Chapter 5

Image Segmentation

John Ashburner & Karl J. Friston
The Wellcome Dept. of Imaging Neuroscience,
12 Queen Square, London WC1N 3BG, UK.

Contents

Abstract

This chapter describes a method of segmenting MR images into different tissue classes, using a modified Gaussian Mixture Model. By knowing the prior spatial probability of each voxel being grey matter, white matter or cerebro-spinal fluid, it is possible to obtain a more robust classification. In addition, a step for correcting intensity non-uniformity is also included, which makes the method more applicable to images corrupted by smooth intensity variations.

5.1 Introduction

Healthy brain tissue can generally be classified into three broad tissue types on the basis of an MR image. These are grey matter (GM), white matter (WM) and cerebro-spinal fluid (CSF). This classification can be performed manually on a good quality T1 image, by simply selecting suitable

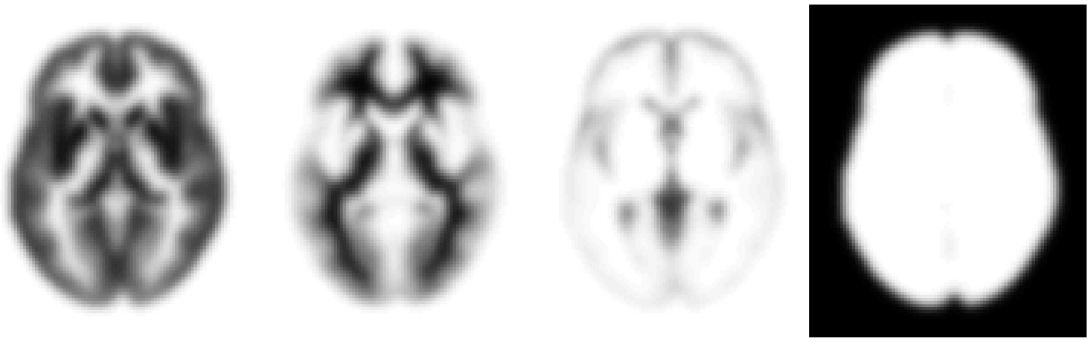

Figure 5.1: The *a priori* probability images of GM, WM, CSF and non-brain tissue. Values range between zero (white) and one (black).

image intensity ranges which encompass most of the voxel intensities of a particular tissue type. However, this manual selection of thresholds is highly subjective.

Some groups have used clustering algorithms to partition MR images into different tissue types, either using images acquired from a single MR sequence, or by combining information from two or more registered images acquired using different scanning sequences or echo times (eg. proton-density and T2-weighted). The approach described here is a version of the 'mixture model' clustering algorithm [8], which has been extended to include spatial maps of prior belonging probabilities, and also a correction for image intensity non-uniformity that arises for many reasons in MR imaging. Because the tissue classification is based on voxel intensities, partitions derived without the correction can be confounded by these smooth intensity variations.

The model assumes that the MR image (or images) consists of a number of distinct tissue types (clusters) from which every voxel has been drawn. The intensities of voxels belonging to each of these clusters conform to a normal distribution, which can be described by a mean, a variance and the number of voxels belonging to the distribution. For multi-spectral data (e.g. simultaneous segmentation of registered T2 and PD images), multivariate normal distributions can be used. In addition, the model has approximate knowledge of the spatial distributions of these clusters, in the form of prior probability images.

Before using the current method for classifying an image, the image has to be in register with the prior probability images. The registration is normally achieved by least squares matching with template images in the same stereotaxic space as the prior probability images. This can be done using nonlinear warping, but the examples provided in this chapter were done using affine registration (see Chapter 3).

One of the greatest problems faced by tissue classification techniques is non-uniformity of the images intensity. Many groups have developed methods for correcting intensity non-uniformities, and the scheme developed here shares common features. There are two basic models describing image noise properties: multiplicative noise and additive noise. The multiplicative model describes images that have noise added before being modulated by the non-uniformity field (i.e., the standard deviation of the noise is multiplied by the modulating field), whereas the additive version models noise that is added after the modulation (standard deviation is constant). The current method uses a multiplicative noise model, which assumes that the errors originate from tissue variability rather than additive Gaussian noise from the scanner. Figure 5.2 illustrates the model used by the classification.

Non-uniformity correction methods all involve estimating a smooth function that modulates

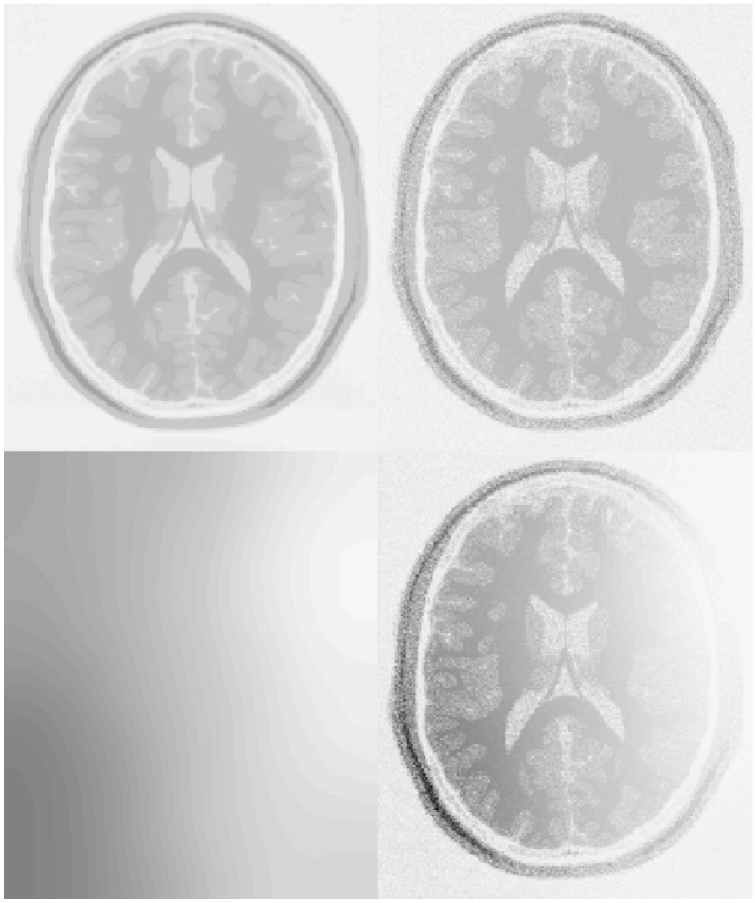

Figure 5.2: The MR images are modeled as a number of distinct clusters (top left), with different levels of Gaussian random noise added to each cluster (top right). The intensity modulation is assumed to be smoothly varying (bottom left), and is applied as a straightforward multiplication of the modulation field with the image (bottom right).

