Adaptive online learning based tissue segmentation of MR brain images

Damkat, C.

Award date:
2007

TU/e technische universiteit eindhoven

Faculty of Electrical Engineering
Section Design Technology For Electronic Systems (ICS/ES)
ICS-ES 885

Master's Thesis

Adaptive online learning based tissue segmentation of MR brain images.

Chris Damkat

Coach:
Supervisor: Prof.dr.ir. G. de Haan
Dr. Ahmet Ekin (Philips Research)
Date: April 2007

Adaptive online learning based tissue segmentation of MR brain images

Master Thesis

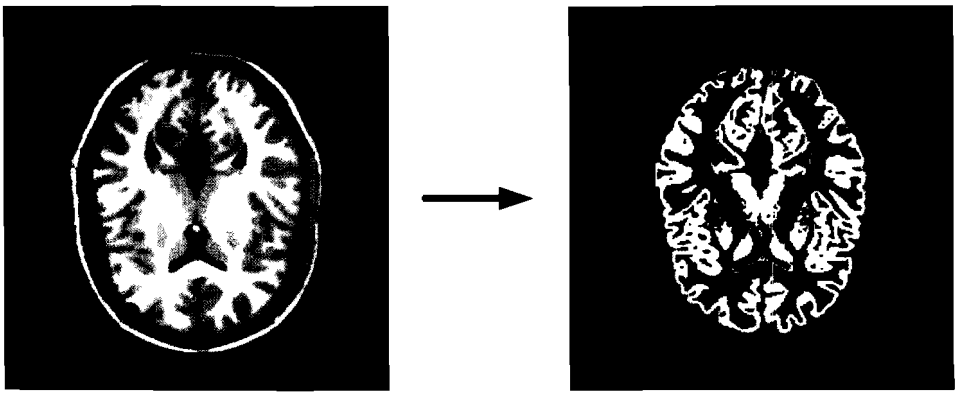

Chris Damkat

Electrical Engineering

Eindhoven University of Technology

Supervisors

Dr. Ahmet Ekin[1]

Prof. Dr. Ir. Gerard de Haan[1,2]

[1]Philips Research Eindhoven, Video Processing & Analysis

[2]Eindhoven University of Technology

March 2007

Preface

This work was performed at Philips Research in the Video Processing & Analysis department, under supervision of Ahmet Ekin and Gerard de Haan. In this department there is much experience with video processing algorithms for television, most importantly with motion estimation. It is interesting to see how this knowledge suits other problems, and specifically in this work how it fits MR image segmentation.

Based on *3D recursive search* motion estimation (3DRS) [DeHaan93] an algorithm has been developed for magnetic resonance image based brain tissue segmentation. Mainly the aspect of so-called online (sample-by-sample) learning was introduced to this problem and was adopted from 3DRS. This resulted in an interesting new algorithm that differs from the common approach. An intermediate paper [Damkat06] was written and presented at the *Symposium on Intelligent Algorithms 2006*[1] (SOIA06). Besides the basic algorithm described in that paper, further developments and validation of the algorithm and an implemented bias correction method are documented in this work.

The author would like to thank Ahmet and Gerard for their support and supervision. Furthermore thanks go to the department and fellow students for a wonderful time at Philips Research.

Chris Damkat,

March 2007.

[1] http://www.extra.research.philips.com/soia/

Abstract

The aging population in the European Union and the US has increased the importance of research in neurodegenerative diseases. Imaging plays an essential role in this endeavor by providing insight to the intricate cellular and inter-cellular processes in living tissues that will otherwise be difficult, or impossible, to gain. Because of the sheer size of the imagery data, the lack of sufficient medical staff, and the inaccuracies resulting from manual processing, automated processing of image-based data to generate quantitative and reproducible results is necessary. To this effect, in this thesis a fully automatic image-processing algorithm for brain tissue segmentation from *magnetic resonance* (MR) images is proposed. Contrary to the present iterative *expectation maximization* (EM) based algorithms, it uses online (sample-by-sample) learning to adapt to the intensity inhomogeneity inherent to MR images. Since the proposed method can adapt to the intensity inhomogeneity online, multiple iterations over the data as in the present algorithms are not necessary, and consequently the processing time is decreased dramatically. The used online learning scheme is based on *learning vector quantization* and is further tailored to the segmentation of MR images by integration of spatial context and the use of a special locality-preserving scanning order of the data. Explorations of various scanning orders and a modification to the learning rule to allow for 3D learning have lead to three variants of the proposed algorithm. These proposed methods are validated by comparing the segmentation masks to basic k-means clustering, and present EM-based methods, namely, FAST and the state-of-the-art EMS, on simulated and real datasets. The proposed methods demonstrated a significant reduction of the processing time (a factor of 20) compared to the EM-based methods. Tests on BrainWeb simulated data showed that segmentation accuracy is comparable to the EM-based methods, however, tests on real data, where the segmentations of EMS were used as ground truth, showed lower accuracy than the EM-based FAST. Moreover, the tests on real data showed that the proposed methods as well as FAST make a significant amount of misclassifications in the so-called deep grey matter areas, which suggests the necessity of a spatial prior atlas as it is used in EMS.

1 Introduction

The significance of effective healthcare solutions has increased dramatically as a result of the aging population worldwide, particularly in North America and in the European Union. Because age is considered as one of the high risk factors for some neurodegenerative diseases, neurology will be immediately affected by this demographical change. The related statistics across the EU and the US also supports this claim. In Europe, the European Brain Council estimates the cost regarding neurological diseases as 139 billion euros per year. In the U.S. alone, 4.5 million people have Alzheimer's disease and by 2050, there will be 11 to 16 million patients.

Neuro-research requires high-resolution imaging of the brain (see Figure 1 for a reference of the brain). One of the best available medical imaging options for the brain is *magnetic resonance imaging* (MRI). MRI is a non-invasive medical imaging modality that captures the images of internal soft tissues at a very high contrast compared to the alternative *computed tomography* (CT), without exposing the body to harmful ionizing radiation. MR images can be used for diagnoses of several brain related diseases, e.g. Multiple Sclerosis, Alzheimer, schizophrenia, epilepsy, Creutzfeldt-Jakob, brain tumors and lesions [VanLeemputPHD01], as well as for other quantitative applications. As most neurological diseases are associated with structural changes in the brain, finding an outline of the brain structures and quantifying the volume of certain structures is invaluable for neuro-research.

This work deals with automated MRI-based brain tissue segmentation/classification, where the objective is to segment the brain image into the major tissue types: grey matter, white matter, and cerebrospinal fluid (see Figure 3). In most cases, the resulting large set of MR images after scanning makes human-based processing impractical and demands for automated solutions. Furthermore besides being very time-consuming and thus expensive, human expert segmentations are subjective, i.e., there is large inter- and intra-observer variability of the obtained segmentations. Thus, an automatic segmentation algorithm, which delivers reproducible, objective results and has computationally attractive complexity, is of high-value for medical applications. Therefore, this is an area of extensive research, as is also indicated by the high number of publications on this subject in the recent years.

As a contribution to this area, a new brain tissue segmentation algorithm is proposed in this work that is based on so-called online (sample-by-sample) learning. Using online learning the tissue class statistics can be deduced from the data and adaptation to the changes of these statistics from one location to another, due to the intensity inhomogeneity inherent to the MR images, is possible. This approach is different from the methods found in the (present) literature, which model the intensity inhomogeneity as a multiplicative bias field and iteratively improve approximations of the bias field and segmentation. Since the proposed method can adapt to the intensity inhomogeneity online, multiple iterations are not needed and the processing time is decreased dramatically. The used online learning scheme is based on *learning vector quantization* (LVQ) [Kohonen86], and is further tailored to the segmentation of MR images by integration of spatial context and the use of a special locality-preserving scanning order of the data. The proposed algorithm is validated by comparing its segmentation masks to basic k-means clustering, and Expectation Maximization based

methods, namely, FAST[2] [Zhang01] and the state-of-the-art EMS[2] [VanLeemput99], on simulated BrainWeb [BrainWeb] and real datasets.

What follows next in this introduction, is a short description of the principles behind the MR images, the problem definition describing the goal and challenges of this research, the main contributions and finally the global outline of this complete work.

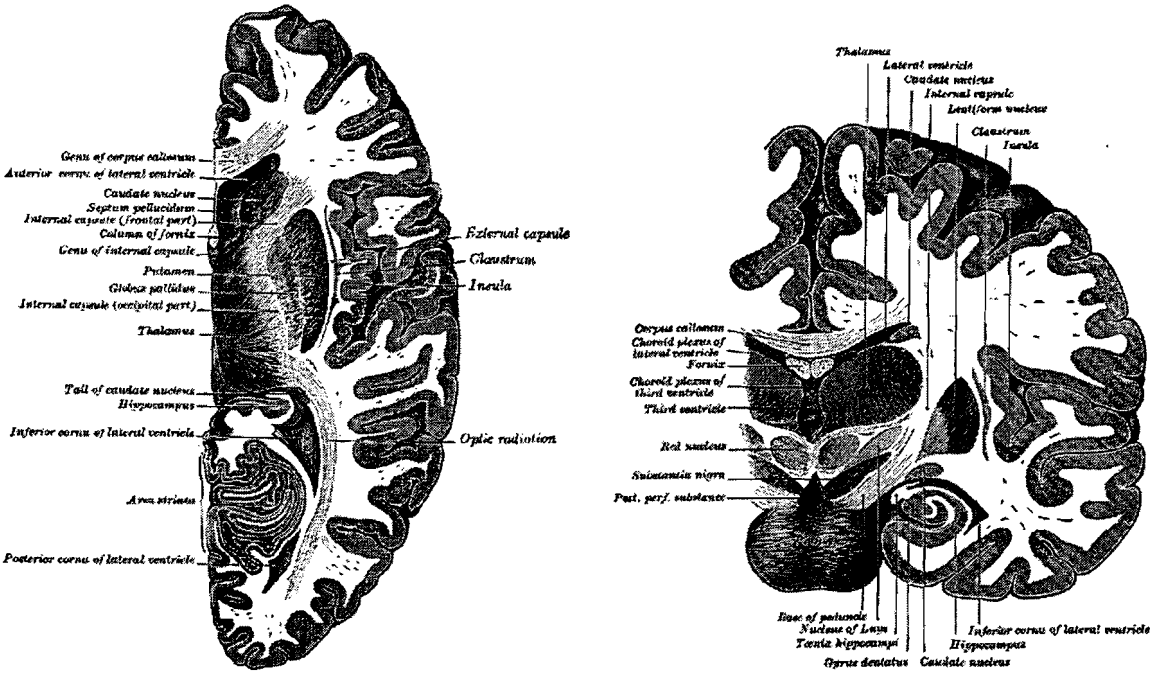

(a) (b)

Figure 1: Anatomical illustrations from 'Gray's Anatomy' [Gray18]. (a) Axial slice and (b) coronal slice showing the brain and its deeper structures.

1.1 MR Images

MRI is an imaging methodology which can be used to visualize the inside of living organisms. The obtained images can be used for medical purposes to diagnose patients without having to penetrate the body, i.e. it is non-invasive. Images can be obtained of internal organs, for example images of the brain. MRI is best suited for brain imaging as it is able to produce contrasts of soft tissues, moreover, compared to a CT scan, the MRI scanner does not use ionizing radiation, which is detrimental to the human health in large doses.

MRI is based on the principle of *nuclear magnetic resonance* (NMR), and for medical purposes it is based on the resonance of the hydrogen nuclei, i.e., single protons. These protons behave as small magnets, with a magnetic moment, i.e., they rotate. In an MRI scanner the subject is placed in a powerful uniform static magnetic field, the so-called B_0 field, which is directed in the z-direction, along the length axis of the patient. The field strength is in the order of 0.3 to 3 Tesla (or even 7 Tesla in research), this is in the order of 100.000 times higher than the Earth's magnetic field. The protons align with this magnetic

[2] FAST = FMRIB Automated Segmentation Tool, EMS = Expectation Maximization Segmentation, FMRIB = Oxford Centre for Functional Magnetic Resonance Imaging of the Brain

field, parallel and anti-parallel, a slightly higher amount (~one in a million) is aligned parallel with the field; consequently, the subject is magnetized in the z-direction, i.e. the longitudinal direction.

This magnetization has a typical 'frequency', called the Larmor frequency, that is dependent on the gyro-magnetic ratio of hydrogen, and linearly dependent on the strength of the magnetic field. For hydrogen in a field with strength 1.5 Tesla the Larmor frequency is 63.864MHz. If an RF-pulse introduces a time-varying magnetic field with this specific frequency perpendicular to the B_0 field, the magnetization is flipped partially or totally (depending on the strength and duration of the pulse) in the xy-plane, i.e. the energy is absorbed and the protons are excited. This time-varying magnetic field is called the B_1 or excitation field. After excitation, the magnetization in the xy-plane rotates with the Larmor frequency, since all nuclei are now in-phase. After the RF-pulse is switched-off the magnetization starts to return to the z-direction, and while this happens the nuclei return the RF-signal, this signal can be received. Basically, the (transversal) magnetization in the xy-plane dies out while the longitudinal magnetization re-establishes its original value; this is due to two relaxation processes.

The first process is the so-called *spin-lattice relaxation*, T1, characterized by the time for the z-component of the magnetization to re-establish its original value. The second process is the *spin-spin relaxation*, T2, and it refers to the disappearance of the xy-component of the magnetization due to de-phasing of the individual magnetization vectors of the nuclei due to spin-spin interaction. The individual returned signals emitted by the nuclei are modulated by properties of the tissue they are in. Finally, the signal is also modulated by the amount of protons in a volume element which gives the *proton density* (PD) and is also tissue specific.

By choosing specific scanning parameters, the so-called T1 or T2 relaxation times or the PD for the tissue in a specific location can be retrieved. As the returned values are tissue specific, a contrast (a greyscale or multi-spectral[3] image) can be obtained. The resulting 3D image is called a volume, consisting of 2D slices, see Figure 2. An element in this volume is called a voxel, which is the equivalent of a pixel in a 2D image.

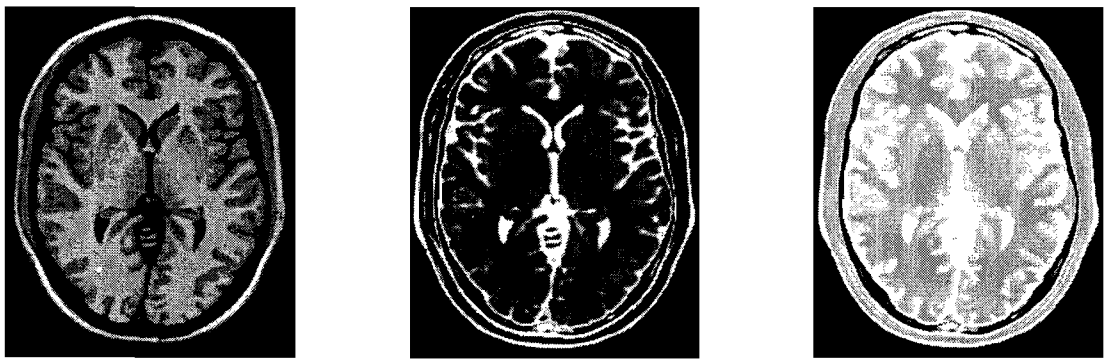

(a) (b) (c)

Figure 2: Example slices of several 'modalities', these are simulated images from the BrainWeb database. (a) T1-weighted image, (b) T2-weighted image, and (c) PD-weighted image.

For a brain scan, the resolution of a volume is typically in the order 256×256×180 voxels, and the dimensions of a voxel are in the order $1\times1\times1\text{mm}^3$, and typically 16 bits per voxel are used to store the intensity (although most formats support 8 bits or 32 bits as well). In Figure 2, examples of a T1-weighted, T2-weighted and PD image are shown. In these images, it can be seen that an MR image presents a contrast where the basic tissue types can be distinguished

[3] By combination of multiple modalities (e.g. T1 and T2 or T2 and PD).

based on the intensity values. Note that the different anatomical structures as indicated in Figure 1 cannot be distinguished by intensity alone. The presented images above give a rather ideal image, which is not encountered in practice. The real MR scans contain noise and intensity inhomogeneity, which is a shading effect, and other distortions, that hamper straightforward tissue segmentation of the brain.

1.2 Problem definition

The brain is composed of several tissue types, namely *white matter* (WM), *grey matter* (GM), *cerebrospinal fluid* (CSF), and in a head scan also skull, muscle, skin[4], and *background* (BG) are included. Finding the outline of the different brain tissues and quantifying the amount of each tissue may be invaluable for the quantification of brain atrophy, which is correlated with the progression of neurodegenerative diseases, or for studying the brain development over time in general. For example in Alzheimer disease, which is associated with neuronal loss, shrinking of the cortex, hippocampus and other structures can be observed while the ventricles enlarge [Fischl02]. Basically the brain shrinks, and the GM amount decreases while the CSF amount increases. If a segmentation of the brain in the basic tissue classes is available, the progress of volumetric changes in these tissues can be monitored.

This main task is illustrated in Figure 3, where we see the input MRI image on the left and the desired segmentation (classification) on the right. The image shown on the left is a T1-weighted image; these images are commonly used for the segmentation as they present a good GM-WM contrast and a good signal-to-noise ratio. On the right a labeled image with the tissue classes CSF, GM, and WM is shown. From this figure we can see that the task is to label every voxel with a tissue class. In the ideal MR, image every tissue class has its own absolute intensity and once the tissue intensities are known, segmentation could be accomplished by simple multi-level thresholding with the appropriate thresholds.

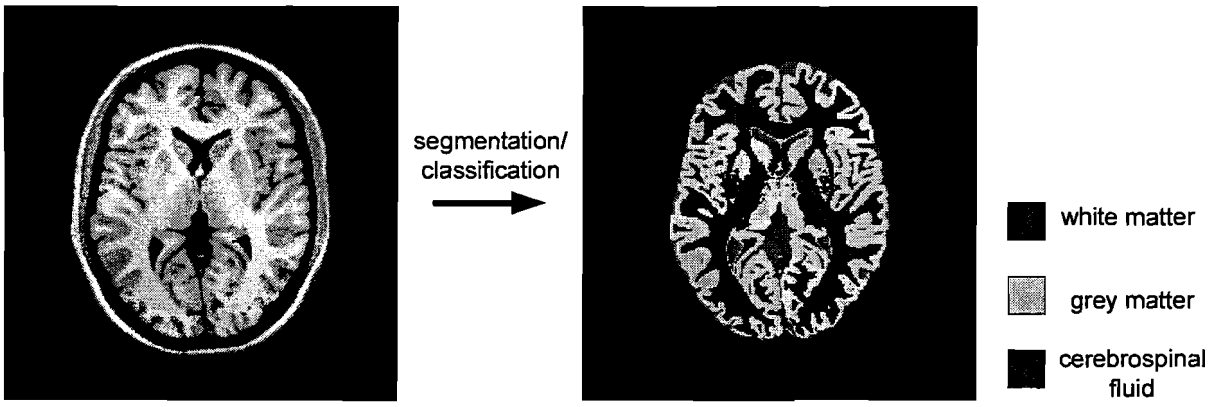

Figure 3: The objective of the algorithm. On the left we see the 'greyscale' input image, on the right the segmented image.

MR-based image segmentation, however, is not straightforward. The intensity values of a real MR image vary with the scanning parameters; hence, the use of a standard thresholding or a non-adaptive algorithm is not possible. Furthermore, MR images are degraded by field inhomogeneity that results from various factors, such as imperfectness of magnetic coils, patient attributes, etc., which cause intra-scan intensity variations of the tissues. This inhomogeneity is generally modeled as a bias field, which is a spatially low-frequency

[4] Skull, muscle, and skin and other non-brain tissues are ignored further in this work as they can be removed with a so-called skull stripping algorithm, such as BET (see Chapter 6).

(slowly varying) multiplicative field [Sled98] and results in smearing of the intensity histogram (see Figure 4); hence, the tissue classes overlap in intensity. In addition to the above, there is also noise, which is assumed to be Gaussian[5] and additive in this work and further increases the tissue class overlap. Then finally, the limited imaging resolution may result in multiple tissues contributing to a single voxel intensity, leading to a blurred appearance of the image. This phenomenon is called the *partial volume* (PV) effect and may require special attention if severe. Figure 4 shows a visual impression of the above-mentioned degradations.

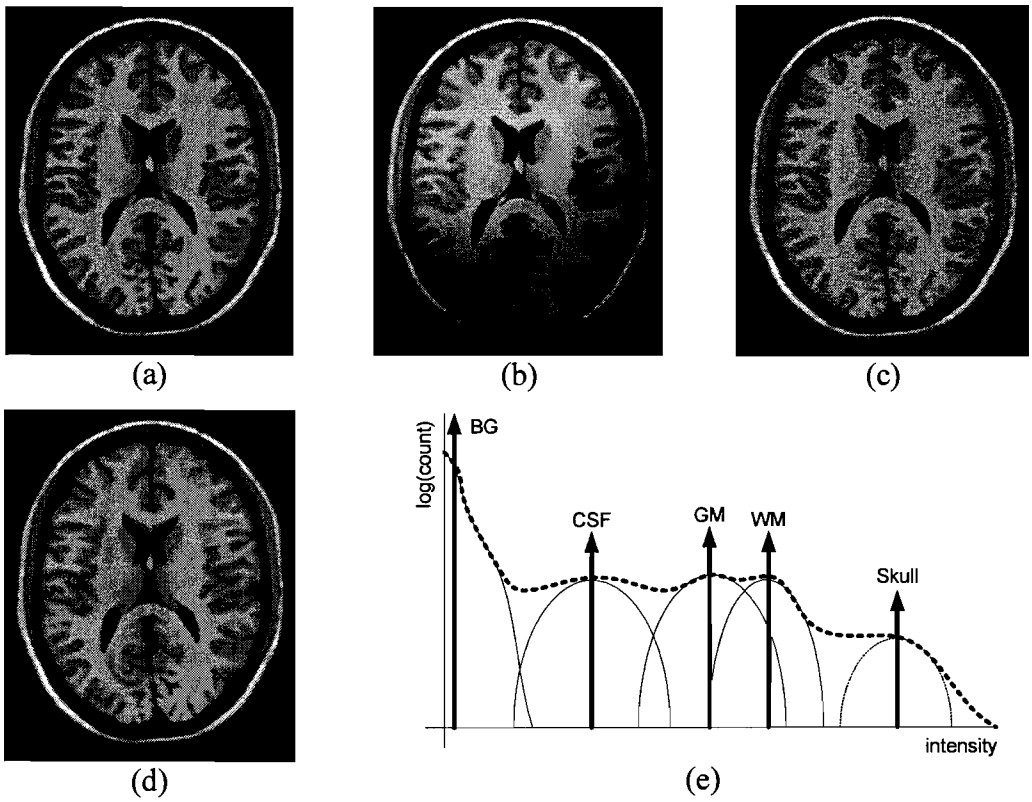

Figure 4: [BrainWeb] simulations illustrating the typical degradations of the ideal image. (a) Ideal image, (b) image with intensity homogeneity, (c) image with noise, and (d) image with partial volume effects due to thicker slices. (e) The intensity histogram where the arrows indicate the ideal histogram, where each tissue has one intensity, and the dotted line indicates the smeared realistic histogram as a result of the degradations.

