AN ACCURATE AND EFFICIENT BAYESIAN METHOD FOR AUTOMATIC SEGMENTATION OF BRAIN MRI

J. L. Marroquín, B. C. Vemuri, S. Botello, F. Calderón and A. Fernández-Bouzas, M. D.

Comunicación Técnica No I-02-08/08-05-2002
(CC/CIMAT)

An accurate and efficient Bayesian method for automatic segmentation of brain MRI

J.L. Marroquin[1], B. C. Vemuri[2], S. Botello[1], F. Calderon[3]
and A. Fernandez–Bouzas, M.D.[4]
[1]Centro de Investigación en Matemáticas
Apdo. Postal 402 Guanajuato, Gto. 36000, Mexico
[2]Dept. of Computer & Information Science & Engineering
University of Florida
Gainesville, Fl. USA. 32611-6120
[3]División de Estudios de Posgrado, Facultad de Ingenieria Electrica
Universidad Michoacana de San Nicolás de Hidalgo.
Santiago Tapia 403, Morelia, Mich., Mexico.
[4]Centro de Neurobiologia
UNAM Campus Juriquilla
Juriquilla, Qro., Mexico.

Abstract

Automatic 3D segmentation of the brain from MR scans is a challenging problem that has received enormous amount of attention lately. Of the techniques reported in literature, very few are fully automatic. In this paper, we present an efficient and accurate, fully automatic 3D segmentation procedure for brain MR scans. It has several salient features namely, (1) instead of a single multiplicative bias field that affects all tissue intensities, separate parametric smooth models are used for the intensity of each class. (2) A brain atlas is used in conjunction with a robust registration procedure to find a non–rigid transformation that maps the standard brain to the specimen to be segmented. This transformation is then used to: segment the brain from non–brain tissue; compute prior probabilities for each class at each voxel location and find an appropriate automatic initialization. (3) Finally, a novel algorithm is presented which is a variant of the EM procedure, that incorporates a fast and accurate way to find optimal segmentations, given the intensity models along with the spatial coherence assumption. Experimental results with both synthetic and real data are included, as well as comparisons of the performance of our algorithm with that of other published methods.

1 Introduction

Image segmentation is a fundamental problem in image processing and computer vision with numerous applications including but not limited to medical image analysis, image compression, etc. Three-dimensional processing and visualization of medical images is a rapidly growing area of research and MRI has provided a means for imaging tissue at very high resolutions providing the desired information for use in fields like reparative surgery, radiotherapy treatment planning, stereotactic neurosurgery and others [24]. In the context of neuro-imaging, 3-D segmentation of white matter (WM), gray matter (GM) and Cerebro–spinal fluid (CSF)

is extremely important for quantitative analysis such as volume measurements. It has already been established that volumetric analysis of different parts of the brain is useful in assessing progress or remission of various diseases e.g., the Alzheimer's [11] and epilepsy [19].

The main difficulties found in the automatic segmentation of MR brain images derive from the fact that image intensities are not necessarily constant for each tissue class. Thus, traditional clustering schemes, such as those reported in [7, 24], or methods based on thresholding [37] do not yield desired results. An additional difficulty stems from the presence of noise in the data, which causes pixel-wise classification methods to give unrealistic results, where the regions corresponding to each tissue class may appear granular, fragmented, or violating anatomical constraints. Thus, it is necessary to devise methods that simultaneously include the estimation of variable intensity models for each class, and prior knowledge about the location and spatial coherence of the corresponding regions.

The most natural framework for the design of such methods is probabilistic, and involves the simultaneous estimation of a discrete field (the tissue class) and a continuous one (the corresponding intensities). Spatial coherence assumptions may be incorporated using a Bayesian approach, in the form of prior Markov Random Field (MRF) models [17, 29], and anatomical constraints about the location of each class may be specified as prior probabilities, obtained from statistical studies [31]. The resulting estimation procedure is iterative, and consists, in general, of 2 steps that are repeated until convergence:

1. Estimate the most likely segmentation, given the intensity models for each class.

2. Estimate the intensity models, given a segmentation,

with an appropriate initialization step.

If a "soft" segmentation is computed in the first step, i.e., if one computes the marginal probabilities for each tissue class, this iterative procedure is equivalent to the Expectation maximization (EM) algorithm [12], and it has been used by a number of researchers [18, 4, 36, 42, 26, 30, 31]. The efficiency of these methods, however, has been limited by several factors. If no spatial coherence assumptions are included, the estimation of the posterior marginal probabilities for each class at each voxel is straightforward and computationally efficient (as in [42]); the quality of this type of segmentation, however, degrades rapidly if noise is present in the data. If, on the other hand, spatial coherence assumptions are included, the exact computation of the optimal segmentation becomes intractable, so that approximations must be made. The most precise are based on Markov Chain Monte Carlo (MCMC) methods, such as the Gibbs Sampler or Metropolis algorithms [17], but are computationally very expensive. Deterministic approximations, such as the ones based on Mean Field (MF) Theory [44, 16, 43, 22] (as used in [31]) are faster, although less precise and more vulnerable to noise [33].

As an alternative, one can estimate a hard segmentation in the first step. Here, one has to choose an appropriate cost function, to get the corresponding optimal estimator (e.g., the Maximum a Posteriori (MAP) or the Maximizer of the Posterior Marginals (MPM) estimators). It was shown in [32], that for low SNR (signal to noise ratio), the MPM criterion yields superior segmentation results whereas for high SNR, both are more or less equivalent.

Another important issue to be taken into consideration is: for a given criterion, what is the best algorithm for finding the optimal estimator? Computational efficiency is a crucial issue that influences the answer to this question. Again, in the presence of immense computational resources, MCMC schemes (such as Simulated Annealing [27] as used in [21] for computing the MAP) are very apt. On the other hand, if computational efficiency is of high priority, it is necessary to resort to viable approximations, such as the ICM algorithm [3] used in [34, 21, 35, 45] for computing the MAP, which is very fast, but highly vulnerable to noise. It is also possible to use the MF or the Gauss-Markov Measure Field (GMMF) [33] approximations for computing the MPM. In [33], it is shown that for low SNR, GMMF is found to be superior to MF, both in terms of accuracy and of computational efficiency. For medium-high SNR MF and GMMF give equivalent results, but GMMF is faster.

The computation of the image intensity model parameters associated with each tissue class for a given segmentation (step 2) is complicated by the fact that both intrascan and interscan intensity inhomogeneities often appear, due to poor radio frequency coil uniformity, operating conditions of the MR equipment, etc. These inhomogeneities are usually modeled as a single multiplicative degradation (bias field) that affects a constant intensity model for each class [42, 26, 35, 30, 31]. Thus, the intensity associated to class k at voxel r is usually modeled as:

$$I_k(r) = \beta(r)\mu_k \tag{1}$$

where the bias field β is assumed to have a slow spatial variation, and μ_k denotes the unknown constant intensity for class k. This intensity model, however, may not be very realistic: there may be variations in the magnetic susceptibility of the tissues; attenuation may be different depending on the location of anatomical structures, etc. To compensate for these effects, in [21] a Parzen-window distribution is used to model the conditional probability of the intensity for each tissue class. This distribution, however, requires of high level user intervention to be defined in each case, and leads to equations that are, in practice, too complicated to be solved exactly.

An additional complication introduced by model (1) is that to make it computationally tractable, the multiplicative degradation must be transformed into an additive one by means of a logarithmic transformation, which has the undesirable effect of altering the intensity distribution, making the tissue separation more difficult, and also, making the additive Gaussian noise assumption invalid for the transformed data. A further difficulty comes from the fact that non–brain tissue may be wrongly classified as gray or white matter in the segmentation step, leading to wrong estimated values for the class intensities. This may be avoided if the brain/non-brain separation is effected prior to segmentation, either manually (as in [35]), or automatically, using for example, active contours [2], active brain templates [35] or registering a standard brain atlas [31].

In summary, while the Bayesian approach to MRI segmentation appears to be very promising, there are several problems which limit its performance and need resolution prior to making it a feasible approach. Below is an itemization of these problems (not necessarily in any specific order).

1. If spatial coherence assumptions are included, an exact computation of the optimal

segmentation at each step becomes intractable and approximations are either computationally expensive or inaccurate.

2. A single bias field affecting all tissue classes equally may not be realistic.

3. A multiplicative bias model may require a logarithmic transformation that distorts the intensity distribution and makes the Gaussian assumption for the noise distribution invalid.

4. Computation of the intensity models for the tissue classes is made rather difficult by the presence of misclassified non-brain tissue.

5. Appropriate initialization of the iterative procedure is difficult to find in a fully automatic way.

The goal of this paper is to present a fully automatic Bayesian–based scheme that overcomes these difficulties, and thus, has superior performance. To achieve this, our method incorporates the following features:

1. Instead of a single multiplicative bias field that affects all tissue intensities, we propose separate parametric, smooth models for the intensity of each class. This may be a more realistic model and avoids the need for a logarithmic transformation and hence the related nonlinear distortions.

2. We propose the use of a brain atlas (as in [31]), together with a robust registration procedure to find a non-rigid transformation that maps the standard brain to the specimen to be segmented. This transformation is then used to segment the brain from non–brain tissue; compute prior probabilities for each class at each voxel location and find an appropriate automatic initialization.

3. We propose a novel variant of the EM algorithm [12] which allows for the use of a fast and accurate way to find optimal segmentations, given the intensity models which incorporate the spatial coherence assumptions.

The rest of this paper is organized as follows: the above described algorithm features are elaborated upon in section 2. Section 3 contains the experimental validation of the complete procedure, and finally some conclusions are drawn in section 4.

2 An Improved Bayesian Approach to MRI Segmentation

In this section, we first present a Bayesian formulation of the segmentation problem followed by a novel, accurate and efficient 3D segmentation algorithm which is dubbed the MPM-MAP algorithm.