the image intensities. If the function is is not forced to be smooth, then it will begin to fit the higher frequency intensity variations due to different tissue types, rather than the low frequency intensity non-uniformity artifact. Spline [17, 13] and polynomial [14, 15] basis functions are widely used for modeling the intensity variation. In these models, the higher frequency intensity variations are restricted by limiting the number of basis functions. In the current method, a Bayesian model is used, where it is assumed that the modulation field ($\mathbf{U}$) has been drawn from a population for which the *a priori* probability distribution is known, thus allowing high frequency variations of the modulation field to be penalized.

5.2 Methods

The explanation of the tissue classification algorithm will be simplified by describing its application to a single two dimensional image. A number of assumptions are made by the classification model. The first is that each of the $I \times J$ voxels of the image ($\mathbf{F}$) has been drawn from a known number (K) of distinct tissue classes (clusters). The distribution of the voxel intensities within each class is normal (or multi-normal for multi-spectral images) and initially unknown. The distribution of voxel intensities within cluster k is described by the number of voxels within the cluster (h_k), the mean for that cluster (v_k), and the variance around that mean (c_k).

Because the images are matched to a particular stereotaxic space, prior probabilities of the voxels belonging to the grey matter (GM), white matter (WM) and cerebro-spinal fluid (CSF) classes are known. This information is in the form of probability images – provided by the Montréal Neurological Institute [5, 4, 6] as part of the ICBM, NIH P-20 project (Principal Investigator John Mazziotta), and derived from scans of 152 young healthy subjects. These probability images contain values in the range of zero to one, representing the prior probability of a voxel being either GM, WM or CSF after an image has been normalized to the same space (see Figure 5.1). The probability of a voxel at co-ordinate i, j belonging to cluster k is denoted by b_{ijk}[1].

The final assumption is that the intensity and noise associated with each voxel in the image has been modulated by multiplication with an unknown smooth scalar field.

There are many unknown parameters to be determined by the classification algorithm, and estimating any of these requires knowledge of the others. Estimating the parameters that describe a cluster (h_k, v_k and c_k) relies on knowing which voxels belong to the cluster, and also the form of the intensity modulating function. Estimating which voxels should be assigned to each cluster requires the cluster parameters to be defined, and also the modulation field. In turn, estimating the modulation field needs the cluster parameters and the belonging probabilities.

The problem requires an iterative algorithm (see Figure 5.3). It begins by assigning starting estimates for the various parameters. The starting estimate for the modulation field is typically uniformly one. Starting estimates for the belonging probabilities of the GM, WM and CSF partitions are based on the prior probability images. Since there are no prior probability maps for background and non-brain tissue clusters, they are estimated by subtracting the prior probabilities for GM, WM and CSF from a map of all ones, and dividing the result equally between the remaining clusters [2].

Each iteration of the algorithm involves estimating the cluster parameters from the non-

[1] Note that ij subscripts are used for voxels rather than the single subscripts used in the previous chapters. This is to facilitate the explanation of how the modulation field is estimated for 2D images as described in Section 5.2.3.

[2] Where identical prior probability maps are used for more than one cluster, the affected cluster parameters need to be modified so that separate clusters can be characterised. This is typically done after the first iteration, by assigning different values for the means uniformly spaced between zero and the intensity of the white matter cluster.

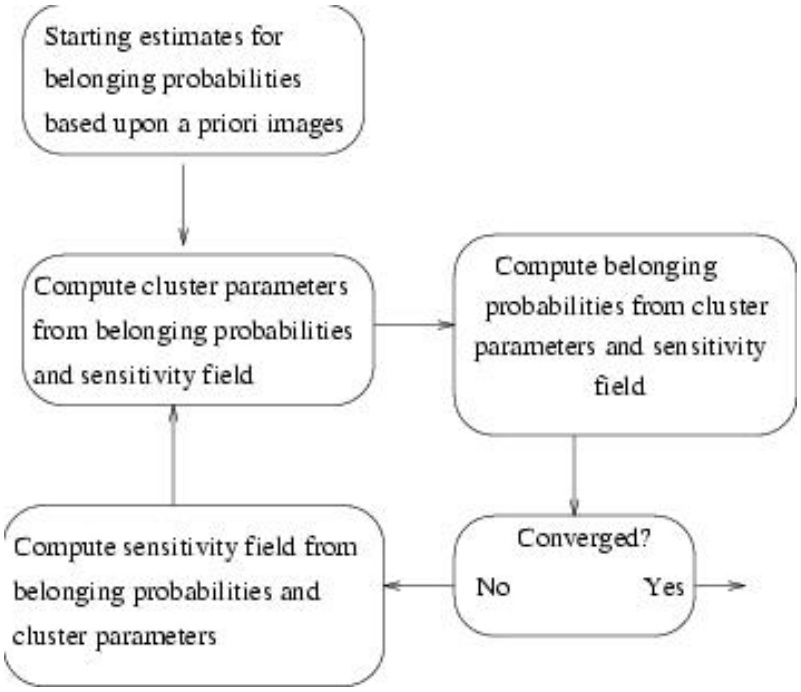

Figure 5.3: A flow diagram for the tissue classification.

uniformity corrected image, assigning belonging probabilities based on the cluster parameters, checking for convergence, and re-estimating and applying the modulation function. With each iteration, the parameters describing the distributions move towards a better fit and the belonging probabilities ($\mathbf{P}$) change slightly to reflect the new distributions. This continues until a convergence criterion is satisfied. The parameters describing clusters with corresponding prior probability images tend to converge more rapidly than the others. This may be partly due to the better starting estimates. The final values for the belonging probabilities are in the range of 0 to 1, although most values tend to stabilize very close to one of the two extremes. The algorithm is in fact an *expectation maximization* (EM) approach, where the *E-step* is the computation of the belonging probabilities, and the *M-step* is the computation of the cluster and non-uniformity correction parameters. The individual steps involved in each iteration are now described in more detail.

