

Segmentation of the Brain from MR Images

Jenny Caesar

LiU-IMT/IM20-EX--05/402--SE

Linköping 2005



Linköpings tekniska högskola
Institutionen för medicinsk teknik

Rapportnr:
LiU-IMT/MI20-EX--05/402--SE

Datum: 2005-06-02

**Svensk
titel**

**Engelsk
titel** Segmentation of the Brain from MR Images

Författare Jenny Caesar

Uppdragsgivare:
KTH,
Avdelningen för Neuronik

Rapporttyp:
Examensarbete

Rapportspråk:
Engelska

Sammanfattning (högst 150 ord).
Abstract (150 words)

KTH, Division of Neuronic Engineering, have a finite element model of the head. However, this model does not contain detailed modeling of the brain. This thesis project consists of finding a method to extract brain tissues from T1-weighted MR images of the head. The method should be automatic to be suitable for patient individual modeling.

A summary of the most common segmentation methods is presented and one of the methods is implemented. The implemented method is based on the assumption that the probability density function (pdf) of an MR image can be described by parametric models. The intensity distribution of each tissue class is modeled as a Gaussian distribution. Thus, the total pdf is a sum of Gaussians. However, the voxel values are also influenced by intensity inhomogeneities, which affect the pdf. The implemented method is based on the expectation-maximization algorithm and it corrects for intensity inhomogeneities. The result from the algorithm is a classification of the voxels. The brain is extracted from the classified voxels using morphological operations.

Nyckelord (högst 8)
Keyword (8 words)

MR images, automatic segmentation, voxel classification, intensity inhomogeneities, the expectation-maximization algorithm

Bibliotekets anteckningar:

Abstract

KTH, Division of Neuronic Engineering, have a finite element model of the head. However, this model does not contain detailed modeling of the brain. This thesis project consists of finding a method to extract brain tissues from T1-weighted MR images of the head. The method should be automatic to be suitable for patient individual modeling.

A summary of the most common segmentation methods is presented and one of the methods is implemented. The implemented method is based on the assumption that the probability density function (pdf) of an MR image can be described by parametric models. The intensity distribution of each tissue class is modeled as a Gaussian distribution. Thus, the total pdf is a sum of Gaussians. However, the voxel values are also influenced by intensity inhomogeneities, which affect the pdf. The implemented method is based on the expectation-maximization algorithm and it corrects for intensity inhomogeneities. The result from the algorithm is a classification of the voxels. The brain is extracted from the classified voxels using morphological operations.

Acknowledgements

Firstly, I would like to thank the staff at KTH, Division of Neuronics, for suggesting this thesis project to me and thus leading me onto the path of medical image processing. Thanks my supervisor at KTH, Johnson Ho, and the rest of the staff for providing a stimulating work environment for me. Thank you also for making me feel welcome and, not to forget, for making me laugh.

Thanks to my opponent Samuel Axelsson for suggesting improvements to my report.

A huge thank you to my supervisor Andreas Wrangsjö! Your ideas and supervision have helped me immensely during the work on this thesis project. Also, your support and encouragement has meant a great deal to me.

Table of Contents

1. Introduction	1
1.1. <i>Background</i>	1
1.2. <i>Assignment</i>	1
1.3. <i>Outline</i>	3
2. Segmentation, Classification, and Labeling	5
3. MR Images	7
3.1. <i>Weighting</i>	7
3.1.1. T1-weighted Images	7
3.1.2. T2-weighted Images	7
3.2. <i>Patient Dependency</i>	8
3.3. <i>Artifacts</i>	8
3.3.1. Intensity Inhomogeneities	8
3.3.2. The Partial Volume Effect	8
4. Segmentation Methods	11
4.1. <i>Features</i>	11
4.2. <i>Manual Segmentation</i>	12
4.3. <i>Thresholding</i>	12
4.4. <i>Atlas-Based Methods</i>	12
4.5. <i>Watershed</i>	13
4.6. <i>Region Growing</i>	13
4.7. <i>Active Contours</i>	14
4.8. <i>Classifiers</i>	14
4.9. <i>Clustering</i>	15
4.10. <i>Markov Random Fields</i>	15
4.11. <i>Fuzzy Connectedness</i>	16
5. Why an Automatic Method?	17
6. The Chosen Method	19
7. Theory for the Implemented Method	21
7.1. <i>Preprocessing</i>	22
7.1.1. Removal of Background Voxels	22
7.2. <i>Modeling the Intensities Emitted by One Tissue</i>	22
7.3. <i>Modeling the Intensity Distribution</i>	23
7.4. <i>Voxel Classification</i>	26
7.4.1. Soft Classification	26
7.4.2. Estimation of the Gaussian Parameters θ	26
7.4.3. Estimation of the Bias Field β	30
7.4.4. The Expectation-Maximization Algorithm	30
7.4.5. Hard Classification	31
7.5. <i>Segmentation of the Brain Using Morphological Operations</i>	32
7.5.1. Dilation	32
7.5.2. Erosion	33
7.5.3. Opening	33
7.5.4. Closing	33

8. Implementation	35
8.1. Preprocessing	35
8.1.1. Removal of Background Voxels	35
8.1.3. Initialization of the EM Algorithm	36
8.2. Voxel Classification	36
8.2.1. Modeling the Bias Field.....	36
8.2.2. Spatial Filtering of the Classifications.....	37
8.3. Segmentation of the Brain Using Morphological Operations.....	38
9. Results.....	39
9.1. Synthetic Images.....	39
9.1.1. Bias Correction.....	40
9.2. MR Images.....	42
9.2.1. Removal of Background Voxels.....	42
9.2.2. Voxel Classification	43
9.2.4. The Number of Classes.....	44
9.2.5. Initialization of the EM Algorithm	46
9.2.6. Bias Correction.....	47
9.2.7. Spatial Filtering of the Soft Classifications	48
9.3. Segmentation of the Brain.....	50
10. Discussion.....	53
10.1. Preprocessing	53
10.1.1. Logarithmic Transformation of the Intensities	53
10.1.2. Initialization of the EM algorithm	53
10.1.3. The Number of Classes	54
10.2. Voxel Classification	54
10.2.1. Spatial Considerations	55
10.2.2. Bias Correction.....	55
10.3. Segmentation of the Brain.....	55
10.4. Extensions of the Method	56
10.5. Validation.....	56
11. Conclusions.....	57
12. References.....	59

1. Introduction

1.1. Background

At KTH, Division of Neuronic Engineering, a finite element model of the head has been developed (Kleiven, 2002). The finite element model is used to model biological tissues to be able to observe how they are affected by for example trauma to the head. The model can thus be used to predict injury by performing numerical calculations of for example pressure and strain in brain tissues.

The finite element model, see Figure 1.1, includes for example white matter, gray matter, cerebrospinal fluid (CSF), bone, major blood vessels, and meninges. However, the model does not have a very detailed geometry.




Figure 1.1. Finite element model of the head.

1.2. Assignment

The aim of in this thesis project is to find and implement a method for segmentation of MR images of the head. The tissues resulting from the segmentation are to be used in a finite element model of the head.

It is outside the scope of this thesis project to find all the tissues in the finite element model. The primary goal is to extract the brain, i.e. white matter, gray matter, and intracerebral CSF, and to detect the convolutions of the brain as well as possible. Thus, the aim is to obtain a more detailed geometrical description of the brain.

The method should be suitable for patient individual modeling, which implies that a method of minimal user interactivity is desirable, i.e. an automatic method.

The images used are T1-weighted gray scale MR images of normal brains. The three-dimensional image volume consists of slices of images. Examples of axial (horizontal) slices from an MR volume are presented in Figure 1.2. The image volume is made up of voxels, which are the three-dimensional equivalent of pixels. The coronal (from the front) and sagittal (from the side) views of the MR volume are shown in Figure 1.3.

The assignment thus consists of finding an appropriate automatic method for extracting the brain tissues from T1-weighted MR images, and of implementing this method.




Figure 1.2. T1-weighted axial image slices of the head.




Figure 1.3. Coronal and sagittal view of an MR volume.

1.3. Outline

Chapters 2-5 contain general background information on segmentation of MR images. Chapter 2 explains the concepts segmentation, classification, and labeling. In chapter 3 characteristics of MR images are described. Chapter 4 gives an overview of the most common segmentation methods and chapter 5 describes why an automatic method is to be preferred.

Chapter 6 gives an introduction to the chosen method, while in chapter 7 the theory behind the method is accounted for. In chapter 8 important details on the implementation of the method are described. The results obtained with the implemented method are presented in chapter 9 and the results and the method are discussed in chapter 10.

2. Segmentation, Classification, and Labeling

Segmentation implies the division of an image into different objects or regions. Which, and how many, these regions are, depend on what the resulting segmentation will be used for. As an example, if the assignment is to make a surface model of a human head it is of no interest to find the contours of the brain. In that case the regions are background and head. In the scope of this thesis the regions are brain, and non-brain, with non-brain being background, CSF, bone, skin, muscle and adipose tissue etc. The brain region consists of the subregions white matter, gray matter, and intracerebral CSF.

The concepts segmentation and classification are often used interchangeably in the literature but there is an actual difference in the meaning of the two words.

Segmentation implies the division of an image into different connected regions that do not overlap. Thus, the union of all the regions is the image itself. A region often has a similar intensity or a distinct boundary (Pham et al, 2000).

Classification is the division of the image but each class does not have to be connected spatially. The voxels are classified as belonging to one of a number of classes. A difficulty of classification may be to determine the number of classes (Pham et al, 2000).

Labeling means that the segmented regions or classes are associated with what they represent – in other words to tell what each region resulting from the segmentation represents (Pham et al, 2000). The result from a segmentation is object 1, 2, 3 etc. The result from a labeling of this segmentation would be to assign for example object 1 = background, object 2 = brain, object 3 = CSF etc.

Labeling can be performed by an expert or a computer program after the segmentation or classification. In classification sometimes labeling can be included in the classification step since it may be known from the start which class represents what (Pham et al, 2000). Also, when using atlas-based methods classification is performed automatically.

3. MR Images

Which segmentation technique to be chosen depends on the properties of the image being segmented. Different imaging techniques bring out different anatomical structures in the head. MR images contain different information than for example CT images. CT images are better at depicting bone, for example fractures. A fracture can be very difficult or impossible to detect in an MR image. On the other hand, MR imaging generate images of high contrast between soft tissues in the body.

This chapter aims to give an understanding of the important characteristics of MR images. It includes a description of some of the acquisition techniques for obtaining images of different contrast between tissues. It also describes certain characteristics of MR images that have to be taken into consideration when segmenting the image. This report does not contain an account of the physics behind MR imaging.

The MR signal emitted by a structure depends on the composition of the tissue and also on the imaging technique, i.e. the operator can choose to bring out certain characteristics of the tissues by using a certain imaging technique.

3.1. Weighting

MR images can be acquired using different techniques. The resulting images highlight different properties of the depicted materials. The most common weightings are T1 and T2, which highlight the properties T1-relaxation and T2-relaxation respectively.

Selection of the most appropriate weighting is important for a successful segmentation. The properties of the tissues that are to be segmented have to be known to make a well-founded decision (Pham et al, 2000).