These degradations hamper the tissue segmentation of the MR brain images. A successful segmentation algorithm should handle all these effects to obtain a valid segmentation. Furthermore, this should be accomplished without user intervention, i.e. unsupervised, to obtain an objective and reproducible result. Therefore, the tissue statistics should be extracted from the data itself. And in the calculation of the statistics and the following classification, inhomogeneity effects, noise, and PV if severe should be well accounted for. The obtained segmentation should be accurate, as medical decisions will be based on it. Finally, the segmentation should be obtained within a reasonable time. Present methods which are publicly available, such as FAST [Zhang01], which is a C implementation of an EM-based method, take about 10 minutes, on a Pentium 4 3GHz with 512 MB of memory, for a dataset with a resolution of 256×256×120 voxels, while more complex methods described in the literature can take up to 30 minutes or more [Awate06]. In this thesis, we aim for a different

[5] The actual noise is Rician, however, Gaussian noise is a good enough approximation for this application, since the signal to noise ratio is relatively high ($\mu >> \sigma$).

approach to reduce the computational complexity and that will show to significantly decrease the necessary processing time.

The proposed algorithm should solve the following challenges:

Challenges:

- It should be robust to the effects inherent to the MR images, being:
    - Inter-scan intensity variations.
    - The intensity inhomogeneity, the intra-scan intensity variations.
    - The partial volume effects, if severe.
    - Noise.
- It should work without supervision and thus obtain objective and reproducible results.
- It should obtain the segmentation in a reasonable amount of time.
- The segmentation should be accurate.

1.3 Main Contributions

This thesis presents a new algorithm for tissue segmentation of MR brain images. The proposed algorithm was influenced by *3D recursive search* (3DRS) motion estimation [DeHaan93], an algorithm for motion estimation in video sequences (for a more detailed description of 3DRS and its relation to the proposed algorithm refer to Appendix B). This influence led to a new algorithm for segmentation of the MR images, which differs from the current algorithms found in the literature. The main point where it differs from the current algorithms is the way it treats the intensity inhomogeneity, namely the proposed algorithm inherently adapts to the inhomogeneity by learning/adapting 'online' (sample-by-sample). Thereby it is not necessary to have multiple iterations over the data to segment the image since it is able to segment the image in one pass over the data. The used online learning scheme is based on LVQ (see Chapter 4), and is further tailored to the segmentation of MR images by integration of spatial context and the use of a special scanning order of the data.

Main contributions:

- Online learning has been introduced to learn the tissue statistics and implicitly track the intensity inhomogeneity inherent to MR images, which results in a significant decrease of computational complexity.
- Different scanning orders have been explored to relax adaptation requirement of the online learning scheme.
- Influence of spatial context has been integrated in the online learning scheme, to add robustness to noise and improve overall stability.
- A modification to the online learning rule has been made to allow for 3D learning of the tissue statistics.
- A separate bias field correction algorithm has been implemented in the initial phase of this research.

1.4 Thesis outline

In Chapter 2 the current methods for handling the bias field and segmentation will be discussed to be able to put the proposed method into perspective. In Chapter 3 an implementation of a bias field correction method is described. Thereafter, in Chapter 4 LVQ is introduced with relevant variants on which the proposed segmentation method is based. In Chapter 5 the proposed method is described with some variants, which will be validated in Chapter 6 by comparing it to base-line methods on simulated and real datasets. Finally, in Chapter 7 the conclusion and discussion are given.

2 Literature on tissue segmentation and bias field correction

This chapter gives an overview of the current methods for bias field correction and segmentation. The description of bias field correction is relevant since the image inhomogeneity is the main artifact hampering the intensity-based segmentation, and the handling of the inhomogeneity plays a central role in the proposed method. Furthermore, there is a strong relation between bias estimation and tissue segmentation, since once the bias is known the segmentation is simplified a lot and conversely once the segmentation is known the bias can be deduced. This fact is used widely in the iterative methods for bias field correction and tissue segmentation. It eventually leads to expectation maximization based segmentation methods integrating bias field correction, classification, tissue class model updates, markov random field parameter estimation, and even registration to an anatomical or tissue atlas in one iterative loop.

2.1 Bias field correction

The intensity inhomogeneity (a.k.a. shading, nonuniformity) (Figure 4(b)) is caused by, among other effects, coil imperfectness, gradient-driven eddy currents, and electrodynamic interactions with the subject, which makes it also patient dependent. Furthermore, its shape is also dependent on the scanning parameters. The inhomogeneity presents no (large) difficulties for visual inspection, however, it does hamper automated analysis profoundly, such as in intensity-based tissue segmentation where the intensities are used to identify tissues and are assumed to be the same throughout the whole volume. Although MR scanners can be improved, the inhomogeneity is an inherent problem in MR imaging. One could measure the inhomogeneity in vivo (during the scan) or by the use of phantom. However, as this approach requires extra scans or hardware, which are not always available or simply impractical, many retrospective (data-driven) methods have been published.

In these methods, the intensity inhomogeneity is generally explicitly modeled as a smooth multiplicative field extending over the entire image [Sled98], although some methods also consider an additive component, such as the method by Likar et al. [Likar01], this is largely ignored.

The model of the bias field is thus described by:

$$I_d(\mathbf{r}) = I_t(\mathbf{r}) \cdot b(\mathbf{r}) + n(\mathbf{r}), \tag{1}$$

where, $\mathbf{r} = [x, y, z]$ is the location, $I_d(\mathbf{r})$ is the observed image, $I_t(\mathbf{r})$ is the true image, $b(\mathbf{r})$ is the bias field, and $n(\mathbf{r})$ is additive noise. Some methods log-transform the data to change the multiplicative bias field to an additive one which simplifies the computation:

$$\log(I_d(\mathbf{r})) = \log(I_t(\mathbf{r})) + \log(b(\mathbf{r})), \tag{2}$$

where the additive noise is ignored.

This low-frequency bias field could be removed by high pass filtering the image. Since the bias field is multiplicative, the image is then first log transformed before the low frequencies can be subtracted, this is called homomorphic filtering. Unfortunately, the frequency spectrum of the brain overlaps that of the bias field, so the estimated bias-field is influenced by the

brain structure as well, and thus by high pass filtering, intensities of structures of the brain are distorted.

Given that the solution is not straightforward, many different methods can be found in the literature. For example: methods using filtering combined with some preprocessing to prevent the distortion of the brain structures [Cohen00; Cheng06], gradient based methods [Garcia05; Luo05; Saha05], methods based on normalization of local histograms [Shattuck01], based on classification/segmentation [VanLeemput99; Madabhushi06; Wells96], and methods based on the global histogram [Sled98; Likar01; Vovk04; Vovk06].

Of all these methods one of the most cited methods is the method of Sled et al. [Sled98], which is a nonparametric method for automatic correction of intensity nonuniformity in MRI data, commonly known as N3. The term 'nonparametric' indicates that it does not rely on a parametric model of the tissue intensities, and thus makes no assumptions and is fully automatic. This iterative method is based on the notion that, due to the intensity variation introduced by the bias field, the intensity distribution of the bias-distorted image is the blurred version of the intensity distribution of the original undistorted image. The idea now is that the distribution of the bias free image can be deduced from the distribution of the image containing a bias field by deconvolving it with the blurring kernel and thus by sharpening (or condensing) the distribution. The mapping of the intensities between these two distributions is applied to the image and spatially smoothed using B-splines to obtain a bias field estimate. This process (Figure 5) is repeated until it converges which requires in the order of 20 iterations. In this process sub-sampled data can be used since the bias field we are trying to estimate is smooth.

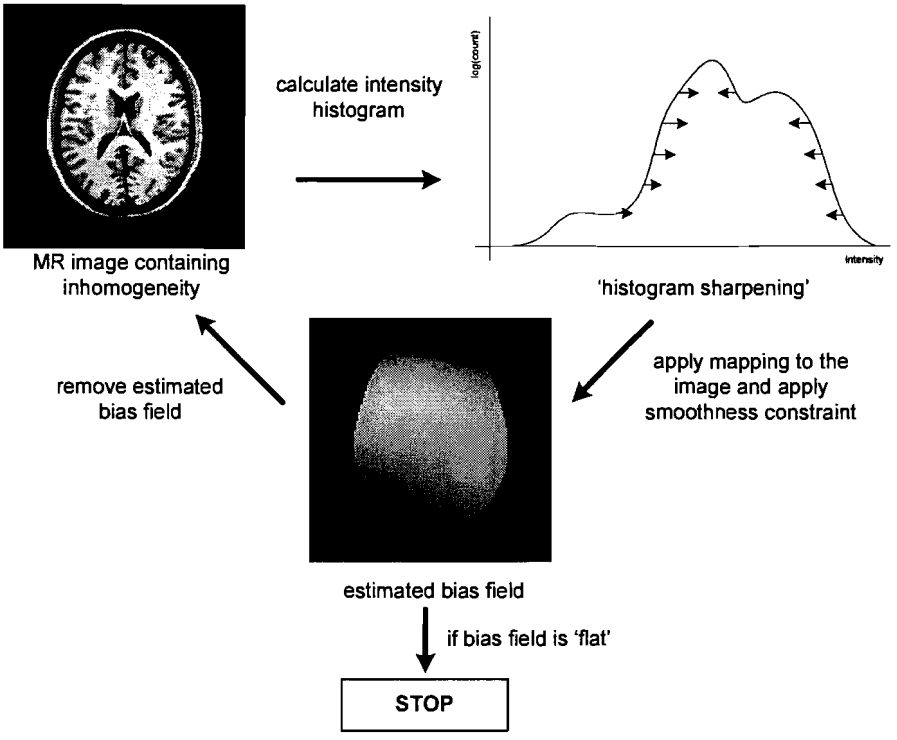

Figure 5: The process used to remove the bias from MR images as in N3 and Vovk's methods.

There are other methods using a similar technique as N3. The method by Likar et al. [Likar01] models the bias field with second order polynomials and the entropy of the image is minimized, i.e. the intensity distribution is condensed, by adjusting the parameters of the polynomial bias-field model. The method by Vovk et al. [Vovk04] calculates 'forces' which will condense the histogram and like in N3 maps these to the image and this estimate is smoothed using a Gaussian kernel. However, in calculating the forces, the second derivative

of the image is taken to create a two-dimensional distribution (feature space) in which there will be more discrimination between tissue types. This is helpful since in the one-dimensional intensity histogram the tissue intensity distributions overlap considerably. In another method by Vovk et al. [Vovk06] the feature space is raised to higher dimensions by combining different modalities (e.g. PD and T2) and the forces are calculated within this space, thereby allowing better discrimination between tissues and reducing overlap, which improves the performance.

In the paper by Sled et al. [Sled98] a limitation hampering all retrospective methods is given, which is that they cannot distinguish between variations caused by the bias and slow natural variations, and thus also remove these slow natural variations. However, especially in high Tesla MRI [Hu04] (3T or higher), the bias is much stronger than the subtle natural variations and renders these insignificant. In the literature, papers giving an overview/review of the methods are available [Sled97; Arnold01; Belaroussi06; Huo06; Vovk07]. From the evaluation by Arnold et al. [Arnold01], where six algorithms are compared, it can be concluded that N3 is robust and delivers the best performance together with the method proposed by Shattuck et al. [Shattuck01]. Although Shattuck's method performs a little better in case of mild bias it performs worse in case of strong bias [Huo06]. Furthermore, given that Shattuck's method uses an explicit brain tissue intensity distribution model for which the parameters are estimated from the global histogram, which makes the method less robust and less generalized, N3 and similar methods can be considered the most valuable.

2.2 Tissue segmentation

In brain tissue segmentation, the objective is to segment a brain MR scan into the basic tissue types, namely the CSF, GM and WM classes, which can be identified by their intensities in the image (see Figure 3). However, given the degradations of the ideal image described in Section 1.2, this is not straightforward. Besides inter-scan intensity variation, the intensity inhomogeneity and noise cause the intensity distributions to overlap, and thus standard histogram based thresholding, with the help of, e.g., k-means clustering (see Appendix A) to obtain these thresholds, leads to misclassifications (Section 6.2).

Therefore, numerous methods, which aim to overcome these degradations, have been published and research in this field is still ongoing. In this work we focus on the classification-based, intensity-driven methods, which segment the brain by classifying (individual) voxels based on their intensities, although extra information is often integrated in the form of spatial context and optionally in the form of an atlas[6]. The intensities can be multi-spectral due to the multiple modalities that can be obtained by changing the parameters of MR scanner. It should be noted that besides the classification-based methods, there are region-/edge-based methods, like [Jimenez-Alaniz06; Kong05]. In such methods, homogeneous regions of voxels are grouped. These region-/edge-based methods are not described here further since most, including the state-of-the-art, methods and our proposed method are intensity-based.

For segmenting an image while handling the bias field degradation, Wells et al. [Wells96] developed one widely cited classification-based method. This is an iterative method based on *expectation maximization* (EM) and includes bias field estimation. In this algorithm, the tissue intensity distributions are modeled by a mixture of Gaussians with parameters mean, variance, and prior probability for each tissue (see Figure 6 for an impression of this mixture model on the lower right). At the start of the algorithm the parameters are all initialized, the mean,

[6] An atlas is a spatial probability map with prior tissue probabilities for each location.

variance, and prior probability are obtained from training, and the estimated bias field is initially flat. Using the tissue class model and the estimated bias field, tissue class probabilities are assigned to all the voxels (E-step) and from this classified volume the estimated bias field can be improved (M-step). The process is repeated for multiple iterations until convergence (typically 5-10 on 1.5 Tesla data).

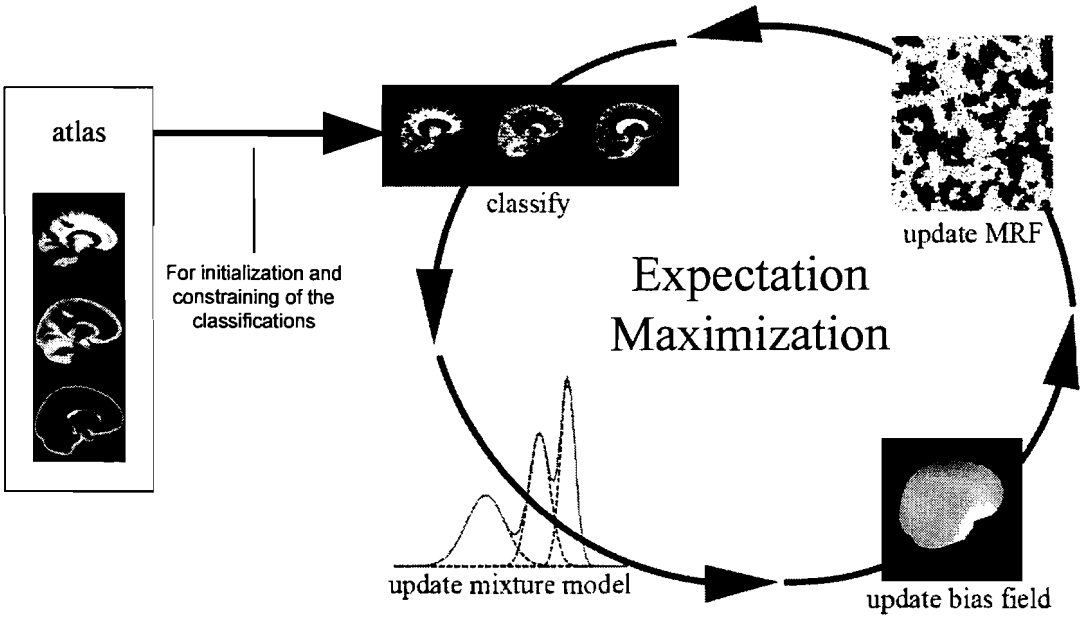

Figure 6: Iterations in the Expectation Maximization method of Van Leemput.

Van Leemput et al. [VanLeemput99] expanded Wells' method by updating the Gaussian mixture model parameters in the iterations, with the use of spatial context via *markov random fields* (MRFs), and integration of atlas information (See Figure 6). This method is considered state-of-the-art in many publications and is often used as a reference. The use of spatial context mainly overcomes the misclassification due to noise in the image since it introduces dependency among neighboring voxels and thus resolves ambiguities (the tissue class overlap due to noise). The atlas provides good initial estimates of the tissue parameters, and during the iterations the atlas is used to constrain the classifications, this improves the robustness of the method to severe bias fields. However, the atlas needs to be spatially registered to the brain prior to the segmentation process, which requires additional computations.

More EM-based methods can be found in the literature [Fischl02; Fischl04; Zhang01; Murgasova06], and similar iterative methods based on fuzzy c-means clustering [Bezdek95], including spatial context models and bias field estimation, are available as well [Liew06; Shen05; Pham99; Cai07]. Although the performance of these iterative methods is good, the main drawback of these iterative methods is that they are time-consuming, since they go through the whole dataset multiple times.

Overviews discussing mainly the methods described above, including the region-based methods, are available in the literature [Liew06; Clarke95; Cuadra05; Pham00; Suri01].

Recent developments are methods combining segmentation and registration to an atlas [Ashburner05; Pohl06] in the iterative loop. In these methods the registration and segmentation are integrated in the iterative EM framework, and not applied sequentially as in Van Leemput's method. Integration of registration in the iterations improves the accuracy of the registration and segmentation, and if an accurately registered atlas contains anatomical labels, anatomical structures can be segmented automatically. For example, in the method by

Pohl et al. [Pohl06] the atlas helps in the segmentation of the thalamus (see Figure 1), which is difficult to segment accurately based on intensity alone.

The proposed method in this thesis handles the inhomogeneity implicitly, while most methods with bias field compensation introduced above model the intensity inhomogeneity explicitly as a multiplicative bias field. Implicit adaptation to the inhomogeneity is achieved by classifying using a tissue intensity model deduced from local sampling of the voxel data, which is valid since the bias field is smooth, and thus the tissue statistics are locally stationary. Other methods handling the bias field implicitly by local sampling are, for example, the methods by Awate et al. [Awate06] and Yan et al. [Yan95]. However, the local samplings in these methods have overlap and thus there is no gain in speed contrary to the proposed method where this overlap is avoided by recursive learning and a special scanning order of the data.

3 Implementation of bias field correction

A new bias field correction method has been implemented; this was done in the early stage of this research to get acquainted with the data and it is a valuable tool for further analysis of MR images. The implementation is similar to the methods proposed by Sled et al. [Sled98] and Vovk et al. [Vovk04], although there are some differences. A detailed description of the method is given in this chapter. The method is validated on some images in Chapter 6, although only qualitatively, since it is not the main subject of this thesis. Furthermore, it is used to compare the performance of the proposed segmentation algorithm before and after bias correction.

3.1 Methodology

As described in Section 2.1, the inhomogeneity in the MR image is commonly modeled as a smooth multiplicative bias field, see Equation (1). In the implemented algorithm we also assume this model and like most methods an additive component of the bias field is ignored. The multiplicative bias field introduces additional variance to the image, which is expressed in the intensity histogram of the image as dispersion, or 'blurring'. Conversely, the histogram of the unknown bias-corrected image can be estimated from the histogram of the image with bias by condensing the histogram. In the implemented method this condensing is performed by application of the mean shift [Cheng95], contrary to Sled and Vovk's methods where other but similar methods are used to condense the histogram. Mean shift is a simple procedure where each point in vector space is mapped to the mean of the data points in its neighborhood. In this method we use the mean shift on the intensity histogram to obtain a mapping between intensities of the distorted image and those of an estimated bias-free image. From this mapping we can calculate intensity-dependent, voxel-wise bias estimates, which can be smoothed to obtain a bias field estimate. This process is repeated until convergence. See Figure 5 for an impression of this algorithm.

Preprocessing

Standardization

Since the global intensity of the MR scan can vary largely, the image is standardized prior to the application of inhomogeneity correction. This standardization is performed by linearly scaling the intensity range of the image such that the 90th percentile of the intensity histogram is set to 400.