5.2.1 Estimating the Cluster Parameters

This stage requires the most recent estimate of the modulation function ($\mathbf{U}$, where u_{ij} is the multiplicative correction at voxel i, j), and the current estimate of the probability of voxel i, j belonging to class k, which is denoted by p_{ijk}. The first step is to compute the number of voxels ($\mathbf{h}$) belonging to each of the K clusters as:

$$h_k = \sum_{i=1}^{I} \sum_{j=1}^{J} p_{ijk} \text{ over } k = 1..K.$$

Mean voxel intensities for each cluster ($\mathbf{v}$) are computed. This step effectively produces a weighted mean of the image voxels, where the weights are the current belonging probability

estimates:

$$v_k = \frac{\sum_{i=1}^{I}\sum_{j=1}^{J} p_{ijk} f_{ij} u_{ij}}{h_k} \text{ over } k = 1..K.$$

Then the variance of each cluster ($\mathbf{c}$) is computed in a similar way to the mean:

$$c_k = \frac{\sum_{i=1}^{I}\sum_{j=1}^{J} p_{ijk} (f_{ij} u_{ij} - v_k)^2}{h_k} \text{ over } k = 1..K.$$

5.2.2 Assigning Belonging Probabilities

The next step is to re-calculate the belonging probabilities. It uses the cluster parameters computed in the previous step, along with the prior probability images and the intensity modulated input image. Bayes rule is used to assign the probability of each voxel belonging to each cluster:

$$p_{ijk} = \frac{r_{ijk} s_{ijk}}{\sum_{l=1}^{K} r_{ijl} s_{ijl}} \text{ over } i = 1..I,\ j = 1..J \text{ and } k = 1..K.$$

where p_{ijk} is the *a posteriori* probability that voxel i, j belongs to cluster k given its intensity of f_{ij}, r_{ijk} is the likelihood of a voxel in cluster k having an intensity of f_{ik}, and s_{ijk} is the *a priori* probability of voxel i, j belonging in cluster k.

The likelihood function is obtained by evaluating the probability density functions for the clusters at each of the voxels:

$$r_{ijk} = \frac{u_{ij}}{(2\pi c_k)^{1/2}} \exp\left(\frac{-(f_{ij} u_{ij} - v_k)^2}{2c_k}\right) \text{ over } i = 1..I,\ j = 1..J \text{ and } k = 1..K.$$

The prior (s_{ijk}) is based on two factors: the number of voxels currently belonging to each cluster (h_k), and the prior probability images derived from a number of images (b_{ijk}). With no knowledge of the spatial prior probability distribution of the clusters or the intensity of a voxel, then the *a priori* probability of any voxel belonging to a particular cluster is proportional to the number of voxels currently included in that cluster. However, with the additional data from the prior probability images, a better estimate for the priors can be obtained:

$$s_{ijk} = \frac{h_k b_{ijk}}{\sum_{l=1}^{I}\sum_{m=1}^{J} b_{lmk}} \text{ over } i = 1..I,\ j = 1..J \text{ and } k = 1..K.$$

Convergence is ascertained by following the log-likelihood function:

$$\sum_{i=1}^{I}\sum_{j=1}^{J} \log\left(\sum_{k=1}^{K} r_{ijk} s_{ijk}\right)$$

The algorithm is terminated when the change in log-likelihood from the previous iteration becomes negligible.

5.2.3 Estimating the Modulation Function

To reduce the number of parameters describing an intensity modulation field, it is modeled by a linear combination of low frequency discrete cosine transform (DCT) basis functions (see

Section **??**), which were chosen because there are no constraints at the boundary. A two (or three) dimensional discrete cosine transform (DCT) is performed as a series of one dimensional transforms, which are simply multiplications with the DCT matrix. The elements of a matrix ($\mathbf{D}$) for computing the first M coefficients of the one dimensional DCT of a vector of length I is given by:

$$\begin{aligned} d_{i1} &= \frac{1}{\sqrt{I}} \; i = 1..I \\ d_{im} &= \sqrt{\frac{2}{I}} \cos\left(\frac{\pi(2i-1)(m-1)}{2I}\right) \; i = 1..I, m = 2..M \end{aligned} \tag{5.1}$$

The matrix notation for computing the first $M \times N$ coefficients of the two dimensional DCT of a modulation field $\mathbf{U}$ is $\mathbf{Q} = \mathbf{D_1}^T \mathbf{U} \mathbf{D_2}$, where the dimensions of the DCT matrices $\mathbf{D_1}$ and $\mathbf{D_2}$ are $I \times M$ and $J \times N$ respectively, and $\mathbf{U}$ is an $I \times J$ matrix. The approximate inverse DCT is computed by $\mathbf{U} \simeq \mathbf{D_1} \mathbf{Q} \mathbf{D_2}^T$. An alternative representation of the two dimensional DCT is obtained by reshaping the $I \times J$ matrix $\mathbf{U}$ so that it is a vector ($\mathbf{u}$). Element $i + (j-1) \times I$ of the vector is then equal to element i, j of the matrix. The two dimensional DCT can then be represented by $\mathbf{q} = \mathbf{D}^T \mathbf{u}$, where $\mathbf{D} = \mathbf{D_2} \otimes \mathbf{D_1}$ (the Kronecker tensor product of $\mathbf{D_2}$ and $\mathbf{D_1}$), and $\mathbf{u} \simeq \mathbf{D}\mathbf{q}$.