3.1.1. T1-weighted Images

The images in this thesis work are T1-images. T1-images show high contrast between tissues having different T1-relaxation times. Tissues with long T1-relaxation time emit little signal and thus they will be dark in the resulting image. In T1-images air, bone and CSF have low intensity, gray matter is dark gray, white matter is light gray, and adipose tissue has high intensity. T1-images have high contrast between white matter and gray matter.

3.1.2. T2-weighted Images

In T2-images, white matter and gray matter are gray and have similar intensities. CSF is bright, while bone, air, and fat appear dark. As opposed to T1-images, T2-images have high contrast between CSF and bone. The contrast between white matter and gray matter is not as good as in T1-images.

3.2. Patient Dependency

Using MR imaging, different patients give rise to images with different image properties. The intensities in the MR images are patient dependant. The gray level of the tissues might for example be very different from person to person. However, MR images share some common characteristics and the general appearance of the histogram is similar for different patients, even though the gray levels might be shifted.

3.3. Artifacts

A variety of artifacts may appear in MR images. Since the artifacts change the appearance of the image they may also affect the performance of a segmentation algorithm. The most important artifacts in image segmentation are intensity inhomogeneities and the partial volume effect.

3.3.1. Intensity Inhomogeneities

Intensity inhomogeneities are not always visible to the human eye, but can nonetheless have negative influence on automatic segmentation. This may manifest itself by for example making intensities in one part of the image brighter or darker than another part. It is often caused by the radio frequency (RF) coils. Different methods exist to compensate for the inhomogeneities. The inhomogeneity is often modeled as a field that varies smoothly over the image. The inhomogeneity field is often thought to be a multiplicative field, which means that the true voxel intensity is multiplied by the value of the field in that voxel. There are methods which remove the inhomogeneities during segmentation. For example Wells et al (1996) and van Leemput et al (1999) alternate estimation of the inhomogeneity field with classification to obtain inhomogeneity corrections.

3.3.2. The Partial Volume Effect

The partial volume effect occurs when a voxel cannot be accurately assigned to one tissue type. This is because the intensity in the voxel originates from more than one tissue. It occurs because one voxel contains many body cells and the signal emitted from these cells make up the detected intensity in this voxel.

The partial volume effect is most apparent at edges between different tissues. It may deteriorate the sharpness of the edges between tissues. The partial volume effect can be a significant problem in brain segmentation since the brain has a complex folded surface.

Another concern is the classification of such voxels, i.e. to which of the tissues should it be assigned. One way of dealing with partial volume effects is by using so called soft segmentation. Soft segmentation, as opposed to hard segmentation, means that a voxel may belong to more than one tissue class. Some methods perform soft segmentation by finding the probabilities that a voxel belongs to different tissues (Pham et al, 2000).

The partial volume effect is caused by the fact that we have a limited resolution in the images. Smaller voxel sizes reduce the partial volume effect since the probability that more than one tissue type is contained in the same voxel is reduced.

4. Segmentation Methods

This chapter aims to give an overview over the most common segmentation techniques. The chapter does not cover all existing segmentation techniques, but hopefully gives an overview over the most common methods that exist today for segmenting MR images.

Different studies have been made in the field of segmentation of MR head images but no universally agreed best method exists. Different kinds of segmentation methods approach the segmentation problem in different ways. Different methods also base the segmentation on different features in the image, e.g. intensity or gradient. The choice of what method that should be chosen should be based on what tissues that are to be segmented, since different methods are better suited for segmenting different tissues.

Segmentation methods can be divided into groups depending on different criteria. One aspect is the level of user interactivity. The groups are then manual methods, which require a high level of user interactivity, computer-aided semiautomatic methods, and completely automatic methods (Shan et al, 2002).

Segmentation methods can also be thought of as being region-based (segmentation) or pixel/voxel-based (classification). Region-based methods find connected regions. A region may have some common property like intensity or texture, or have well defined boundaries, or a characteristic shape. In voxel-based segmentation, or classification, individual voxels are assigned to different objects. Some voxel-based methods only take the current voxel value into account, while other methods also include the influence of voxels in a neighborhood (Pham, 2000). Voxel-based segmentation methods can be more susceptible to intensity inhomogeneities because they often are based only on the intensity values and not on other information like gradient (Pham et al, 2000).

4.1. Features

The segmentation of an image is based on features that are extracted from the image. The most obvious feature that can be found in an image is the intensity. In a single image, as in this thesis project, the intensity is one feature. For multispectral images, e.g. both T1 and T2, the intensity is several features. This represents that each voxel has several different intensities, i.e. one intensity per image.

Other features can be derived from the intensity values. The voxel values in a neighborhood of a voxel or the average intensity in the neighborhood can be useful if one wants to include contextual information to the feature, i.e. that the classification of a voxel also depends on the neighboring voxels. Gradient images are often used because they give specific information about where discontinuities in the image intensity exist. These discontinuities are often boundaries between tissues.

Features can be used alone or as a combination of several features. For example, a combination of intensity and gradient might be used, or a combination of intensity and average neighborhood intensity. The choice of what features to use is important since the features determine the outcome of the segmentation.

The methods described in this chapter can be used alone but are also used in various combinations to obtain better segmentation results.

4.2. Manual Segmentation

The most commonly used, and conceptually simplest, method is manual segmentation. Manual segmentation implies that a user decides the segmentation, for example by drawing the contours of the different objects directly in the image. The user thus decides what image voxels to assign to the different objects. This requires that an expert performs the segmentation, i.e. someone who has extensive knowledge about the anatomy of the regions being segmented. The performance will depend on the complexity of the shapes being segmented. It might for example be complicated to delineate the contours of the convolutions of the brain accurately. Manual segmentation is almost always a region-based segmentation technique.

4.3. Thresholding

Thresholding is a voxel-based segmentation method. By thresholding an image the image is transformed into a binary image. Voxels with intensities lower than the threshold value are set to zero and voxels with intensities higher than the threshold value are set to one.

The thresholds should be chosen so that they separate tissues from each other. Thresholds can be found by various methods. They can be chosen manually by a user or in a more automatic manner by a computer program. Thresholding is often based on the image histogram. The histogram shows how the intensities in the image are distributed. Thresholds are often chosen as the minima in the histogram. In Brummer et al (1993) the choice of thresholds is based on the assumption that the intensities of a particular tissue has a Rayleigh distribution. Shan et al (2002) base the thresholds on the assumption that the intensities of a tissue have a Gaussian distribution. The distribution is fitted to the image histogram and the thresholds are chosen at appropriate distances from the distribution mean value.

4.4. Atlas-Based Methods

Atlas-based segmentation is based on that an atlas of the tissues exists, i.e. the atlas is a map of where different tissues are located in the image. This map is registered to the medical image being segmented. After registration, the segmentation and labeling of the image is completed.

The difficulty in atlas-based segmentation is the registration of the image, i.e. to fit the map to the image being segmented. The result of atlas-based methods depends on the registration method.

Also probabilistic atlases can be used in segmentation. A probabilistic atlas consists of probabilities of certain tissues occurring at certain positions.

The performance of atlas-based methods, but also methods based on probabilistic atlases, depend on the quality of the atlas being used. Anatomic variability in the tissue being

segmented, limit the usefulness of atlases. If the tissue varies a lot between individuals, as is the case with the brain, it can be difficult to obtain a good match between the atlas and the image.

4.5. Watershed

The watershed transform is another region-based segmentation approach. The performance of this algorithm does not depend on the complexity of the shapes that are segmented. To get an understanding of how watershed segmentation works it can help to see the image as a landscape with valleys, hills, and plateaus. The intensity value in the image is then proportional to the altitude. Thus high intensity values are peaks etc. Then imagine that rain falls on the landscape and creates pools of water. Where two pools meet there is a watershed that separates the two pools. The watersheds thus separate different objects from each other. The number of objects resulting from the segmentation depends on how many local minima that exist in the image. In other words, water starts to collect at each local minimum and each local minimum will give rise to one object in the segmentation result. Often, the watershed transform is performed on the gradient image (Grau et al, 2004).

However, watershed segmentation has several drawbacks when used on medical images. Firstly, oversegmentation is often a problem. Oversegmentation means that too many objects have been segmented, i.e. the image is divided into too many different regions. It occurs due to too many local minima in the image. The watershed transform is also sensitive to noise, and may have difficulties in finding thin structures, and in finding regions which are separated with a boundary of lower contrast than other boundaries in the proximity. For example the gray matter-white matter boundary has lower contrast than the gray matter-CSF boundary which might cause problems (Grau et al, 2004).

However, refinements of the watershed segmentation algorithm have been made to improve the performance. For example, the concept of markers has been introduced to remove oversegmentation. Grau et al (2004) have developed a method that is automatic and performs well for white matter-gray matter segmentation.

4.6. Region Growing

Region growing aims at finding regions that share some common characteristic feature. First a seed point is selected. The neighboring voxels are then compared to the seed voxel and added to the region if they fulfill some criteria of similarity. The neighboring voxels to this region are then investigated and compared until the growing stops. The stopping criteria could for example be relative intensity value to the seed point or gradient value. The choice of seed points and criteria of similarity affect the final outcome of a region growing algorithm (Clarke et al, 1995).

If partial volume effects are present or if there are small connections between objects, it can cause the objects to combine into one larger region. Noise can also affect the segmentation negatively (Pham et al, 2000).

4.7. Active Contours

Active contours is a method to find the contours of objects. It can be performed in 2D, also called snakes, or in 3D which is called active surfaces. The idea is to place a contour, or snake, in the image. This snake is then supposed to find the contours of the searched object in an automatic manner. The snake is affected by different forces so that it changes its shape to fit the contour of the object that one wants to find. To picture the snake one can think of a rubber band that changes its shape and size to fit the contour of an object.

In the basic snake model there are internal and external forces that affect the snake. The internal forces tries to make the contour smaller and the external forces counteracts the internal force. The external forces are due to the image itself. It is often the gradient image. How much the different forces affect the snake determines how flexible the snake is. For example if it should be able to delineate the convolutions it has to be very flexible.

The problem in active contours is to find the optimal parametric contour $c(s)$.

$$c(s) = (x(s), y(s))$$

$$s \in [0,1]$$

There are internal and external forces affecting every point on the contour. The so called energy of a point on the contour can be written as:

$$E(c(s)) = E_i(c(s)) + E_e(c(s))$$

E_i is the energy due to the internal forces and E_e is the energy due to the external forces. The total energy of the contour is:

$$E = \int_0^1 E(c(s)) ds = \int_0^1 (E_i(c(s)) + E_e(c(s))) ds$$

The optimal snake is then found by minimizing the total energy of the contour, i.e.:

$$\hat{c}(s) = \arg \min_{c(s)} (E)$$

When using active contours, the contour changes its shape until a local minimum of the energy function E is reached (Denzler & Niemann, 1999).

The difficulty with active contours is that the choice of the parameters, which decide how much the different forces should affect the contours, is not always straightforward. Fine tuning by the user is often needed to get the best results on different sets of images.