2.1 Bayesian Formulation of the Segmentation Problem

In a Bayesian approach to segmentation, one specifies the distribution of the observation noise, and introduces the prior constraint about the spatial coherence of the support regions in the form of a prior probability distribution on the set of possible segmentations. To describe this approach, we introduce the following notation: let L denote the voxel lattice, and $\{g(r), r \in L\}$ denote the observed MR volume; let $\{\Phi(r;\theta_k), k = 1, ..., K\}$ denote a set of models characterized by the parameter vector $\theta = (\theta_1, ..., \theta_K)$, which describe the variation of the value of a property f over L. Note that f may represent a scalar valued data representing the intensity values in a 3D volume, or it may be a vector valued function defined on a 3D image grid (e.g., for multimodal data). Each model k is supposed to be valid in a region $R_k \subseteq L$, so that the value of the field f at pixel r is $f(r) = \sum_{k=1}^{K} \Phi(r;\theta_k) b_k(r)$, where $b_k(r)$ is the indicator function for region R_k, i.e., $b_k(r) = 1$ if $r \in R_k$, and $b_k(r) = 0$, otherwise. Assuming a Gaussian observation noise model, the likelihood of the data given the models and the support regions is :

$$P(g|b,\theta) = \frac{1}{Z_L} \Pi_{r \in L} \Pi_{k=1}^{K} l_k(r)^{b_k(r)}$$

where Z_L is a constant and

$$\begin{aligned} l_k(r) &= \Pr\left(g(r)|\theta, b_k(r) = 1\right) \\ &= \sqrt{\frac{\gamma}{\pi}} \exp\left[-\gamma |g(r) - \Phi(r,\theta_k)|^2\right] \end{aligned} \tag{2}$$

where γ is a parameter that depends on the noise variance (i.e., we are assuming that the noise is independent of the tissue class). One may also have prior probabilities that specify the probability of the voxel at location r belonging to each of the classes. We will denote these probabilities by $q_k(r) = \Pr(\text{voxel } r \in \text{class } k)$. The constraint for the spatial coherence of the support regions $\{R_1, ..., R_K\}$ may be expressed, in probabilistic terms, in the form of a prior Markov Random Field (MRF) model on the indicator variables b; in particular, one may assume a prior generalized Ising (Potts) model [32], where the neighborhood N_r of each site $r \in L$ is formed only by its nearest neighbors (in a 2-D image there are four such neighbors, and in a 3-D volume image there are 6 in the simplest neighborhood system, which is the one we use in all the experiments reported in this paper). This gives the prior Gibbs distribution: $P_b(b) = \frac{1}{Z_b} \exp\left[\beta \sum_{<r,s>} V(b(r), b(s))\right]$. Where, Z_b is a normalizing constant, β is a positive parameter controlling the granularity of the regions, the summation is performed over all nearest–neighbor pairs of sites in L and the Ising potential V_I is given by,

$$V_I\left(b(r), b(s)\right) = \begin{cases} -1 & \text{, if } \sum_{k=1}^{K} b_k(r) b_k(s) = 1 \\ 1 & \text{, otherwise} \end{cases} \tag{3}$$

(note that $b_k(r) \in \{0, 1\}$ and that $\sum_{k=1}^{K} b_k(r) = 1$ for all $r \in L$).

The posterior distribution is computed using Bayes rule and is of the form: $P(b,\theta|g) = \frac{1}{Z}\exp\left[-U(b,\theta)\right]$, where Z is a normalizing constant and

$$\begin{aligned} U(b,\theta) &= -\sum_{r\in L}\sum_{k=1}^{K}\left[b_k(r)\log l_k(r) + \log q_k(r)\right] \\ &\quad +\beta\sum_{<r,s>} V_I(b(r),b(s)) + \log P(\theta) \end{aligned} \tag{4}$$

where $l_k(r)$ is given by Eq. (2), and $P(\theta)$ represents the prior distribution of the parameters that define each intensity model.

In the classical Expectation Maximization (EM) approach [12], the b variables are considered "missing data" and U is minimized in a 2–step procedure consisting of an (E) step in which one first computes the expected value $< b >$ of the b variables, given an estimate of the model parameters θ, while in the second (M) step, one minimizes the U function, where the b variables are replaced by their expected value. The drawback of this approach lies in the computational complexity associated with the computation of $< b >$: since the variables are binary, their expected value is equal to the posterior marginal probabilities given by

$$\pi_k(r) = \Pr\left(b_k(r)=1|\theta,g\right) = \sum_{b:b_k(r)=1}\frac{1}{Z}\exp\left[-U(b,\theta)\right] \tag{5}$$

It is clear that the summation in (5) has too many terms to be exactly computed, so it has to be approximated. There are several methods that have been proposed for estimating the posterior marginals of discrete–valued MRF's: stochastic approaches based on the construction of regular Markov chains [17, 32], and deterministic approaches based on the mean–field approximation [16, 44] in which the estimated marginals are obtained as solutions of large systems of coupled non–linear equations. These approaches, however, are computationally very expensive. Other methods [43, 22] that also involve mean field approximations using different cost functions require solution to a large coupled system of nonlinear equations but possibly lead to more accurate – than the approximation outlined in [44] – approximation of the marginals. In this paper, we propose a novel approach with which one can get fast, high quality estimators solving a set of decoupled, *linear* systems of equations. It is based on the fact that, given a Gibbs distribution with energy of the form:

$$U(b) = \sum_{r\in L} V_0(b(r)) + \beta\sum_{<r,s>} V_I(b(r),b(s)) \tag{6}$$

where the b's are random $K-$vectors satisfying: $b_k(r)\in\{0,1\}$ and $\sum_{k=1}^{K} b_k(r) = 1$ for all $r\in L$, and V_I is an Ising potential, then, the marginals $\pi_k(r) = \Pr(b_k(r)=1)$ may be approximated by the expected value of random vectors p that admit a Gibbsian model with energy of the form:

$$U_p(p) = \sum_{r\in L}|p(r)-\widehat{p}(r)|^2 + \lambda\sum_{<r,s>}|p(r)-p(s)|^2 \tag{7}$$

where $\widehat{p}_k(r) = \frac{1}{Z}\exp\left[-V_0(e_k)\right]$ with Z a normalizing constant and where e_k is a unit vector with a 1 in the k^{th} position. (see [33] for details). Therefore, an approximation to the marginals may be constructed by minimizing (7), that is, by solving K sets of decoupled linear equations, that result from setting the partial derivatives of U_p, with respect to the variables $p_k(r)$, equal to zero. These equations are of the form:

$$(1 + \lambda|N_r|)\, p_k(r) + \lambda \sum_{s \in N_r \cap L} p_k(r) = \hat{p}_k(r) \tag{8}$$

where N_r denotes the neighborhood of voxel r (i.e., its 6 nearest neighbors for interior voxels). The EM algorithm, however, may be quite sensitive to errors in the estimated marginals (see section 2.2). Therefore, if this fast approximation procedure is to be used, it is necessary to modify the algorithm to make it more robust with respect to this kind of errors. This modified algorithm will be described in the next subsection.

2.2 The MPM-MAP Algorithm

The algorithm we propose for the simultaneous estimation of b and θ is based on Bayesian estimation theory. In this approach, we consider b and θ as random vectors whose optimal estimators are to be found via the minimization of the expected value of an appropriate cost function, taken with respect to the posterior distribution. The cost function we propose is:

$$C(\widehat{b}, \widehat{\theta}, b, \theta) = 1 - \delta(\theta - \widehat{\theta}) + \frac{1}{|L|} \sum_{r \in L} \left[1 - \delta\left(b(r) - \widehat{b}(r)\right)\right] \tag{9}$$

where the δ functions equal 1 if their argument is the 0 vector, and equal 0 otherwise. The first term requires that the estimated parameter vector $\widehat{\theta}$ be — on the average — a perfect estimator, while the second term requires that the estimated indicator functions $\widehat{b}$ for the support regions, minimize the expected number of segmentation errors. Denoting $Q(\widehat{b}, \widehat{\theta}) = E\left[C(\widehat{b}, \widehat{\theta}, b, \theta)\right]$, the optimal estimators $(\widehat{b}^*, \widehat{\theta}^*)$ are, therefore,

$$(\widehat{b}^*, \widehat{\theta}^*) = \arg\min_{\widehat{b}, \widehat{\theta}} Q(\widehat{b}, \widehat{\theta})$$

To minimize Q we propose a 2-step procedure in which $Q(\widehat{b}, \widehat{\theta})$ is minimized with respect to $\widehat{b}$ for a given $\widehat{\theta}$ in a first step, and then minimized with respect to $\widehat{\theta}$, keeping the optimal $\widehat{b}$ fixed, in the second step. To derive the implementation of the first step, we take the following into consideration: suppose that $\widehat{\theta} = \overline{\theta}$ is given. The optimal estimator for $\widehat{b}$ is found by minimizing the expected value of the second term of Eq. (9):

$$\begin{aligned} & \sum_{b} \sum_{r \in L} \left[1 - \delta\left(b(r) - \widehat{b}(r)\right)\right] P(b, \overline{\theta}|g) \qquad (10) \\ = \;& |L| - \sum_{r \in L} \sum_{k=1}^{K} \pi_k(r) \widehat{b}_k(r) \end{aligned}$$

where

$$\pi_k(r) = \sum_{b:b_k(r)=1} P(b, \overline{\theta}|g) \tag{11}$$

is the posterior marginal probability for the support region k at pixel r. Expression (10) is minimized by setting $\widehat{b} = \overline{b}$, where $\overline{b}_k(r) = 1$, if $\pi_k(r) > \pi_{k'}(r)$ for $k' \neq k$,and $\overline{b}_k(r) = 0$, otherwise. This estimator is called the Maximizer of the Posterior Marginals or MPM estimator [32] for b given $\overline{\theta}$.