The sensitivity correction field is computed by re-estimating the coefficients ($\mathbf{q}$) of the DCT basis functions such that the product of the likelihood and a prior probability of the parameters is increased. This can be formulated as an iteration of a Gauss-Newton optimisation algorithm (compare with Section **??**):

$$\mathbf{q}^{(n+1)} = \left(\mathbf{C_0}^{-1} + \mathbf{A}\right)^{-1} \left(\mathbf{C_0}^{-1} \mathbf{q_0} + \mathbf{A}\mathbf{q}^{(n)} - \mathbf{b}\right) \tag{5.2}$$

where $\mathbf{q_0}$ and $\mathbf{C_0}$ are the means and covariance matrices describing the *a priori* probability distribution of the coefficients. Vector $\mathbf{b}$ contains the first derivatives of the log-likelihood cost function with respect to the basis function coefficients, and matrix $\mathbf{A}$ contains the second derivatives of the log-likelihood. These can be constructed efficiently using the properties of Kronecker tensor products (see Figure **??** in Chapter 3):

$$\begin{aligned} b_{l_1} &= \sum_{j=1}^{J} d_{2jn_1} \sum_{i=1}^{I} d_{1im_1} \left(-u_{ij}^{-1} + f_{ij} \sum_{k=1}^{K} \frac{p_{ijk}(f_{ij}u_{ij} - v_k)}{c_k}\right) \\ A_{l_1 l_2} &= \sum_{j=1}^{J} d_{2jn_1} d_{2jn_2} \sum_{i=1}^{I} d_{1im_1} d_{1im_2} \left(u_{ij}^{-2} + f_{ij}^2 \sum_{k=1}^{K} \frac{p_{ijk}}{c_k}\right) \end{aligned}$$

where $l_1 = m_1 + M(n_1 - 1)$ and $l_2 = m_2 + M(n_2 - 1)$.

Once the coefficients have been re-estimated, then the modulation field $\mathbf{U}$ can be computed from the estimated coefficients ($\mathbf{Q}$) and the basis functions ($\mathbf{D_1}$ and $\mathbf{D_2}$).

$$u_{ij} = \sum_{n=1}^{N} \sum_{m=1}^{M} d_{2jn} q_{mn} d_{1im} \text{ over } i = 1..I \text{ and } j = 1..J.$$

The Prior Probability Distribution

In Eqn. 5.2, $\mathbf{q_0}$ and $\mathbf{C_0}$ represent a multi-normal *a priori* probability distribution for the basis function coefficients. The mean of the prior probability distribution is such that it would generate a field that is uniformly one. For this, all the elements of the mean vector are set to zero, apart from the first element that is set to $\sqrt{IJ}$.

The covariance matrix $\mathbf{C_0}$ is such that $(\mathbf{q} - \mathbf{q_0})^T \mathbf{C_0}^{-1} (\mathbf{q} - \mathbf{q_0})$ produces an "energy" term that penalizes modulation fields that would be unlikely *a priori*. There are many possible forms

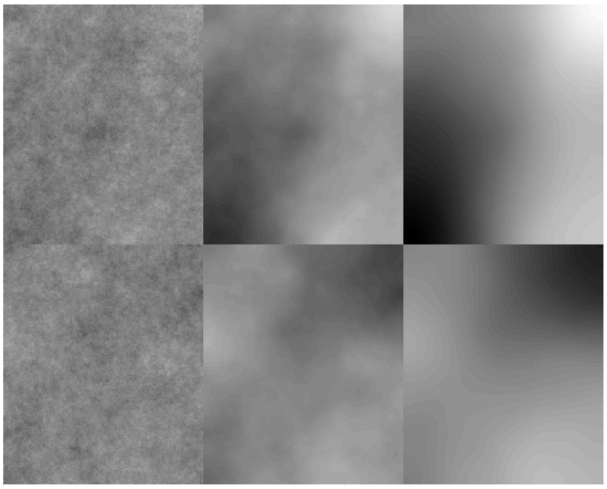

Figure 5.4: Randomly generated modulation fields generated using the membrane energy cost function (left), the bending energy cost function (center) and the squares of the third derivatives (right). These can be referred to as 1st, 2nd and 3rd order regularization.

for this penalty function (see Section **??**). Some widely used simple penalty functions include the "membrane energy" and the "bending energy", which (in three dimensions) have the forms $h = \sum_i \sum_{j=1}^{3} \lambda \left(\frac{\partial u(\mathbf{x}_i)}{\partial x_{ji}} \right)^2$ and $h = \sum_i \sum_{j=1}^{3} \sum_{k=1}^{3} \lambda \left(\frac{\partial^2 u(\mathbf{x}_i)}{\partial x_{ji} \partial x_{ki}} \right)^2$ respectively. In these formulae, $\frac{\partial u(\mathbf{x}_i)}{\partial x_{ji}}$ is the gradient of the modulating function at the ith voxel in the jth orthogonal direction, and λ is a user assigned constant. However, for the purpose of modulating the images, a smoother cost function is used that is based on the squares of the third derivatives (third order regularization):

$$h = \sum_i \sum_{j=1}^{3} \sum_{k=1}^{3} \sum_{l=1}^{3} \lambda \left(\frac{\partial^3 u(\mathbf{x}_i)}{\partial x_{ji} \partial x_{ki} \partial x_{li}} \right)^2$$

This model was chosen because it produces slowly varying modulation fields that can represent the variety of non-uniformity effects that are likely to be encountered in MR images (see Figure 5.4). In two dimensions it can be computed from:

$$\begin{aligned} \mathbf{C_0}^{-1} = \quad & \lambda \left(\dddot{\mathbf{D}}_{\mathbf{2}}{}^T \dddot{\mathbf{D}}_{\mathbf{2}} \right) \otimes \left(\mathbf{D_1}{}^T \mathbf{D_1} \right) + 3\lambda \left(\ddot{\mathbf{D}}_{\mathbf{2}}{}^T \ddot{\mathbf{D}}_{\mathbf{2}} \right) \otimes \left(\dot{\mathbf{D}}_{\mathbf{1}}{}^T \dot{\mathbf{D}}_{\mathbf{1}} \right) + \\ & 3\lambda \left(\dot{\mathbf{D}}_{\mathbf{2}}{}^T \dot{\mathbf{D}}_{\mathbf{2}} \right) \otimes \left(\ddot{\mathbf{D}}_{\mathbf{1}}{}^T \ddot{\mathbf{D}}_{\mathbf{1}} \right) + \lambda \left(\mathbf{D_2}{}^T \mathbf{D_2} \right) \otimes \left(\dddot{\mathbf{D}}_{\mathbf{1}}{}^T \dddot{\mathbf{D}}_{\mathbf{1}} \right) \end{aligned}$$

where the notation $\dot{\mathbf{D}}_{\mathbf{1}}$, $\ddot{\mathbf{D}}_{\mathbf{1}}$ and $\dddot{\mathbf{D}}_{\mathbf{1}}$ refer to the first, second and third derivatives (by differentiating Eqn. 5.1 with respect to i) of $\mathbf{D_1}$, and λ is a user specified hyper-parameter.