4.8. Classifiers

A classifier is a voxel-based classification method. It assigns voxels to different classes depending on how well the voxels agree with the features of a set of voxels that are

characteristic for that class. This is a supervised method since the class specific voxels have to be chosen before the segmentation is performed. A common classifier is the k-nearest-neighbor classifier (kNN). The voxel that is to be classified is compared to the training data and the k number of data that are closest to the features of the voxel determine what class the voxel is assigned to. Another common classifier is Bayes' classifier which assumes that the distribution of the voxel intensities have a parametric intensity distribution. A common assumption about the distribution in MR images is that the voxel intensities of a certain tissue are normally distributed around a tissue-specific mean intensity. Since an image contains several tissues the intensity distribution of the entire image, i.e. the histogram, can be thought of as being a mixture of Gaussians distributions. The parameters of the Gaussians are determined by training data that are selected in advance. Wells et al (1996) use Bayes' classifier in a method which also includes inhomogeneity corrections. Bayes' classifier can also perform soft segmentations by allowing voxels to be assigned to more than one Gaussian (Pham et al, 2000).

4.9. Clustering

Clustering is very similar to classifiers only that it is an unsupervised method. That means that no training data is used. Instead algorithms that interleave classification and estimation of class specific properties are used. Common algorithms for clustering are K-means (ISODATA), fuzzy c-means, and the Expectation-Maximization (EM) algorithm. The class specific property in K-means is a mean value that is determined iteratively for each class. Then the voxels are assigned to the mean which they are closest to. Fuzzy c-means is a generalization of K-means where soft segmentations are allowed (Pham et al, 2000). The EM algorithm, used by for example van Leemput et al (1999), is the clustering equivalent to Bayes' classifier. It alternates soft classification and estimation of the Gaussian parameters (Pham et al, 2000). van Leemput et al (1999) have also included inhomogeneity correction and Markov random fields in their method.

One difficulty with clustering is the choice of initialization of the algorithms. The iterations have to start with either an estimate of the classifications or of the class specific properties. Often the initialization affects how well the clustering algorithm performs (Pham et al, 2000).

4.10. Markov Random Fields

One disadvantage of classifiers and clustering is that they do not include contextual information, i.e. the classification of a voxel is totally independent of all the other voxels. One way of dealing with this is by using Markov random fields (MRFs). In MRFs the classification of a voxel depends on its neighboring voxels. An example of this is that the voxel being classified has a high probability of belonging to the same class as its neighboring voxels, or that certain tissues never are in contact with each other, like for example skin and brain. MRFs have been used with the assumption that the underlying tissue distributions are parametric (Zhang et al, 2001), i.e. describable by for example Gaussians, and also by nonparametric intensity distributions (Held et al, 1997). Since the brain has a complex boundary it is preferable not to include too many neighbors in the MRF model (Held et al, 1997). MRFs are often less susceptible to noise than classifiers and clustering methods due to the inclusion of contextual information (Pham et al, 2000). A difficulty is that it may be

problematic to choose the dependencies of neighboring voxels, i.e. how much the neighborhood should affect the classification (Pham et al, 2000).

4.11. Fuzzy Connectedness

Even though an object consists of voxels that have different intensities, the appearance to a viewer is that they form one object. Thus, the images are fuzzy and Udupa & Saha (2003) suggest that this can be taken care of by means of fuzzy connectedness. Fuzzy connectedness is related to region growing in that it is a fuzzy variant of region growing (Pham et al, 2000).

Fuzzy connectedness finds how strongly voxels in the image are connected to each other. The connections between every pair of voxels in the image are calculated and from that the strongest connection between two voxels is found. Fuzzy connectedness consists of finding the strength of the connection between every pair of voxels. Affinity is a local relationship of two voxels. It is determined by how far apart two voxels are, on their surrounding intensities, and on object specific intensities. Fuzzy connectedness is the global fuzzy relationship between two voxels. It is the strongest of all the paths between two voxels (Udupa & Saha, 2003).

Fuzzy connectedness is a large optimization problem, but computations have become much faster with dynamic programming (Udupa & Saha, 2003).

5. Why an Automatic Method?

This chapter aims to give an understanding why it is desirable to use an automatic segmentation method instead of a method that requires user interactivity.

Different segmentation methods require different levels of user interactivity. There are manual methods, which require a high level of user interactivity, computer-aided semiautomatic methods, and completely automatic methods (Shan et al, 2002). The aim in this thesis is to implement a method that needs minimal user interactivity. This section presents arguments to why an automatic method should be chosen in favor of a manual method.

Methods that need user interactivity often suffer from intra- and interobserver variability (Clarke et al, 1995). Interobserver variability means that the segmentation result will not be objective, i.e. different users will make different choices that affect the segmentation process. Intraobserver variability means that the same user will make different choices on different occasions that will affect the performance of the segmentation algorithm.

In manual segmentation the user has to take a great part in the segmentation by for example drawing the contours of the brain on the images, i.e. the user chooses where the contours are located. Obviously, this kind of approach leads to segmentation results that depend very much on the user. In other words, different users will choose different contours and thus there is a great risk for interobserver variability. Manual methods may also suffer from intraobserver variability.

In semiautomatic methods the user interactivity is not as extensive as in manual methods. The user might for example be required to choose some seed points or set some parameter values. The results of such segmentations can also vary from time to time and from user to user. The use of an automatic method, however, removes these risks. An automatic method applied on the same image several times will produce the same result every time.

Also, manual and semiautomatic methods often require a trained expert to perform the segmentations. That is someone who has extensive knowledge of the anatomy being segmented and the segmentation technique being used.

Another drawback of manual delineation of brain contours is that it is a difficult task to perform due to the complexity of the shape of the brain. It may be an impossible task even for a trained expert to perform good segmentations of the convolutions of the brain.

A very important drawback of manual methods is that they are labor intense. Especially if there are many slices and sets of images, it will take considerable time for the user to perform the segmentation (Shan et al, 2002). An automatic method on the other hand only takes up computer power and does not require the user to be there during the segmentation procedure.

A drawback of automatic methods is that they do not always produce meaningful segmentations, i.e. the segmented regions may not correspond to different tissues (Clarke et al, 1995). This is due to the fact that the method cannot make use of all the information that exists about the image, such as anatomical knowledge. It can also be difficult to find

completely automatic initializations to the automatic methods that work well for different sets of images.

6. The Chosen Method

As Pham et al (2000) mention, the choice of the best segmentation technique is not an easy task. Considerations should be taken to for example which tissues that are to be found and what kind of images that are available.

When segmenting an image it is advantageous to make use of as much previous knowledge as possible about the image and the structures that are depicted. This can be properties of the imaging modality, and spatial knowledge about the structures, e.g. the brain is in the middle part of the head and it is surrounded by CSF and bone.

The method that has been implemented in this thesis project is based on a clustering technique which uses an EM classifier. EM classifiers have shown good results by others, e.g. Wells et al (1996), van Leemput et al (1999), and Kapur et al (1996). The method uses the assumption that voxel intensities emitted from a specific tissue is distributed around a specific tissue mean value.

Since intensity inhomogeneities are very common in MR images, it is important to choose an algorithm that is able to compensate for this. Wells et al (1996) and van Leemput et al (1999) have incorporated inhomogeneity correction in the EM algorithm.

The method was also chosen due to the fact that it is an automatic method. Many of the segmentation techniques described earlier require some level of user interactivity. However, several of the methods described earlier may work for this particular task, but one method had to be chosen and this method seemed most suitable.

7. Theory for the Implemented Method

This chapter explains the theory behind the method that has been implemented in this thesis project. The method is based on a statistical model of how the intensities in MR images are distributed. It is made up of:

1. Preprocessing
2. Voxel classification
3. Segmentation of the brain using morphological operations

The method used in this thesis is based on the assumption that the voxel intensities in the image can be described by parametric models. This is the same as saying that the probability density function (pdf) can be described by parametric models. Alternatively, the pdf could be described by the histogram which is a nonparametric way of describing the pdf. The histogram of a T1-weighted MR image of the head is presented in Figure 7.1.




Figure 7.1: Histogram of a T1-image of the head.

7.1. Preprocessing

7.1.1. Removal of Background Voxels

The first high peak in the histogram in Figure 7.1 corresponds to the background noise. The background, which is the surroundings of the head, does not emit any signal. Hence these intensities are only due to measurement noise. Even though the surroundings of the head do not emit any signal, they will give rise to Rayleigh distributed noise of low intensity (Brummer, 1993; van Leemput et al, 1999), see Figure 7.2. The background voxels thus affect the image histogram and make the distributions of the tissues not as distinct. The number of background voxels is quite large. Hence, if most of these voxels were to be removed from the image the relative amount of head voxels would be higher (Brummer 1993). This improves the appearance of the histogram and may improve the segmentation. Since the background voxel intensities only arise due to noise, they are not affected by intensity inhomogeneities (van Leemput et al, 1999).




Figure 7.2: The background voxels have a Rayleigh distribution.

7.2. Modeling the Intensities Emitted by One Tissue

In this section a model for the voxel intensities is set up according to the reasoning in Prima et al (2001). The model shows how the voxel intensities depend on the tissue type they belong to, on noise, and on intensity inhomogeneities.

In MR images a first assumption is often made that each tissue type emits a specific intensity, i.e. gives rise to a specific voxel intensity. This is however not quite true due to several

reasons. First there is biological noise due to the fact that tissues are not perfectly homogenous – they contain for example smaller structures and blood vessels. Due to the biological noise, the emitted intensities from a tissue are not the same, but distributed around a tissue specific mean value. With this way of thinking, the signal x_i in voxel i belonging to tissue Γ_i can be expressed as:

$$x_i = m_{\Gamma_i} + n_i^{bio}$$

where m_{Γ_i} is the mean intensity that tissue Γ_i emits and n_i^{bio} is the biological noise in voxel i .

Noise may also appear in the signal due to the nature of MR images, i.e. due to the way the image is acquired. In MR images it is also very common with artifacts such as intensity inhomogeneities. Even though an inhomogeneity field may not affect the perception of the image at a visual inspection, it may affect or even ruin automatic segmentation. The inhomogeneity field is often modeled as a multiplicative field. The bias field in voxel i is b_i . In the model below, measurement noise, n^{mea} , is also included in the model:

$$x_i = b_i(m_{\Gamma_i} + n_i^{bio}) + n^{mea}$$

From this general model different researchers have made different simplifications and assumptions before using the model in segmentation algorithms. A common assumption is to model the intensity x_i in voxel i belonging to tissue Γ_i as normally distributed around $b_i \cdot m_{\Gamma_i}$, i.e.:

$$x_i = b_i \cdot m_{\Gamma_i} + n_i^{\Gamma_i}$$

where $n_i^{\Gamma_i}$ is white Gaussian noise depending on the class Γ_i .

A slightly different approach, used by for example Wells et al (1996) and van Leemput et al (1999), is to perform a logarithmic transformation of the voxel intensities to make the bias field additive instead of multiplicative. In this model the values of the log-intensities are assumed to be normally distributed. Hence $n_i^{\Gamma_i}$ is white Gaussian noise depending on the class Γ_i .

$$y_i = \log_e x_i = \log_e(b_i \cdot m_{\Gamma_i}) + n_i^{\Gamma_i} = \log_e b_i + \log_e(m_{\Gamma_i}) + n_i^{\Gamma_i} = \beta_i + \mu_{\Gamma_i} + n_i^{\Gamma_i}$$

This means that the log-intensity y_i in voxel i belonging to tissue Γ_i is normally distributed around $\beta_i + \mu_{\Gamma_i}$.

The implemented method estimates an additive bias field and thus the model using a logarithmic transformation of the intensities is used.

