(3):238–251, 1997.

[14] K. Held, E. R. Kops, B. J. Krause, W. M. Wells, R. Kikinis, and H. Muller-Gartner. Markov random field segmentation of brain MR images. *IEEE Trans. Medical Imaging*, 16(6):878–886, 1997.

[15] E. F. Jackson, P. A. Narayana, J. S. Wolinsky, and T. J. Doyle. Accuracy and reproducibility in volumetric analysis of multiple sclerosis lesions. *J. Comput. Assist. Tomogr.*, 17:200–205, 1993.

[16] T. Kapur, E. L. Grimson, R. Kikinis, and W. M. Wells. Enhanced spatial priors for segmentation of magnetic resonance imagery. *Lecture Notes in Computer Science*, 1496:148–157, 1998.

[17] R. Kikinis, M. Shenton, F. A. Jolesz, G. Gerig, J. Martin, M. Anderson, and D. Metcalf. Quantitative analysis of brain and cerebrospinal fluid spaces with MR imaging. *J. Magn. Res. Imag.*, pages 619–629, 1992.

[18] A. Kundu. Local segmentation of biomedical images. *Comput. Med. Imaging Graph*, 14:173–183, 1990.

[19] A. Kundu, Y. He, and P. Bahl. Recognition of handwritten word: First and second order hidden Markov model based approach. *Pattern Recognition*, 22(3):283–297, 1989.

[20] K. V. Leemput, F. Maes, D. Vandermeulen, A. Colchester, and P. Suetens. Automated segmentation of MS lesions from multi-channel MR images. *Lecture Notes in Computer Science*, 1679:11–21, 1999.

[21] T. Lei and W. Sewchand. Statistical approach to x-ray ct imaging and its applications in image analysis–part ii: A new stochastic model-based image segmentation technique for x-ray ct image. *IEEE Trans. Medical Imaging*, 11(1):62–69, 1992.

[22] S. Z. Li. *Markov Random Field Modeling in Computer Vision*. Springer-Verlag, 1995.

[23] K. O. Lim and A. Pfefferbaum. Segmentation of MR brain images into cerebrospinal fluid spaces, with and gray matter. *J. Computer Assisted Tomography*, 13(4):588–593, 1989.

[24] J. R. Mitchell, S. J. Karlik, D. H. Lee, and A. Fenster. Computer-assisted identification and quantification of multiple sclerosis lesions in MR imaging volumes in the brain. *J. Magnetic Resonance Imaging*, 4:197–208, 1994.

[25] N. Otsu. A threshold selection method from gray-level histogram. *IEEE Trans. Syst. Man Cybern.*, 9(1):62–66, 1979.

[26] L. R. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.

[27] J. C. Rajapakse and F. Kruggel. Segmentation of MR images with intensity inhomogeneities. *Image and Vision Computing*, 16(3):165–180, 1998.

[28] R. A. Robb. *Visualization Methods for Analysis of Multimodality Images*, chapter Functional Neuroimaging: Technical Foundations, pages 181–190. Orlando, FL: Academic Press, 1994.

[29] P. Schroeter, J.-M. Vesin, T. Langenberger, and R. Meuli. Robust parameter estimation of intensity distributions for brain magnetic resonance images. *IEEE Trans. Medical Imaging*, 17(2):172–186, 1998.

[30] J.A. Vlontzos and S.Y. Kung. Hidden Markov models for character recognition. *IEEE Trans. Image Processing*, 1(4):539–543, 1992.

[31] W. M. Wells, E. L. Grimson, R. Kikinis, and F. A. Jolesz. Adaptive segmentation of MRI data. *IEEE Trans. Medical Imaging*, 15(4):429–442, 1996.

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

[33] J. Zhang. The mean field theory in EM procedures for blind markov random field image restoration. *IEEE Trans. Image Processing*, 2(1):27–40, 1993.

[34] S. C. Zhu and A. Yuille. Region competition: Unifying snakes, region growing, and Bayes/MDL for multiband image segmentation. *IEEE Trans. Pattern Anal. Machine Intell.*, 18(9):884–900, 1996.