5.3 Examples

Figure 5.5 shows a single sagittal slice through six T1-weighted images. The initial registration to the prior probability images was via the 12-parameter affine transformation described in Section

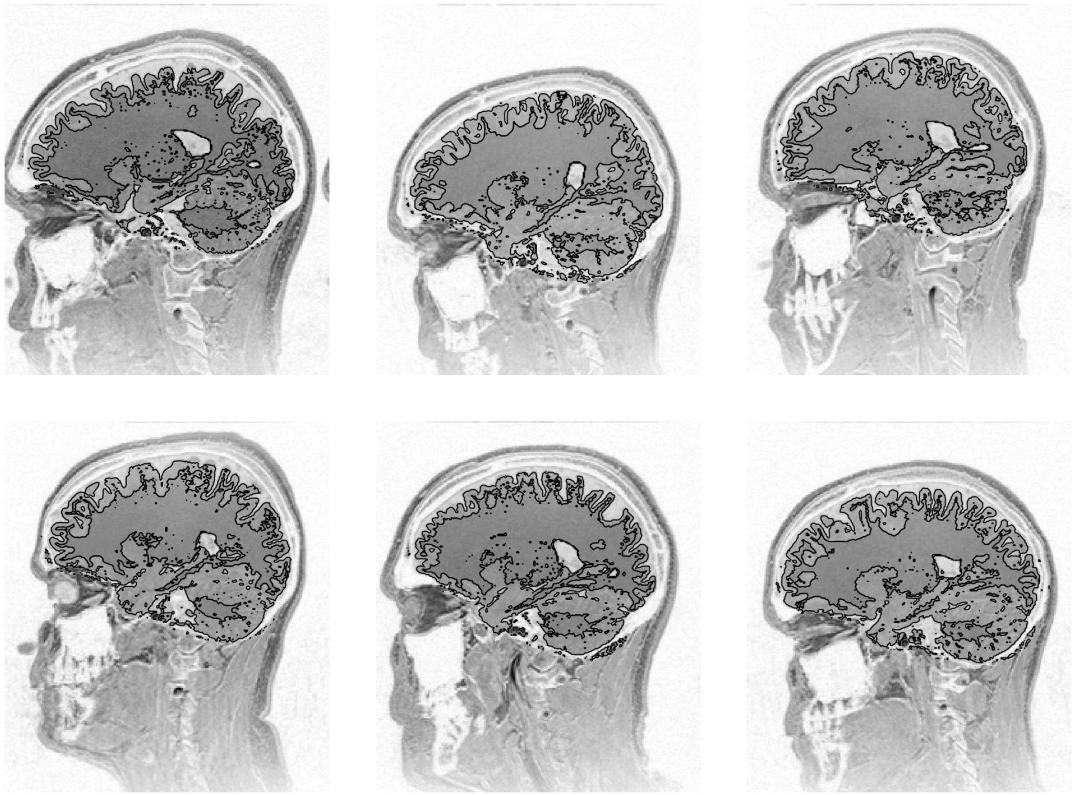

Figure 5.5: A single sagittal slice through six T1-weighted images (2 Tesla scanner, with an MPRAGE sequence, 12^o tip angle, 9.7ms repeat time, 4ms echo time and 0.6ms inversion time). Contours of extracted grey and white matter are shown superimposed on the images.

??. The images were automatically classified using the method described here, and contours of extracted grey and white matter are shown superimposed on the images.

Tissue classification methods are often evaluated using simulated images generated by the BrainWeb simulator [2, 9, 3]. It is then possible to compare the classified images with ground truth images of grey and white matter using the κ statistic (a measure of inter-rater agreement):

$$\kappa = \frac{p_o - p_e}{1 - p_e}$$

where p_o is the observed proportion of agreement, and p_e is the expected proportion of agreements by chance. If there are N observations in K categories, the observed proportional agreement is:

$$p_o = \sum_{k=1}^{K} f_{kk}/N$$

where f_{kk} is the number of agreements for the kth category. The expected proportion of agreements is given by:

$$p_e = \sum_{k=1}^{K} r_k c_k / N^2$$

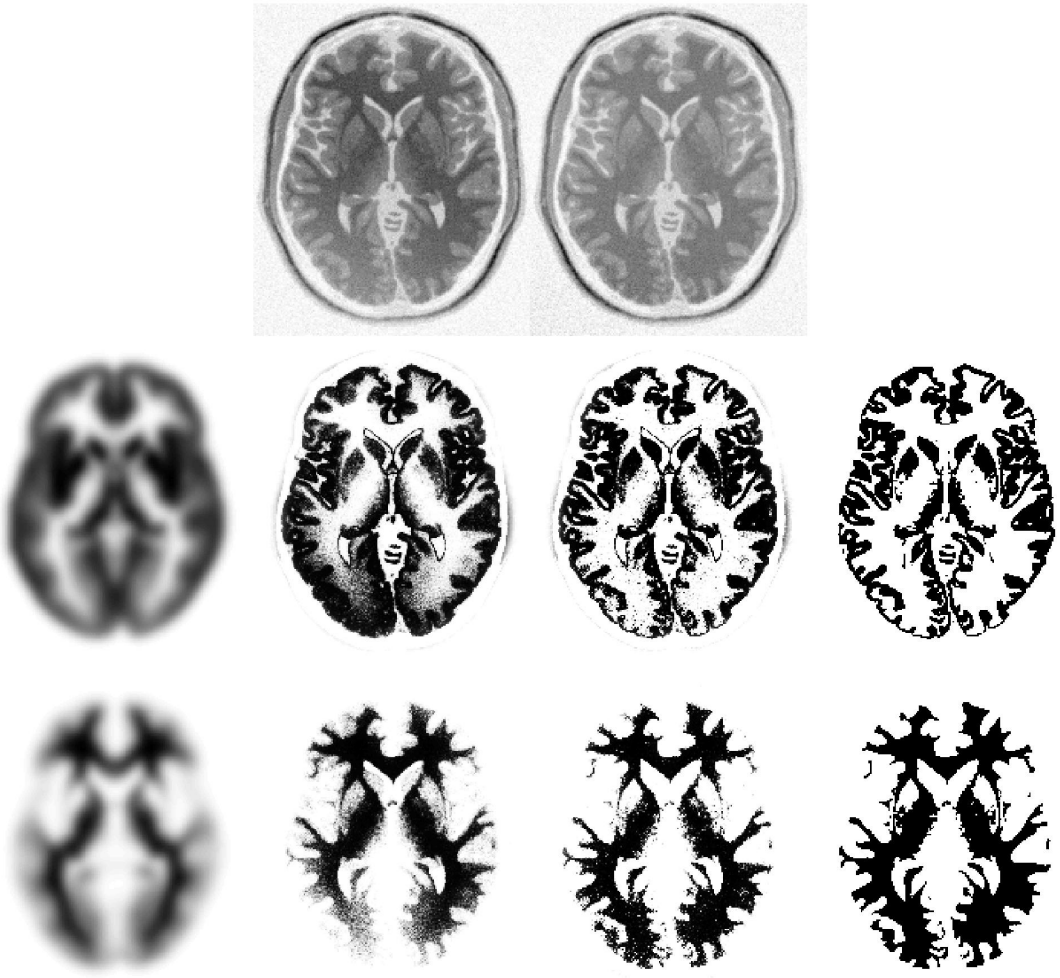

Figure 5.6: The classification of the simulated BrainWeb image. The top row shows the original simulated T1-weighted MR image with 100% non-uniformity, and the non-uniformity corrected version. From left to right, the middle row shows the *a priori* spatial distribution of grey matter used for the classification, grey matter extracted without non-uniformity correction, grey matter extracted with non-uniformity correction and the "true" distribution of grey matter (from which the simulated images were derived). The bottom row is the same as the middle, except that it shows white matter rather than grey. Without non-uniformity correction, the intensity variation causes some of the white matter in posterior areas to be classified as grey. This was also very apparent in the cerebellum because of the intensity variation in the inferior-superior direction.

where r_k and c_k are the total number of voxels in the kth class for both the "true" and estimated partitions.

The classification of a single plane of the simulated T1 weighted BrainWeb image with 100% non-uniformity is illustrated in Figure 5.6. It should be noted that no pre-processing to remove scalp or other non-brain tissue was performed on the image. In theory, the tissue classification method should produce slightly better results if this non-brain tissue is excluded from the computations. As the algorithm stands, a small amount of non-brain tissue remains in the grey matter partition, which has arisen from voxels that lie close to grey matter and have similar intensities.

5.4 Discussion

The current segmentation method is fairly robust and accurate for high quality T1 weighted images, but is not beyond improvement. Currently, each voxel is assigned a probability of belonging to a particular tissue class based only on its intensity and information from the prior probability images. There is a great deal of other knowledge that could be incorporated into the classification. For example, if all a voxel's neighbors are grey matter, then there is a high probability that it should also be grey matter. Other researchers have successfully used Markov random field models to include this information in a tissue classification model [17, 16, 15, 11, 18]. Another very simple prior, that can be incorporated, is the relative intensity of the different tissue types [7]. For example, when segmenting a T1 weighted image, it is known that the white matter should have a higher intensity than the grey matter, which in turn should be more intense than the CSF. When computing the means for each cluster, this prior information could sensibly be used to bias the estimates.

In order to function properly, the classification method requires good contrast between the different tissue types. However, many central grey matter structures have image intensities that are almost indistinguishable from that of white matter, so the tissue classification is not always very accurate in these regions. Another related problem is that of partial volume. Because the model assumes that all voxels contain only one tissue type, the voxels that contain a mixture of tissues may not be modeled correctly. In particular, those voxels at the interface between white matter and ventricles will often appear as grey matter. This can be seen to a small extent in Figures 5.5 and 5.6. Each voxel is assumed to be of only one tissue type, and not a combination of different tissues, so the model's assumptions are violated when voxels contain signal from more than one tissue type. This problem is greatest when the voxel dimensions are large, or if the images have been smoothed, and is illustrated using simulated data in Figure 5.7. The effect of partial volume is that it causes the distributions of the intensities to deviate from normal. Some authors have developed more complex models than mixtures of Gaussians to describe the intensity distributions of the classes [1]. A more recent commonly adopted approach involves modeling separate classes of partial volumed voxels [10, 11, 12].

In order for the Bayesian classification to work properly, an image volume must be in register with a set of prior probability images used to instate the priors. Figure 5.8 shows the effects of mis-registration on the accuracy of segmentation. This figure also gives an indication of how far a brain can deviate from the normal population of brains (that constitute the prior probability images) in order for it to be segmented adequately. Clearly, if the brain cannot be well registered with the probability images, then the segmentation will not be as accurate. This fact also has implications for severely abnormal brains, as they are more difficult to register with images that represent the prior probabilities of voxels belonging to different classes. Segmenting such abnormal brains can be a problem for the algorithm, as the prior probability images are based on normal healthy brains. The profile in Figure 5.8 depends on the smoothness or resolution of the prior probability images. By not smoothing the prior probability images, the segmentation would be