7.3. Modeling the Intensity Distribution

In the previous section a model for the voxel intensities emitted by a tissue is described. To extend this, a parametric model for how the intensities in the image are distributed is set up, based on the tissue model. As is mentioned previously, each tissue emits a signal which is

normally distributed around a specific tissue mean, but distorted by a bias field. This is the same as saying that the probability that class Γ_i has generated voxel value y_i at position i that is affected by the bias field β_i , is a Gaussian distribution centered around the mean value $\mu_{\Gamma_i} + \beta_i$:

$$p(y_i | \Gamma_i, \theta_{\Gamma_i}, \beta_i) = G(y_i - \mu_{\Gamma_i} - \beta_i, \sigma_{\Gamma_i}) = G(y_i - (\mu_{\Gamma_i} + \beta_i), \sigma_{\Gamma_i}) \quad (7.1)$$

where $G(y_i - \mu_{\Gamma_i} - \beta_i, \sigma_{\Gamma_i}) = \frac{1}{\sigma_{\Gamma_i} \sqrt{2\pi}} e^{-(y_i - \mu_{\Gamma_i} - \beta_i)^2 / 2\sigma_{\Gamma_i}^2}$

and $\theta_{\Gamma_i} = \{ \mu_{\Gamma_i}, \sigma_{\Gamma_i} \}$, i.e. the Gaussian parameters,

and β_i is the bias field in voxel i .

Another way to explain this is that after the voxel intensity y_i has been corrected for by the bias field β_i in that voxel, the probability that a certain tissue class has emitted this bias corrected intensity $y_i - \beta_i$ is a normal distribution centered around the tissue mean μ_{Γ_i} , i.e.:

$$p(y_i | \Gamma_i, \theta_{\Gamma_i}, \beta_i) = G((y_i - \beta_i) - \mu_{\Gamma_i}, \sigma_{\Gamma_i})$$

Since the bias corrected intensity distribution of each tissue class is modeled by a normal distribution centered around the tissue mean, there are as many normal distributions as there are classes. Hence each voxel value has one probability of belonging to each of the classes.

However, it is useful to have an expression for the probability of a voxel intensity regardless of its class. This probability is describes by the pdf of the voxel intensities. Hence, what is wanted is the marginal probability for the voxel intensity, ignoring information about the class. The marginal probability for an event A , to ignore information about a second event B , is found from the joint probability of these events. It is found by integrating or summing the joint probability over the ignored event B , i.e. the marginal probability for A is:

$$P(A) = \sum_B P(A, B)$$

An expression for the probability of intensity y_i independent of class Γ_i is sought. Hence the expression for the marginal probability becomes:

$$p(y_i | \theta, \beta_i) = \sum_{\Gamma_i} p(y_i, \Gamma_i | \theta, \beta_i)$$

where θ is the parameters for all the Gaussian distributions.

The problem is that there is no expression for the joint probability of y_i and Γ_i given. However, the joint probability for two events A and B can also be expressed as:

$$P(A, B) = P(A | B) \cdot P(B) = P(B | A) \cdot P(A) \quad (7.2)$$

and hence:

$$P(A) = \sum_B P(A, B) = \sum_B P(A | B) \cdot P(B)$$

which gives:

$$p(y_i | \theta, \beta_i) = \sum_{\Gamma_i} p(y_i, \Gamma_i | \theta, \beta_i) = \sum_{\Gamma_i} p(y_i | \Gamma_i, \theta, \beta_i) \cdot p(\Gamma_i | \theta, \beta_i)$$

The first factor in the sum is the Gaussian distribution for a class, as expressed in (7.1). The second factor $p(\Gamma_i | \theta, \beta_i)$ can be simplified, since the probability of a certain tissue does not depend on θ and β_i . This is because θ and β_i are parameters in the model for the distribution of voxel intensities, and the probability of a certain tissue class is in reality independent of the model. The probability of a tissue class is actually decided by the anatomy of the patient, i.e. on what the relative amount of this tissue that is found in this particular patient. Hence:

$$p(\Gamma_i | \theta, \beta_i) = p(\Gamma_i) \quad (7.3)$$

which gives:

$$p(y_i | \theta, \beta_i) = \sum_{\Gamma_i} p(y_i | \Gamma_i, \theta, \beta_i) \cdot p(\Gamma_i) \quad (7.4)$$

Thus, the probability of a voxel is a weighted sum of normal distributions that is also affected by a bias field. This probability is the sum of the voxel's probabilities of belonging to each of the classes, since these class probabilities for the bias corrected voxel values are a Gaussian distribution each. The factors of the sum are weighted by the so called class prior. A prior is known or assumed information about an event before the event has occurred. In this case it is the probability of a tissue, i.e. the relative amount of a tissue.

This is an expression for the pdf for one voxel. An expression for the whole image is obtained by assuming that the voxel intensities are statistically independent of each other, as in Wells et al (1996) and van Leemput et al (1999). The pdf for the whole image, y , is then the product of the pdfs for all the voxels, using the rule for statistically independent samples:

$$P(A, B) = P(A) \cdot P(B) \quad , A \text{ and } B \text{ statistically independent}$$

which gives:

$$p(y | \theta, \beta) = p(y_1, y_2, \dots, y_N | \theta, \beta) = \prod_i p(y_i | \theta, \beta_i) \quad (7.5)$$

The assumption that the voxels are independent of each other can be thought of being rather erroneous, since it is obvious that the intensity in one voxel depends on the voxels surrounding it. However, by making this assumption, the only thing that occurs is that a lot of spatial information that exists in the image is not used. To include spatial considerations, one solution can be to use a Markov random field model.

7.4. Voxel Classification

The classification problem is to, given a voxel intensity, find the tissue type that this voxel most probably belongs to. What is wanted is thus the probabilities that each voxel belongs to each of the classes. This is the probabilities that voxel value y_i at position i belongs to the different classes Γ_i . Hence, there are as many probabilities as there are classes. These probabilities, that represent a soft classification of the tissues, are what are wanted from the voxel classification step. However, to solve this problem it has to be sorted out what is given in this problem and which the unknown variables are.

The only given information in this problem are the voxel intensities. Also known is the assumption that the pdf of the intensities are a sum of Gaussian distributions affected by a bias field.

7.4.1. Soft Classification

However, what is wanted are the classifications, i.e. the probabilities that each voxel belongs to each of the classes. An expression for this can be obtained from the pdf for voxel intensity (7.4) using Bayes' rule (7.6). By rearranging the expression for joint probabilities (7.2), Bayes' rule can be derived:

$$P(B | A) = \frac{P(A | B) \cdot P(B)}{P(A)} \quad (7.6)$$

Setting $A = y_i$ and $B = \Gamma_i$, the expression for the classifications of the voxels becomes:

$$p(\Gamma_i | y_i, \theta, \beta_i) = \frac{p(y_i | \Gamma_i, \theta, \beta_i) \cdot p(\Gamma_i | \theta, \beta_i)}{p(y_i | \theta, \beta_i)}$$

Using (7.3) and (7.4) in this expression gives:

$$p(\Gamma_i | y_i, \theta, \beta_i) = \frac{p(y_i | \Gamma_i, \theta, \beta_i) \cdot p(\Gamma_i | \theta, \beta_i)}{p(y_i | \theta, \beta_i)} = \frac{p(y_i | \Gamma_i, \theta, \beta_i) \cdot p(\Gamma_i)}{\sum_{\Gamma_i} p(y_i | \Gamma_i, \theta, \beta_i) \cdot p(\Gamma_i)} \quad (7.7)$$

However, the problem is that to be able to perform these classifications the Gaussian parameters, the bias field, and the voxel values have to be known. This is not the case, but expressions for estimating the parameters θ , and the bias field β are needed.

7.4.2. Estimation of the Gaussian Parameters θ

The parameters can be estimated as in Wells et al (1996) by manually selecting areas of voxels belonging to each tissue. The mean and standard deviation are then calculated from the intensity distribution of the voxels in this selected region. However, this method is user dependent, and the aim of this thesis is to find an automatic method.

An expression for the Gaussian parameters θ can be derived in the same way as the expression for the soft classification of the voxels, i.e. by using Bayes' rule (7.6). However, the classification is only based on the voxel value in the voxel being classified, and the estimation of θ should be based on all the voxel values in the image. Hence, Bayes' rule (7.6) is used on the pdf for the entire image (7.5), which gives:

$$p(\theta | y, \beta) = \frac{p(y | \theta, \beta) \cdot p(\theta | \beta)}{p(y)}$$

Since the best estimate for θ is sought, the aim is to maximize this probability. Hence, the denominator $p(y)$ is a normalizing constant which can be ignored. Also, since the parameters θ are independent of the bias field, the expression can be simplified and the estimation $\hat{\theta}$ becomes:

$$\hat{\theta} = \arg \max_{\theta} (p(\theta | y, \beta)) = \arg \max_{\theta} (p(y | \theta, \beta) \cdot p(\theta)) \quad (7.8)$$

This is a maximum a posteriori (MAP) estimation of θ . It requires that a prior $p(\theta)$ for the parameters is known. The MAP estimation can thus be used if there is knowledge about the probabilities for various parameter values.

However, there are cases when there is no knowledge about the prior, or when it might be best not to make any assumptions about the prior, i.e. if it is undesirable to give certain parameter values a higher probability of being chosen. Then another similar method called maximum likelihood (ML) estimation can be used, as is done by van Leemput et al (1999). It does not include any assumptions about the prior.

The likelihood function is a conditional probability that is a function of the second "given" argument with the first argument held fixed, i.e.:

$$L(B | A) = p(A | B)$$

In this problem it will be:

$$L(\theta | y, \beta) = p(y | \theta, \beta)$$

Once again, the parameters $\hat{\theta}$ that maximize the likelihood function are wanted. This means that the parameters θ that maximize the pdf for the image are sought:

$$\hat{\theta} = \arg \max_{\theta} (L(\theta | y, \beta)) = \arg \max_{\theta} (p(y | \theta, \beta)) \quad (7.9)$$

As can be seen, this expression is very similar to the MAP expression (7.8). The only difference being the prior.

Since a maximum is searched, the derivative of the ML, or MAP, expression is equal to zero in this point. Hence, the estimates $\hat{\theta}$ is found by differentiating the ML, or MAP, expression with respect to the searched parameters and setting the partial derivatives equal to zero. The calculations are shown for the ML estimation, but the MAP calculations are similar.

Using the expression for the pdf of the image (7.5) in the ML expression (7.9) gives:

$$\hat{\theta} = \arg \max_{\theta} (p(y | \theta, \beta)) = \arg \max_{\theta} \left(\prod_i p(y_i | \theta, \beta_i) \right)$$

To make the calculations simpler, a logarithmic transformation is performed. This makes the formula for the pdf of the image additive instead of multiplicative:

$$\log_e (p(y | \theta, \beta)) = \log_e \left(\prod_i p(y_i | \theta, \beta_i) \right) = \sum_i \log_e (p(y_i | \theta, \beta_i))$$

Since all the values for $p(y_i | \theta, \beta_i)$ are in the interval $[0, 1]$, the logarithmic values will be in the interval $]-\infty, 0]$. The new ML estimation for θ is thus:

$$\hat{\theta} = \arg \min_{\theta} \left(-\sum_i \log_e (p(y_i | \theta, \beta_i)) \right).$$

Using this reasoning an estimate of the mean is found by setting the partial derivative of the minimization expression equal to zero.