Masking non-brain areas

The skull and other outliers, which might influence the bias estimation negatively, are masked by selecting the larger smooth areas using a simple morphological operator[7]. The smooth areas that remain are mainly the brain tissues. The morphological operation is applied per slice and defined as follows:

$$I_{preprocessed}(x,y) = \begin{cases} 0, & \text{if } \mathrm{diff}(x,y) > T_{smooth} \\ I(x,y), & \text{otherwise} \end{cases} \quad (3)$$

where

[7] Alternatively the non-brain areas can be removed using a skull-stripping algorithm.

$$\text{diff}(x,y) = \left(\begin{array}{l} (|\, I(x+d,y) - I(x,y)\,| + |\, I(x-d,y) - I(x,y)\,| \\ +|\, I(x,y+d) - I(x,y)\,| + |\, I(x,y-d) - I(x,y)\,| \end{array} \right) \Big/ I(x,y) \quad (4)$$

where $I(x,y)$ and $I_{preprocessed}(x,y)$ are the input and pre-processed slice, d is set to be 5mm and the smoothness threshold T_{smooth} is set to 1. An example T1 input slice and its result can be seen in Figure 7, where mainly only the brain tissues remain. The pre-processed slices will be used for the estimation of the bias field, and the estimated bias field is applied to the original input image.

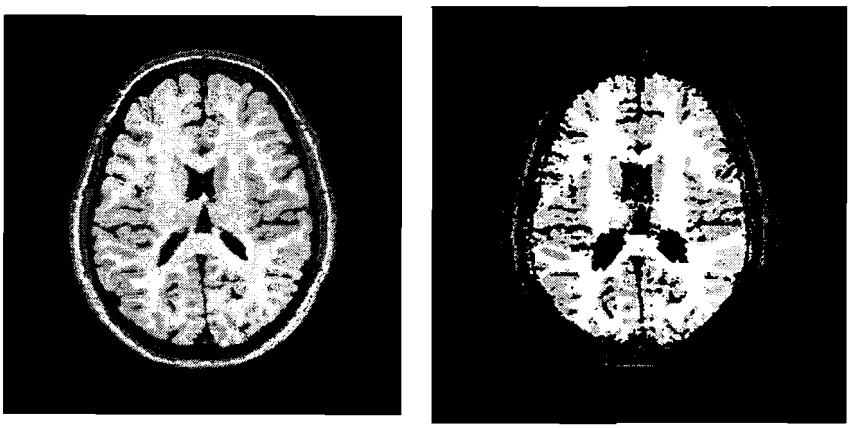

Figure 7: Example of preprocessing on a T1 image.

Intensity-dependent bias estimates

The intensity-dependent bias estimates are calculated from the intensity histogram of the whole preprocessed volume using mean-shift mapping. We define the mean-shift mapping in the intensity histogram as:

$$f_c(f_d) = \frac{\sum_{i=-s/2}^{s/2} H_{distorted}(f_d + i) \cdot (f_d + i)}{\sum_{i=-s/2}^{s/2} H_{distorted}(f_d + i)} \quad (5)$$

where $H_{distorted}(i)$ is the histogram of preprocessed input volume, f_d is the distorted intensity, $f_c(f_d)$ is the estimated corrected intensity for f_d, and s is the size of the mean-shift window which has been set 64 in the implementation. Thus f_c is the mean intensity in the neighborhood of intensity f_d in the intensity histogram. From this intensity mapping, $f_d \rightarrow f_c(f_d)$, an intensity-dependent bias estimate can be calculated:

$$b(f_d) = f_d / f_c(f_d) \quad (6)$$

From the formula we can see that mapping in the background, where values are close to 0, will be inaccurate due to lower signal to noise ratio and quantization. So, these inaccurate estimates will have a large negative effect on the bias-field estimation, if not excluded. Therefore, we exclude the background intensities and also outliers on the higher end of the intensity range from the bias field estimation:

$$b'(f_d) := \begin{cases} 1, & f_d < T_{low} \\ f_d / f_c(f_d), & T_{low} \le f_d < T_{high} \\ 1, & T_{high} \le f_d \end{cases} \quad (7)$$

In case of the brain, and given the correct thresholds T_{low} and T_{high}, the selected intensity range represents the brain tissue. In practice, the thresholds were set to 100 and the maximum of the intensity range. The obtained intensity-dependent bias estimates can be applied to the image to obtain voxel-wise, intensity-dependent bias estimates, i.e. $b'(I_{preprocessed}(\mathbf{r}))$.

Smoothing

Of the voxel-wise, intensity-dependent bias estimates, $b'(I_{preprocessed}(\mathbf{r}))$, some are erroneous since there is overlap between the individual tissue intensity distributions in the histogram and thus the mapping for some voxels will be incorrect. However, it is assumed that the majority of the estimates is good. By spatially smoothing this initial estimate with a Gaussian kernel, to apply the smoothness constraint of the bias field, a good approximation, $b_{smooth}(\mathbf{r})$, of the true bias field can obtained.

Smoothing is performed in two dimensions within a slice, i.e. no inter slice-smoothing of the bias estimates is performed. Smoothing only within the slices, and not in axial direction, allows the algorithm to correct for rapid interslice intensity variations that can occur in some scanning protocols. The used Gaussian smoothing kernel has a standard deviation of 20mm and was cut off at the 10% points to save calculations. The resulting bias field estimate can be removed from the input image by dividing the image by it.

Multiple iterations

The estimated bias field does not correct the intensity inhomogeneity at once. Therefore, the procedure of calculating the bias estimates from the histogram, smoothing, and division is repeated for multiple iterations, until convergence. However to improve the performance, two extra steps are added to the procedure. First, to speed up the convergence, the shift calculated by the mean shift is amplified:

$$f'_c(f_d) = A \cdot (f_c(f_d) - f_d) + f_d \quad (8)$$

where A, the amplification, has a value of 5.0 in the implementation.

Secondly, to prevent drifting of the average image intensity, the average shift produced by the mean-shift mapping is compensated for:

$$f''_c(f_d) = f'_c(f_d) + \frac{\sum_{i=T_{low}}^{T_{high}} H_{distorted}(i)(f'_c(i) - i)}{\sum_{i=T_{low}}^{T_{high}} H_{distorted}(i)} \quad (9)$$

And $f''_c(f_d)$ replaces $f_c(f_d)$ in Equation (7).

Finally, these iterations are repeated until the energy of the estimated bias field in an iteration goes below a certain threshold or a maximum number of iterations is reached. The energy E is calculated using the following equation:

$$E = \frac{\sum_{\bar{x}} | b_{smooth}(\mathbf{r}) - 1 |}{N} \tag{10}$$

where $b_{smooth}(\mathbf{r})$ is the smoothed bias field estimate and N is number of voxels in the volume. The threshold to terminate the iterations based on the energy has been set to 0.001 and the maximum number of iterations to 20.

Sub-sampling

Since the bias field is smooth it can be calculated on a sub-sampled version of the input image, this significantly reduces the processing time of the algorithm (this is also applied by Sled et al.). In this implementation the image is sub-sampled in the x and y direction to a lower working resolution while the number of slices is maintained. After the algorithm is converged and the final bias field is calculated, this final bias field is up-scaled and applied to the original input image. In the implementation the standard working resolution was chosen to be 128×128 pixels in a slice and the number of slices is maintained.

Pseudo-code

For final clarity the algorithm is described in pseudo-code in Table 1.

input: MR image distorted by intensity inhomogeneity
Step 1. standardize the intensity range
Step 2. sub-sample the image to the working resolution
Step 3. mask the non-brain brain areas by using Eq. (3)
repeat
    Step 4. calculate the intensity histogram
    Step 5. calculate voxel-wise intensity-dependent bias estimates
    Step 6. spatially smooth the bias estimates to obtain an estimated bias field
    Step 7. remove the bias field by division
    Step 8. update the total bias field
until (bias field energy < threshold) **or** (maximum number of iterations is reached)
Step 9. upscale the total bias field
Step 10. remove the bias field from the original input image by division
output: corrected MR Image and estimated bias field

Table 1: Pseudo-code of the implemented bias-correction method.

4 Learning Vector Quantization

As *learning vector quantization* (LVQ) is the learning scheme on which the proposed segmentation method is based and since it essential to understanding of the method an introduction to LVQ is given in this chapter. LVQ [Kohonen86] belongs to the so-called *competitive learning* methods of which other well-known examples are k-means clustering [Bezdek95] (see Appendix A) and the *self-organizing map* (SOM) [Kohonen90].

LVQ is a learning scheme that can be used to obtain a classifier or quantizer, where classification/quantization is based on a set of features, an input vector, e.g. the intensity of a voxel. The obtained classifier is similar to one obtained through the well-known k-means clustering, where k classes are represented by the mean of the samples assigned to them, the centroid (reference vector), and each sample is assigned to the class with the closest centroid calculated in the previous iteration, in this way the within-class variance is iteratively minimized (see Appendix A). As a result, one obtains a meaningful representation of the data by a relatively small number of classes, which represent the clusters in the data.

The main difference between k-means and LVQ is that in k-means the centroids are updated after classifying the whole dataset, while in LVQ the reference vectors are updated after each sample. Thus in LVQ the samples are processed one-by-one, this is so-called online learning.

In the next section the basics are explained, and then some relevant modifications of LVQ are given after which LVQ is related to the proposed method.

4.1 Basics

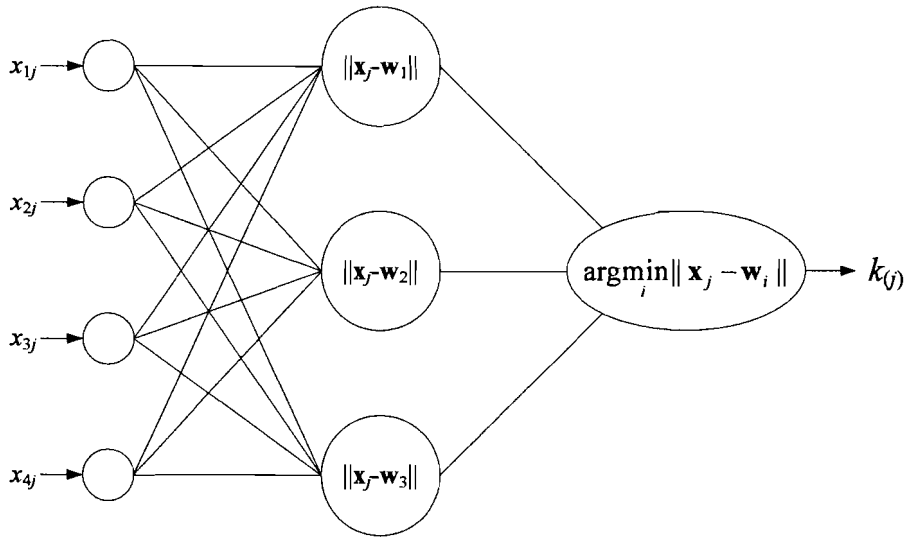

Figure 8: An example LVQ network. $\mathbf{x}_j=\{x_{1j}, x_{2j}, x_{3j}, x_{4j}, \ldots, x_{nj}\}$.

The LVQ algorithm can be thought of as a neural network where each class is represented by a neuron[8] with a reference vector. Figure 8 depicts an exemplary LVQ network where the input layer has four[9] dimensions and the competition layer has three neurons with reference

[8] Classes can also be represented by a group of neurons; however, this is ignored here since it is not used in our algorithm.

[9] Please note that in the present application of the proposed algorithm only one input dimension (the T1 image), a scalar, is used. In our application more input dimensions could be obtained from different scans of the same brain using different scanning parameters, i.e. other modalities.

vectors, $\mathbf{w}_i$, for class i. The reference vectors represent three classes and the number of classes is chosen in advance. For each input vector $\mathbf{x}_j$, the neuron whose reference vector is closest to the input vector is assigned as the 'winner'. In general, the euclidean distance measure is used:

$$k_{(j)} = \arg\min_i \| \mathbf{x}_j - \mathbf{w}_i \| \tag{11}$$

where $\mathbf{x}_j$ is the input vector, $\mathbf{w}_i$ is the reference vector of neuron i, representing class i, and k is the winning neuron/winner class for sample j.

To obtain a network with reference vectors for valuable classification, the network should be trained, i.e. it should 'learn'. This learning is achieved by feeding the network with samples from the dataset and updating the reference vectors after each sample with the following recursive formula [Wu03]:

$$\mathbf{w}_i(t) = \mathbf{w}_i(t-1) + \alpha(t) \cdot h_{i,j,k} \cdot (\mathbf{x}_j - \mathbf{w}_i(t-1)) \tag{12}$$

where $h_{i,j,k}$ is a neighborhood function and $\alpha(t)$ is a learning rate, which is usually chosen as a function decreasing with time.

If we have a training set of samples with their classification available, labeled by a human supervisor, we can perform so-called supervised learning. For supervised LVQ, LVQ1 [Kohonen90], the neighborhood function is defined as:

$$h_{i,j,k} = \begin{cases} 1 & \text{if } i = k \text{ and } \mathbf{x}_j \text{ is classified correctly} \\ -1 & \text{if } i = k \text{ and } \mathbf{x}_j \text{ is classified incorrectly} \\ 0 & \text{if } i \neq k \end{cases} \tag{13}$$

After the training phase is completed, the trained classifier can be used to classify new input samples.

For unsupervised LVQ (ULVQ), where there are no labeled samples available for training, the neighborhood function is defined as:

$$h_{i,j,k} = \begin{cases} 1 & \text{if } i = k \\ 0 & \text{if } i \neq k \end{cases} \tag{14}$$

These neighborhood definitions lead to a crisp update method, in which only the reference vector of the winner neuron, $\mathbf{w}_k$, is updated; this scheme is called 'winner takes all'.

When using ULVQ with a decreasing learning rate, $\alpha(t)$, convergence is implied and the resulting reference vectors are much like those obtained through k-means clustering. Namely, the reference vectors automatically drift towards the dense clusters in the input data and represent these clusters. In our application these clusters are the CSF, GM, and WM intensity distributions.

For final clarity, in Table 2 the pseudo code for the training phase of ULVQ is given.

initialize reference vectors

$t = 0$

repeat

take an input sample j

find the winner k using (11)

update the reference vector using (12) and (14)

j = random sample number, $t = t + 1$

until ($\alpha(t)$ is equal to 0) **or** (the maximum number of iterations is reached)

Table 2: Pseudo-code for training using ULVQ.

4.2 Modifications

In the following, two modifications to the basic ULVQ algorithm are described, which are known in the literature and relevant to our application. These modifications are the introduction of fuzziness and a Gaussian kernel.

4.2.1 Fuzzy LVQ

Fuzziness increases the performance of LVQ by preventing the algorithm from getting stuck to local minima[10], as it does in fuzzy c-means [Bezdek95]. Like in fuzzy c-means one can assign each sample to all classes to a different degree. If we redefine the neighborhood function we can introduce this fuzziness in the update scheme. Consequently, we update not only the winner but also the other reference vectors and thereby abandon the 'winner-takes-all' scheme, which is prone to local minima.

Fuzzy membership is described with the following membership function:

$$\mu_i(\mathbf{x}_j) = \frac{\| \mathbf{x}_j - \mathbf{w}_i \|^{-2/(m-1)}}{\sum_{i'=1}^{c} \| \mathbf{x}_j - \mathbf{w}_{i'} \|^{-2/(m-1)}} \tag{15}$$

which is the same as the membership function in fuzzy c-means, where c is the number of classes, and m is the degree of fuzziness, typically m=2 in fuzzy c-means, and for the limit $m \to 1$ it leads to crisp (hard) classification (see Figure 9). This membership is based on the relative inverse (squared) distance of a sample to all reference vectors.

The membership degree is then used in the neighborhood function in the following way [Bezdek95; Wu03]:

$$h_{i,j,k} = (\mu_i(\mathbf{x}_j))^m \tag{16}$$

see Figure 9 for graphical impression.

[10] It can converge into a set of neuron weights, which represent the data badly.

4.2.2 Kernel LVQ

An LVQ method described in the literature using kernels in the update function is that of Wu et al. [Wu06]. The kernel is typically defined as a Gaussian and makes the algorithm more robust to outliers and noise.

The update function (12) is modified by adding an Gaussian kernel weight to it, which weighs the contribution of the sample to the update of the reference vector based on its distance to the reference vector, assuming samples with a large distance from the reference vector are noise or outliers and should be ignored. Thus, the update function becomes:

$$\mathbf{w}_i(t) = \mathbf{w}_i(t-1) + \alpha(t) \cdot h_{i,j,k} \cdot \exp\{-\frac{\| \mathbf{x}_j - \mathbf{w}_i(t-1) \|^2}{\beta(t)}\} \cdot (\mathbf{x}_j - \mathbf{w}_i(t-1))$$

with

$$\beta(t) = \frac{\sum_{i=1}^{c} \| \mathbf{w}_i(t-1) - \overline{\mathbf{w}}(t-1) \|^2}{c} \tag{17}$$

where $h_{i,j,k}$ is defined as in ULVQ (14). For adaptation to the scale of the dataset, $\beta(t)$ is introduced, which is the variance of the reference vectors, where c is the number of classes, and $\overline{\mathbf{w}}(t-1)$ is the average of all reference vectors. Again, see Figure 9 for a graphical impression.

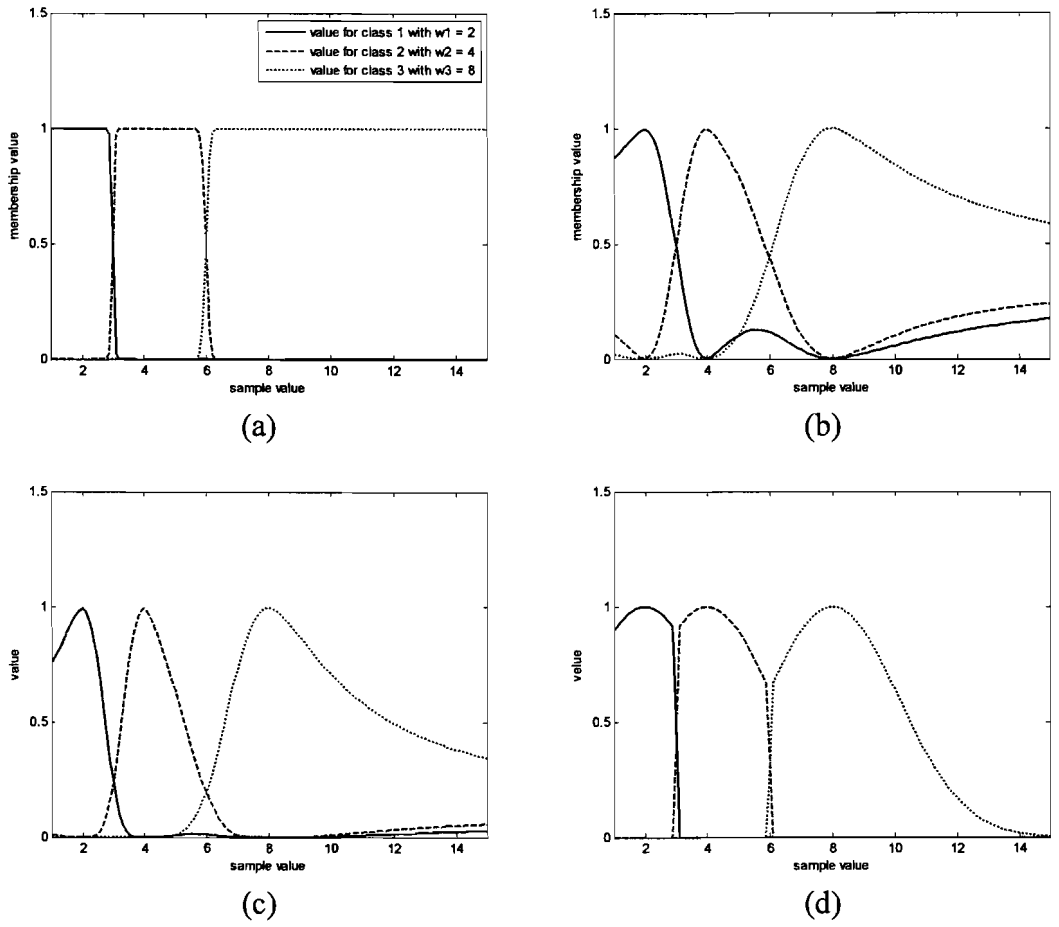

Figure 9: Graphical representation of some of the discussed functions for three classes with reference vectors with values 2, 4, and 8. (a) Fuzzy membership function (15) with m = 1.1 and in (b) with m = 2. (c) is neighborhood function (16) with m = 2, and (d) crisp neighborhood multiplied with kernels as in (17).

4.3 Use in the proposed method

To improve the robustness of our method both the Gaussian kernel and the fuzzy membership are used.

The Gaussian kernel improves robustness to noise and outliers. Using a Gaussian kernel seems essential in our application, especially with the use of online learning, where the kernel helps to stabilize the algorithm. The kernel creates the ability to ignore outliers, which are not modeled, and noise. In our case outliers can be artifacts in the scan, lesions, or non brain tissue (skull, skin, bone, etc.) if they were not removed in advance. Without using the kernel, the outlier intensities would have a large influence on the behavior of the algorithm and would destabilize it.

Fuzziness has proven its ability to avoid getting stuck to local minima, since all the reference vectors are updated, contrary to the winner-takes-all scheme, where some reference vectors can be completely ignored [Karayiannis99]. Besides this valuable feature it appears to be a good model for the partial volume effects, given the results of the fuzzy c-means based segmentation methods mentioned in Section 2.2. Namely, the fuzzy membership of voxels to different tissue classes models the mixture of pure tissue types present in one voxel, and might thus even be essential to incorporate.

5 An adaptive online learning based tissue segmentation method

Using LVQ (Chapter 4) we can train a classifier to classify the voxels of a brain image into the tissue classes based on their intensity. If we use a LVQ network with 4 classes, one for BG, CSF, GM, and WM, and train it with samples from the input image the reference vectors will move towards the clusters in the data which are the BG, CSF, GM and WM intensity distributions. If, after training, all voxels are assigned to the class with the closest reference vector (i.e. the winner class) we obtain a brain image segmented into the tissue classes. This segmentation is basically a multilevel thresholding procedure, which does not include intensity inhomogeneity compensation and spatial context (such as MRFs) for noise compensation. If the intensity inhomogeneity and/or noise are significant, this leads to an unacceptable amount of misclassifications. Methods similar to the process described above, where the training and classification phase are separated, are described in [Karayiannis99; Alirezaie95].