To minimize Q with respect to $\widehat{\theta}$ for a fixed $\widehat{b} = \overline{b}$ one needs to consider only the expected value of the first term of (9), so that the optimal (Maximum a Posteriori or MAP) estimator for θ is found by minimizing $U(\overline{b}, \theta)$ (Eq. (4)) with respect to θ. The complete algorithm is therefore:

Algorithm 1

1. Compute an initial estimate $\overline{b} = b^{(0)}$ for the segmentation indicator variables and set $t = 0$;
2. Until convergence do steps (3-4):
3. MAP-step: compute $\theta^{(t+1)} = \arg\min_\theta U(\overline{b}, \theta)$;
4. MPM-step:
   (a) Compute the estimators $p_k(r)$ for the posterior marginals $\pi_k(r)$, $r \in L$, $k = 1, ...K$ solving the linear system (8).
   (b) Set $\overline{b}_k(r) = 1$, if $p_k(r) > p_{k'}(r)$ for $k' \neq k$,and $\overline{b}_k(r) = 0$, otherwise; set $t := t + 1$;

The convergence of the algorithm follows from the facts that Q is bounded below (it is always non–negative) and that it cannot increase in one full iteration; it could converge to a limit cycle, but in practice it is found to converge to a fixed point. *Note that this algorithm is more robust than EM [12] with respect to errors in the estimation of the posterior marginals* $\{\pi_k(r)\}$, *because only the location of the mode of each distribution* $\pi(r)$ *is of consequence.* This is illustrated in Fig. 1, where we show the solutions found by the two procedures in the estimation of piecewise constant models (i.e., $\Phi(r, \theta_k) = \theta_k$) for the gray-level segmentation of an image. To simplify the computations, we assumed no spatial coherence term (i.e., $\beta = 0$ in Eq. (6)), so that the marginals can be exactly computed: they are equal to the normalized likelihood, which for the assumed Gaussian model is:

$$\pi_k(r) = \frac{1}{Z} \exp\left[\frac{(g(r) - \theta_k)^2}{2\sigma^2}\right]$$

where Z is a normalizing constant. We assumed that the σ parameter is not known, and studied the convergence behavior as its value is varied between 0 and 1 (its true value is 0.2). The initial models were computed using the windows shown in the upper left panel. EM always finds three models; for low values of σ, two of these models have the same value for

θ — which is very close to the true value $\theta_1 = 1.0$ — for intermediate σ values they have different θ values, and for high σ values they collapse into a wrong intermediate value of $= 1.5$. In contrast, MPM-MAP is practically insensitive to these variations, and always finds 2 models with the correct values $\theta_1 = 1$ and $\theta_2 = 2$.

2.3 Parametric Models of the Intensities in Each Tissue Class

To complete the specification of our approach, it is necessary to define the parametric models that represent the spatial variation of the intensity within each tissue class. These intensity variations are due to several causes: the primary one is connected with variations in the sensitivity of the receptor coils, and is well explained by the multiplicative bias model of Eq. (1). There are other factors, however, for which this model is not so appropriate [38]; these include: induced currents; non–uniform excitation; variation of electromagnetic properties of biological tissues, etc. For this reason, in this paper we use a more general approach in which the variable intensity for each class is separately modeled. The choice of the appropriate parametric models is important: polynomial models are usually too rigid (if their degree is low like in planar surfaces) or too flexible (for high degrees), so that a single model may fit more than one class. What is needed, therefore, are models that can account for smoothly varying intensities and are easy to compute given the support. Here, we use spline models with a Gibbsian prior $P(\theta)$ that imposes a controlled smoothness constraint. This is equivalent to a finite element approximation to a membrane spline. In particular, we define the model $\Phi(r, \theta_k)$ as:

$$\Phi(r, \theta_k) = \sum_{j=1}^{m} N_j(r)\theta_{kj} \tag{12}$$

where $\{N_j\}$ are the classical trilinear interpolation functions used for the 8-node Lagrangian element[46], and the parameters θ_{kr} correspond to the height of the membrane at the m nodes of the finite element mesh.

The smoothness constraint is specified in the form of a Gibbs distribution:

$$P(\theta_k) = \frac{1}{Z} \exp\left[-\eta \int_L |\nabla\Phi(r, \theta_k)|^2 dr\right] \tag{13}$$

where Z is a normalizing constant and η is a parameter. Note that using (12), the integral in (13) may be evaluated as a quadratic form in the parameter vector θ_k. The value for these parameters is found in the MAP step by minimization of:

$$\begin{aligned} U_\theta(\theta_k) &= \sum_{r\in L} b_k(r) \left(\Phi(r, \theta_k) - g(r)\right)^2 \\ &\quad +\eta \int_L |\nabla\Phi(r, \theta_k)\,(r)\,|^2 dr \end{aligned} \tag{14}$$

This scheme allows one to model arbitrary smooth shapes, while retaining sufficient control on the smoothness of each model (given by the positive parameter η). Substituting (12) into (14) and setting the partial derivatives of U_θ with respect to the parameter vector θ_k equal to

zero, one obtains a system of *linear* equations whose solution gives the optimal parameters. In this way, *the whole procedure involves only the solution of linear systems (in both the MAP and MPM steps) and is, therefore, computationally efficient.*

Note that these membrane models have more degrees of freedom than the standard multiplicative model (1), which may be considered a particular case; if the data is in fact consistent with (1), then one should have that, for all class pairs i, j, $\Phi(r, \theta_i) \approx R_{ij}\Phi(r, \theta_j)$, where $R_{ij} = \mu_i/\mu_j$ is a constant, for all voxels r that are close to the regions of validity of models i and j (i.e., close to the boundary between classes i and j). This fact may thus be used to test experimentally the validity of model (1), which we do in section 4.

In the case of multimodal data sets, i.e., when several pulse sequences are used, the use of decoupled membrane models presents an additional advantage; since some pulse sequences are more sensitive to magnetic inhomogeneities than others, fitting a separate membrane for each tissue class and each pulse sequence permits a more accurate modeling that takes this effect into account.

2.4 Prior Class Probabilities and Automatic Initialization

To improve the performance of a segmentation method, we should take advantage of domain–dependent relevant information. In the particular case of brain MRI, we consider the fact that the location of the main anatomical structures is approximately the same in most subjects, if the brain is brought to a standard position and scale. Therefore, it should be possible to incorporate into the segmentation procedure, prior information about the approximate location of the different tissue classes. This information may then be used to provide an initial approximate segmentation to initialize the algorithm (step 1 in section 2.2) and also to estimate the prior class probabilities for each voxel ($q_k(r)$ in Eq. (4)). The class location information may be obtained from statistical studies, in which hand–segmented MRI's for a given population of subjects are brought to a standard pose, so that the frequencies with which each tissue class appears at each standard voxel location can be computed. One such study, based on a population of 151 normal subjects, is available in [1], and is the one we used in our experiments. To incorporate this information in our procedure, it is necessary to find a transformation that maps the standard anatomical model (to which the intensity distribution is referred) into the specimen to be segmented. A robust and precise registration procedure is crucial for the success of the complete algorithm; we use a 2–step registration scheme: in the first step, an affine transformation is found (of the form $T_{aff}[r] = Ar + b$, where A is a 3×3 matrix and b is a 3–vector), so that it minimizes the sum of square differences between the intensities derived from the transformed anatomical model and the specimen. For this, we use a multi–scale Broyden–Fletcher–Goldfarb–Shanno (BFGS) minimization method [13]. In the second step, a non–rigid transformation that refines the affine map found in the first step is computed using a level–set based Partial Differential Equation (PDE) scheme described in [41]. The governing equation of the non-rigid registration method is given by

$$\frac{\partial V}{\partial t} = (I_2(\mathbf{x}) - I_1(V(\mathbf{x}))\frac{\nabla I_1(V(\mathbf{x}))}{\|\nabla I_1(V(\mathbf{x}))\|} \tag{15}$$

where $V(\mathbf{x}) = (x + u, y + v, z + w)$, (u, v, w) defines the displacement field at each point on the image grid $\mathbf{x} = (x, y, z)$. This is a nonlinear hyperbolic partial differential equation characterizing the non-rigid deformation field $V(\mathbf{x})$. We solve this equation using minmod finite difference techniques for automatically determining the time step (see [41] for details). Most often than not, the computed non-rigid deformation is not smooth, and one way to obtain a smooth deformation field is to apply to the vector field obtained from the above equation at each iteration a linear smoothing operator S, e.g., a Gaussian filter with a small fixed variance.

There are other schemes that have been used to do atlas–based segmentation. In Chen et.al., [5], an atlas-based segmentation scheme is described that uses expert input to define the atlas and then warp it to the subject brain MRI to segment the subject brain which is then followed by morphometrics. Their algorithm is almost identical to the work reported in Dawant et.al., [10] and Thirion [39]. These methods have been compared with the technique in Vemuri et.al., [41] – used here – and the latter has been shown to outperform these techniques in computational efficiency and accuracy. In [25], atlas-based segmentation was achieved using the so called fluid-flow model introduced by Chritensen et.al., [6]. A more general and mathematically thorough treatment of the non-rigid registration which subsumes the fluid-flow methods was presented in Trouve [40]; this last method, however, has not been used in atlas-based segmentation applications.

Once the complete (affine+non-rigid) transformation T is computed, it is used for the following purposes:

1. To map a mask that separates brain from non–brain tissue. Note that since a "hard" brain segmentation is performed, it is unnecessary to include an "outlier class" for non–brain tissue, as in [31] or [20].

2. To map the anatomical model to get an initial segmentation.

3. To map the empirical frequencies f_k obtained from the statistical study, from which the prior probabilities may be estimated.

We have found that it is not convenient to make the prior probabilities equal to the empirical frequencies as in [31], since class distributions may have wide variations in populations which are different from the one on which the study is based. Therefore, it is convenient to mitigate their effect on the prior probability estimates by contaminating them with a uniform distribution, so that the prior probability field is estimated as:

$$q_k(r) = \alpha f_k(r) + (1 - \alpha)\frac{1}{K} \tag{16}$$

where K is the number of classes considered and $\alpha \in [0, 1]$ is a parameter that weights the influence of the empirical frequencies.