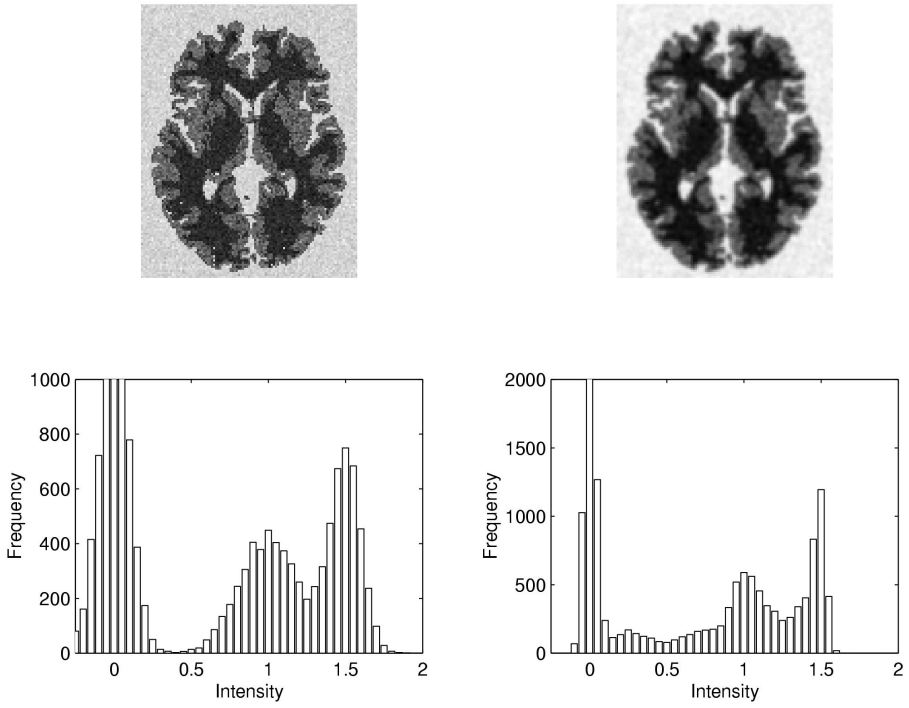

Figure 5.7: Simulated data showing the effects of partial volume on the intensity histograms. On the upper left is a simulated image consisting of three distinct clusters. The intensity histogram of this image is shown on the lower left and consists of three Gaussian distributions. The image at the top right is the simulated image after a small amount of smoothing. The corresponding intensity histogram no longer shows three distinct Gaussian distributions.

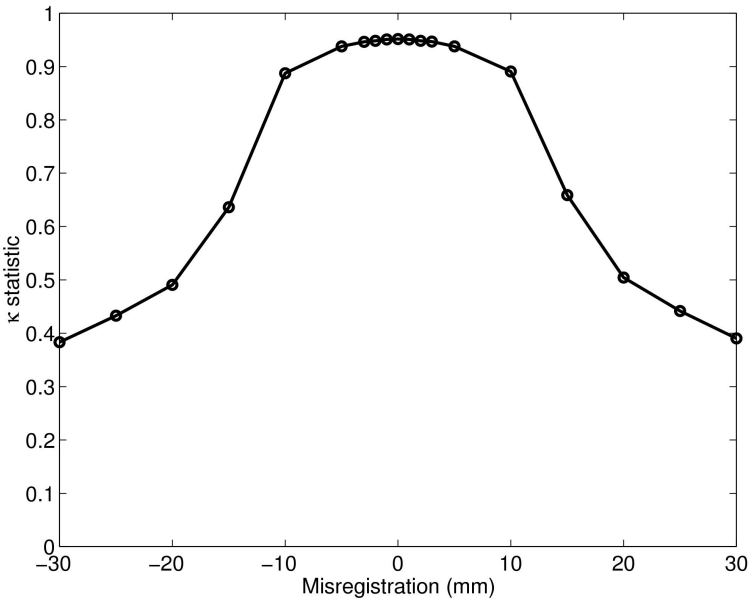

Figure 5.8: Segmentation accuracy with respect to misregistration with the prior probability images.

optimal for normal, young and healthy brains. However, these images may need to be smoother in order to encompass more variability when patient data are to be processed.

As an example, consider a subject with very large ventricles. CSF may appear where the priors suggest that tissue should always be WM. These CSF voxels are forced to be misclassified as WM, and the intensities of these voxels are incorporated into the computation of the WM means and variances. This results in the WM being characterized by a very broad distribution, so the algorithm is unable to distinguish it from any other tissue. For young healthy subjects, the classification is normally good, but caution is required when the method is used for severely pathological brains.

MR images are normally reconstructed by taking the modulus of complex images. Normally distributed complex values are not normally distributed when the magnitude is taken. Instead, they obey a Rician distribution. This means that any clusters representing the background are not well modeled by a single Gaussian, but it makes very little difference for most of the other clusters.

The segmentation is normally run on unprocessed brain images, where non-brain tissue is not first removed. This results in a small amount of non-brain tissue being classified as brain. However, by using morphological operations on the extracted GM and WM segments, it is possible to remove most of this extra tissue. The procedure begins by eroding the extracted WM image, so that any small specs of misclassified WM are removed. This is followed by conditionally dilating the eroded WM, such that dilation can only occur where GM and WM were present in the original extracted segments. Although some non-brain structures (such as part of the sagittal sinus) may remain after this processing, most non-brain tissue is removed. Figure 5.9 shows how the GM and WM partitions can be cleaned up using this procedure, and surface rendered images of brains automatically extracted this way are shown in Figure **??**.