The proposed segmentation method in this thesis is also LVQ-based, however, contrary to the above-mentioned LVQ methods there is no real separation between the training and classification phase whereby intensity inhomogeneity compensation is enabled. Namely, in the method proposed here we continue learning during classification, i.e. $\alpha(t)>0$ during the classification. By sample-by-sample learning during the classification the drifting tissue intensities (due to the inhomogeneity) can be tracked, and if the scanning order of the data allows this, the bias field is compensated for implicitly. In this way no modeling of the bias field is needed and we can classify the data in one pass. This gives a significant reduction in processing time, compared to present well-known EM-based methods for MR brain image segmentation described in Section 2.2. These methods overcome the intensity inhomogeneity in the data by iteratively classifying the data and updating an explicit model of the bias field. As this means classifying the whole dataset multiple times and each time calculating a bias field using low-pass filtering or polynomials it is computationally expensive.

The proposed method is based on the LVQ scheme with both fuzziness and a Gaussian kernel as described in Chapter 4. However, additionally we need a special scanning order to be able to track the intensity variations due to the bias in the 3 dimensional data. Therefore, several so-called locality-preserving scanning orders are proposed in the next section. Furthermore, we also make a modification to the update rule to allow influence of spatial context in classification of a voxel for handling noise and improving spatial consistency of the produced segmentations. Finally, there is an option to abandon the inherently 1D update rule of LVQ and adopt a new 3D learning rule where the scanning order is of less importance. The flowchart of the proposed algorithm is depicted in Figure 10.

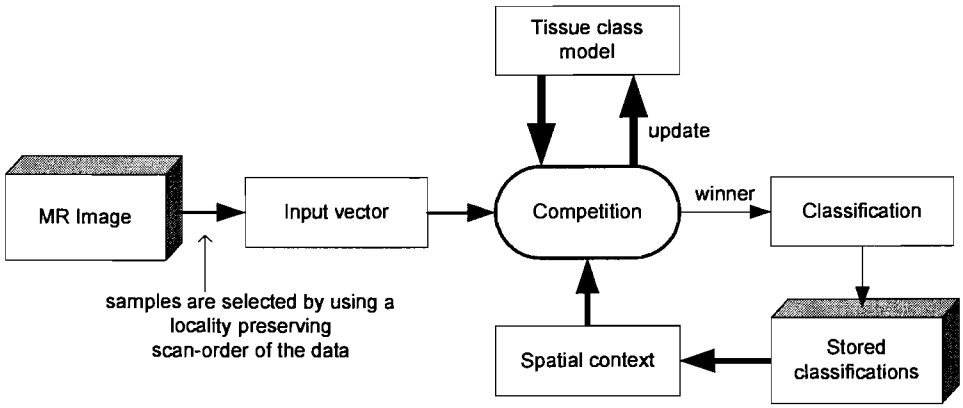

Figure 10: Flowchart of the algorithm. The thick lines indicate multiple vectors while thin lines indicate a single vector.

The different scanning orders and the 3D learning rule lead to three versions of the proposed algorithm, namely two based on 1D learning with different scanning patterns, the raster and hybrid scan, and one based on 3D learning.

5.1 Scanning (1D-learning)

Since the LVQ scheme is inherently 1D, i.e. it learns by feeding samples sequentially, and the input MR volume is 3D, a special 1D scanning order of this 3D data should be defined. In order to be able to track the drift of the intensities due to the bias there should be no jumps in the data observed along the 1D trail. Since as long as we do not jump location in the MR volume, the present local model (reference vectors) is valid for the next (neighboring) sample and only needs a small adjustment based on that next sample. Thus the scanning pattern should have what is called a locality-preserving behavior. In this section two scanning patterns are proposed, the raster scan and a hybrid scan based on the Hilbert space-filling curve.

5.1.1 Raster

The most straightforward way of scanning the volume is to do it slice-by-slice and line-by-line. However to preserve the locality of the data along the 1D trail, we alternate direction for every line and every slice, i.e. we always move to a neighboring voxel and never jump through the data.

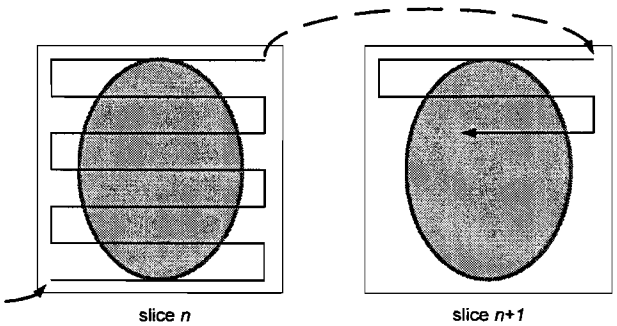

Figure 11: Raster scanning; meandering through lines and slices. The scale of the scanning pattern is exaggerated for display.

Figure 11 shows in which order the samples are classified within a slice, once the slice is fully classified we simply move up one voxel from the present location, assuming we are classifying the volume from bottom to top, and scan this next slice.

Scanning directions were chosen in this specific order after examining the bias field found in the MR datasets used in this thesis. Namely, it was observed that the speed at which the bias field typically changes is lowest in the horizontal direction, higher in the vertical direction, and even higher the z-direction, although this can be different for other datasets.

5.1.2 Hybrid

Theoretically the best locality-preserving pattern is defined by the Hilbert space-filling curve, see Figure 12. The Hilbert curve is defined for squares of size $2^n \times 2^n, n \in \mathbf{N}$, and can be computed by recursion, namely the curve of order n consists of 4 rotated curves of order n-1. Although the curve is also defined for higher dimensions, we limit ourselves to the 2D version, since we also want to use spatial context information calculated in the previous slice, as described in the next section.

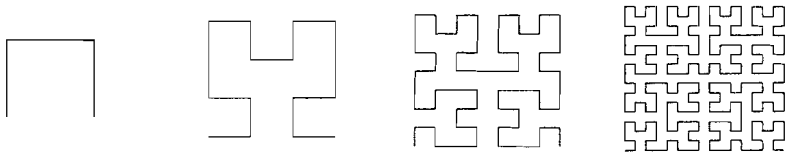

Figure 12: Hilbert curves of order 1, 2, 3, and 4.

Compared to the raster scan, the Hilbert curve is not biased towards a specific direction and therefore progresses slower through the data whereby there is more 'time' to adapt to the drifting intensities. However, in our application, specifically in brain scans, the Hilbert scan is not suited. This is since the hidden assumption, when using the Hilbert scan for our application, is that the data covers the whole square and the inhomogeneity can be tracked everywhere in this square. Since the brain only covers a circular segment of the square, using the Hilbert scan would introduce jumps in the 1D observed data. Namely, due to the random like pattern, the curve can exit the brain in a certain position and enter at a position far from where it exited, and then the locality preserving behavior does not hold. These jumps disable the ability of the online learning scheme to track the inhomogeneity.

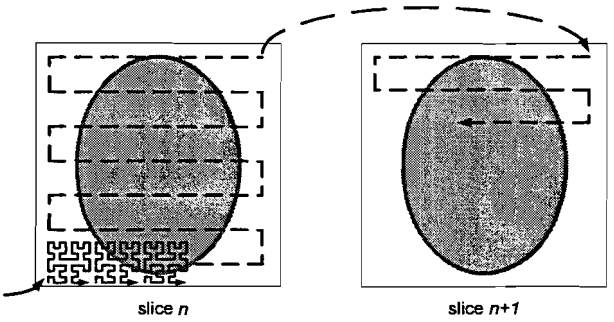

Figure 13: Hybrid scanning; meandering through lines and slices. The scale of the scanning pattern is exaggerated for display.

Therefore a hybrid scanning pattern is proposed, which is a combination of the raster and Hilbert scanning. There, small squares are scanned using the Hilbert curve and these squares are concatenated in a raster pattern, see Figure 13. In this pattern the jumps in the observed data are constraint, i.e. it will always reenter the brain close to where is exited. The hybrid method is a balanced version of the Hilbert and raster scans and thereby it preserves the locality and prevents large jumps in the observations.

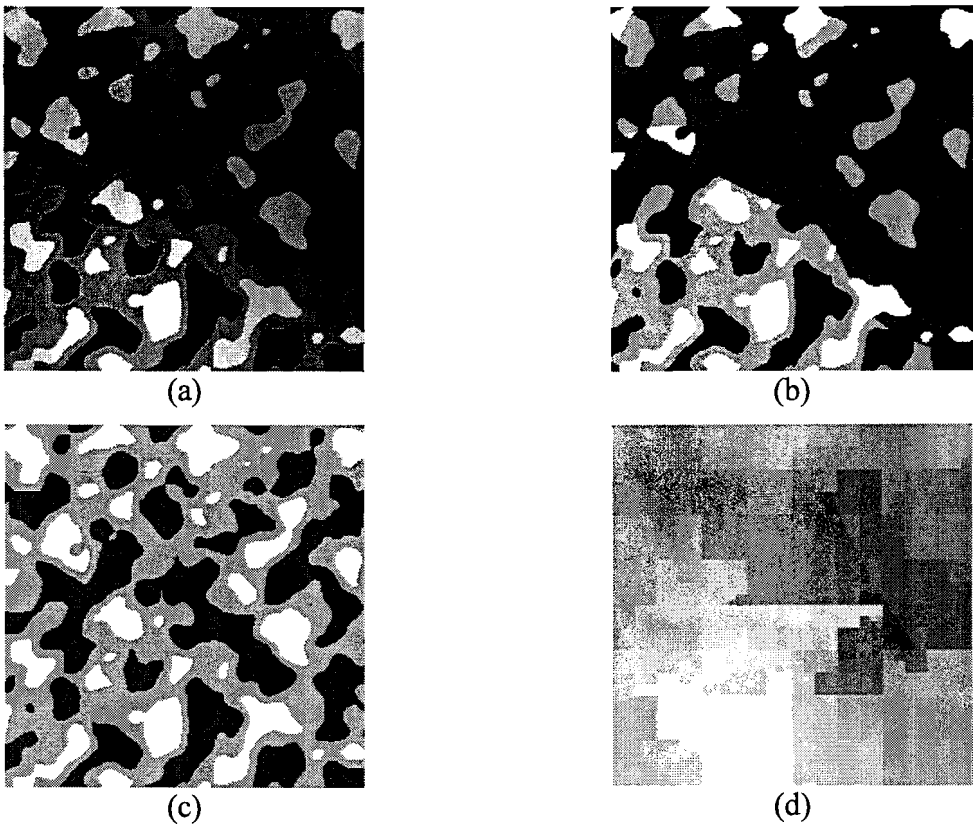

Figure 14: Example of proposed method combined with Hilbert scanning. (a) The input image with four classes and inhomogeneity. (b) the segmentation using k-means based thresholding, and (c) the result using the proposed LVQ based method. (d) shows the reference vector of the brightest class for each location.

Although the Hilbert scan is not used on the brain data, it is interesting to illustrate the performance of the proposed segmentation method using the Hilbert scan on a synthetic image[11], see Figure 14. This synthetic image does match the square of the Hilbert curve and the inhomogeneity can be tracked everywhere. For this image an 8^{th} order Hilbert curve was used, i.e. 256×256. The image contains four classes and has strong inhomogeneity degradation. Figure 14 (d) illustrates the tracking behavior of the proposed method.

5.2 Spatial context

In addition to the special scanning order, the influence of spatial context is added to the algorithm. The spatial context is used to benefit from the fact that segments are homogeneous connected regions spanning many voxels and thus there is correlation between the classes of neighboring voxels. That is to say, if neighboring voxels are of class x then the current voxel is probably of class x as well. Using this spatial context improves the algorithm's noise robustness and overall stability. The noise robustness is improved since by using the spatial context the ambiguity due to the noise induced tissue intensity distribution overlap can be resolved, by biasing the decision for classification to the class most found in the neighborhood. Furthermore, the overall stability is improved since spatial consistency of the tissue regions is imposed.

[11] The synthetic image was manually created with a graphics editor.

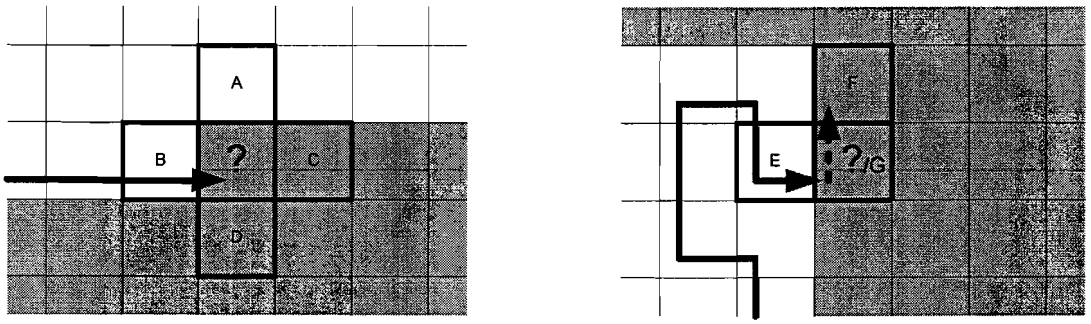

Figure 15: Used spatial context, i.e. neighborhood classifications. (a) for raster scanning, (b) for hybrid scanning.

In Figure 15, the used spatial context is shown. The arrow indicates the current scanning direction, the '?' indicates the current voxel which should be classified, the white voxels have already been classified, and the classification of the grey voxels is unavailable in the current slice. However, the classification of these grey voxels is available in the previous slice[12]. If we assume that the slices are thin (approximately 1mm) then the classifications can be taken from the previous slice instead, since it will be highly correlated.

As the available classifications depend on the scanning pattern which is used, two contexts are defined: in case of raster the neighboring classification of voxels A, B, C, and D can be taken, and in case of the hybrid scan, due to its random like pattern, neighbors E, F, and G are taken along the Hilbert curve.

5.3 Update rules

We base our method on LVQ with fuzziness and a kernel. Additionally, the spatial context influence should be integrated in the LVQ scheme. The influence of the spatial context is expressed by adjusting the 'sensitivity' of the neurons based on the neighborhood classifications (A, B, C, and D or E, F, and G) of the voxel to be classified. We model this sensitivity by adjusting the width of the Gaussian kernel:

$$\mathbf{w}_i(t) = \mathbf{w}_i(t-1) + \alpha \cdot (\mu_i(\mathbf{x}_j))^m \cdot \exp\{-\frac{\| \mathbf{x}_j - \mathbf{w}_i(t-1) \|^2}{2(s_i\sigma)^2}\} \cdot (\mathbf{x}_j - \mathbf{w}_i(t-1)) \qquad (18)$$

where α is now constant since we keep learning during classification, σ is a constant and determines the default width of the kernels, and s_i, the class dependent sensitivity is defined as:

$$s_i = \beta^{\#_i} \qquad (19)$$

with $\#_i$ is the number of voxels in the neighborhood belonging the class i and β is a constant, which controls the influence of the neighborhood. Note that hereby the $\beta(t)$ from Section 4.2.2 is replaced and the width of the kernel is dependent on the spatial context and not on the variance of the reference vectors.

Finally, the winner is defined as the neuron k having the maximum response:

---

[12] Note that this gives a problem at initialization. However, this is solved by the introduction of an initialization phase, see Section 6.2.2. Furthermore, using the context in this way prevents the need for iterative methods which are used for MRF such as the commonly used iterated conditional modes (ICM) algorithm [Besag86; Shattuck01].

$$k_{(j)} = \arg\max_i \exp\{-\frac{\| \mathbf{x}_j - \mathbf{w}_i(t-1) \|^2}{2(s_i\sigma)^2}\} \tag{20}$$

5.4 3D-learning

Instead of using the traditional 1D-learning which requires special scanning orders we could also use 3D-learning, where the recursion is three dimensional. This is achieved by storing the model, i.e. all reference vectors, for all visited locations. Now by using the raster scanning described above we can obtain a 3D recursion by averaging the model with the models obtained in the previous slice and previous line, which were obtained using a different scanning direction and thus different recursion direction.

We achieve this by not only updating the model with the sample we classify but by also updating the model with the earlier obtained models in the neighborhood, so additionally to the usual update step we have an extra model update being:

$$\mathbf{w}_i(t') = \frac{\mathbf{w}_i(t-1) + \delta_{i,l(t_A)} \cdot \mathbf{w}_i(t_A) + \alpha_z(\delta_{i,l(t_C)} \cdot \mathbf{w}_i(t_C) + \delta_{i,l(t_D)} \cdot \mathbf{w}_i(t_D))}{1 + \delta_{i,l(t_A)} + \alpha_z(\delta_{i,l(t_C)} + \delta_{i,l(t_D)})} \tag{21}$$

Where t_A, t_C, and t_D (note $t_B = t\text{-}1$) indicate the times when the locations that are defined as the spatial context in Figure 15(a) were visited, $l(t)$ is the class label assigned to voxel which was visited at time t, and $\delta_{a,b}$ is 1 if a equals b and 0 otherwise, α_z is used to weight the contribution of the models found in the previous slice, which might be less accurate due to thicker slices or abrupt intensity variation between slices.

Using the delta functions assures that the model we use from the neighborhood was updated recently as well, this is necessary since the classes only appear in parts of the image. If we would average the model with models in the neighborhood everywhere and neglect the delta functions we would allow our model to be averaged with possibly worse models.

After the additional step (21) for 3D learning, the default rules (18), (19) and (20) are used, however with t-1 replaced by t'.

Although 3D-learning improves the learning capability, there is a drawback, namely since the models have to be stored in all locations the amount of memory needed for the algorithm increases.

5.5 Discussion

In this chapter a new LVQ-based segmentation has been proposed. The main properties are the use of online learning in combination with a special scanning order to adapt to the drifting intensities and the integration of the spatial context for noise robustness and spatial consistency.

Using the learning rate, α, the adaptivity of the methods can be controlled. Setting the learning rate too low can cause the algorithm to be unable to keep up with the changing intensities. Setting the learning rate too high leads to instabilities since the model will then only be based on a short history of measurements.

α_z is used in 3D-learning and controls to adaptivity in the axial direction. Choosing a high value allows only smooth variations in the axial direction to be tracked. Choosing a low value allows for adaptation to rapid intensity changes in the axial direction, however, at the sacrifice of stability.

σ controls the default width of the kernels. Choosing a high value allows more influence of outliers. On the other hand, a low value can make the algorithm lose its tracking ability since then even changes due to the drifting intensities will be considered outliers.

With β the influence of the neighborhood classifications on the current classification can be controlled. Setting β too low allows noise to pop through in the segmentations. Furthermore, if β is too low, the consistency of the segments is not imposed. This especially decreases the performance of the 1D learning based methods where consistency needs to be implied along the axial direction since it is not implied in the continuity of the model in this direction as in 3D learning. On the other hand, setting β too high leads to loss of detail, since fine detail can be mistaken for noise.

With the proposed methods the feasibility of using online learning for inhomogeneity adaptation in MRI segmentation can be assessed. Therefore, in the next chapter, the methods will be validated and from there conclusions can be drawn for further improvements and/or modifications to the presented methods.

6 Validation

The results of the proposed algorithm should aid in a medical decision-making process and therefore these results should be accurate. Consequently, thorough validation of a proposed medical image-processing algorithm is needed before it would ever be fully trusted in medical decision-making. In most cases the validation of medical image processing algorithms is not straightforward, given the nature of the medical images with large natural variations, variations and distortions due to the imaging equipment, and absence of a ground truth. Jannin et al. [Jannin06] state that validation is rarely the main objective of traditional papers in medical image processing, furthermore, there has been very little work in the standardization of the validation processes and this is primarily caused by the diversity of problems and approaches in medical image processing. Fortunately, for brain tissue segmentation there has been fine work done at the McConnell Brain Imaging Centre with the creation of the BrainWeb database [BrainWeb], which contains simulated MR brain scans for which the ground-truth tissue maps are available. However, it is found in this chapter that good performance for segmentation on the simulated brain scans does not validate the performance on real brain scans.

Firstly, the correction performance of the implemented bias correction will be qualitatively validated in the next section. Secondly, in Section 6.2, the proposed tissue segmentation method is validated extensively by comparing it to several known methods on simulated data from BrainWeb and 3 Tesla real data from the Leiden University Medical Center.

6.1 Bias field correction

As the implemented bias field correction method (Chapter 3) is based on a proven principle and it is not the main focus of this thesis only qualitative validation is given. The method has been tested on brain scans and visual inspection shows good results, i.e. the shading effect disappears. The processing of these volumes takes in the order of 1-2 minutes on a 3GHz Pentium 4 with 512 MB RAM using a C implementation of the algorithm.

In Figure 16, the result on a simulated brain image containing intensity inhomogeneity is shown. This volume also contains, next to the smooth bias field, rapid interslice intensity variations, clearly visible in the sagittal view. Since the implemented method smoothens only within the slice it allows for adaptation to these rapid intensity variations. The final corrected images show homogeneous tissue intensities throughout the whole volume, indicating that the bias has been completely removed.

In Figure 17, results on real data are shown. The T1 image has a resolution of 256×256×120 voxels and the T2 images a resolution of 1024×1024×40 voxels and both are from the same patient. Especially for the T2 images the downscaling to the 128×128 working resolution in implemented method gives a dramatic speedup. These images were produced by a 3 Tesla scanner, and contain more severe intensity variations within the slice compared to the simulated data. The corrected images again show a homogeneous image where the bias has been removed. Especially the correction on T2 image allows for more details to be visible in central area in this specific contrast setting. The method was tested on 22 T1 and T2 images and shows consistent results.