2.5 Complete Segmentation Procedure

In summary, the complete segmentation procedure we are proposing consists of the following steps:

1. Find the complete transformation T that maps the standardized (atlas) space into the specimen space using the procedure of section 2.4. This procedure consists of the following steps: let I_R denote the reference (atlas) volume, and I_S the specimen.

   a) Find the optimal affine transformation parameters A^*, b^*:

   $$(A^*, b^*) = \arg\min_{A,b} \sum_{r \in L} (I_R(Ar + b) - I_S(r))^2$$

   using a multiscale BFGS minimization algorithm.

   b) Find the affine–transformed reference image $\hat{I}_R(r) = I_R(A^* r + b^*)$

   c) Set the initial non–rigid deformation field $V^{(0)}(r) = r$

   d) Iterate until convergence:

   $$\begin{aligned} \tilde{V}(r) &= V^{(t)}(r) + (I_S(r) - \hat{I}_R(V^{(t)}(r))) \frac{\nabla \hat{I}_R(V^{(t)}(r))}{|\nabla \hat{I}_R(V^{(t)}|} \\ V^{(t+1)} &= S(\tilde{V}) \end{aligned}$$

   where S is a smoothing operator.

   Let V^* be the fixed point of this iterative procedure, and let I be an image in the atlas (reference) space. The corresponding transformed image $\hat{I}$ in the specimen space is found as:

   $$\hat{I}(r) = \tilde{I}(V^*(r)) \tag{17}$$

   with $\tilde{I}(r) = I(A^* r + b*)$.

2. To peel the brain, obtain from the mask I_M that separates the brain parenchyma from non–brain tissue in the atlas, the transformed mask $\hat{I}_M$ using (17) and apply it to the specimen.

3. Obtain the transformed empirical frequencies for each tissue class $\hat{f}_k$, $k = 1, 2, 3$, from the reference ones using (17), and compute the prior class probability field using (16).

4. Obtain the transformed segmentation $\hat{b}$ from the atlas anatomical model using (17) and set the initial segmentation $b^{(0)} = \hat{b}$.

5. Apply the MAP–MPM algorithm (Algorithm 1 of section 2.2) until convergence.

In the case of multi–band images (e.g., T1, T2 and PD weighted MRI), the individual images have to be registered separately, and a different membrane model has to be adjusted for each band in the MAP step in a decoupled manner. One finds, however, a unique segmentation in each MPM step.

This procedure has a number of parameters: γ in equation (2); λ in Eq. (7); η in Eq.(13) and α in Eq.(16). The best value for these parameters has to be hand–picked using a trial

and error procedure on a test image. Once these values are found, however, they may be used in different data sets without need of further tuning, since the performance of the method is quite insensitive to their precise value (see section 3). The values we have used in *all* the experiments reported in the next section, both for synthetic and real data, are: $\gamma = 1$; $\lambda = 0.1$; $\eta = 1000$ and $\alpha = 0.1$. The finite element mesh size used in all our experiments to be described in the next section is $(32, 32, 8)$.

3 Experimental Validation

To validate the performance of our method, we performed 2 sets of experiments, one on simulated MR and another on real MR brain scans. In the first set, we used the simulated MR images of the head generated using the BrainWeb MR simulator [8, 28, 9]. Since in this case, the anatomical model (ground truth) available, it is possible to obtain a quantitative assessment of the performance of the algorithm, under different conditions. It is important to note that the original MRI simulator described in [28] had the limitation of producing sharp tissue boundaries in simulated images, due to the discrete nature of the labelled data set (phantom), and thus, Partial Volume Effects (PVE) were not appropriately modeled. The version we are using, however, includes the "fuzzy phantom" developed in [9], where instead of assigning a voxel to a given tissue class, one has a 3-D tissue model, with one volume per class, where voxel intensity represents the fraction of tissue (between 0 and 1) within the voxel. The MR simulator described in [28] is used to predict image contrast, and the phantom is used to map tissue intensities into images. The simulated volumes available in [8], therefore, account for the effects of partial volume averaging, noise and spatial inhomogeneities. We considered the following cases:

1. T1 weighted images with 1 through 9 % noise levels and no spatial inhomogeneity.
2. Same as (1), with 40 % spatial inhomogeneity (i.e., with intensity variations of up to 40 % for each tissue class).
3. Multi–band (T1, T2 and PD weighted) data with 1 through 9 % noise and 40 % spatial inhomogeneity .

The technical data pertaining to these simulations are shown in table 1. Since these experiments were also reported in [31], in order for the results to be meaningfully compared, we used the performance index reported therein. This index is defined as:

$$I_1(k) = \frac{2V_{p\cap g}(k)}{V_p(k) + V_g(k)} \tag{18}$$

where $V_{p\cap g}(k)$ denotes the number of voxels classified by both the proposed method and the ground truth as class k, and $V_p(k)$ and $V_g(k)$ represent the number of voxels classified as class k by the proposed method alone and by the ground truth, respectively. Note that this index will take the value 1 if the proposed method coincides with the ground truth,

and will decrease as the quality of the segmentation deteriorates. The results are shown in Figures (2) and (3) respectively. As one can see, the proposed method shows an excellent and stable performance, outperforming the best results (to our knowledge) reported to date in literature. A sample of the segmentation results is presented in Figure (4). The figure depicts corresponding slices from the ground truth (panel d), automatic segmentation (panel e) and segmentation errors, as well as the final estimated intensities, using separate membrane models for each tissue class and each pulse sequence (panels g-i).

It is interesting to note that the benefits of using multi–band images over using only T1 weighted data are relatively small, while the increase in computational load is significant, which suggests that, for segmentation purposes, it might often be more convenient to use only single band, say T1 weighted images.

We also performed experiments with these simulated images to test the sensitivity of the method with respect to the precise setting of the control parameters. Specifically, we investigate the influence of:

1. The weight of prior probabilities (i.e., α in Eq. (16)).
2. The weight of the spatial interacion term (λ in Eq. (7)).
3. The rigidity of the membrane intensity models (η in Eq. (13)).
4. The value of the noise parameter γ.

For the test volume, we take the "worst case" simulation (9 % noise and 40 % spatial inhomogeneities). We do not test here the factors related to the registration with an atlas, since in this case the atlas and simulated images are already registered. The results are presented in table 2; in each case, when a parameter value was varied, the others were kept fixed at their default values. As one can see, the performance is relatively stable, which is the reason why one can use the same (default) parameter values for all the experiments reported here.

We also performed experiments to test the effect of replacing the MPM–MAP algorithm by the standard EM procedure, with the Mean Field algorithm [44, 16, 43, 22] for the E step, using the same test volume. In this case we find practically the same performance (indices equal to 0.885 for CSF, 0.87 for GM and 0.89 for WM), but the processing time increases from 248 sec. (for MPM–MAP) to 466 sec. (for EM).

Showing good performance with simulated data, however, is not sufficient to validate a segmentation procedure. It is also very important to test it with real images and compare it with other published methods. This comparison is difficult to do, because most of the published methods work with different data sets. To overcome this problem, we use the 20 normal MR brain data made publicly available on the world wide web by the Center for Morphometric Analysis at Massachusetts General Hospital (MGH). In addition, manual expert segmentations and performance results from five automatic segmentation methods are also provided at this site [23], making it convenient to compare our results with those reported by the five automatic methods. The 20 coronal scans of this data set were chosen

because they have been used in published volumetric studies in the past, and because they have various levels of difficulty; the worst scans have low contrast and relatively large spatial inhomogeneities. Hence, these test images permit a standardized mechanism for testing the sensitivity of a proposed segmentation method to signal to noise ratio, contrast to noise ratio, shape complexity, degree of partial volume effect, etc.

The repository also contains performance indices that measure the amount of overlap between the expert hand–guided segmentation and a collection of automatic methods. The index in this case is the Tanimoto coefficient [14], and is defined as:

$$I_2(k) = \frac{V_{p\cap g}(k)}{V_{p\cup g}(k)}$$

where $V_{p\cap g}(k)$ denotes the number of voxels classified as class k by both the proposed method and the expert (taken in this case as ground truth), and $V_{p\cup g}(k)$ denotes the number of voxels classified as class k by either the proposed method or the expert. It should be noted that for a given automatic classification and ground truth, one always has that $I_2(k) \leq I_1(k)$ (although both indices range from 0 to 1), where I_1 is given by Eq. (18).

We applied our method to the twenty complete head MR scans in this data set, *using the same parameter values in all cases*, which were also the same values used to segment the simulated data, i.e., $\gamma = 1$, $\lambda = 0.01$, $\eta = 1000$ and $\alpha = 0.1$, so that it may, in effect, be considered an automatic procedure. The results are summarized in table 3 and figure (5). The dimensions of the image stacks were $256 \times 256 \times 64$ voxels, and the average total processing time (including registration for peeling the skull and non-brain material and segmentation) was 29 minutes on a single processor of an SGI ONYX machine. As one can see, these results are significantly better than all other reported methods (see [35] for details on these other methods). Figure (6) depicts a sample of our segmentation results, together with the corresponding expert (ground truth) segmentation. We used an itensity value of $255, 170$ and 85 to depict the WM, GM and CSF pixels respectively.

A final validation was performed by computing the performance indices I_2, comparing a manual segmentation performed by an expert radiologist (Dr. Fernandez–Bouzas, M.D.) with the automatic segmentation, for a MR volume not included in the MGH data base. The index values were 0.82 for WM and 0.84 for GM.

An interesting question is to quantify the relative importance of each feature in the overall efficiency of the method. To answer this, we performed a set of experiments to investigate the influence of each one of the following factors:

1. Use of spatially varying prior class probabilities (i.e., using $\alpha > 0$ in Eq. (16)).

2. Use of an anatomical model registered with the specimen to be segmented to initialize the method vs. using a more standard initialization procedure (e.g., equispaced constant intensities as initial models).