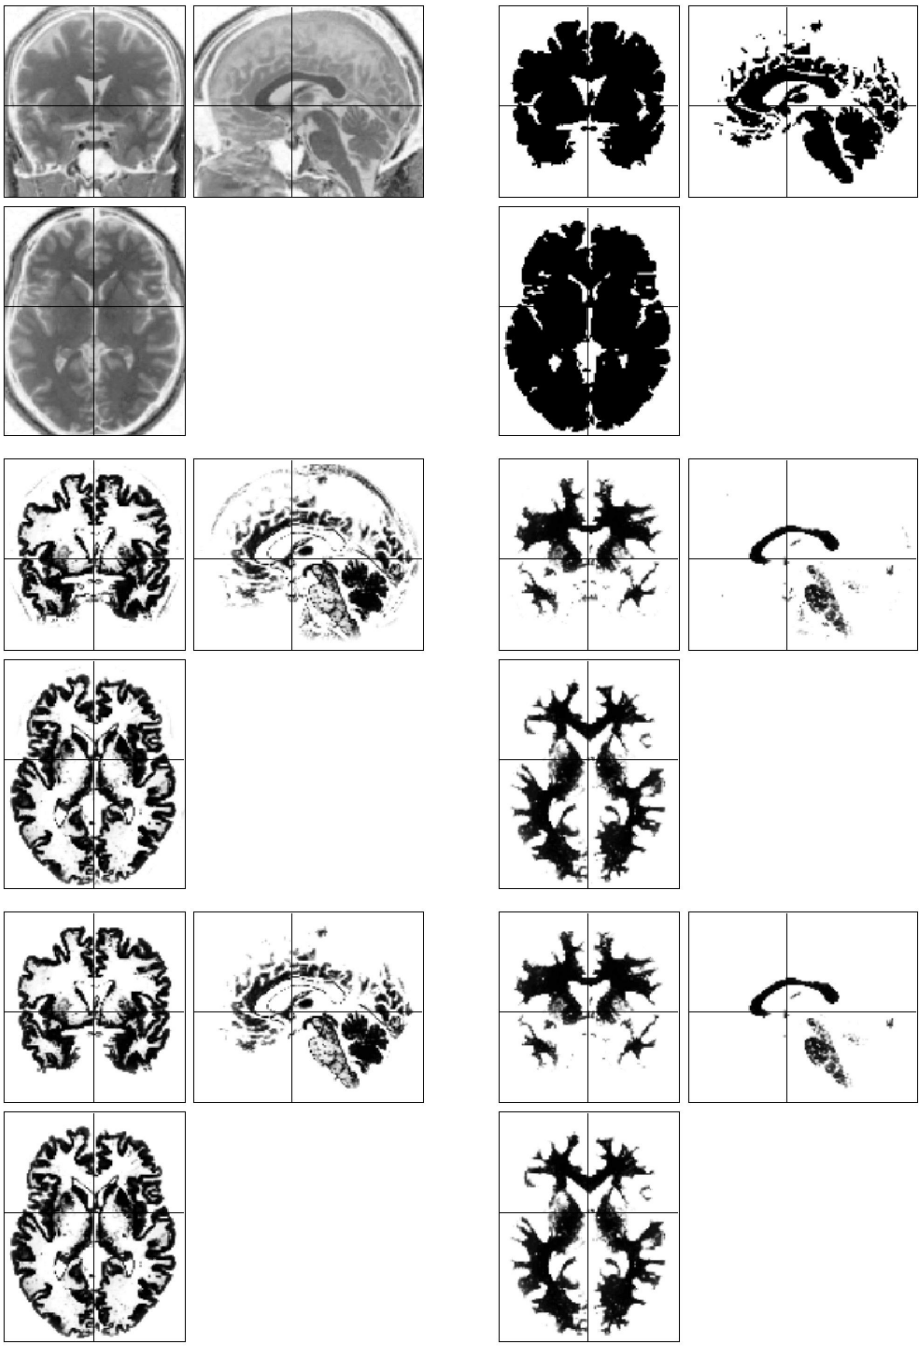

Figure 5.9: Example of automatically cleaned up segmented images. The top row shows the original T1 weighted MR image, next to an automatically generated mask of brain derived from the initial grey and white matter partitions. The second row shows the initial extracted grey and white matter. The bottom row shows the grey and white matter partitions after cleaning up by multiplying with the brain mask.

Bibliography

[1] E. Bullmore, M. Brammer, G. Rouleau, B. Everitt, A. Simmons, T. Sharma, S. Frangou, R. Murray, and G. Dunn. Computerized brain tissue classification of magnetic resonance images: A new approach to the problem of partial volume artifact. *NeuroImage*, 2:133–147, 1995.

[2] C.A. Cocosco, V. Kollokian, R.K.-S. Kwan, and A.C. Evans. Brainweb: Online interface to a 3D MRI simulated brain database. *NeuroImage*, 5(4):S425, 1997.

[3] 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 Transactions on Medical Imaging*, 17(3):463–468, 1998.

[4] 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 *Proc. IEEE-Nuclear Science Symposium and Medical Imaging Conference*, pages 1813–1817, 1993.

[5] A. C. Evans, D. L. Collins, and B. Milner. An MRI-based stereotactic atlas from 250 young normal subjects. *Society of Neuroscience Abstracts*, 18:408, 1992.

[6] A. C. Evans, M. Kamber, D. L. Collins, and D. Macdonald. An MRI-based probabilistic atlas of neuroanatomy. In S. Shorvon, D. Fish, F. Andermann, G. M. Bydder, and Stefan H, editors, *Magnetic Resonance Scanning and Epilepsy*, volume 264 of *NATO ASI Series A, Life Sciences*, pages 263–274. Plenum Press, 1994.

[7] B. Fischl, D. H. Salat, E. Busa, M. Albert, M. Dieterich, C. Haselgrove, A. van der Kouwe, R. Killiany, D. Kennedy, S. Klaveness, A. Montillo, N. Makris, B. Rosen, and A. M. Dale. Whole brain segmentation: Automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33:341–355, 2002.

[8] J. A. Hartigan. *Clustering Algorithms*, pages 113–129. John Wiley & Sons, Inc., New York, 1975.

[9] R. K.-S. Kwan, A. C. Evans, and G. B. Pike. An extensible MRI simulator for post-processing evaluation. In *Proc. Visualization in Biomedical Computing*, pages 135–140, 1996.

[10] D. H. Laidlaw, K. W. Fleischer, and A. H. Barr. Partial-volume bayesian classification of material mixtures in MR volume data using voxel histograms. *IEEE Transactions on Medical Imaging*, 17(1):74–86, 1998.

[11] S. Ruan, C. Jaggi, J. Xue, J. Fadili, and D. Bloyet. Brain tissue classification of magnetic resonance images using partial volume modeling. *IEEE Transactions on Medical Imaging*, 19(12):1179–1187, 2000.

[12] D. W. Shattuck, S. R. Sandor-Leahy, K. A. Schaper, D. A. Rottenberg, and R. M. Leahy. Magnetic resonance image tissue classification using a partial volume model. *NeuroImage*, 13(5):856–876, 2001.

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

[14] K. Van Leemput, F. Maes, D. Vandermeulen, and P. Suetens. Automated model-based bias field correction of MR images of the brain. *IEEE Transactions on Medical Imaging*, 18(10):885–896, 1999.

[15] K. Van Leemput, F. Maes, D. Vandermeulen, and P. Suetens. Automated model-based tissue classification of MR images of the brain. *IEEE Transactions on Medical Imaging*, 18(10):897–908, 1999.

[16] D. Vandermeulen, X. Descombes, P. Suetens, and G. Marchal. Unsupervised regularized classification of multi-spectral MRI. In *Proc. Visualization in Biomedical Computing*, pages 229–234, 1996.

[17] M. X. H. Yan and J. S. Karp. An adaptive baysian approach to three-dimensional MR brain segmentation. In Y. Bizais, C. Barillot, and R. Di Paola, editors, *Proc. Information Processing in Medical Imaging*, pages 201–213, Dordrecht, The Netherlands, 1995. Kluwer Academic Publishers.

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