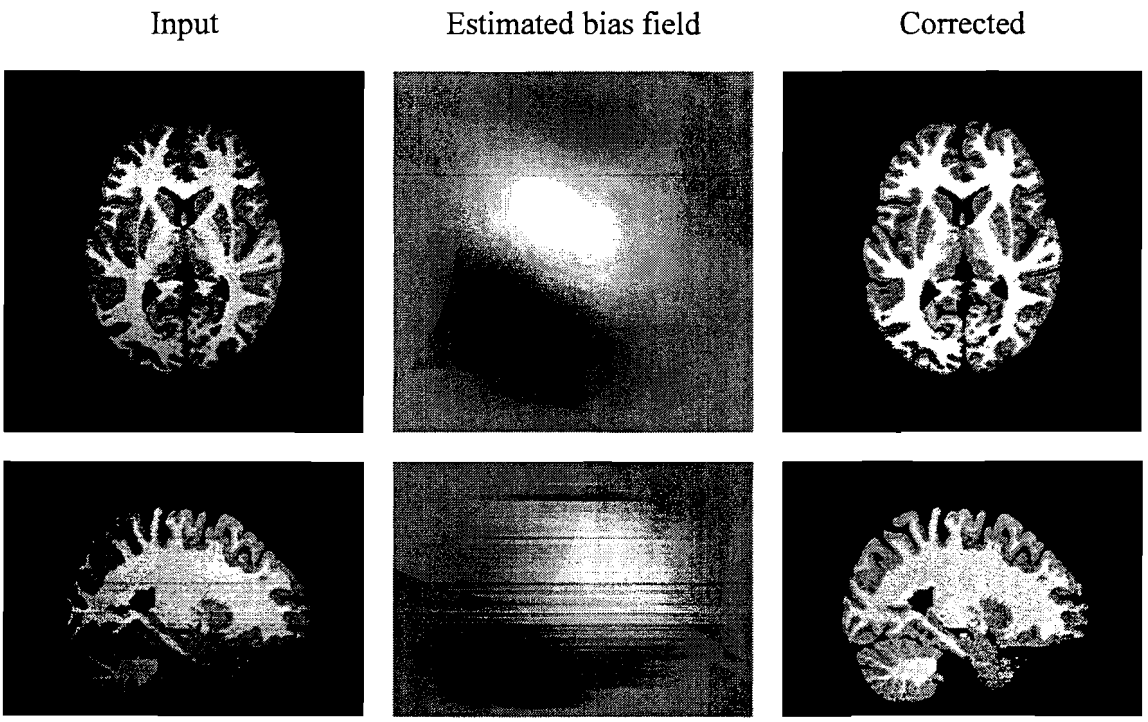

Figure 16: Results of implemented bias correction method on simulated data. On the left the input images, in the middle the estimated bias field, on the right the corrected images. Top row axial view, bottom row sagittal view showing the rapid interslice intensity variations.

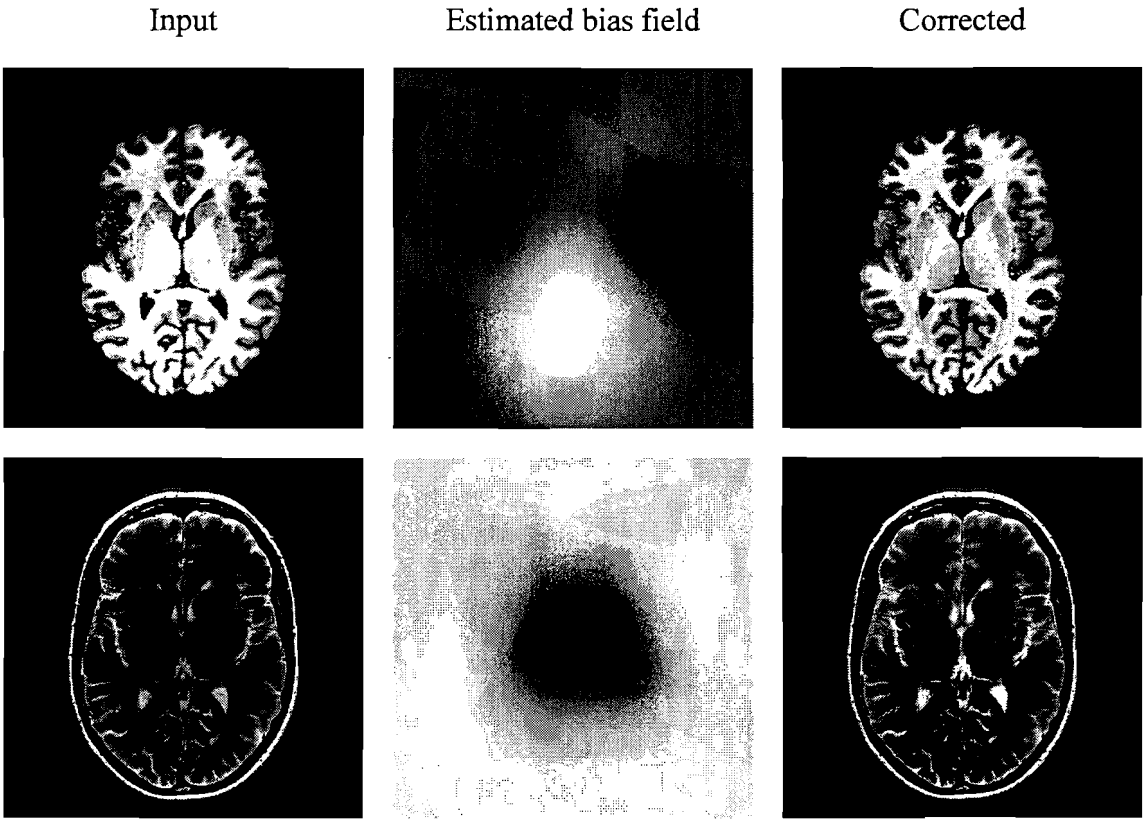

Figure 17: Results of implemented bias correction method on real data. On the left the input images, in the middle the estimated bias field, on the right the corrected images. Top row T1-weighted image, bottom row T2-weighted image.

6.2 Tissue segmentation

In this section we compare the methods proposed in Chapter 5 to several widely known and available methods. Namely, the proposed methods are compared to k-means based thresholding (Appendix A), FAST (FMRIB's[13] Automated Segmentation Tool) [Zhang01], and EMS (Expectation Maximization Segmentation) [VanLeemput99]. FAST is an EM based method with Gaussian mixture modeling, bias field compensation and MRF modeling. EMS is similar to FAST in the mentioned features however it has additional information in the form of atlas which contains tissue priors for each location, see Chapter 2. FAST and EMS are freely available through the Internet[14] and were used with their standard settings. K-means based segmentation has been implemented by the author in MATLAB.

We compare these methods on simulated data from BrainWeb [BrainWeb] and real 3 Tesla data from the Leiden University Medical Center.

6.2.1 Methodology

Validation of the segmented images presents a difficulty since there is no ground-truth at hand with which the produced segmentations can be compared. As ground truth a manually segmented image by a human-expert could be used, however there is inter- and intra-observer variability, and obtaining it is very time-consuming task. Therefore the BrainWeb database has been created, which contains simulated MRI brain images for which the ground truth is known. The simulated images are created from segmented images of real subjects, which were obtained through an extensive process [Aubert-Broche06NI; Aubert-Broche06MI]. These simulated images model the relaxation times of the tissues, the inhomogeneity, partial volume effects, and noise, and the amount of these can be controlled independently. However, there is no ground truth available for the real MR dataset. Therefore, we take the results of EMS as baseline in these experiments, since these seem closest to the truth after visual inspection by the author.

To compare the quality of the segmentations as produced by the automatic methods, we measure their agreement with the ground truth using the so-called Dice coefficient [Dice45]:

$$D_i = \frac{2 \cdot \Omega^i_{g \cap a}}{\Omega^i_g + \Omega^i_a} \tag{22}$$

where D_i is the Dice coefficient for class i, with Ω^i_g and Ω^i_a the volume of the segmentation masks for class i of the ground truth and the automated method respectively, and $\Omega^i_{g \cap a}$ the volume of their intersection. The measure gives a value between 0, meaning no overlap of the segments, and 1, meaning perfect overlap. This measure is a commonly used measure in the brain segmentation field. Besides this measure, segmentation results are depicted for visual inspection.

Prior to segmentation the skull and other non-brain tissue were stripped from all the images using BET2 (Brain Extraction Tool) of the FSL toolkit[14]. After skull stripping only brain tissue is left and all remaining voxels are CSF, GM, WM or BG (background).

---

[13] Oxford Centre for Functional Magnetic Resonance Imaging of the Brain.

[14] FAST is part of the FSL Toolkit of Oxford University, http://www.fmrib.ox.ac.uk/fsl/. EMS is available at Medical Image Computing, Leuven, Belgium, http://www.medicalimagecomputing.com/downloads/ems.php.

6.2.2 Initialization and settings for the proposed methods

Before using the proposed methods, the intensity range of the image is standardized by scaling it so that the 90^{th} percentile of the intensity histogram is at 400. On the resulting brain images we classify using four classes, WM, GM, CSF and BG. The reference vectors of the four classes are initialized by equidistantly distributing them over the range 0-400. With the initial reference vectors we start with refining them on the middle four slices using the proposed method, which is the training phase. At the end of this training we assume that the reference vectors are correct for the location where the training stopped. From there we classify the upper and lower half of the brain by scanning upwards and downwards.

For the hybrid scanning method we use Hilbert squares of 16×16 pixels. For 1D-learning the learning rate α was set to a constant 0.005, for 3D-learning α was set to 0.05 and α_z to 0.2. The width of the kernels σ was set to 20, and the spatial context parameter β was to 1.3 for all three methods. Finally, the fuzziness parameter m was set to the commonly used value of 2. These parameters were determined experimentally. A deeper study of the algorithm's behavior depending on these parameters could be for future research, for now we aim at validation of the produced segmentations and comparing the performance depending on the scanning pattern and 1D versus 3D learning.

6.2.3 Results

Processing time

It was suggested that the use of online learning would give a significant reduction in computational complexity. For the experiments a 3GHz Pentium 4 with 512 MB RAM running Windows XP was used, the processing times on this machine clearly support this claim. However, it should be noted that the implementation of the publicly available methods as well as the implementation of the proposed methods could almost certainly be optimized further.

The processing time of FAST, which is a C implementation, on the used data was approximately 10 minutes per volume for the standard setting of 8 iterations. Since EMS requires additional registration of the brain to an atlas containing the priors, it requires more time. Furthermore, EMS is available only as a MATLAB-based implementation, although some critical functions are implemented as dynamically linked libraries (dll's). The registration step takes approximately 8 minutes and segmentation 16 minutes, thus together 24 minutes per volume.

The proposed methods were implemented in C and require in the order of 30 seconds per volume for the used datasets. This is a dramatic speedup compared to the EM based methods. There is no significant difference between the proposed methods, although 3D learning is more complex and requires slightly more time in the order of a few seconds. 3D learning is more complex since it has an extra update rule and requires buffers for the storage of the class model for each location.

Finally, there is no reference for the k-means method since it has been implemented in MATLAB. However, since it ignores spatial information and is thus purely histogram-driven, an implementation in C would probably be even faster than the proposed methods.

Simulated data

The simulated BrainWeb images are provided with their ground truth and the noise levels and inhomogeneity of the simulated brain images can be controlled. The simulated brain images

have a resolution of 256×256×181 voxels, where the voxels have a dimension $1\times1\times1mm^3$. For these tests the T1-weighted images were used, where the voxel intensities simulate the tissues T1 relaxation times with a precision of 16 bits.

Table 3 and Figure 18 show the results of comparing the methods on the standard BrainWeb volume where various amounts of inhomogeneity were introduced. In this specific test the inhomogeneity was set to 0%, 20% and 40% while the noise level of 3% was constant. Where for a 20% level, the introduced multiplicative field has a range of values of 0.90-1.10 over the brain area, and the percentage of noise indicates the ratio between the standard deviation of the noise and the intensity of the white matter tissue. The noise level was set 3% since this was found comparable or higher than the noise levels found in the real data.

Inhomogeneity [%]	class	k-means	EMS	FAST	raster	Hybrid	3D-learning
0%	CSF	0.9526	0.9475	0.8942	0.9351	0.9313	0.9388
	GM	0.9474	0.9379	0.9310	0.9374	0.9346	0.9445
	WM	0.9524	0.9345	0.9607	0.9463	0.9442	0.9525
20%	CSF	0.9422	0.9476	0.8931	0.9350	0.9309	0.9375
	GM	0.9263	0.9378	0.9313	0.9368	0.9338	0.9433
	WM	0.9290	0.9347	0.9614	0.9455	0.9433	0.9508
40%	CSF	0.9258	0.9478	0.8933	0.9335	0.9296	0.9344
	GM	0.8921	0.9373	0.9291	0.9343	0.9313	0.9387
	WM	0.8913	0.934	0.9583	0.9427	0.9405	0.9451

Table 3: Results on standard BrainWeb volume, with various amounts of inhomogeneity and 3% noise. Segmentation agreement using Dice coefficient.

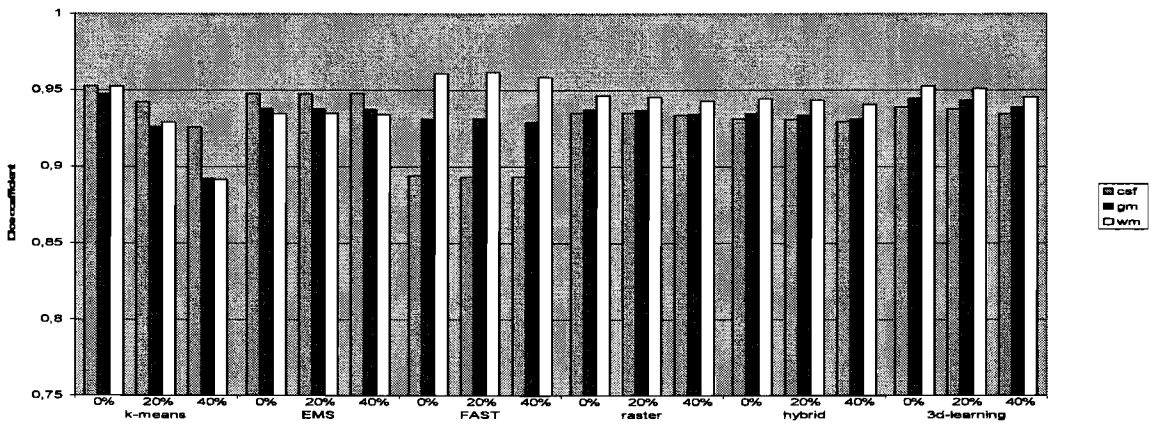

Figure 18: Results on standard BrainWeb volume, with various amounts of inhomogeneity and 3% noise. Segmentation agreement using Dice coefficient.

From the Table 3 we can clearly see that the performance of k-means decreases as the inhomogeneity increases, this is as expected since this method has no bias field compensation. It should also be noted that it has the best performance in case there is no inhomogeneity at all. The remaining methods all have bias field compensation and give an almost constant performance invariant of the amount of bias. However, there are differences between the methods, and especially FAST stands out. FAST has a lower performance in CSF classification and a higher performance in WM classification. Furthermore, it is interesting to

note that the performance of the three proposed methods is similar on this data, with 3D learning slightly better.

Subject	class	k-means	EMS	FAST	Raster	hybrid	3D-learning
50	CSF	0,8187	0.8135	0,7936	0,8144	0,8125	0,8154
	GM	0,8818	0.8817	0,8809	0,8895	0,8878	0,8895
	WM	0,8989	0.8974	0,9069	0,9064	0,9048	0,9048
51	CSF	0,7971	0.8006	0,7707	0,7924	0,7906	0,7939
	GM	0,8928	0.8961	0,8959	0,9032	0,9015	0,9031
	WM	0,8986	0.9010	0,9128	0,9118	0,9101	0,9100
4 with field A	CSF	0,8238	0.8321	0,8062	0,8292	0,8268	0,8296
	GM	0,8142	0.8852	0,8738	0,8903	0,8883	0,8908
	WM	0,8048	0.8964	0,8981	0,9045	0,9024	0,9024
5 with field A	CSF	0,8214	0.8357	0,8029	0,8326	0,8296	0,8337
	GM	0,8089	0.8896	0,8710	0,8946	0,8918	0,8947
	WM	0,7956	0.8960	0,8810	0,9033	0,9004	0,8993
38 with field B	CSF	0,7878	0.8123	0,7987	0,8210	0,8186	0,8241
	GM	0,7709	0.8864	0,8619	0,8887	0,8895	0,8957
	WM	0,7444	0.8938	0,8613	0,8844	0,8866	0,8868
41 with field B	CSF	0,8010	0.8290	0,8114	0,8331	0,8312	0,8351
	GM	0,7851	0.8939	0,8769	0,8949	0,8943	0,9001
	WM	0,7661	0.9036	0,8856	0,8972	0,8973	0,9017
45 with field C	CSF	0,8224	0.8433	0,8163	0,8422	0,8393	0,8423
	GM	0,8264	0.8946	0,8934	0,8967	0,8951	0,9013
	WM	0,8199	0.9041	0,9133	0,9042	0,9037	0,9105
46 with field C	CSF	0,8265	0.8358	0,8192	0,8345	0,8323	0,8348
	GM	0,8164	0.8901	0,8872	0,8864	0,8848	0,8952
	WM	0,8063	0.9021	0,9114	0,8941	0,8935	0,9076
Mean	CSF	0,8123	0.8253	0,8024	0,8249	0,8226	0,8261
	GM	0,8246	0.8897	0,8801	0,8930	0,8916	0,8963
	WM	0,8168	0.8993	0,8963	0,9007	0,8999	0,9029
Standard deviation	CSF	0,0147	0.0147	0,0154	0,0157	0,0154	0,0153
	GM	0,0427	0.0050	0,0116	0,0054	0,0053	0,0049
	WM	0,0560	0.0039	0,0188	0,0085	0,0073	0,0076

Table 4: Results on simulated data with various bias fields. Segmentation agreement using Dice coefficient.

The BrainWeb database was recently expanded with more simulated patients and also various bias fields are available separately. Using this data we can test the performance of the methods under patient and bias field variability. The noise and inhomogeneity specifications of the new patients are not given at present. However, inspection shows that the noise level is 3-5% and the bias field is absent while there are some rapid interslice intensity variations. Using these patients and the separately available bias fields A, B, and C a set of new test subjects has been created. This set includes two patients without additional inhomogeneity, two with field A, two with field B, and two with field C. The introduced inhomogeneity was scaled to have a value of approximately 40%.

The results of these tests are given in Table 4 and Figure 19. Again, as we saw in the previous test, the k-means method is most affected by the inhomogeneity and FAST has the lowest performance in CSF classification. From visual inspection of the data this appears to be mainly due to overestimation in size of the CSF segments. Furthermore, if we ignore k-means, FAST has the highest standard variation in performance in this test, indicating that it is less stable than the other methods. On this data set there is again no real difference among the proposed methods.

For all methods the performance of CSF classification is lower compared to Table 3, this is since in this test BET2 was used for skull stripping while in the previous test the ground truth was used as mask. Since the area between the skull and the brain is filled with CSF some of the CSF is usually removed together with the skull, and consequently this affects the volume of the CSF. This is something that should always be taken into account when accurate measurement of CSF volume is required. Furthermore, the overall score is lower than in Table 3, this can be attributed for a part to the fact that there is a mismatch between the resolution of the simulated brain scans and the provided ground truth. Therefore, the ground truth was downscaled and this seems to have introduced some interpolation errors to which the segmentation comparisons are very sensitive. Nevertheless, the relative comparisons are still valid.

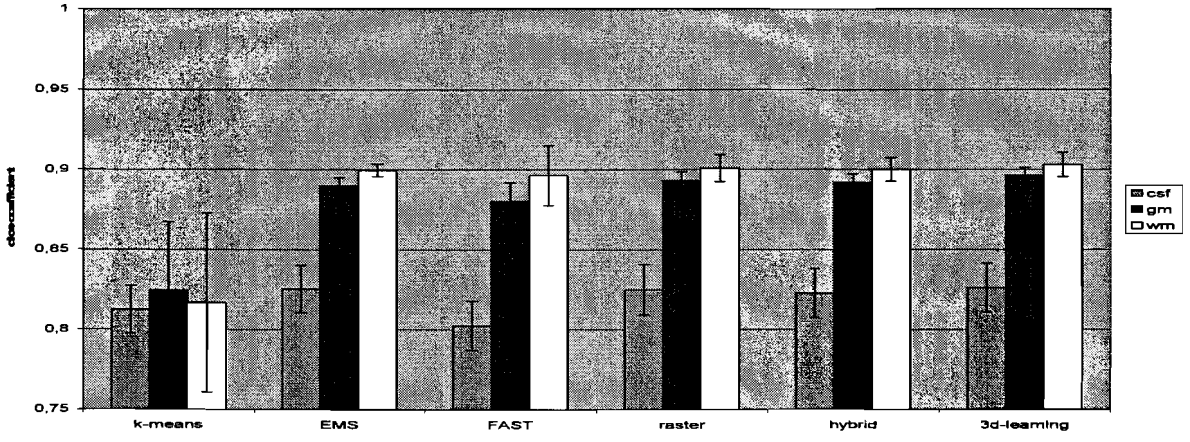

Figure 19: Mean and standard deviation of the results shown in Table 4.

Since the Dice coefficient gives only a simplified impression of the performance, the produced segmentations for subject 46 with field C are shown in Figure 20 and Figure 21. Here we see axial and sagittal slices of a simulated MR image, the ground truth and the segmentations of all methods. Clearly, k-means is affected enormously by the inhomogeneity and the produced segmentation contains many misclassifications. The segmentations of the remaining methods show far better results where the main artifact is the loss of some detail.