3. Use of a prior spatial coherence (MRF) model (i.e., using $\lambda > 0$ in Eq.(7)).

4. Use of spatially varying (membrane) intensity models, instead of constant intensity models.

5. Use of the non–parametric, level-set-based refinement after the multi-scale affine registration.

We applied our segmentation procedure by varying these factors, one at a time, on the 20 brains of the Mass. General data base. The average performance indices obtained in each case, together with the average time required for the segmentation (MPM-MAP phase only, excluding the registration step) are shown in table 4.

From these results one can draw the following conclusions:

1. The most important feature of the method is the inclusion of prior probabilities derived from statistical studies. This emphasizes the importance of performing such studies, as well as of developing more precise and efficient registration procedures that map individual brains into a standard space.

2. Using the deformed anatomical model instead of equi-spaced constant intensity models to initialize the method has a significant effect on the performance, and makes the method significantly faster. The importance of this initialization procedure is even greater if multi–band (T1, T2 and PD weighted) data are used, since in this case one cannot simply use equi-spaced intensities, but one must estimate the multi-band models based on randomly placed windows, which may have a greater adverse effect on the performance.

3. The use of membranes instead of constant intensity models has a significant effect on the average performance, which is consistent with the fact that significant spatial inhomogeneities (bias) were present in several brains of the data set.

4. The use of the non–parametric, level–set based refinement of the multi–scale affine registration has little effect on the average performance for the test data. However, application of this non-rigid registration of priors might show helpful in cases with strong morphological differences. The average registration time is 5.4 minutes for the affine registration, and 4.3 minutes for the refinement. Since the refinement step does not take too much time, it is definitely worth doing to obtain better performance.

5. The inclusion of a prior MRF to model spatial interactions had, by itself, only a minor effect on the performance, due to the fact that spatial inhomogeneities played a more significant role than noise in these data. To corroborate this, we evaluated this effect on the performance indices obtained with the synthetic MRI images available from [15]. In this case, the performance indices for both gray and white matter showed a 6 % increase when spatial interactions were included, for data with 9 % noise, but showed practically no change when the noise level was below 3 %.

Finally, we note that although these experiments give an indication of the effect of each feature used by our method, there is, without doubt, a synergistic effect between them, which explains the effectiveness of our procedure when they are used together.

4 Discussion and Conclusions

The Bayesian estimation framework is very convenient for MRI segmentation, because it permits the inclusion of explicit models for the spatial coherence and location of the different tissues, as well as for the spatial inhomogeneities of the corresponding image intensities caused by the bias field in the magnetic field of the MR instrument. For Bayesian-based methods to be effective, however, it is crucial to have: realistic spatially varying models for the intensity of each class; efficient and accurate ways for computing optimal segmentations given these intensity models and robust and precise registration procedures, so that prior information about the location of anatomical structures may be incorporated.

We have presented here a fully automatic method that incorporates these features, and therefore depicted a competitive performance with existing methods reported in literature. The method has a number of control parameters that need to be adjusted; its performance, however, is not too sensitive to their precise setting; thus, it is possible to find suitable values using images with known ground truth, and then use these default values to segment new real data. This is the approach we followed here: the *same* parameter values were used for *all* the experiments presented here, both for simulated and real images.

An important issue is related to the fact that the method we are proposing produces a "hard" segmentation, which means that each voxel is assigned to only one tissue class, whereas in reality, many voxels may have a mixture of 2 tissue classes (the so called Partial Volume Effect or PVE). One would hope, however, that since the method is not biased towards any particular class, these effects would tend to cancel out when total volumes are computed. To test this hypothesis, we computed the error in the total volume using the segmentation produced by our method, with respect to the true volume (which takes into account the PVE) of the anatomical model of the Brainweb MR simulator [8, 28, 9]. We found that for the case with 40 % of spatial inhomogeneities, and with up to 5 % noise one gets less than 1 % error in the total volume for GM and less than 2 % error for WM. We think these error rates are acceptable: in the MGH data base, it is reported that if one uses index I_2 to measure the agreement between hand–guided segmentations performed by 2 different experts, one gets values of 0.83 for WM and 0.87 for GM. This means that the total computed volumes for the corresponding tissues could differ by up to 17 % (resp. 13 %). Compared with these figures the errors induced by the PVE do not seem significant, at least for this task.

Another important question is related to the adequacy of a single multiplicative bias (Eq. (1)) for modeling the intensity variations of all tissue classes. This may be tested by computing the ratio of the estimated individual class intensities (using membrane models, which include the multiplicative bias model as a particular case; see section 2.3) for voxels that lie in the neighborhood of the corresponding inter–class boundary; we consider specifically gray and white matter (classes 2 and 3). If the multiplicative bias model is correct, one should have that $\Phi(r, \theta_k) \approx \mu_k \beta(r)$ and hence,

$$R_{23}(r) = \frac{\Phi(r, \theta_2)}{\Phi(r, \theta_3)} \approx \frac{\mu_2}{\mu_3} = \text{ constant}$$

One may compute the maximum absolute deviation:

$$D_{max} = 100 \times \max_{r \in N_{23}} \frac{|R_{23}(r) - \bar{R}_{23}|}{\bar{R}_{23}}$$

where $\bar{R}_{23}$ is the average of $R_{23}(r)$ taken over N_{23}, which denotes a tubular neighborhood, 2 voxels wide, around the boundary between classes 2 and 3. For synthetic MRI (from the Brainweb simulator), one has that, even in the case of 40 % spatial inhomogeneity and 9 % noise, D_{max} is small (less than 6 %). This is not surprising, since the bias was artificially introduced in the simulation using the multiplicative model. For real images (i.e., for the MGH data base), however, one finds values of D_{max} as high as 30 %, which indicates that factors different from sensitivity variations in the reception coil may have a significant effect, so that the multiplicative bias model may not be adequate, and it may be preferable to use individual bias models for each class, as is done in this paper.

Acknowledgments. J.L. Marroquin and S. Botello were supported in part by grant 34575-A from Conacyt, Mexico. B.C. Vemuri was in part supported by the grants NSF IIS9811042 and NIH RO1-RR13197 respectively. The authors are grateful to Dr. T. Harmony for many useful discussions, and for her advice on the experimental validation of the method.

References

[1] J. Ashburner, K. Friston, A. Holmes, and J.B. Poline. Statistical parametric mapping. Technical report, The Wellcome Dept. Cognitive Neurology, Univ. College, London, U.K. Available: http://www.fil.ion.ucl.ac.uk/spm/, 2000.

[2] M.S. Atkins and B.T. Mackiewich. Fully automatic segmentation of the brain in mri. *IEEE Trans. on Medical Imaging*, 17(1):98–107, 1998.

[3] J. Besag. Spatial interaction and statistical analysis of lattice systems. *Journal of Royal Statistical Society B*, 36:192–236, 1974.

[4] M. E. Brummer. Automatic detection of brain contours in mri data sets. *IEEE Trans. on Med. Imaging*, 12:153–166, 1993.

[5] M. Chen, T. Kanade, H. A. Rowley, and D. Pomerleau. Quantitative study of brain anatomy. In *IEEE Workshop on Biomedical Image Analysis*, pages 84–92, 1998.

[6] G. E. Christensen, M. I. Miller, and M. Vannier. Individualizing neuroanatomical atlases using a massively parallel computer. *IEEE Computer*, 1(29):32–38, 1996.

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

[8] C.A. Cocosco, V. Kollokian, R.K. Kwan, and A.C. Evans. Brainweb: Online interface to a 3d mri simulated brain database. *Neuroimage*, 5(4):part2/4, S425 (images available on www.bic.mni.mcgill.ca /brainweb/), 1997.

[9] D.L. Collins, A.P. Zijdenbos, V. Kollokian, J.G. Sled, N.J. Kabani, C.J. Holmes, and A.C. Evans. Design and construction of a realistic digital brain phantom. *IEEE Trans. on Med. Imaging*, 17(3):463–468, 1998.

[10] B. M. Dawant, S. L. Hartmann, J. P. Thirion, F. Maes, D. Vandermeulen, and P. Demaerel. Automatic 3-d segmentation of internal structures of the head in mr images using a combination of similarity and free-form transformations: Part i, methodology and validation on normal subjects. *IEEE Trans. on Medical Imaging*, 18(10):909–916, 1999.

[11] C. DeCarli, J. Moisog, D. G. M. Murphy, D. Teichberg, S. I. Rapoport, and B. Horwitz. Method of quantification of brain, ventricular, and subarachnoid csf volumes for mri images. *Journal of Computerized Tomography*, 16(2):274–284, 1992.

[12] A.P. Dempster, N.M. Laird, and D.B. Rubin. Maximum likehood from imcomplete data via the em algorithm. *J. Roy. Stat. Soc. B*, 39:1–38, 1977.

[13] J.E. Dennis and R.B. Schnabel. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. SIAM Classics in Applied Mathematics, Philadelphia, PA, 1996.

[14] R.O. Duda, P.E. Hart, and D. G. Stork. *Pattern Classification*. Wiley Interscience, New York, 2001.

[15] A.C. Evans, D.L. Collins, S.R. Mills, E.D. Brown, R.L. Kelly, and T.M. Peters. 3d statistical neuroanatomical models from 305 mri volumes. In *Proceedings IEEE Nuclear Sc. Symp., Medical Imaging Conference*, pages 1813–1817, 1993.

[16] G. Geiger and F. Girosi. Parallel and deterministic algorithms for mrfs: Surface reconstruction and integration. *IEEE Trans. on PAMI*, 12:401–412, 1991.

[17] S. Geman and D. Geman. Stochastic relaxation, gibbs distributions and the bayesian restoration of images. *IEEE Trans. PAMI*, 6:721–741, 1984.

[18] G. Gerig, J. Martin, R. Kikinis, O. Kubler, M. Shenton, and F. Jolesz. Automatic segmentation of dual-echo mr head data. In *XII Intl. Conf. on IPMI*, pages 175–187, 1991.

[19] R. L. Gilmore, M. D. Childress, C. M. Leonard, R. Quisling, S. Roper, and S. Eisehschenk. Hippocampal volumetrics in patients with temporal lobe and extemporal lobe epilepsy. *Neurology*, 44(A350), 1994.