Finding $\hat{\mu}_{\Gamma_i}$:

$$\begin{aligned} \frac{\partial}{\partial \mu_{\Gamma_i}} - \sum_i \log_e (p(y_i | \theta, \beta_i)) &= \\ &= \frac{\partial}{\partial \mu_{\Gamma_i}} - \sum_i \log_e \left(\sum_{\Gamma_i} p(y_i | \Gamma_i, \theta, \beta_i) \cdot p(\Gamma_i) \right) = \\ &= - \sum_i 2(y_i - \mu_{\Gamma_i} - \beta_i) / \sigma_{\Gamma_i}^2 \cdot \frac{p(y_i | \Gamma_i, \theta, \beta_i) \cdot p(\Gamma_i)}{\sum_{\Gamma_i} p(y_i | \Gamma_i, \theta, \beta_i) \cdot p(\Gamma_i)} \end{aligned}$$

Then Bayes' rule (7.6) is used, in the opposite direction as when finding the expression for the classifications (7.7), which gives:

$$- \sum_i 2(y_i - \mu_{\Gamma_i} - \beta_i) / \sigma_{\Gamma_i}^2 \cdot \frac{p(y_i | \Gamma_i, \theta, \beta_i) \cdot p(\Gamma_i)}{\sum_{\Gamma_i} p(y_i | \Gamma_i, \theta, \beta_i) \cdot p(\Gamma_i)} = - \sum_i \frac{2(y_i - \mu_{\Gamma_i} - \beta_i)}{\sigma_{\Gamma_i}^2} \cdot p(\Gamma_i | y_i, \theta, \beta_i)$$

Since a minimum is searched this expression is set to zero, which gives:

$$\begin{aligned} \sum_i (y_i - \mu_{\Gamma_i} - \beta_i) \cdot p(\Gamma_i | y_i, \theta, \beta_i) &= \\ &= \sum_i (y_i - \beta_i) \cdot p(\Gamma_i | y_i, \theta, \beta_i) - \mu_{\Gamma_i} \cdot \sum_i p(\Gamma_i | y_i, \theta, \beta_i) = 0 \end{aligned}$$

Rearranging the terms gives an expression for the maximum likelihood estimate of $\hat{\mu}_{\Gamma_i}$:

$$\mu_{\Gamma_i} = \frac{\sum_i (y_i - \beta_i) \cdot p(\Gamma_i | y_i, \theta, \beta_i)}{\sum_i p(\Gamma_i | y_i, \theta, \beta_i)} \quad (7.10)$$

The same approach is used to find an expression for $\hat{\sigma}_{\Gamma_i}$:

$$\begin{aligned} \frac{\partial}{\partial \sigma_{\Gamma_i}} - \sum_i \log_e(p(y_i | \theta, \beta_i)) &= \\ &= - \sum_i \left(\frac{(y_i - \mu_{\Gamma_i} - \beta_i)^2}{\sigma_{\Gamma_i}^3} - \frac{1}{\sigma_{\Gamma_i}} \right) \frac{p(y_i | \Gamma_i, \theta, \beta_i) \cdot p(\Gamma_i)}{\sum_{\Gamma_i} p(y_i | \Gamma_i, \theta, \beta_i) \cdot p(\Gamma_i)} = \\ &= - \sum_i \left(\frac{(y_i - \mu_{\Gamma_i} - \beta_i)^2}{\sigma_{\Gamma_i}^3} - \frac{1}{\sigma_{\Gamma_i}} \right) p(\Gamma_i | y_i, \theta, \beta_i) = \\ &= \sum_i \frac{(y_i - \mu_{\Gamma_i} - \beta_i)^2}{\sigma_{\Gamma_i}^3} \cdot p(\Gamma_i | y_i, \theta, \beta_i) - \sum_i \frac{1}{\sigma_{\Gamma_i}} \cdot p(\Gamma_i | y_i, \theta, \beta_i) \end{aligned}$$

Setting this to zero:

$$\frac{1}{\sigma_{\Gamma_i}^3} \cdot \sum_i (y_i - \mu_{\Gamma_i} - \beta_i)^2 \cdot p(\Gamma_i | y_i, \theta, \beta_i) - \frac{1}{\sigma_{\Gamma_i}} \cdot \sum_i p(\Gamma_i | y_i, \theta, \beta_i) = 0$$

Rearranging the expression gives:

$$\begin{aligned} \frac{1}{\sigma_{\Gamma_i}^3} \cdot \sum_i (y_i - \mu_{\Gamma_i} - \beta_i)^2 \cdot p(\Gamma_i | y_i, \theta, \beta_i) &= \frac{1}{\sigma_{\Gamma_i}} \cdot \sum_i p(\Gamma_i | y_i, \theta, \beta_i) \\ \frac{1}{\sigma_{\Gamma_i}^2} \cdot \sum_i (y_i - \mu_{\Gamma_i} - \beta_i)^2 \cdot p(\Gamma_i | y_i, \theta, \beta_i) &= \sum_i p(\Gamma_i | y_i, \theta, \beta_i) \end{aligned}$$

which leads to the following expression for the maximum likelihood estimate of $\hat{\sigma}_{\Gamma_i}$:

$$\sigma_{\Gamma_i} = \sqrt{\frac{\sum_i (y_i - \mu_{\Gamma_i} - \beta_i)^2 \cdot p(\Gamma_i | y_i, \theta, \beta_i)}{\sum_i p(\Gamma_i | y_i, \theta, \beta_i)}} \quad (7.11)$$

However, these estimates require the knowledge of both the bias field and the classifications!

7.4.3. Estimation of the Bias Field β

Finding an expression for β , the same reasoning can be used as when finding the estimate for θ – either a MAP estimation or an ML estimation can be made. The MAP estimation becomes:

$$\hat{\beta} = \arg \max_{\beta} (p(\beta | y, \theta)) = \arg \max_{\beta} (p(y | \theta, \beta) \cdot p(\beta)) \quad (7.12)$$

and the ML estimation:

$$\hat{\beta} = \arg \max_{\beta} (p(y | \theta, \beta)) \quad (7.13)$$

For example, Wells et al (1996) use a MAP formulation for estimating the bias field. Their expression looks like the one above, except that θ is given and excluded from the expression. They model the prior for the bias field as a Gaussian centered on zero.

van Leemput et al (1999) use an ML estimation to obtain estimates for the bias field. They use a parametric model for the bias field, modeling it as a linear combination of smooth basis functions. The ML estimation becomes an estimation of the parameters of the basis functions.

The estimates of β will depend on the other searched variables, i.e. the classifications and the Gaussian parameters. Since all the searched variables depend on the values of the other variables, it implies the need for an iterative algorithm which estimates one of the unknown parameters while assuming the others as known. One such algorithm is the expectation-maximization (EM) algorithm.

7.4.4. The Expectation-Maximization Algorithm

The expectation-maximization (EM) algorithm has been used by for example Wells et al (1996) and van Leemput et al (1999). It is an algorithm that finds the ML or MAP estimates for some parameters which depend on observed information, but also on some missing information. In this problem the observed information is the voxel values, the missing information is the soft classifications of the voxels, and the parameters are the Gaussian parameters θ , and the bias field β . The algorithm alternates between two steps – the E-step and the M-step.

In the E-step the algorithm estimates the missing information, i.e. the classifications of the voxels, basing this estimation on estimates of the parameters θ and β . The expression for soft classification (7.7) is then used.

The M-step then finds the ML or MAP estimates of the parameters, basing this estimation on the estimates of the soft classification. The estimation of β is then made using (7.12) or (7.13), and then the Gaussian parameters θ are estimated using (7.10) and (7.11). The algorithm is then iterated until some criterion for convergence is reached.

For each iteration, the prior on tissue class, $p(\Gamma_i)$, can also be estimated. It is the probability of a certain tissue. Hence the sum of the probabilities of the tissues is equal to one:

$$\sum_{\Gamma_i} p(\Gamma_i) = 1$$

This parameter thus describes what proportion of the total amount of voxels that belong to tissue Γ_i . In the E-step the probability that each of the voxels belongs to tissue Γ_i is estimated. The sum of these probabilities is thus the probability of tissue Γ_i :

$$p(\Gamma_i) = \sum_i p(\Gamma_i | y_i, \theta, \beta_i)$$

However, to be able to use this algorithm it has to be initialized. It can either be initialized at the E-step, which requires an initial estimate of the parameters θ and β , or it can be initialized at the M-step, which requires an initial estimation of the soft classification. Good initialization of the EM algorithm is crucial since the algorithm only converges to a local minimum.

Initialization at the E-step implies that an estimate of the distribution parameters and the bias field must be supplied. An estimate can often be obtained from the histogram, but it often requires some level of user interactivity. Zhang et al (2001) have developed a method for automatic initialization at the E-step.

For initialization at the M-step an estimate of the voxel classifications is required. These estimates can be obtained with automatic methods from a probability map. A probability map contains the posterior probabilities of different tissues at all locations in the volume. There is thus one probability per tissue for each location. A probability map is an estimate of probability and location, usually calculated from a number of different images (van Leemput et al, 1999).

There might be difficulties with the registration of the map, i.e. to associate the values in the probability map to the voxels in the image being segmented. Also the map may not reflect the biological diversity that exists and this can cause segmentation errors for tissues that vary a lot from the map.

7.4.5. Hard Classification

One of the results from the EM algorithm is the soft classification of the voxels, i.e.

$$p(\Gamma_i | y_i, \theta, \beta_i)$$

To proceed, a hard segmentation of the voxels is wanted. To obtain the hard classification, each voxel is assigned to the class to which the voxel has the highest probability of belonging to, i.e.:

$$\hat{\Gamma}_i = \arg \max_{\Gamma_i} (p(\Gamma_i | y_i, \theta, \beta_i))$$

This is a MAP classification of the voxels, since a prior was used in the soft classification.

7.5. Segmentation of the Brain Using Morphological Operations

The brain cannot directly be separated from the classification of the voxels by for example finding the largest connected component in the image. This is because other tissues, e.g. muscle, are classified as brain tissue and there are thin connections between the brain and these other tissues. However, after the classification of the voxels has been performed, segmentation of the brain from the rest of the tissues is a rather simple problem, which can be performed using morphological operations.

Morphological operations are simple operations used to process the binary images. They manipulate the shape of the objects in the image. The most common morphological operations are dilation, erosion, opening, and closing.

7.5.1. Dilation

Dilation implies that object areas are expanded along the border to the background. This means that background voxels closer than a given distance, r , to an object voxel are converted into object voxels. An example of dilation is presented in Figure 7.3.




Figure 7.3. Example of dilation; Original image and dilated image.

Conditional dilation implies that background voxels are only converted to object voxels depending on whether they are object voxels in a mask image. An example of conditional dilation is presented in Figure 7.4. There are also conditional variants for erosion, opening and closing.



Figure 7.4. Example of conditional dilation; Original image, mask image, and conditionally dilated image.

7.5.2. Erosion

Erosion implies that object areas are shrunk along the border to the background. This means that all object voxels closer than a given distance, r , to a background voxel are converted into background voxels. An example of erosion is presented in Figure 7.5.