In Figure 20 we can see a difference between the three proposed methods. In the lower left of the brain, the area indicated with the arrow, some of the details are lost in the segmentations of the proposed methods due to the darkening of the image is this area. This effect is strongest in the 1D-learning with raster scanning, as it progresses quickly in the horizontal direction it can not adapt to rapid change of intensity in this direction. Improvement can be seen in the hybrid method and even more in the 3D-learning method, although only the EM methods preserve the details in this slice fully. This indicates that the online adaptivity is not high enough in the proposed methods, raising the learning rate might improve the results although higher learning rates lead to instabilities. In Figure 21 we can see some loss of fine detail in the cerebellum (indicated with the arrow). This is mainly due to the spatial context/MRFs which are known to remove small details, since small details are easily mistaken for noise.

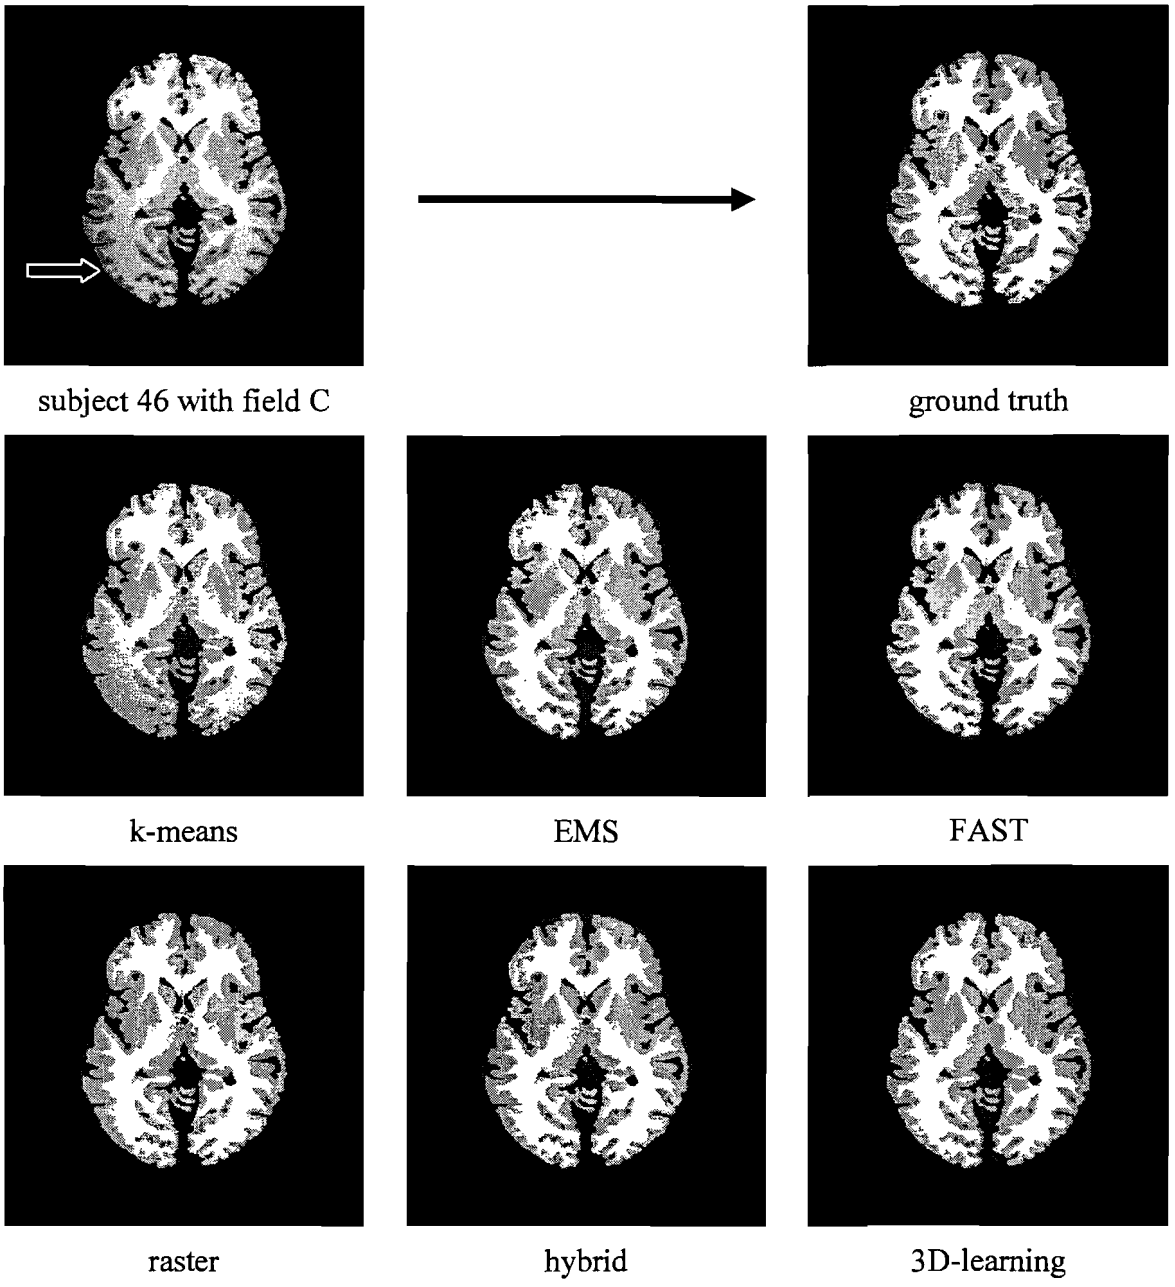

Figure 20: An axial slice of the segmentation of the different methods on subject 46 with field C.

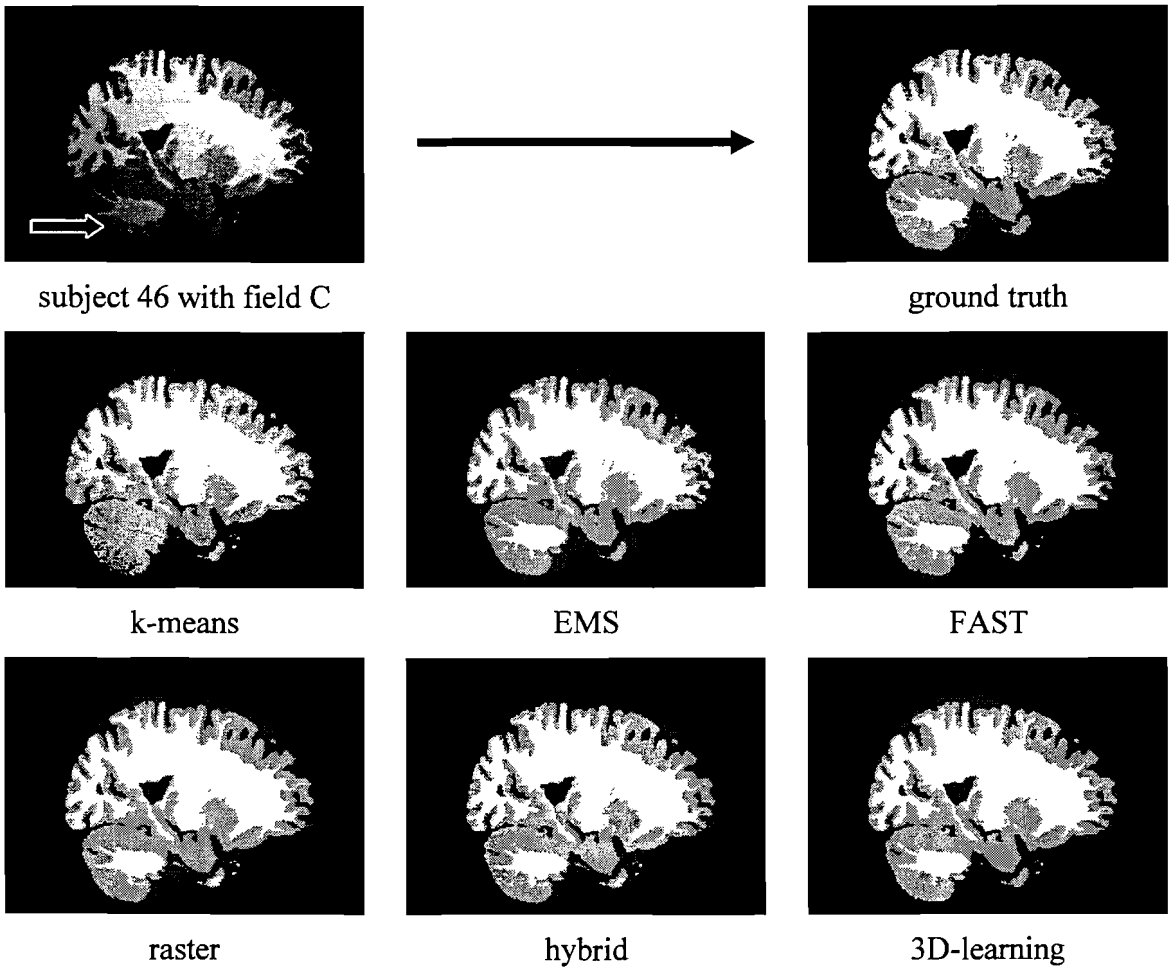

Figure 21: An sagittal slice of the segmentation of the different methods on subject 46 with field C.

Real data

In these tests we compare the segmentation results on real data. The real data has been provided by the Leiden University Medical Center. The used T1-weighted images have a resolution of 256×256×120 voxels, where the voxels have a dimension of $0.88 \times 0.88 \times 1.20 mm^3$ and the T1 relaxation time is represented with 16 bits. These scans were produced by a scanner with a field-strength of 3 Tesla. As the intensity inhomogeneity increases with the field strength it is severe in these images. Furthermore, since there is no ground truth available for these images, the results of EMS are used as ground truth. That EMS can be used as ground truth has been decided by the author after visual inspection of the results.

Table 5 and Figure 22 show us the results of a comparison using the Dice coefficient and EMS as ground truth on 5 real subjects. K-means again gives the lowest mean performance and highest standard deviation since it has no bias compensation. The best overall performing method is FAST, although it shows low scores on CSF classification like it did on the simulated data. The higher performance of FAST suggests that the iterative EM approach with Gaussian mixture and explicit bias field modeling do pay off for real data, which is found to be more complex than simulated data, i.e. there is more natural tissue variation and severe inhomogeneity. One could suggest that by increasing the learning rate in the 3D-

learning method the results would improve further; however, experiments show that this is not the case.

subject	class	k-means	FAST	raster	hybrid	3D-learning	bias correction, 3D-learning
R1	CSF	0,8388	0,8249	0,8727	0,8671	0,8539	0,8920
	GM	0,8116	0,9037	0,8507	0,8463	0,8749	0,8747
	WM	0,8242	0,9195	0,8499	0,8480	0,8928	0,8731
R2	CSF	0,8226	0,8697	0,8931	0,8866	0,8753	0,9035
	GM	0,8347	0,9291	0,8621	0,8595	0,8954	0,8749
	WM	0,8340	0,9302	0,8401	0,8399	0,8883	0,8519
R3	CSF	0,6865	0,7792	0,8073	0,8019	0,7655	0,8222
	GM	0,7832	0,9101	0,8535	0,8442	0,8753	0,8590
	WM	0,8147	0,9326	0,8568	0,8492	0,8980	0,8620
R4	CSF	0,8407	0,8340	0,8548	0,8451	0,8049	0,8563
	WM	0,7781	0,8829	0,8458	0,8385	0,8716	0,8514
	GM	0,8166	0,9285	0,8635	0,8605	0,9048	0,8694
R5	CSF	0,8182	0,8315	0,8246	0,8188	0,7774	0,8174
	GM	0,8138	0,9024	0,8575	0,8547	0,8745	0,8638
	WM	0,8399	0,9323	0,8618	0,8616	0,8982	0,8737
Mean	CSF	0,8014	0,8279	0,8505	0,8439	0,8154	0,8583
	WM	0,8043	0,9056	0,8539	0,8486	0,8783	0,8648
	GM	0,8259	0,9286	0,8544	0,8518	0,8964	0,8660
Standard deviation	CSF	0,0650	0,0323	0,0349	0,0345	0,0477	0,0392
	GM	0,0234	0,0166	0,0062	0,0084	0,0096	0,0102
	WM	0,0109	0,0054	0,0096	0,0091	0,0062	0,0092

Table 5: Results on real data with EMS as ground truth. Numbers indicate the segmentation agreement with the Dice coefficient.

Of the proposed methods 3D-learning is clearly the winner, considering the GM-WM classification, and stands out more than it did on the simulated data. This is due to the higher adaptivity and higher stability, which are accomplished by the higher learning rate and averaging of the model in all directions. However, 3D-learning scores lower on the CSF classification. This can be understood by looking at the sagittal view of an example segmentation on subject R4 depicted in Figure 24. There we can see that contrary to the 1D-learning methods misclassifications are made in the cerebellum (indicated with the arrow). There, due to the inhomogeneity, the image darkens dramatically. And since in 3D-learning the change of the intensity model in the axial direction is constraint, i.e. it takes reference vectors from the previous slice, this quick change cannot be tracked. By adjusting α_z this might be improved, however stability will be sacrificed. Since the 1D-learning methods do not have this constraint they can adapt to this rapid change. This explanation is supported by the results of first applying the implemented bias correction method before using the 3D-

learning method, then we see that the CSF classification performance improves. However, the performance of the GM and WM classification decreases. This also caused by misclassifications in the cerebellum, because of the higher brightness of this area after bias correction the 3D-learning method seems to favor WM over GM, like the 1D-learning methods also do.

Interestingly, in Figure 23, which shows an axial view of the segmentations, we can see disagreement in the central area (indicated with the box) between EMS and all the other methods. Here areas which are classified as GM by EMS are classified as WM by the other methods. These areas are the so-called deep grey matter, comprising the caudate nucleus, putamen, and thalamus (see Figure 1), which are considered to belong to the GM class. However, these appear to be brighter than the cortical GM and since they are closer to the WM class in intensity they are considered to belong to the WM class by the other methods. EMS can resolve this issue since it uses an atlas which gives priors which indicate a high probability of GM in those areas. It is clear that it is the atlas in EMS that contributes to the better classification of GM in this area since FAST which is essentially the same as EMS apart from the atlas makes the same misclassifications as the proposed methods. Furthermore, after bias correction the 3D-learning method still makes the same misclassifications, which shows it is not related to shortcoming of bias compensation.

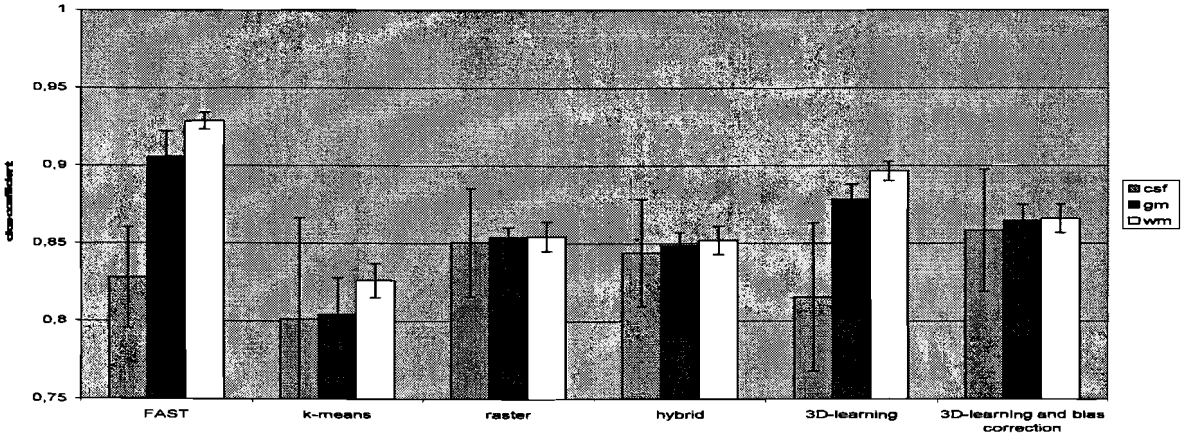

Figure 22: Mean and standard deviation of the results shown in Table 5.

The misclassification of deep grey matter in the real data compared to the simulated data, where this did not happen, indicate that the simulated data is different. Namely, the simulated data does not model the natural variation of the tissue intensities, this also stated in Aubert-Broche et al. [Aubert-BrocheNI06]. This can be seen when visually comparing the simulated images with the real images, in the simulated images the tissue classes themselves are homogeneous, discrete apart from the noise, while in the real images there is more variation in the tissue intensities. Thus, validation on simulated images does not validate for real data.

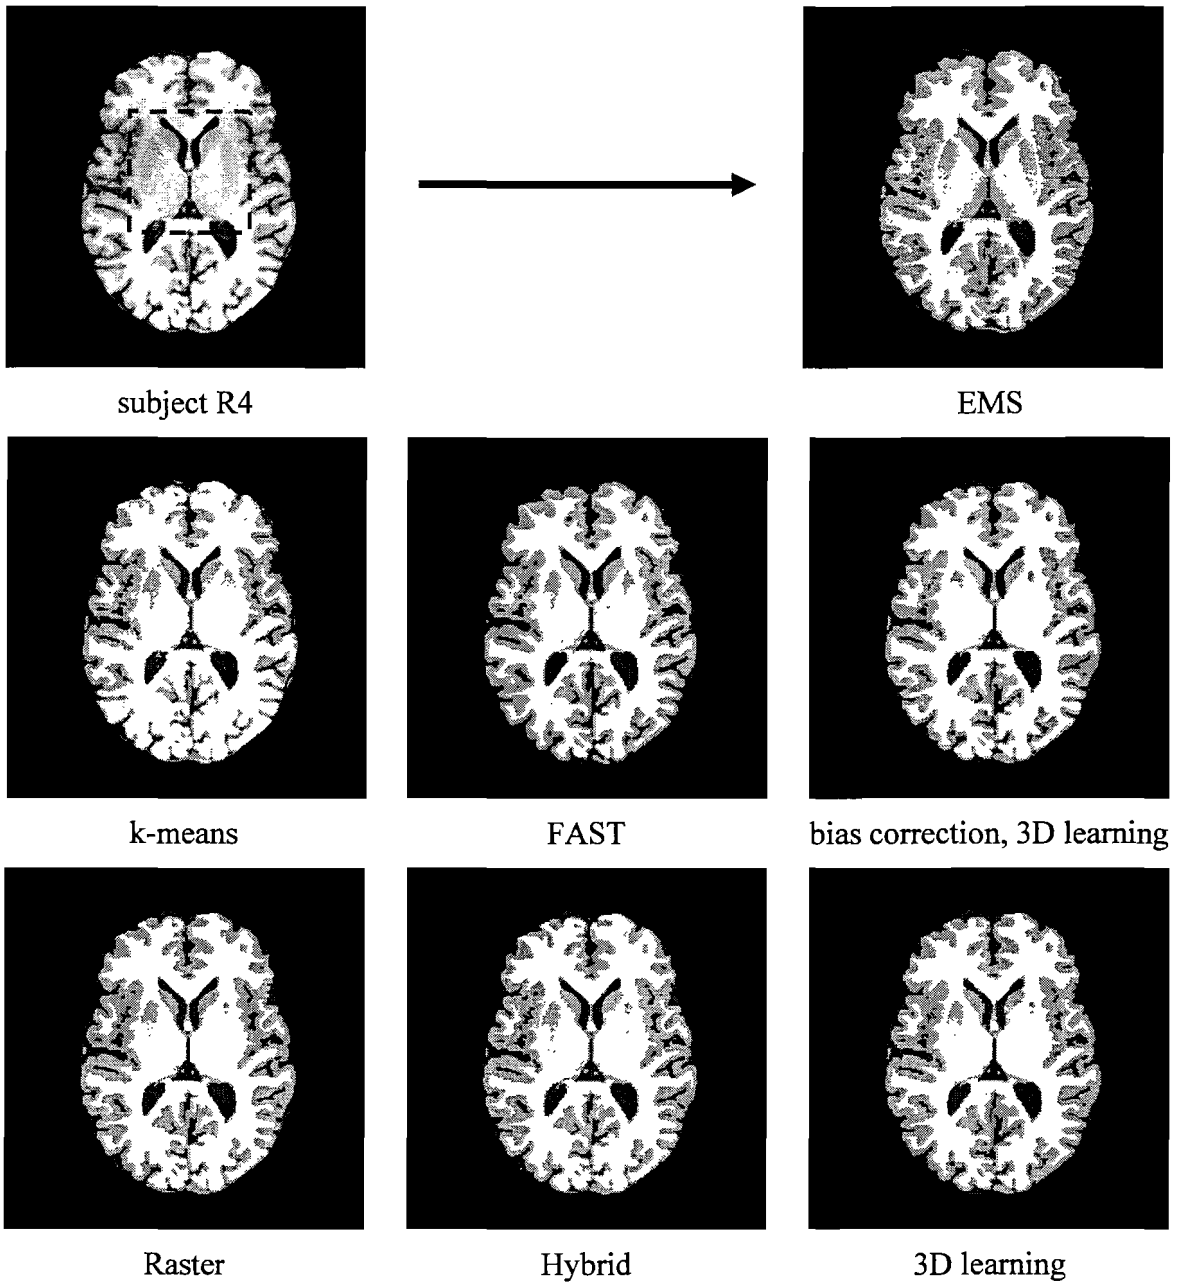

Figure 23: An axial slice of the segmentation of the different methods on subject R4, mainly showing the misclassification in the deep grey matter structures.