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

[21] K. Held, E.R. Kopsa, B.J. Krause, W.M. Wells, R. Kikinis, and H.W. M ller-Gartner. Markov random field segmentation of brain mr images. *IEEE Trans. on Med. Imaging*, 16(6):878–886, 1997.

[22] T. Hofmann, J. Puzicha, and J. M. Buhmann. Unsupervised texture segmentation in a deterministic annealing framework. *IEEE Trans. on PAMI*, 20(8):803–818, 1998.

[23] ISBR. Internet brain segmentation repository. Technical report, Massachusetts General Hospital, Center for Morphometric Analysis. Available: http://neuro-www.mgh.harvard.edu/cma/ibsr/, 2000.

[24] M. Joliot and B. M. Majoyer. Three dimensional segmentation and interpolation of magnetic resonance brain images. *IEEE Trans. on Medical Imaging*, 12(2):269–277, 1993.

[25] S. Joshi, A. Banerjee, G. E. Christensen, J. G. Csenansky, J. W. Haller, M. I. Miller, and L. Wang. Gaussian random fields on sub-manifolds for characterizing brain surfaces. In *Proc. of XVth Intl Conf. on Info. Processing in Medical Imaging*, pages 381–386. Springer, 1997.

[26] T. Kapur, W.E.L.Grimson, W. M. Wells III, and R. Kikinis. Segmentation of brain tissue from magnetic resonance images. *Med. Image Aanalysis*, 1(2):109–127, 1996.

[27] S. Kirkpatrick, C.D. Gelatt, and M.P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.

[28] R.K.S. Kwan, A.C. Evans, and G.B. Pike. An extensible mri simulator for post–processing evaluation. In *Visualization in Biomedical Computing (VBC'96). Lecture Notes in Computer Science, Vol. 1131*, pages 135–140. Springer Verlag, 1996.

[29] R. Leahy, T. Herbert, and R. Lee. Application of markov random fields in medical imaging. In *IPMI*, pages 1–14, 1989.

[30] K. V. Leemput, F. Maes, D. Vandermeulen, and P. Suetens. Automated model-based bias field correction of mr images of the brain. *IEEE Trans. on Med. Imaging*, 18(10):885–896, 1999.

[31] K. V. Leemput, F. Maes, D. Vandermeulen, and P. Suetens. Automated model-based tissue classification of mr images of the brain. *IEEE Trans. on Med. Imaging*, 18(10):897–908, 1999.

[32] J. Marroquin, S. Mitter, and T. Poggio. Probabilistic solution of ill–posed problems in computational vision. *J. Am. Stat. Asoc.*, 82:76–89, 1987.

[33] J. Marroquin, F. Velasco, M. Rivera, and M. Nakamura. Gauss-markov measure field models for low-level vision. *IEEE Trans. PAMI*, 23:337–348, 2001.

[34] T. N. Pappas. An adaptive clustering algorithm for image segmentation. *IEEE transactions on Signal Processing*, 40(4):901–914, 1992.

[35] J.C. Rajapakse and F. Krugge. Segmentation of mr images with intensity inhomogeneities. *Image and Vision Computing*, 16:165–180, 1998.

[36] S. Sandor and R. Leahy. A 3d morphological algorithm for automated labeling of the cortex in magnetic resonance brain images. In *AAAI Spring Symp. on Applications of Compu. Vision in Medical Image Processing*, 1994.

[37] T. Schiemann, U. Tiede, and K. H. Hohne. Segmentation of visible human for high quality volume-based visualization. *Medical Image Analysis*, 1(4):263–270, 1996.

[38] J.G. Sled, A.P. Zijdenbos, and A.C. Evans. A non–parametric method for automatic correction of intensity nonuniformity in mri data. *IEEE Trans. on Medical Imaging*, 17(1):87–97, 1998.

[39] J. P. Thirion. Image matching as a diffusion process: an analogy with maxwell's demons. *Medical Image Analysis*, 2(3):243–260, 1998.

[40] A. Trouve. Diffeomorphisms groups and pattern matching in image analysis. *International Journal of Computer Vision*, 28(3):213–221, 1998.

[41] B. C. Vemuri, J. Ye, Y. Chenand, and C. M. Leonard. A level-set based approach to image registration. In *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pages 86–93, June 2000.

[42] W.M. Wells, W.E.L.Grimson, R. Kikinis, and F. A. Jolesz. Adaptive segmentation of mri data. *IEEE Trans. on Med. Imaging*, 15:429–442, Aug. 1996.

[43] A.L. Yuille and Kosowsky. Statistical physics algorithms that converge. *Neural Computation*, 6(3):341–356, 1994.

[44] J. Zhang. Mean field theory in em procedures for mrf's. *IEEE Trans. on Sig. Proc.*, 40:2570–2583, 1992.

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

[46] O.C. Zienkiewicz and R.L. Taylor. *The Finite Element Method*. McGraw Hill, 1989.

Pulse Sequence	T1	T2	PD
Scan Technique	Spoiled Flash	Dual Echo Spin Echo, Late Echo	Dual Echo Spin Echo, Early Echo
Rep. Time TR (ms)	18	3300	3300
Echo Time TE (ms)	10	35, 120	35, 120
Flip Angle (deg.)	30	90	90

Table 1: Technical data for the simulated MRI used in the experiments

	CSF	GRAY	WHITE
Default	.8899	.8742	.8894
$\lambda = 0.0$	.8863	.8693	.8849
$\lambda = 0.005$	.8854	.8718	.8868
$\lambda = .1$	.8846	.8693	.8877
$\eta = 100$	.8867	.8747	.8919
$\eta = 10000$	.8798	.8595	.8742
$\eta = \infty$	.8759	.8545	.8698
$\alpha = 0$	.8759	.8428	.8788
$\alpha = .05$	.8846	.8907	.8854
$\alpha = .3$	.8771	.8826	.8852
$\alpha = .4$	.8523	.8871	.8753
$\alpha = .5$	.8174	.8790	.8616
$\alpha = .6$	.7662	.8675	.8467
$\gamma = 0.5$	.8900	.8683	.8809
$\gamma = 1.5$	.8866	.8662	.8872

Table 2: Sensitivity of the performance indices to variations in the value of the control parameters for simulated MRI.

Method	GM	WM
adaptive MAP	0.564	0.567
biased MAP	0.558	0.562
fuzzy c-means	0.473	0.567
Maximum a Posteriori Probability (MAP)	0.550	0.554
tree-structure k-means	0.477	0.571
Maximum-Likelihood	0.535	0.551
MPM-MAP	0.662	0.683
Manual (4 brains averaged over 2 experts)	0.876	0.832

Table 3: Average Overlap (performance indices) between manually-guided segmentations and various methods for the 20 brain scans of the MGH data base.

	CSF	GM	WM	Seg. Time (min.)
Complete Procedure	.227	.662	.683	19.2
No prior Probs.	.168	.579	.654	19.3
Std. Initialization	.211	.648	.644	32.8
No prior MRF	.227	.659	.680	19.2
Constant intensity mod.	.220	.645	.650	10.3
No refinement in regist.	.248	.648	.684	18.3

Table 4: Effect of different features of our algorithm on the average performance indices and on the average processing time for the 20 brain scans of the MGH data base.

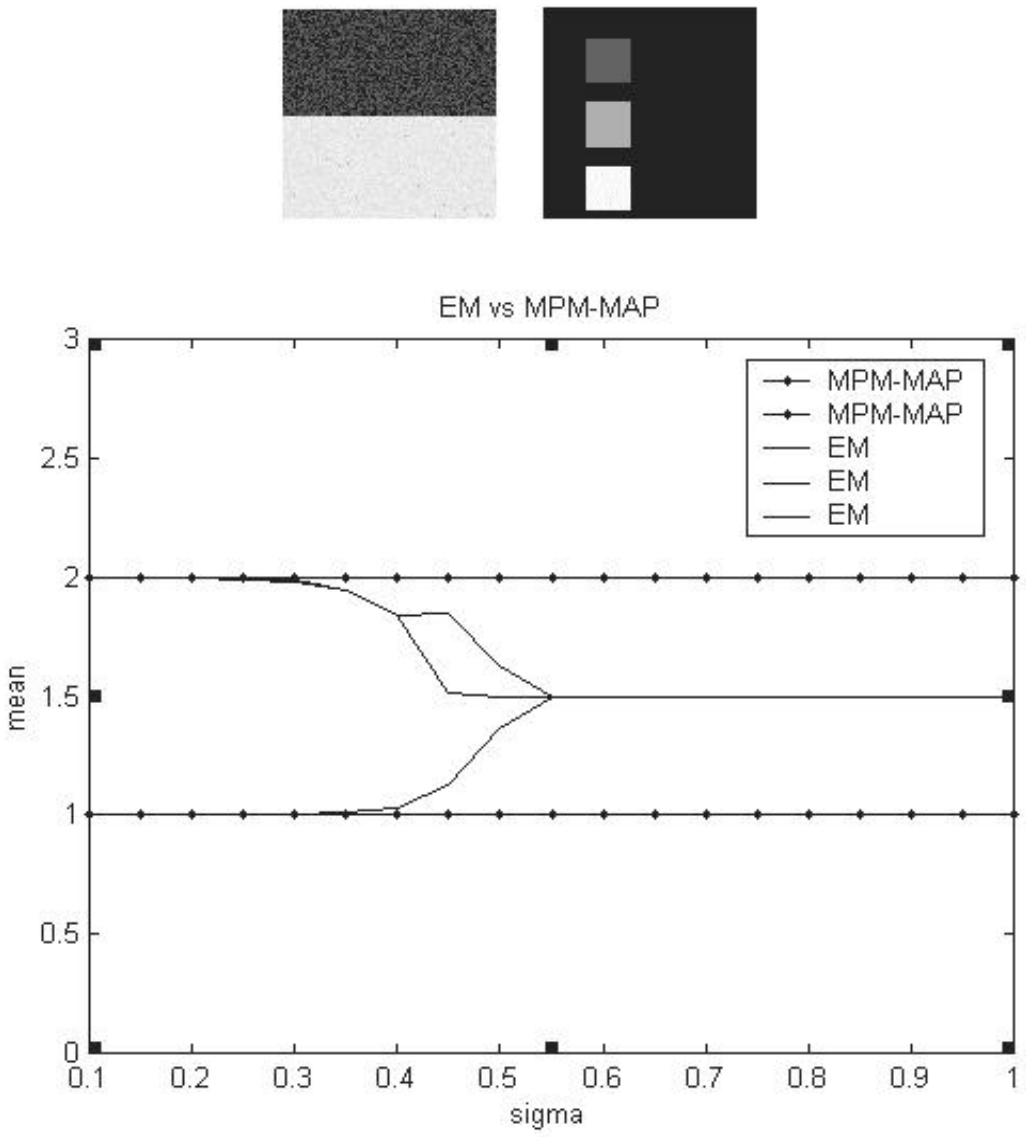

Figure 1: Comparison of EM and MPM-MAP algorithms with respect to errors in the estimation of the posterior marginals