Figure 7.5. Example of erosion; Original image and eroded image.

7.5.3. Opening

Opening is erosion followed by dilation. The effect of opening is that thin connections between objects are removed. By having the same distance r for both erosion and dilation, most of the original image is preserved. An example of opening is presented in Figure 7.6.




Figure 7.6. Example of opening; Original image and opened image.

7.5.4. Closing

Closing is dilation followed by erosion. The effect of closing is that cracks and holes in the objects are filled. By having the same distance r for both dilation and erosion, most of the original image is preserved. An example of closing is presented in Figure 7.7.




Figure 7.7. Example of closing; Original image and closed image.

8. Implementation

This chapter contains details on the implementation. The algorithm is developed in Matlab. It consists of:

1. Preprocessing
2. Voxel classification
3. Segmentation of the brain using morphological operations

8.1. Preprocessing

In preprocessing all the steps before the EM algorithm are included. Finding estimates for the initialization of the algorithm is also regarded as preprocessing.

8.1.1. Removal of Background Voxels

The background voxels have been removed by creating a mask of the head and disregarding voxels that are not included in the mask, see Figure 8.1. The mask is created by first thresholding the image. The threshold is chosen as the midpoint between the two largest local maxima in the histogram. Morphological closing is then performed on the mask to get the mask to include most of the CSF in the brain. After these operations the image might still consist of several objects. The head is then found by finding the largest connected component in the image. To avoid that low intensity voxels inside the head are not included in the mask, flood fill operations from the corners of the image are performed. The voxels that are not filled after these operations are also assigned to the mask. The process of removing background voxels is visualized in Figure 8.1. All of the following segmentation steps are performed on the masked image.

As an alternative, the background removal could have been made using active contours.

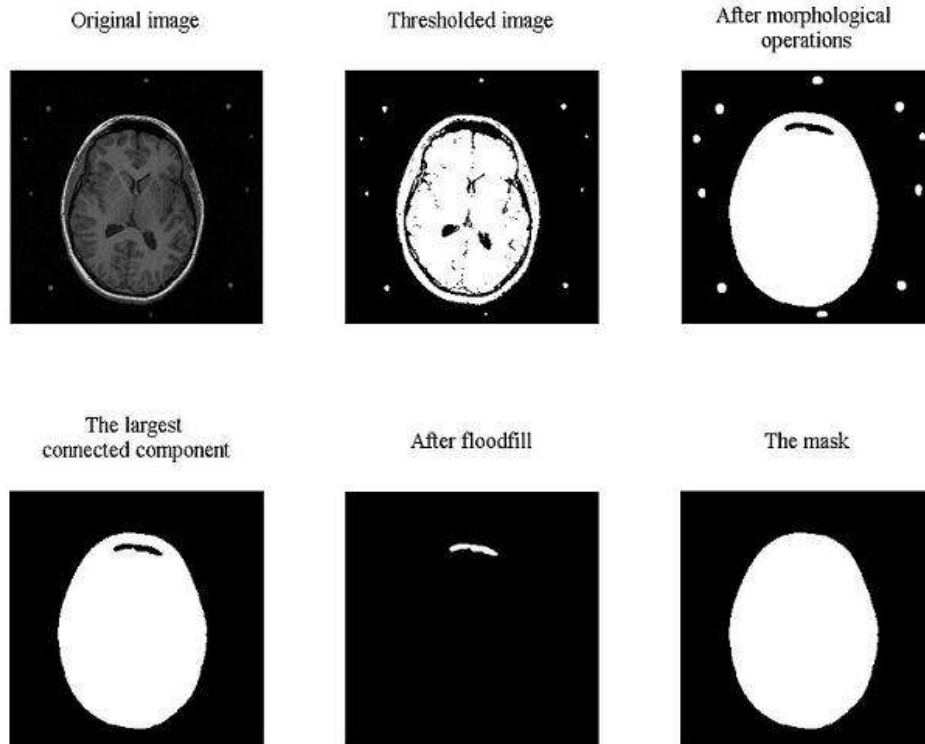


Figure 8.1. Background removal.

8.1.3. Initialization of the EM Algorithm

The EM algorithm is initialized on the E-step. Hence an estimate of the Gaussian parameters and the bias field must be provided. The bias field is initialized to zero, i.e. making the initial assumption that there is no bias field in the image. The means of the Gaussian distributions are estimated manually by inspection of the histogram. The Gaussian standard deviations are set equally for all classes. The tissue prior is simply divided equally over the classes, giving each class equal initial probability.

8.2. Voxel Classification

The EM algorithm is implemented according to the description in the previous chapter. The only exception is a slightly different model for the bias field.

8.2.1. Modeling the Bias Field

To obtain an estimate of the bias field in one voxel, the expression for the pdf of one voxel (7.4) is used together with Bayes' rule (7.6), which gives:

$$\hat{\beta}_i = \arg \max_{\beta_i} (p(\beta_i | y_i, \Gamma_i, \theta)) = \arg \max_{\beta_i} \frac{p(y_i | \Gamma_i, \theta, \beta_i) \cdot p(\beta_i | \Gamma_i, \theta)}{p(y_i | \theta, \beta_i)}$$

As is mentioned previously, simplifications can be made since the bias field is independent of the Gaussian parameters and the tissue class. Also, since a maximum estimate is wanted, the denominator is just a normalizing constant which can be ignored. Hence:

$$\hat{\beta}_i = \arg \max_{\beta_i} (p(y_i | \Gamma_i, \theta, \beta_i) \cdot p(\beta_i)) \quad (8.1)$$

Originally this expression will give one contribution to the bias field from each of the tissue classes. However, to simplify, a MAP estimate of the most probable class is performed before estimating the bias field:

$$\hat{\beta}_i = \arg \max_{\beta_i} (p(y_i | \hat{\Gamma}_i, \theta, \beta_i) \cdot p(\beta_i))$$

The estimate for the bias field in a voxel is the maximum of a normal distribution, since:

$$\hat{\beta}_i = \arg \max_{\beta_i} (p(y_i | \hat{\Gamma}_i, \theta, \beta_i) \cdot p(\beta_i)) = \arg \max_{\beta_i} (G(y_i - \mu_{\Gamma_i} - \beta_i, \sigma_{\Gamma_i}) \cdot p(\beta_i))$$

The maximum of a normal distribution is simply the mean value of the distribution, i.e. when:

$$y_i = \mu_{\Gamma_i} + \beta_i$$

Rearranging this expression gives an estimate of the bias field in voxel i :

$$\hat{\beta}_i = y_i - \mu_{\Gamma_i}$$

The bias field in a voxel is thus approximated as the difference between the intensity value in that voxel and the mean of the tissue that this voxel belongs to.

However, the bias field is not a local feature since it varies slowly across the image. Thus, a large spatial low pass filter is applied to the bias field estimates to make the bias field slowly varying across the image. This filter is applied slice by slice, since intensity offsets are common between slices.

To avoid that the bias field introduces a DC level in the image, the prior for the bias field is chosen as a zero mean uniform distribution.

8.2.2. Spatial Filtering of the Classifications

Since there are no spatial considerations taken when classifying the voxels using the EM algorithm, the result may become rather granular. Because the resulting brain from this segmentation is to be used in a finite element model, it may be desirable to make the classes appear less granular. To obtain the smoothing of classes, a local spatial low pass filter can be applied to the soft classification images, before the hard segmentation is performed.

8.3. Segmentation of the Brain Using Morphological Operations

Morphological operations are applied to the resulting classification in accordance with Kapur et al (1996). First a morphological erosion is applied because the thin connections between the brain and the other tissues have to be removed. In this new binary image, the brain is found as the largest connected component in the image. Conditional dilation is then performed of the largest connected component to obtain the final segmentation result.

9. Results

This chapter presents the obtained results when applying the implemented algorithm on T1-images and synthetic images.

9.1. Synthetic Images

The voxel classification has been applied to synthetic images affected by an additive bias field. The original image is affected by white Gaussian noise, and the second image is also affected by an additive bias field, see Figure 9.1. The bias field makes the image brighter to the left and darker to the right. Note that the images consist of 320×320 pixels.




Figure 9.1. The original image and the image affected by an additive bias field.

The histograms for these images are shown in Figure 9.2. It can be seen that the distributions are not as clear in the histogram for the image affected by the bias field.

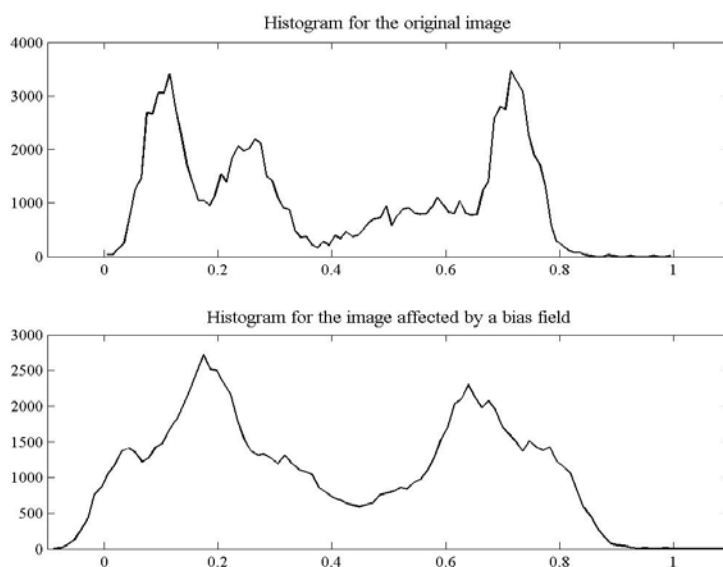


Figure 9.2. Histogram for the original image and for the image affected by a bias field.

9.1.1. Bias Correction

The image has been classified with and without bias correction. The resulting Gaussian distributions, Figure 9.3 and 9.4, give a hint that the classifications will be better with bias correction. The dotted lines are the Gaussian distributions and the thick line is the sum of these distributions. The thinner line is a histogram of the bias corrected intensity values. As can be seen, the bias correction makes the intensity distribution look more like a sum of four Gaussians.

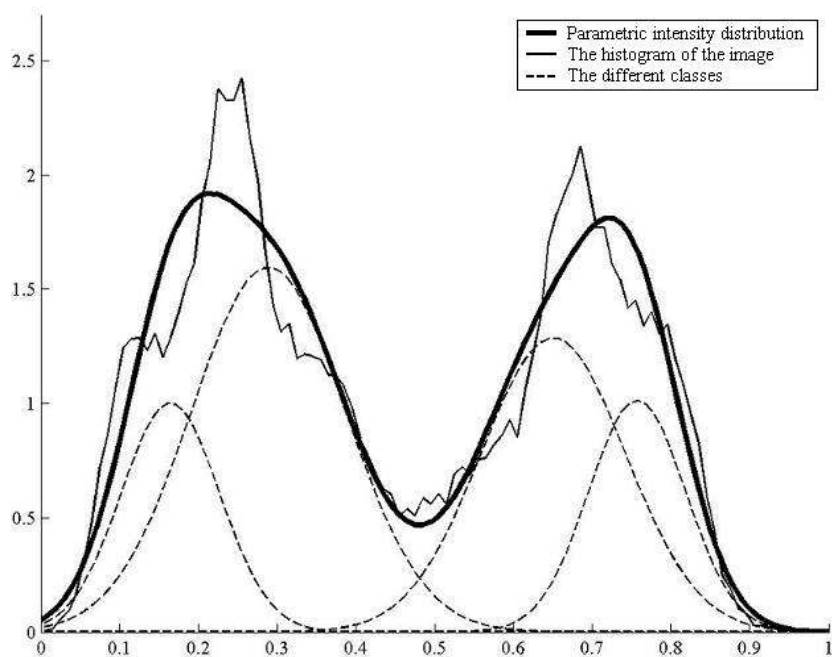


Figure 9.3. The Gaussian approximation without bias correction.