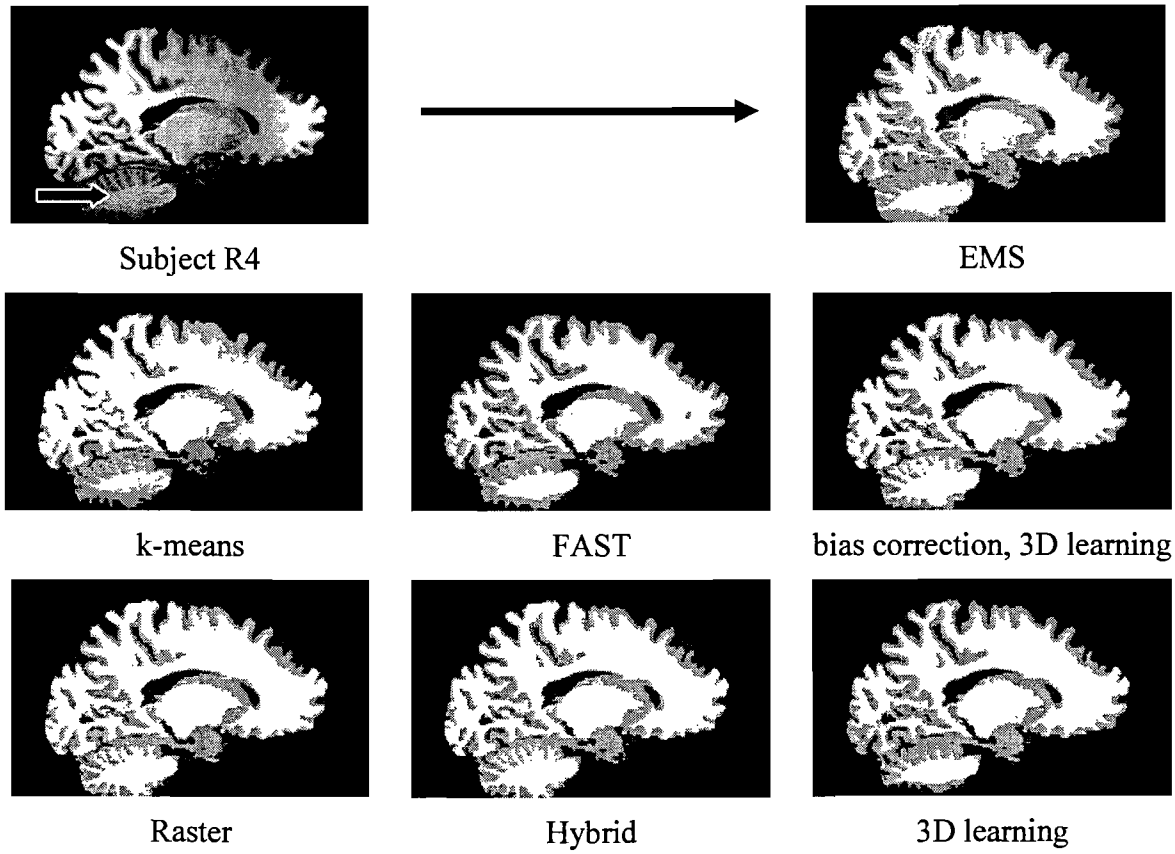

Figure 24: An sagittal slice of the segmentation of the different methods on subject R4, mainly showing the misclassification in the deep grey matter structures.

Additional inspection

As the deep grey matter areas have an intensity between the GM and WM intensities it could be suggested that they can be modeled by fuzziness. However, the study of the fuzzy membership of the voxels in the proposed methods indicates that this does not overcome the problem. To illustrate this, the fuzzy segmentation of the voxels as produced by the 3D learning method is depicted in Figure 25 where it is compared with the probabilistic segmentation of EMS. The fuzzy segmentation is produced by storing the fuzzy membership (15) for all classes for each location during the run of the algorithm. We can see that in the fuzzy segmentations some more details are visible compared to the hard segmentations. However, for example the thalamus (indicated with the arrow) can still not be distinguished in the fuzzy segmentation of 3D learning, i.e. it is classified as 100% WM. Study of the stored reference vectors depicted in Figure 26 shows this is caused by the fact that the relatively large deep grey matter areas without clear boundaries are mistaken for inhomogeneity; i.e. the smooth natural variations are mistaken for inhomogeneity. Therefore, it seems an atlas is essential for good deep grey matter segmentation.

Furthermore, Figure 26 allows us to compare the behavior of the three proposed methods. Like the fuzzy segmentations, these images were created by storing the reference vectors of the winner class for each location[15]. Contrary to the similar segmentations produced by the methods there are clear differences in the evolution of the reference vectors. As the raster scanning progresses fast in the horizontal direction there is almost no adaptation to changes in

[15] Note that the storage of *all* reference vectors is integrated in the 3D-learning method by definition.

this direction, consequently, the reference vector approximates the average of a tissue class in a horizontal line. More adaptation in the horizontal direction is possible with the hybrid scan, however, the results appear blocky and the smoothness constraint of the bias field is not completely maintained. On the other hand, 3D learning allows for the best adaptation in all directions while preserving the smoothness constraint since it is implied by averaging with rule (21). If we combine this observation with the results in the previous tests on real and simulated data it is clear that 3D learning does pay off.

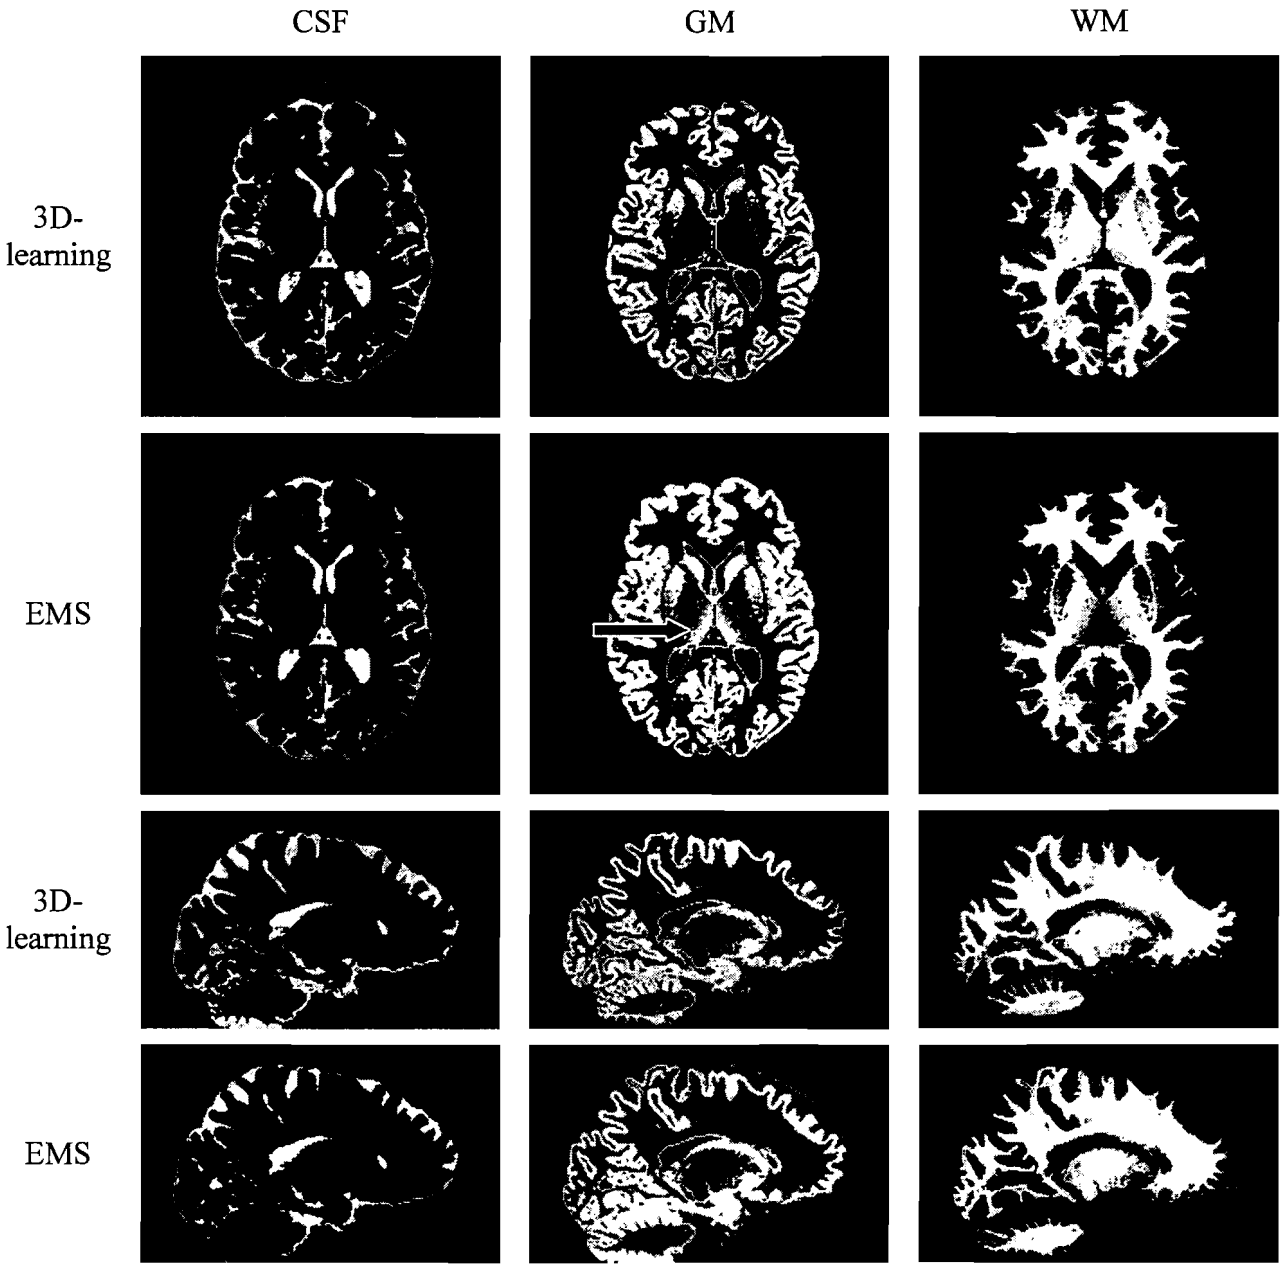

Figure 25: Fuzzy segmentation of the 3D learning method compared to the probabilistic segmentation of EMS.

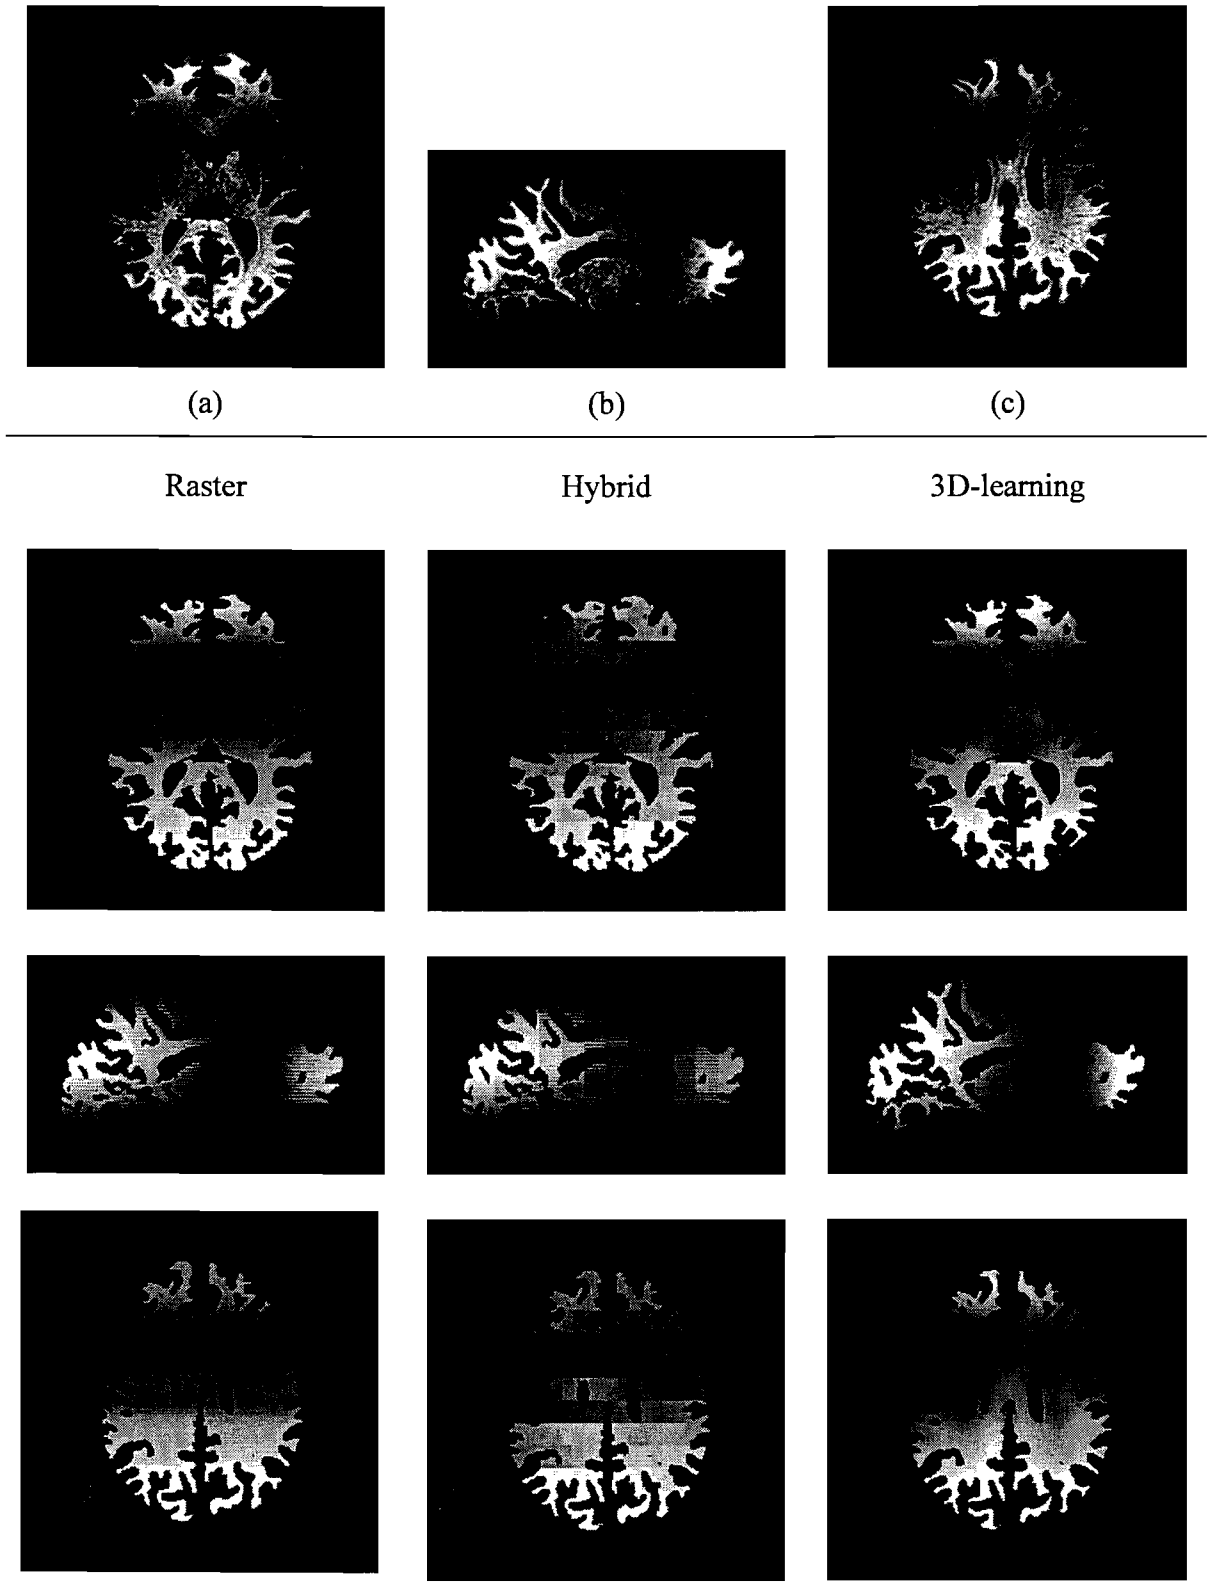

Figure 26: Top row shows subject R4 with contrast enhanced to show the inhomogeneity, (a) and (b) axial slices, (b) sagittal slice. Remaining images show the reference vectors of the winner class for all locations of the three proposed methods. Note that due to the high contrast setting only the reference vectors of the WM class are visible.

7 Conclusion and discussion

A new brain tissue segmentation method for MR images has been proposed. The method is LVQ-based and uses online learning for the handling of the intensity inhomogeneity. Together with an efficient integration of spatial context this leads to an efficient algorithm for segmentation. A significant reduction of computation time (a factor of approximately 20) compared to the EM-based methods has been demonstrated with this method. Tests on simulated data showed segmentation accuracy comparable with these EM-based methods. However, tests on real data where EMS was used as ground truth showed lower performance than the EM-based FAST.

It should be noted that the focus in the literature has not been on speed, since the accuracy of the methods is still the main subject. Therefore, it is not known how far the EM-based methods could be optimized in computational sense. Perhaps sub-sampling, algorithmic improvements, and code optimization can drastically speed up these methods and take away the need for an alternative method if speed is an issue.

The attempt to use the Hilbert space-filling curve exposed a weakness of the use of online learning for compensation for intensity inhomogeneity. Namely, we assume that the drifting intensities can be tracked everywhere. However, this is not true for all data and this leads to gaps in the observed data along the scanning path, where the intensities can not be tracked since the model is not updated. Consequently, after a gap the model might mismatch the real intensities and lead to misclassifications. The EM methods, however, are not vulnerable to gaps.

The tests showed that there is difference between the proposed methods. Hybrid scanning does improve the classification compared to raster scanning if there is strong intensity drift in the scanning direction. However, in the overall Dice coefficients on the test set no real differences are observed and raster appears slightly better. Although, the overall validation showed clearly that 3D learning does improve the inhomogeneity compensation and stability, and thus the classification, and therefore it is the preferred method.

Further study of the LVQ scheme with its variants might improve the performance. For example the tracking ability of the online learning scheme might be improved by explicitly modeling the multiplicative effect of the bias field in the update rule. Meaning that if one class increases in intensity, probably the other classes should do so as well. This can be accomplished by the use of a different neighborhood function where the multiplicative effect is modeled by scaling all the reference vectors accordingly. However, the use of this rule would not allow for the relative tissue intensities to be learned. Possibly, a combination of the present rule with the rule suggested here could be considered. Karayiannis et al. [Karayiannis99] discusses the use of different definitions of the membership and neighborhood functions and applies these to MR brain segmentation, however in the traditional sense where the learning and classification phase are separated. Furthermore, the higher modeling of the tissue distributions as Gaussian mixtures (as is used in the EM methods) might be necessary for the real data, where more natural variations are present. In tests, however, it was found out that integrating the variance of the tissue intensity distributions in the proposed method led to instabilities.

An important observation is that validation on the BrainWeb simulated data does not validate the performance of a method on real data. This is due to the fact that the simulated brains do not fully model the real data, namely, the natural variations of the tissue classes are not modeled in the simulated data, as is also stated in Aubert-Broche et al. [Aubert-BrocheNI06].

These natural tissue variations lead to misclassifications in the deep grey matter areas of real data, by the proposed methods as well as by the EM-based FAST. Only EMS can classify these areas correctly by the use of the prior atlas. There are also papers which address the classification of the deep grey matter structures specifically [Murgasova06; Pohl06; Fischl02].

It could be suggested that these areas can also be identified by fuzzy membership as they appear as large partial volume areas. However, it was shown that fuzziness alone does not overcome this problem, since the relatively large deep grey matter structures without clear boundaries are mistaken for intensity inhomogeneity. Finally, introducing extra classes which model the mixed CSF-GM and GM-WM classes could be suggested. However, it was found in experiments that this destabilizes the algorithm. Namely, if a higher number classes are present, tracking ability of the intensity inhomogeneity decreases, since often the classifier will hop to another class instead of tracking the inhomogeneity.

It is expected that integration of atlas information in the proposed method will improve the segmentation accuracy. However, atlas registration is currently a time consuming task and if no quick method is available, the speed advantage of the proposed method is less pronounced. Furthermore, if atlas information is needed for the classification, the EM-based methods integrating the atlas registration (see Section 2.2) in the iterative loop become of more interest. These methods allow for more accurate registration and can provide labeling of structures as well. Alternatively, improvement of deep grey matter classification might be accomplished by the use of multiple modalities or a special scanning protocol [Prince95; Fischl04] where the GM and WM classes are separated better, and thus atlas information is not needed.

Although the performance is not highly accurate on real data, the method could for example be used in cortical grey matter segmentation, CSF segmentation, bias estimation/correction, or lesion detection. Furthermore, there might be other domains where the proposed method can be applied, for example in microscopy where often illumination is non-uniform.

References

[Alirezaie95] Alirezaie, J., Jernigan, M.E., and Nahmias, C., Neural network based segmentation of magnetic resonance images of the brain, (1995) IEEE Nuclear Science Symposium & Medical Imaging Conference 3, pp. 1397-1401.

[Arnold01] Arnold J.B., Liow J., Schaper K.A., Stern J.J., Sled J.G., Shattuck D.W., Worth A.J., Cohen M.S., Leahy R.M., Mazziotta J.C., and Rottenberg D.A., Qualitative and Quantitative Evaluation of Six Algorithms for Correcting Intensity Nonuniformity Effects, NeuroImage, Vol. 13, No. 5. (May 2001), pp. 931-943.

[Ashburner05] Ashburner, J. and Friston, K.J., Unified segmentation, NeuroImage, Volume 26, Issue 3, , 1 July 2005, Pages 839-851.

[Aubert-Broche06MI] Aubert-Broche, B., Griffin, M., Pike, G.B., Evans, A.C., and Collins, D.L., Twenty New Digital Brain Phantoms for Creation of Validation Image Data Bases, Medical Imaging, IEEE Transactions on , vol.25, no.11pp. 1410- 1416, Nov. 2006

[Aubert-Broche06NI] Aubert-Broche, B., Evans, A.C., and Collins, L., A new improved version of the realistic digital brain phantom, NeuroImage, Volume 32, Issue 1, , 1 August 2006, Pages 138-145.