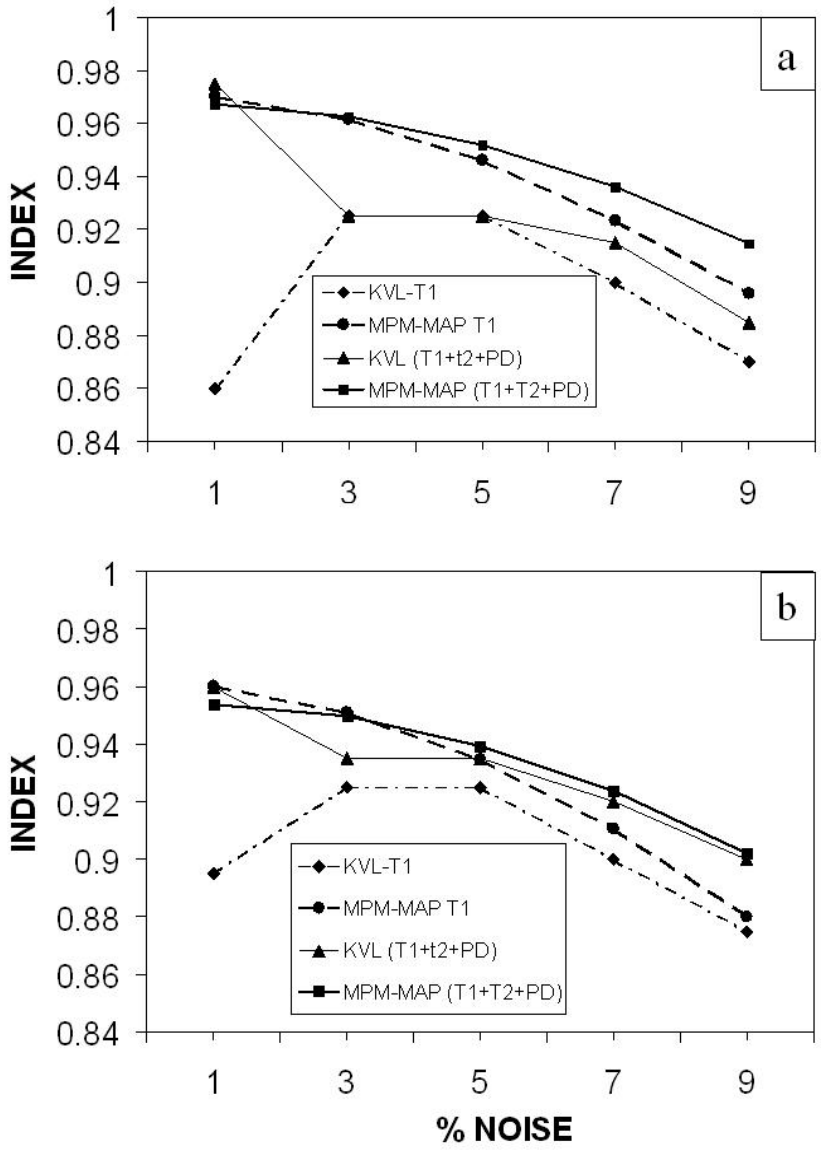

Figure 2: Performance indices for the segmentation of MRI simulated data (from the brainweb [8]) with 0% spatial non-homogeneities, using our algorithm and the one in [31] (labeled KVL) (a) for WM and (b) GM

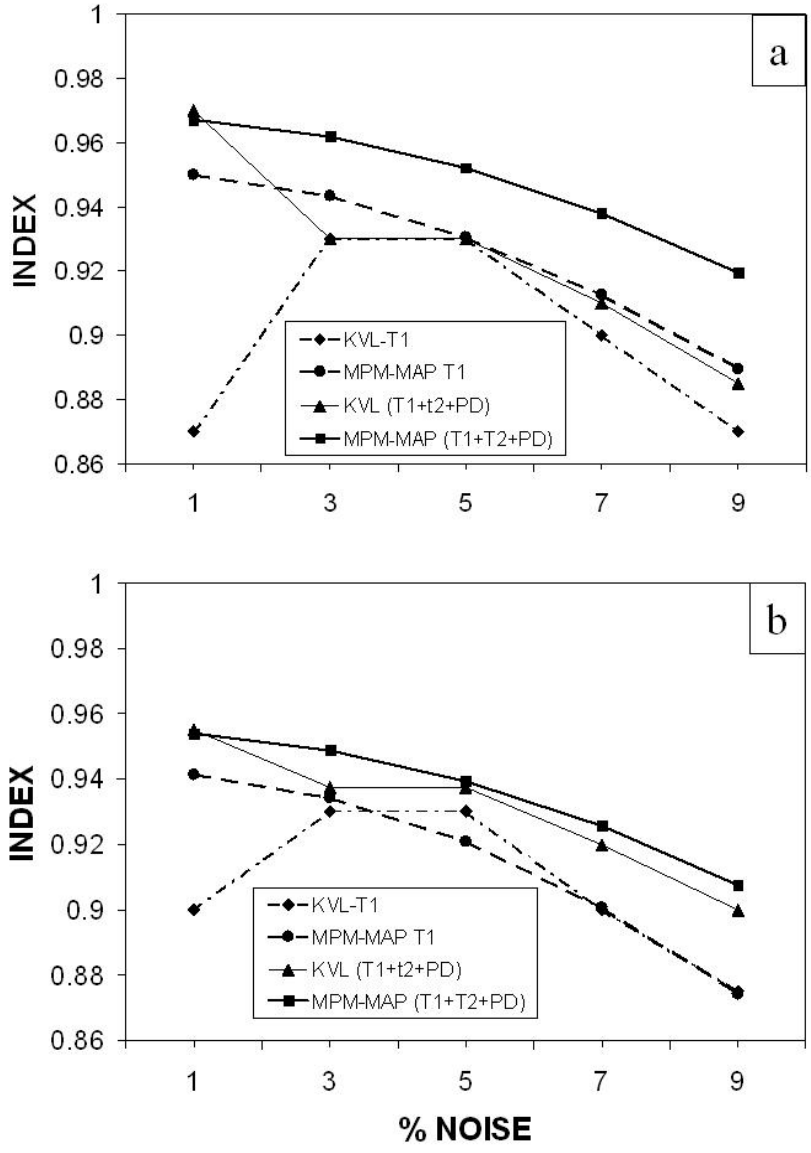

Figure 3: Performance indices for the segmentation of MRI simulated data (from the brainweb [8]) with 40% spatial non-homogeneities, using our algorithm and the one in [31] (labeled KVL) (a) for WM and (b) GM