Figure 9.4. The Gaussian approximation with bias correction.

If no bias correction is performed, there are a lot of errors in the classifications, see Figure 9.5.




Figure 9.5. The true classifications and the classifications without bias correction.

When bias correction is applied, the classifications improve significantly, see Figure 9.6.




Figure 9.6. The true classifications and the classifications with bias correction.

The bias field that the algorithm has estimated is shown in Figure 9.7 next to the true bias field.




Figure 9.7. The true bias field and the estimated bias field.

9.2. MR Images

This section presents the results after the algorithm has been applied to T1-weighted images of the head.

9.2.1. Removal of Background Voxels

The background removal removes most of the background voxels. The MR image with and without background is presented in Figure 9.8.




Figure 9.8. The image before the background is removed and the image without background.

The removal of the background voxels gives a more balanced histogram, see Figure 9.9.

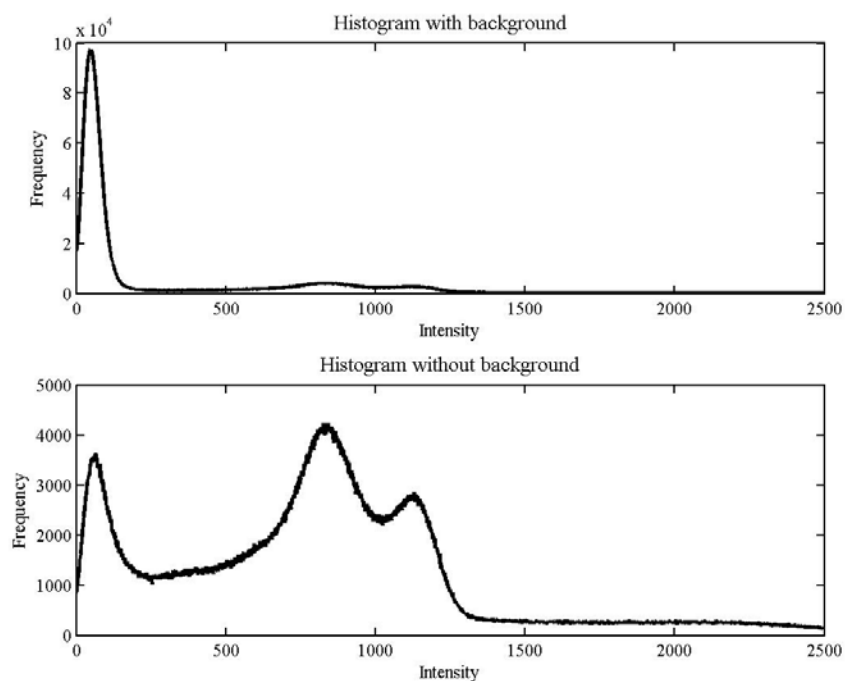


Figure 9.9: The histogram with and without background voxels.

9.2.2. Voxel Classification

The result of the EM algorithm is a mixture of Gaussians approximating the pdf for the bias corrected voxel intensities, see Figure 9.10. The dotted lines are the Gaussian distributions and the thick line is the sum of these distributions. The thinner line is a histogram of the bias corrected intensity values.




Figure 9.10: The resulting Gaussians.

The other result from the EM algorithm is the soft classification of the voxels. Each voxel is then assigned to the class that it most probably belongs to. An example of the original image and the resulting classification is shown in Figure 9.11. The different gray levels in the second image represent different classes. In this image there are six classes. From the darkest to the brightest, the classes represent mainly the following tissues:

- Background
- Bone
- CSF
- Gray matter
- White matter
- Adipose tissue

Other tissues in the head are classified as either of these classes, e.g. muscle tissue is classified as gray matter.




Figure 9.11. The original image and the result of the classification.

9.2.4. The Number of Classes

The number of classes must be chosen as input to the EM algorithm. The algorithm has been tested with different number of classes. The class for background voxels is not included in the EM algorithm since these voxels have been identified in the preprocessing step. Figure 9.13 shows the original image and the resulting classifications using four, five, and six tissue classes respectively. Using four tissue classes, CSF and gray matter are assigned to the same class. The difference between five and six tissue classes is not very significant.




Figure 9.13. The original image and the classifications using four, five, and six tissue classes respectively.

9.2.5. Initialization of the EM Algorithm

The resulting classification depends on the initialization of the algorithm. Figure 9.14 shows how different the results from the same image can be when different initialization is used. For the first image, the mean values are based on the appearance of the histogram, and in the second image they are distributed equally between the minimum and maximum intensity. The other initialization parameters are the same for both. The failure of the second classification can also be seen in the resulting Gaussians in Figure 9.15.




Figure 9.14. The resulting classifications using different initializations.




Figure 9.15. The Gaussian mixture when poor initialization is used.

9.2.6. Bias Correction

The algorithm with and without bias correction is applied to T1-weighted images. It can be seen that the MR image in Figure 9.16 is affected by intensity inhomogeneities. The bottom part is brighter, while the upper left corner is darker.




Figure 9.16. An MR image affected by intensity inhomogeneities (BrainWeb; Cocosco et al, 1997; Kwan et al, 1999; Kwan et al, 1996; Collins et al, 1998).

The image is classified using the algorithm without bias correction, see Figure 9.17. It can be seen that the classification is not perfect. In the encircled parts white matter is classified as gray matter.




Figure 9.17. The original image and the image classified without bias correction.

The image is also classified using the algorithm with bias correction, see Figure 9.18. The classifications have improved compared to the classifications without bias correction. As can be seen, the results have improved significantly in the encircled region.




Figure 9.18. The original image and the image with bias correction.

9.2.7. Spatial Filtering of the Soft Classifications

The spatial filtering of the soft classifications before the hard classification reduces the granularity in the image. This effect is shown in Figure 9.19 and 9.20 for two different MR images.




Figure 9.19. The classifications without and with filtering of the soft classification before hard classification.




Figure 9.20. The classifications without and with filtering of the soft classification before hard classification.

Due to the filtering of the soft classifications some misclassifications are corrected. For example, in the second image in Figure 9.20 some white matter voxels are classified as adipose tissue, but when spatial filtering is performed these voxels are correctly classified, see Figure 9.21.




Figure 9.21 The filtering of the soft classifications correct some misclassified voxels.

9.3. Segmentation of the Brain

The brain class, i.e. the classes for white matter and gray matter, is shown in Figure 9.22 and 9.23. As can be seen in the images, many other tissues are included in the brain class, like for example meninges and muscle tissue. There are thin connections between the brain and other tissues.




Figure 9.22. Axial slices of the brain class.




Figure 9.23. Coronal and sagittal views of the brain class.

After morphological operations have been applied on the brain class only the brain remains. The resulting segmentations of the images above are presented in Figure 9.24 and 9.25, and a 3D-visualization of the resulting brain is presented in Figure 9.27. However, the morphological segmentation may cause not only non-brain tissue but also brain tissue to become disconnected from the brain. This is seen in Figure 9.26.




Figure 9.24. Axial slices of the resulting brain.




Figure 9.25: Coronal and sagittal views of the resulting brain.




Figure 9.26. Some brain voxels are removed by the morphological operations.




Figure 9.27. 3D-visualization of the resulting brain.

10. Discussion

The performance of the method is discussed in this chapter and some improvements are suggested.

The implemented method is applied to T1-weighted images of the head. The brain is segmented from the other tissues of the head in an automatic manner. The convolutions of the brain are detected and white matter, gray matter, and CSF are separated. The method works on images independent of their orientation (axial, coronal, or sagittal).

10.1. Preprocessing

Preprocessing is a crucial step for good convergence of the EM algorithm. Even though the main focus of this thesis work has not been on this aspect of the segmentation process, some problems have been found and are mentioned in this section. They deal with the logarithmic transformation of the voxel intensities and the initialization of the algorithm.

10.1.1. Logarithmic Transformation of the Intensities

Intensity inhomogeneities are often modeled as a multiplicative field, but the EM algorithm estimates an additive bias field. To obtain an additive bias field, a logarithmic transformation should be performed on the voxel intensities. However, good results have been obtained on the original intensities. This may partly be explained by the fact that gray matter and white matter have similar intensities and thus the multiplicative effect of the inhomogeneity field is not as large. If it is important that tissues having significantly different intensities, e.g. bone, also should be classified correctly, it may prove more important to perform the logarithmic transformation before classification of the voxels.

The intensities of one class are often modeled as having a Gaussian distribution. However, in the model using logarithmic transformation of the intensities, the distribution model changes. To make the log-model resemble the original model better, it may work better if skewed distributions are introduced in the log-model, instead of Gaussian distributions.

10.1.2. Initialization of the EM algorithm

Since the EM algorithm only converges to a local minimum, the performance depends on the initialization, which is shown in the results. This shows that the initialization is important for the results.

Also, the initialization is not completely automatic, which is desired. To get the best and the most robust performance of the algorithm, efforts should be made to find better and more automatic initialization. It may be possible to find an automatic initialization at the E-step by using the cumulative histogram, see Figure 10.1. If the relative amount of each tissue type is known, the means and standard deviations can be calculated.

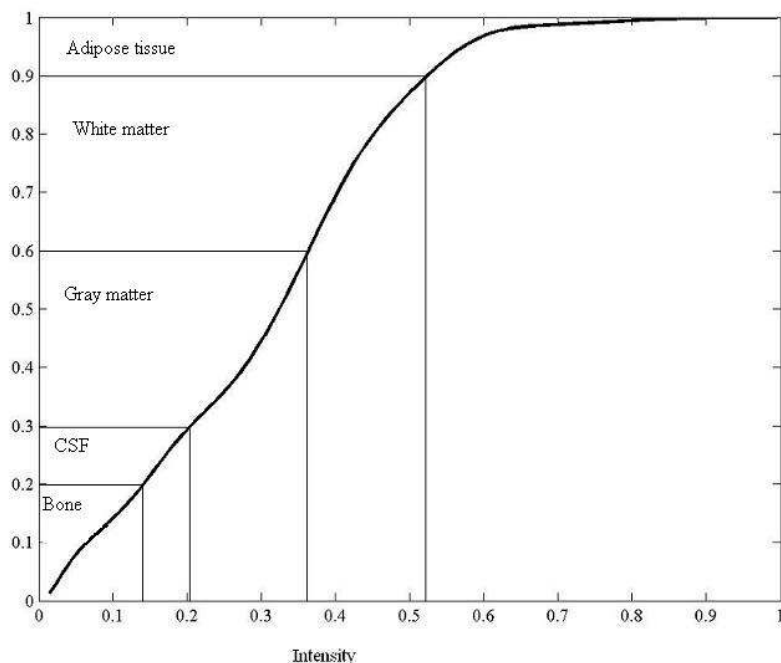


Figure 10.1. Suggestion of how to use the cumulative histogram to calculate the distribution parameters.