[Awate06] Awate, S.P., Tasdizen, T., Foster, N. and Whitaker, R.T., Adaptive Markov modeling for mutual-information-based, unsupervised MRI brain-tissue classification, Medical Image Analysis, Volume 10, Issue 5, The Eighth International Conference on Medical Imaging and Computer Assisted Intervention - MICCAI 2005, October 2006, Pages 726-739.

[Belaroussi06] Belaroussi, B., Milles, J., Carme, S., Zhu, Y.M., and Benoit-Cattin, H., Intensity non-uniformity correction in MRI: Existing methods and their validation, Medical Image Analysis, Volume 10, Issue 2, April 2006, Pages 234-246

[Besag86] Besag, J., On the Statistical Analysis of Dirty Pictures, Journal of the Royal Statistical Society. Series B (Methodological), Vol. 48, No. 3 (1986), pp. 259-302

[Bezdek95] Bezdek, J.C., and Pal, N.R., Two soft relatives of learning vector quantization, (1995) Neural Networks, 8 (5), pp. 729-743.

[BrainWeb] Cocosco, C.A., Kollokian, V., Kwan, R.K.-S., and Evans, A.C., BrainWeb: Online Interface to a 3D MRI Simulated Brain Database, NeuroImage, vol.5, no.4, part 2/4, S425, 1997 -- Proceedings of 3-rd International Conference on Functional Mapping of the Human Brain, Copenhagen, May 1997., http://www.bic.mni.mcgill.ca/brainweb/

[Cai07] Cai, W., Chen, S., and Zhang, D., Fast and robust fuzzy c-means clustering algorithms incorporating local information for image segmentation, Pattern Recognition, Volume 40, Issue 3, March 2007, Pages 825-838.

[Cheng06] Cheng, H., and Huang, F., Magnetic resonance imaging image intensity correction with extrapolation and adaptive smoothing, (2006) Magnetic Resonance in Medicine 55 (4), pp. 959-966.

[Cheng95] Cheng, Y., Mean shift, mode seeking, and clustering, Pattern Analysis and Machine Intelligence, IEEE Transactions on, Vol.17, Iss.8, Aug 1995, Pages: 790-799

[Clarke95] Clarke, L.P., and Hall, L.O., MRI segmentation: methods and applications., Magn Reson Imag 13(3): 343-368 (1995)

[Cohen00] Cohen, M.S., Dubois, R.M., and Zeineh, M.M., Rapid and effective correction of RF inhomogeneity for high field magnetic resonance imaging, Human Brain Mapping, vol. 10, no. 4, pp. 204-211, 2000.

[Cuadra05] Cuadra, M.B., Cammoun, L., Butz, T., Cuisenaire, O., and Thiran, J.P., Comparison and validation of tissue modelization and statistical classification methods in T1-weighted MR brain images, Medical Imaging, IEEE Transactions on, vol. 24, number 12, 1548-1565, 2005

[Damkat06] Damkat, C., Ekin, A., and de Haan, G., Adaptive segmentation of brain MRI using context and online learning, Proc. of Symposium on Intelligent Algorithms 2006 (SOIA06), 121-133

[DeHaan93] de Haan, G., Biezen, P.W.A.C., Huijgen, and H., Ojo, O.A., True-motion estimation with 3-D recursive search block matching, Circuits and Systems for Video Technology, IEEE Transactions on , vol.3, no.5pp.368-379, 388, Oct 1993

[Dice45] Dice, L.R., Measures of the Amount of Ecologic Association Between Species, Ecology, Vol. 26, No. 3 (Jul., 1945), pp. 297-302

[Fischl02] Fischl, B., Salat, D.H. , Busa, E., Albert, M., Dieterich, M., Haselgrove, C., van der Kouwe, A., Killiany, R., Kennedy, and D., Klaveness, S., Whole Brain Segmentation: Automated Labeling of Neuroanatomical Structures in the Human Brain, Neuron, Volume 33, Issue 3, 31 January 2002, Pages 341-355.

[Fischl04] Fischl, B., Salat, D.H., van der Kouwe, A.J.W., Makris, N., Segonne, F., Quinn, B.T., and Dale, A.M., Sequence-independent segmentation of magnetic resonance images, NeuroImage, Volume 23, Supplement 1, Mathematics in Brain Imaging, 2004, Pages S69-S84.

[Garcia05] Garcia, M., Fernandez, E., Grana, M., and Torrealdea, F. J., A gradient descent MRI illumination correction algorithm, In: Proceedings Computational Intelligence and Bioinspired Systems, 8th International Work Conference on Artificial Neural Networks, Barcelona, Spain, June 8-10, 2005. Berlin/Heidelberg: Springer 2005. Lecture Notes in Computer Science Vol. 3512. 2005: 913-20

[Gray18] Gray, H., Anatomy of the human body, 20th ed., Philadelphia: Lea & Febiger, 1918, ISBN: 1-58734-102-6, On-line edition: Bartleby.com, 2000. www.bartleby.com/107/.

[Hu04] Hu, X., and Norris, D.G., Advances in high-field magnetic resonance imaging, (2004) Annual Review of Biomedical Engineering, 6, pp. 157-184.

[Huo06] Hou, Z., A Review on MR Image Intensity Inhomogeneity Correction, International Journal of Biomedical Imaging, vol. 2006, Article ID 49515, 11 pages, 2006.

[Jannin06] Jannin, P., Krupinski, E., and Warfield, S., Guest Editorial Validation in Medical Image Processing, Medical Imaging, IEEE Transactions on , vol.25, no.11pp. 1405-1409, Nov. 2006

[Jimenez-Alaniz06] Jimenez-Alaniz, J. R., Medina-Banuelos, V., and Yanez-Suarez, O., Data-driven brain MRI segmentation supported on edge confidence and a priori tissue information, IEEE Transactions on Medical Imaging. Jan. 2006, 25(1): 74-83

[Karayiannis99] Karayiannis, N. B., and Pai, P.-I., Segmentation of magnetic resonance images using fuzzy algorithms for learning vector quantization, (1999) IEEE Transactions on Medical Imaging 18 (2), pp. 172-180.

[Kohonen86] Kohonen, T., Learning Vector Quantization, Helsinki University of Technology, Laboratory of Computer and Information Science, Report TKK-F-A-601, 1986

[Kohonen90] Kohonen, T., The self-organizing map, Proceedings of the IEEE, Volume 78, Issue 9, Sept. 1990 Page(s): 1464 - 1480

[Kong05] Kong, J., Zhang, J., Lu, Y., Wang, J., and Zhou, Y., A novel approach for adaptive unsupervised segmentation of MRI brain images, (2005) Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 3789 LNAI, pp. 918-927.

[Liew06] Liew, A.W.C., and Yan, H., Current methods in the automatic tissue segmentation of 3D magnetic resonance brain images, Current medical imaging reviews 2 (1): 91-103 FEB 2006

[Likar01] Likar, B., Viergever, M.A., and Pernus, F., Retrospective correction of MR intensity inhomogeneity by information minimization, Medical Imaging, IEEE Transactions on , vol.20, no.12pp.1398-1410, Dec 2001

[Luo05] Luo, J., Zhu, Y., Clarysse, P., and Magnin, I., Correction of bias field in MR images using singularity function analysis, IEEE Transactions on Medical Imaging (ISSN: 02780062), date: 2005, volume: 24, issue: 8, startpage: 1067

[MacQueen67] MacQueen, J. B., Some Methods for classification and Analysis of Multivariate Observations, Proceedings of 5-th Berkeley Symposium on Mathematical Statistics and Probability, Berkeley, University of California Press, 1:281-297, 1967

[Madabhushi06] Madabhushi, A., Udupa, J.K., and Souza, A., Generalized scale: Theory, algorithms, and application to image inhomogeneity correction, (2006) Computer Vision and Image Understanding 101 (2), pp. 100-121.

[Murgasova06] Murgasova, M., Dyet, L., Edwards, D., Rutherford, M., Hajnal, J.V., and Rueckert, D., Segmentation of Brain MRI in Young Children, Medical Image Computing and Computer-Assisted Intervention, MICCAI 2006, Pages 687- 694

[Pham00] Pham, D.L., X, C., and Prince, J.L., Current method in medical image segmentation, Annual Review of Biomedical Engineering 2000 2, 315-337

[Pham99] Pham, D.L., and Prince, J.L., Adaptive fuzzy segmentation of magnetic resonance images, (1999) IEEE Transactions on Medical Imaging 18 (9), pp. 737-752.

[Pohl06] Pohl, K.M., Fisher, J., Grimson, W.E.L., Kikinis, R., and Wells, W.M., A Bayesian model for joint segmentation and registration, NeuroImage, Volume 31, Issue 1, , 15 May 2006, Pages 228-239.

[Prince95] Prince, J.L., Tan, Q., and Pham, D.L., Optimization of MR pulse sequences for Bayesian image segmentation, Med. Phys. 22 (1995) (10), pp. 1651–1656.

[Saha05] Saha, P. K., A new non-parametric method for image intensity inhomogeneity correction using a non-uniform gradient filter and path integrals, Proceedings of the SPIE The International Society for Optical Engineering. 2005, 5747(1): 1544-53

[Shattuck01] Shattuck, D.W., Sandor-Leahy, S.R., Schaper, K.A., Rottenberg, D.A., and Leahy, R.M., Magnetic resonance image tissue classification using a partial volume model, NeuroImage 13 (2001) (5), pp. 856-876

[Shen05] Shen, S., Sandham, W., Granat, M., and Sterr, A., MRI fuzzy segmentation of brain tissue using neighborhood attraction with neural-network optimization, IEEE Transactions on Information Technology in Biomedicine 9 (3), pp. 459-467, 2005

[Sled97] Sled, J.G., Zijdenbos, A.P., and Evans, A.C., A comparison of retrospective intensity non-uniformity correction methods for MRI, in Information Processing in Medical Imaging, Lecture Notes in Computer Science vol. 1230, Springer, pp. 459-464, 1997.

[Sled98] Sled, J. G., Zijdenbos, A. P., and Evans, A. C., A non-parametric method for automatic correction of intensity non-uniformity in MRI data, IEEE Transactions on Medical Imaging, vol. 17, pp. 87-97, February 1998.

[Suri01] Suri, J. S., Singh, S. and Reden, L., Computer Vision and Pattern Recognition Techniques for 2-D and 3-D MR Cerebral Cortical Segmentation: A State-of-the-Art Review, Accepted for Publication In: Inter. Journal of Pattern Analysis and Applications, 2001

[VanLeemput99] Van Leemput, K. L., Maes, F., Vandermeulen, D., and Suetens, P., Automated model-based tissue classification of MR images of the brain, IEEE Trans. Med. Imag., vol. 18, no. 10, pp. 897-908, Oct. 1999.

[VanLeemputPHD01] Van Leemput, K., Quantitative analysis of signal abnormalities in MR imaging for multiple sclerosis and creutzfeldt-jakob disease, K.U.Leuven, May 2001, Leuven, Belgium

[Vovk04] Vovk, U., Pernuš, F., and Likar, B., MRI intensity inhomogeneity correction by combining intensity and spatial information, Physics in Medicine and Biology, vol. 49, no. 17, pp. 4119-4133, 2004.

[Vovk06] Vovk, U., Pernuš, F. and Likar, B., Intensity inhomogeneity correction of multispectral MR images, NeuroImage, Volume 32, Issue 1, 1 August 2006, Pages 54-61

[Vovk07] Vovk, U., Pernuš, F. and Likar, B., A Review of Methods for Correction of Intensity Inhomogeneity in MRI, IEEE Trans. Med. Imag., vol. 26, no. 3, pp. 405-421, March 2007.

[Wells96] Wells, W. M., Grimson, W. E. L., Kikinis, R., and Jolesz, F. A., Adaptive segmentation of MRI data, IEEE Trans. Med. Imag., vol. 15, no. 4, pp. 429-442, Aug. 1996.

[Wu03] Wu, K.-L., and Yang, M.-S., A fuzzy-soft learning vector quantization, Neurocomputing, Volume 55, Issues 3-4, October 2003, Pages 681-697

[Wu06] Wu, K.-L., and Yang, M.-S., Alternative learning vector quantization, (2006) Pattern Recognition 39 (3), pp. 351-362.

[Yan95] Yan, M.X.H., and Karp, J.S., Segmentation of 3D brain MR using an adaptive K-means clustering algorithm, Nuclear Science Symposium and Medical Imaging Conference, 1994., 1994 IEEE Conference Record , vol.4, no.pp.1529-1533 vol.4, 30 Oct-5 Nov 1994

[Zhang01] Zhang, Y., Brady, and M., Smith, S., Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm, Medical Imaging, IEEE Transactions on, vol.20, no.1pp.45-57, Jan 2001

Appendix A: k-means clustering

K-means clustering [MacQueen67] is a parametric data-clustering algorithm. Given a vector space with data points, which form multiple clusters, the k-means algorithm automatically partitions this vector space into classes representing these clusters. In k-means each class is represented only by its mean, the *centroid*, and the number of classes, k, is fixed a priori. After initialization, the algorithm works fully automatic through iterative minimization of the within-class variance. The following equation expresses this minimization objective:

$$J = \sum_{j=1}^{k} \sum_{i \in S_j} \| \mathbf{x}_i - \mathbf{c}_j \|^2$$

this is the so-called objective function, which should be minimized, S_j are the samples assigned to class j, $\mathbf{x}_i$ are the samples, and $\mathbf{c}_j$ are the class centroids.

This objective function can be minimized by iteratively:

- assigning each sample to the class with the nearest centroid,
- updating the class centroids by calculating the mean of the samples assigned to each class.

The centroids should be initialized; this can be done randomly or uniformly over the sample space, or via some smarter initialization algorithm. Alternatively, the assigned classes of the samples can be initialized randomly and the algorithm can be started from the second step. Since the initialization is not trivial, especially if the vector space is higher dimensional, and the resulting clustering is dependent on this initialization, since it can converge into local minima, the use of k-means should be carefully considered.

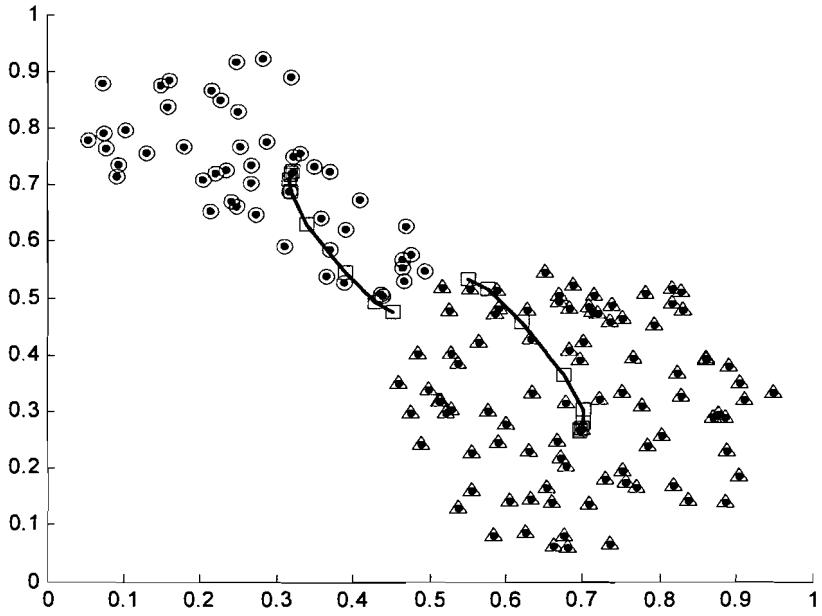

Figure 27: Result of the k-means algorithm on a 2D dataset, with $k = 2$.

In Figure 27, a result of the k-means algorithm is shown. This is the result of the algorithm on a two-dimensional dataset with k, the number classes, set to 2. The values of the centroids

during the iteration steps are indicated with the squares, starting from the middle, and the followed trail during the iterations is indicated with a thick line. The final classification for the two classes is indicated with the triangles and circles.

Appendix B: 3D Recursive Search Motion Estimation and its relation to the proposed method

Given the success of the *3D recursive search* (3DRS) algorithm in motion estimation [DeHaan93], where motion is estimated on block-basis for objects in a video sequence, and recently in other video related applications, 3DRS has been used as an inspiration for an algorithm for the segmentation of a MRI image, which can be thought of as a sequence of frames. Although there are some fundamental differences between motion estimation and image segmentation, a 3DRS-based approach resulted in an interesting new algorithm for this specific segmentation problem.

In motion estimation the video frame is (usually) divided into small blocks (e.g. 8x8 pixels). The objective is to find motion vectors for each of these blocks. These motion vectors should reflect the motion of the objects in video sequence, i.e. the displacement of them between the current and previous frame, on block resolution. The simplest way to find these motion vectors is to compare each block with displaced blocks in the previous frame in a predefined search range. The displacement vector resulting in the best match, the *motion vector*, is then assigned to the block. This method takes a lot of comparisons for each block (in the order of 64x32, the size of the search range, block-comparisons), which is computationally expensive. Furthermore, the results do not portray the real object motion, due to local ambiguity, which is caused by repetition in the texture, total lack of texture, or change in brightness. 3DRS solves this problem by only testing a small set of candidates for each block. These candidates are obtained from the previously calculated motion vectors in the neighborhood. These candidates are correlated, since objects are assumed to be bigger than blocks and they have inertia, and thus provide good predictions for the motion vector of the current block, while implying spatial (and temporal) consistency.

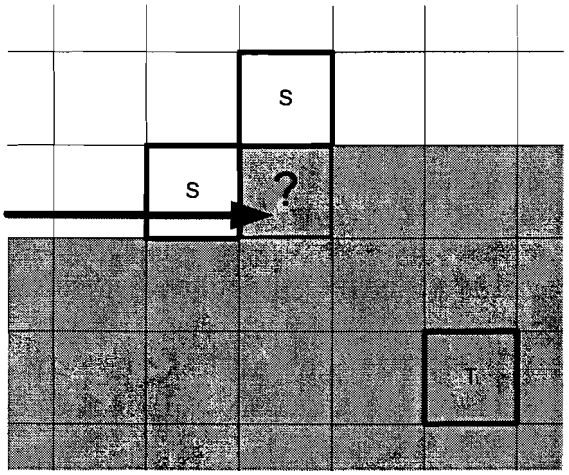

Figure 28: Candidates for the estimation of '?', S: Spatial candidates, T: Temporal candidate. The grey area is not processed yet in the current frame.

In Figure 28, we can see where the candidates for the estimation of block '?' are taken from[16]. The blocks are handled one by one in a scanning manner, as indicated by the arrow. Some candidates can be taken from the neighborhood in the same frame since there the motion vectors are already known, these are denoted with the 'S', the so-called spatial candidates. However, in the grey area, which has not been processed yet, the motion vectors are not known, but they were known for these locations in the estimation of the previous frame, from where they can be taken. These candidates are called temporal candidates, denoted with 'T'. Since the process needs to initiate and adapt, there are also update candidates, which are spatial candidates with a noise vector added to them. All these candidates are tested, some with penalties to prefer certain candidates, and the winner is assigned to the current block. Finally, by alternating the scanning direction each line (meandering), convergence is achieved faster. The algorithm steps are given in pseudo-code in Table 6. For a more detailed description of 3DRS please refer to [DeHaan93].

```
start
        take candidates from the neighborhood
        generate extra update candidates by adding noise to some candidates
        test all the candidates on the current block with penalties for some
        candidates
        assign the winner to the current block
        go to the next block in the scanning order
goto start
```

Table 6: Pseudo-code for the 3DRS motion estimation algorithm.

Mapping this algorithm to segmentation is not straightforward since there are fundamental differences between motion estimation and intensity-based classification, which is what we are doing in the segmentation of the MR brain images. First of all, in motion estimation only the motion vector is unknown, while in segmentation both the class model and the classifications are unknown and these are mutually dependent. Furthermore, in motion estimation the motion of the current block identifies itself by the displacement vector giving the best block-match, i.e. the motion can be locally observed, although this can be ambiguous and is not always true. In intensity-based classification the class of a current sample (e.g. pixel intensity) can only be identified in the context of the (local) intensity distribution/histogram, as belonging to a cluster. From this histogram class statistics can be deduced, such as class mean, variance, etc. Contrary to motion vectors, the histogram cannot be observed locally since it is calculated from a set of samples. Finally, the number of classes in intensity-based classification cannot be deduced explicitly and is thus usually set to a predefined number. While in motion estimation, where the motion identifies itself, the number of classes is not an issue.

By adopting aspects of 3DRS, which is inherently an efficient algorithm, it was expected that this would lead to a computationally efficient algorithm for segmentation of brain MRI. The aspects of 3DRS that were adopted in the MR segmentation algorithm are the way we use spatial context by using the earlier calculated classifications from the neighborhood, and the sample-by-sample model updating, optionally also with earlier calculated models from the neighborhood (3D-learning). The purpose of the spatial context is to imply spatial consistency

[16] The shown set up is the one typically used in 3DRS, although others can be used

in the obtained segments and introduce robustness to noise influence by biasing the classification of the current sample towards the classes found in the neighborhood. This resembles so-called *markov random fields* (MRFs) that are widely used in image segmentation. These neighborhood classifications can be obtained in a similar way as the candidates in 3DRS motion estimation if we interpret the MR image as a sequence of slices. Using the context in this way prevents the need for iterative methods that are used for MRF, such as the commonly used *iterated conditional modes* (ICM) algorithm [Besag86; Shattuck01]. Sample-by-sample updating of model parameters was found to resemble online learning. This was found to be applicable particularly to MR image segmentation since it enables a new approach to handling the intensity inhomogeneity inherent to MR images, and this is in a more efficient way than it is handled in the present known algorithms. This eventually led to an online competitive learning scheme based on *learning vector quantization* with integration of spatial context and a special scanning order of the data for MRI segmentation.