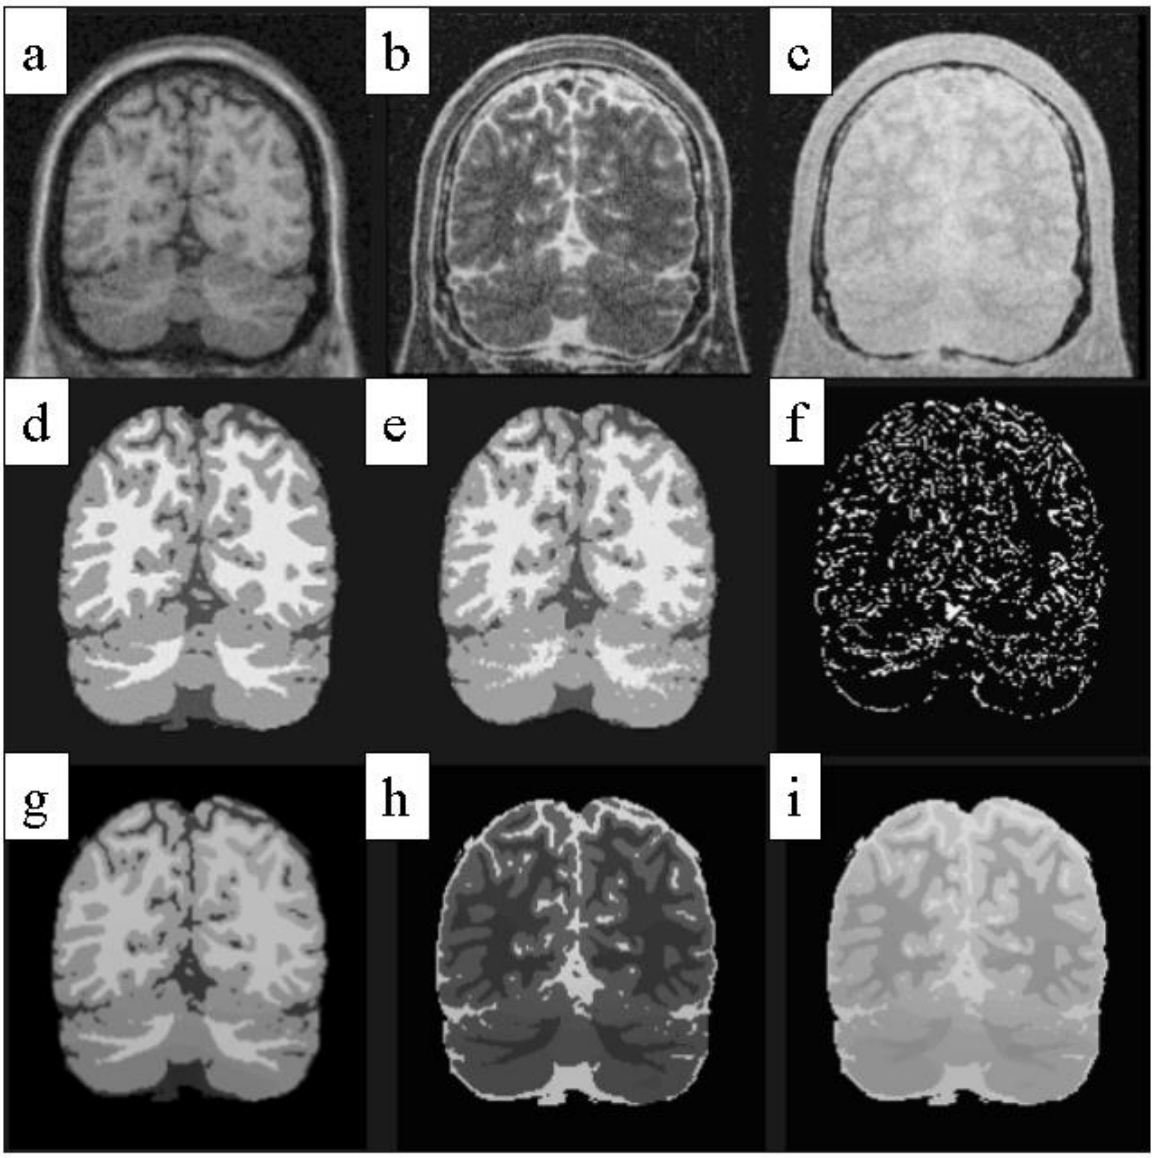

Figure 4: (a-c): T1, T2 and PD weighted simulated images from Brainweb database, with 9% noise and 40% spatial inhomogeneity. (d) Anatomical model (ground truth). (e) MPM-MAP segmentation (f) MPM-MAP segmentation errors (white pixels) (g-i) Estimated intensities (membrane models) for the images in panels (a-c)

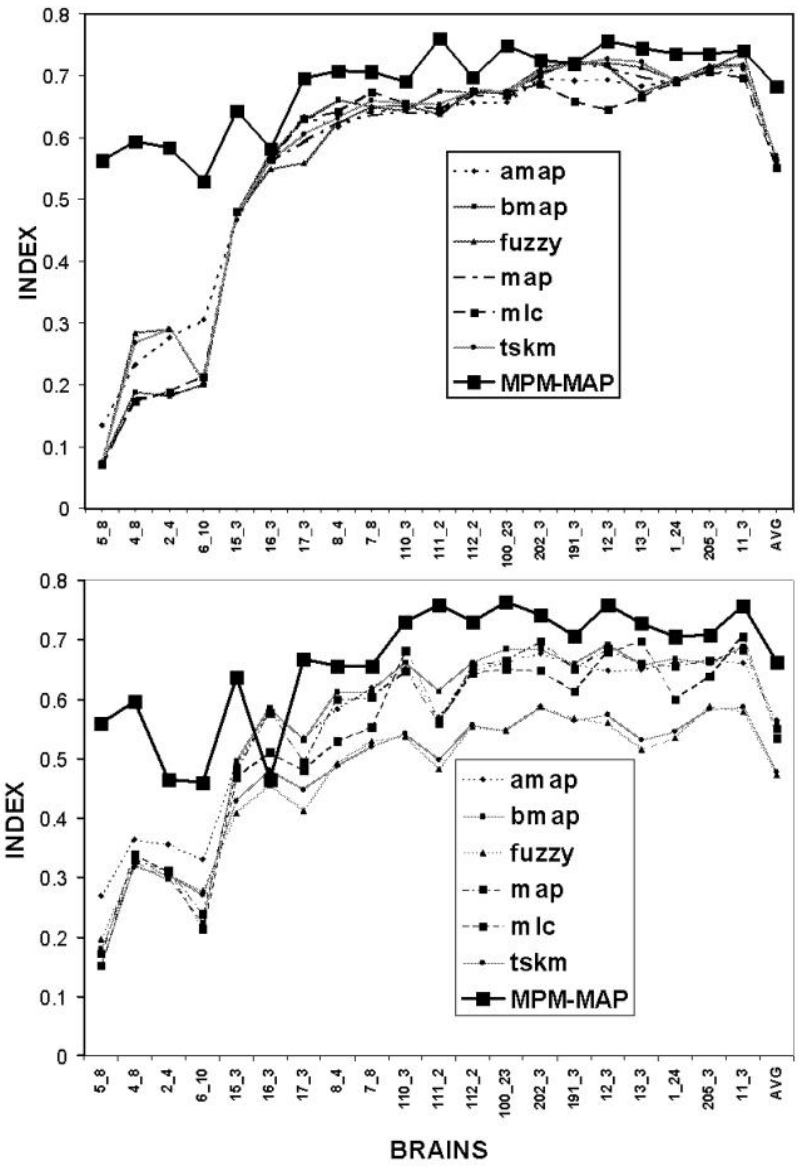

Figure 5: Performance indices for the individual brains of the MGH database for different segmentation methods. The bold line corresponds to our method. Above: WM; below: GM.

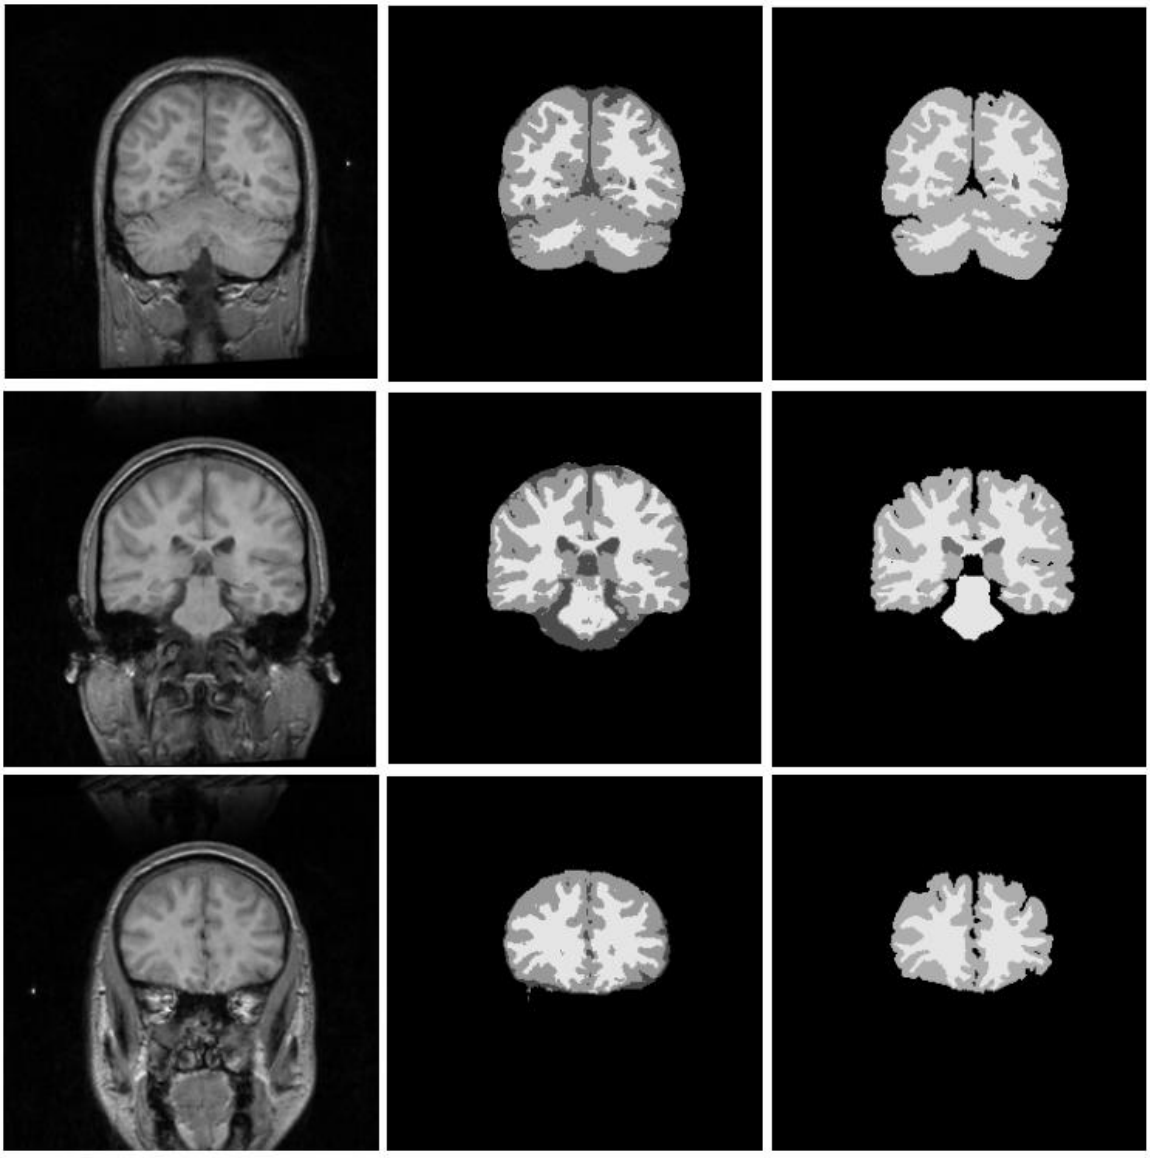

Figure 6: Sample images from the MGH database. On each row: left: original image; center: MPM-MAP automatic segmentation; right: expert hand-guided segmentation (see text for details).