The EM algorithm can also be initialized at the M-step instead of the E-step. Automatic initialization may then be performed, as described in van Leemput et al (1999), by using probability maps of the various tissues. However, this introduces a registration problem.

10.1.3. The Number of Classes

The algorithm has been initialized using different number of classes. The results are affected by the number of classes used. Four tissue classes and one background class is not enough since gray matter and CSF are classified as the same class. Thus the tissues in the brain cannot be separated from each other and the brain cannot be separated from the extracerebral CSF. Also there are thicker connections between the brain tissue and the muscle tissue. Using six tissue classes as opposed to five does not improve the classifications and is thus unnecessary. Five tissue classes and one background classify the image into different tissues. The brain tissues are separated from each other and there are only thin connections connecting the brain tissues to the muscle tissue.

The parametric model used is never going to be an exact representation of the true intensity distribution. It may be argued that using more classes can make the model more exact. One idea is to introduce new classes to represent partial volume voxels.

10.2. Voxel Classification

In the voxel classification the voxels are divided into different classes. However, in the EM algorithm no spatial relationships are included. It is also important to correct for image inhomogeneities.

10.2.1. Spatial Considerations

Since the implemented EM algorithm does not take any spatial considerations into account, a perfect segmentation is impossible to achieve, since the Gaussian distributions are overlapping. Some voxels will thus always be assigned to the wrong class. Also the resulting classifications may appear granular. With the knowledge that tissues are not granular but continuous, spatial dependency should be included in the method. One way of doing this is by the use of a Markov random field model. It is possible to incorporate MRFs in the EM algorithm, see for example van Leemput et al (1999).

In this thesis work spatial considerations have been included using spatial low pass filters. This introduces spatial dependency between adjacent voxels. By the low pass filtering, the neighbors of a voxel affect the classification in that voxel. The spatial filtering has shown to remove some of the granularity and also correct some of the misclassified voxels. However, it has not been validated that the results are improved with this method. The spatial filtering could introduce errors in the classification. Nonetheless, the classifications become less granular, which may be important for the finite element modeling of the tissues.

10.2.2. Bias Correction

Both the results from the synthetic image and the MR images show that the bias correction improves the classifications. The bias estimation could be improved to make the EM algorithm perform even better.

There may be errors in the estimated bias field due to the partial volume effect. This is because these voxels belong to more than one class. The bias correction may be improved by removing partial volume voxels from the calculations of the bias field. This can be done by removing the voxels that have high values in the gradient image, since the gradient image often shows the boundaries between tissues.

The bias estimate may also be improved by weighting the contribution of the bias field in a voxel by the inverse of the variance of the class that the voxel belongs to. This is because voxels belonging to a narrower distribution give a better estimate of the bias field (van Leemput et al, 1999).

Another concern is that the filter size may not be optimal. However, different sizes have been tried and the one giving the best results was chosen.

The algorithm has been applied to the original intensities. To better compensate for the multiplicative effect of the bias field, the algorithm should be performed on the log-intensities.

10.3. Segmentation of the Brain

The brain is segmented from the other tissues using morphological operations. The segmentation is automatic, using the same morphological operations for the tested images. By using morphological operations brain voxels may be removed from the brain due to the fact

that morphological operations are simple operations. However, this loss of voxels is not very significant and does not affect the end results notably.

The use of morphological operations has given good results on the segmented images. However, there might be images which have thicker connections between the brain and the other tissues. Then morphological operations may not be enough to separate the brain from the other tissues. Active contours may be a solution to this problem (Kapur et al, 1996).

10.4. Extensions of the Method

The existing method consists of voxel classification and brain segmentation. Some improvements to the method have already been suggested in the discussion. This section introduces some suggestions to possible extensions of the method.

The resulting tissues are to be used in a finite element model of the brain. However, the complete finite element model should include more than these three tissues. The cranium and extracerebral CSF are included in the finite element model. The method could be extended to segment these tissues too. This can be done from the resulting classifications. However, the contrast between bone and CSF may be higher for other MR image weightings. If wanted other tissues like muscle and adipose tissue could also be segmented from the classifications to be included in the finite element model.

This method has been applied to T1-weighted images of the head. However, an investigation of the most appropriate weighting might have been useful. According to Pham et al (2000) the extraction of the brain volume is particularly difficult for T1-images. This is due to that non-brain tissues around the brain may have similar intensities. However, T1-images are often used due to the fact that they have high resolution even without longer acquisition times (Pham et al, 2000). Since the finite element model includes the cranium and extracerebral CSF, an extension of the method to extract bone and extracerebral CSF is wanted. Thus T1-images might not be the best choice, since bone and CSF may have similar intensities. To separate CSF from bone T2-images might be better since they have high contrast between also bone and CSF, as well as between CSF and brain matter. However, the contrast between gray matter and white matter is better in T1-images. Thus, a combination of T1- and T2-images might be appropriate. For example is the cerebrospinal fluid the parts which are bright in T2-images and dark in T1-images (Clarke et al, 1995). It is possible to extend the implemented method to include multichannel data, e.g. both T1- and T2-images. However, using multichannel data introduces a registration problem to the method.

10.5. Validation

The method should be applied to more images to see that the performance is acceptable on many images. Validation of the method can be done on both synthetic and real MR images. Applying the image on synthetic MR images, e.g. the images on BrainWeb (BrainWeb; Cocosco et al, 1997; Kwan et al, 1999; Kwan et al, 1996; Collins et al, 1998) can give a measure of the accuracy of the method. It is also important to validate the method on real MR images for example by comparisons with hand segmented images.

11. Conclusions

The implemented method segments the brain tissues from the other tissues of the head in an automatic manner. The convolutions of the brain are detected and white matter, gray matter, and CSF are separated. The method compensates for intensity inhomogeneities. However, improvements can be made to the algorithm to make it more robust and automatic. The initialization of the algorithm must be improved to make the method completely automatic. The bias field correction could also be improved to obtain better results. To obtain connected tissues it is important to include spatial considerations in the model. The method should also be evaluated using both synthetic and real MR images.

12. References

- BrainWeb, <http://www.bic.mni.mcgill.ca/brainweb/> (Acc: 2005-05-12)
- Brummer, M.E., Mersereau R.M., Eisner R.L., Lewine R.R.J. (1993), "Automatic Detection of Brain Contours in MRI Data Sets", *IEEE Transactions on Medical Imaging*, Vol. 12, No. 2, pp.153-166
- Clarke L.P., Velthuizen R.P., Camacho M.A., Heine J.J., Vaidyanathan M., Hall L.O., Thatcher R.W., Silbiger M.L. (1995), "MRI Segmentation: Methods and Applications", *Magnetic Resonance Imaging*, Vol. 13, No. 3, pp.343-368
- Cocosco C.A., Kollokian V., Kwan R.K.-S., Evans A.C. (1997) : "BrainWeb: Online Interface to a 3D MRI Simulated Brain Database", *NeuroImage*, Vol. 5, No. 4, part 2/4, S425
- Collins D.L., Zijdenbos A.P., Kollokian V., Sled J.G., Kabani N.J., Holmes C.J., Evans A.C. (1998): "Design and Construction of a Realistic Digital Brain Phantom", *IEEE Transactions on Medical Imaging*, Vol. 17, No. 3, pp.463-468
- Denzler J. & Niemann H (1999), "Active Rays: Polar-transformed Active Contours for Real-Time Contour Tracking", *Real-Time Imaging*, Vol. 5, No. 3 pp.203-213.
- Grau V., Mewes A.U.J., Alcañiz M., Kikinis R., Warfield S.K. (2004), "Improved Watershed Transform for Medical Image Segmentation Using Prior Information", *IEEE Transactions on Medical Imaging*, Vol. 23, No. 4, pp.447-458
- Held K., Rota Kops E., Krause B.J., Wells III W.M., Kikinis R., Müller-Gärtner H-W. (1997), "Markov Random Field Segmentation of Brain MR Images", *IEEE Transactions on Medical Imaging*, Vol. 16, No. 6, pp.878-886
- Kapur T., Grimson W.E.L., Wells III W.M., Kikinis R. (1996), "Segmentation of brain tissue from magnetic resonance images", *Med Image Anal*, Vol. 1, No. 2, pp.109-127.
- Kleiven S. (2002), "Finite Element Modeling of the Human Head", Doctoral Thesis, Department of Aeronautics, KTH, Report 2002-9, Stockholm, Sweden.
- Kwan R.K.-S., Evans A.C., Pike G.B. (1999): "MRI simulation-based evaluation of image-processing and classification methods", *IEEE Transactions on Medical Imaging*, Vol. 18, No. 11, pp.1085-1097
- Kwan R.K.-S., Evans A.C., Pike G.B. (1996): "An Extensible MRI Simulator for Post-Processing Evaluation", *Visualization in Biomedical Computing (VBC'96)*. Lecture Notes in Computer Science, Vol. 1131. Springer-Verlag, 1996. pp. 135-140.
- Pham D., Xu C., Prince J.L. (2000), "Current methods in medical image segmentation," *Annual Review of Biomedical Engineering*, Vol. 2, pp. 315-337
<http://medic.rad.jhmi.edu/dpham/p124r.pdf> (Acc: 2005-01-18).

Prima S., Ayache N., Barrick T., Roberts N. (2001), "Maximum Likelihood Estimation of the Bias Field in MR Brain Images: Investigating Different Modelings of the Imaging Process", In W.J. Niessen and M.A. Viergever, editors, 4th Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI'01), volume 2208 of LNCS, Utrecht, The Netherlands, pages 811-819, October 2001.
<ftp://ftp-sop.inria.fr/epidaure/Publications/Prima/22080811.pdf> (Acc: 2005-01-19).

Shan Z.Y., Yue G.H., Liu J.Z. (2002), "Automated Histogram-Based Brain Segmentation in T1-Weighted Three-Dimensional Magnetic Resonance Head Images", *NeuroImage*, Vol. 17, No. 4, pp.1587-1598

Udupa J.K. & Saha P.K. (2003), "Fuzzy Connectedness and Image Segmentation", *Proceedings of the IEEE*, Vol. 91, No. 10, pp.1649-1669

van Leemput K., Maes F., Vandermeulen D., Suetens P. (1999), "Automated Model-Based Bias Field Correction of MR Images of the Brain", *IEEE Transactions on Medical Imaging*, Vol. 18, No. 10, pp.885-896

van Leemput K., Maes F., Vandermeulen D., Suetens P. (1999), "Automated Model-Based Tissue Classification of MR Images of the Brain", *IEEE Transactions on Medical Imaging*, Vol. 18, No. 10, pp.897-908

Wells III W.M., Grimson W.E.L., Kikinis R., Jolesz F.A. (1996), "Adaptive Segmentation of MRI Data", *IEEE Transactions on Medical Imaging*, Vol. 15, No. 4, pp.429-442

Zhang Y., Brady M., Smith S. (2001), "Segmentation of Brain MR Images Through a Hidden Markov Random Field Model and the Expectation-Maximization Algorithm", *IEEE Transactions on Medical Imaging*, Vol. 20, No. 1, pp.45-57