

Prior Information for Brain Parcellation

by

Kilian Maria Pohl

Dipl.-Math. techn., University of Karlsruhe, Germany (1999)

Submitted to the Department of Electrical Engineering and Computer
Science

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2005

© Massachusetts Institute of Technology 2005. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
May 19, 2005

Certified by
W. Eric L. Grimson
Bernard Gordon Professor of Medical Engineering
Thesis Supervisor

Certified by
Ron Kikinis
Professor of Radiology, Harvard Medical School
Thesis Supervisor

Accepted by
Arthur C. Smith
Chairman, Department Committee on Graduate Students

Prior Information for Brain Parcellation

by

Kilian Maria Pohl

Submitted to the Department of Electrical Engineering and Computer Science
on May 19, 2005, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Electrical Engineering and Computer Science

Abstract

To better understand brain disease, many neuroscientists study anatomical differences between normal and diseased subjects. Frequently, they analyze medical images to locate brain structures influenced by disease. Many of these structures have weakly visible boundaries so that standard image analysis algorithms perform poorly. Instead, neuroscientists rely on manual procedures, which are time consuming and increase risks related to inter- and intra-observer reliability [53]. In order to automate this task, we develop an algorithm that robustly segments brain structures.

We model the segmentation problem in a Bayesian framework, which is applicable to a variety of problems. This framework employs anatomical prior information in order to simplify the detection process. In this thesis, we experiment with different types of prior information such as spatial priors, shape models, and trees describing hierarchical anatomical relationships. We pose a maximum *a posteriori* probability estimation problem to find the optimal solution within our framework. From the estimation problem we derive an instance of the Expectation Maximization algorithm, which uses an initial imperfect estimate to converge to a good approximation.

The resulting implementation is tested on a variety of studies, ranging from the segmentation of the brain into the three major brain tissue classes, to the parcellation of anatomical structures with weakly visible boundaries such as the thalamus or superior temporal gyrus. In general, our new method performs significantly better than other standard automatic segmentation techniques. The improvement is due primarily to the seamless integration of medical image artifact correction, alignment of the prior information to the subject, detection of the shape of anatomical structures, and representation of the anatomical relationships in a hierarchical tree.

Thesis Supervisor: W. Eric L. Grimson

Title: Bernard Gordon Professor of Medical Engineering

Thesis Supervisor: Ron Kikinis

Title: Professor of Radiology, Harvard Medical School

Readers: Alan Willsky

William M. Wells III

Acknowledgments

I always saw this document as the final step in a longer journey, which I am sad to have completed. For somebody, who was almost banned from school 20 years ago because of a lack of intellectual capabilities, these last four years were quite a treat as I have enjoyed the freedom to explore. Looking back, this journey was so beautiful because of the many wonderful personal interactions I had.

First, I would like to thank my committee. I am in deep gratitude to Prof. Eric Grimson for giving me the opportunity to do a Ph.D. at this institution. I admire Prof. Alan Willsky's enthusiasm for the advancement in the field of sciences. I am thankful to Prof. Sandy Wells for teaching me so much about the field. I am in great debt to Prof. Ron Kikinis whose support has carried me through many valleys.

This thesis would be incomplete without the many collaborations I had. I want to thank Prof. Simon Warfield for introducing me to the field of segmentation and Sylvain Bouix for familiarizing the Schizophrenia group with my segmentation method. I would like to thank the Schizophrenia group for providing me with the necessary data. I am grateful to Steve Pieper for introducing my technology to a larger audience and John Fischer for being my cheerleader and ear during this journey.

As I am not a native English speaker, I am especially in debt to the many readers of this document which include Ali Rahimi, Bruce Fischl, Bryan Russell, Corey Kemper, Eric Cosman, Koen Van Leemput, Lauren O'Donnell, Lilla Zollei, and Polina Golland.

Finally, I would like to acknowledge the people that have carried me through this journey in my personal life. First, I would like to thank my two sisters, Prisca and Anna, for always supporting me throughout my life. I would like to thank my parents for showing me where dreams can take you. I would like to thank the Krushnics for providing me with shelter in so many ways. My deep admiration goes towards Monica, without whom I would never have been accepted to this program. She always showed me that getting a Ph.D. is a piece of cake compared to the problems people face in the real world. Finally, I would like to thank Frottee who was always at my side and whose calm playful way was a true inspiration to me. I deeply miss her.

Contents

1	Introduction	17
1.1	Prior Information	20
1.2	Registration	22
1.3	Segmentation	23
1.3.1	Modeling Image Artifacts	24
1.3.2	Including Prior Information	24
1.4	Objective and Contribution of the Thesis	27
1.4.1	Combine Registration and Segmentation	28
1.4.2	Model Shape Constraints in Statistical Framework	29
1.4.3	Represent Hierarchical Relationships in a Tree	30
1.5	Overview of Validation Techniques	31
1.6	Organization of the Thesis	32
2	The Expectation-Maximization Algorithm	33
2.1	Background	34
2.2	Deriving the EM Algorithm	36
2.2.1	EM: A Lower Bound Optimization Method	37
2.2.2	Properties of the Lower Bound	39
2.2.3	Summary	42
2.3	Examples of EM-Approaches	42
2.3.1	Applying EM to a Multiple Clustering Example	42
2.3.2	Image Inhomogeneity Correction of MR Images as Maximum Likelihood Problem	46

2.3.3	Image Inhomogeneity Correction of MR Images as a Maximum <i>a Posteriori</i> Probability Estimation Problem	52
2.4	Discussion	56
3	Integration of Spatial Prior into EM Framework	59
3.1	Defining Spatial Priors of Anatomical Structure	60
3.2	Separate Framework for Registration and Segmentation	64
3.2.1	A Sequential Approach	64
3.2.2	Validation	69
3.3	Integrated Registration Framework	75
3.3.1	Derivation of the EM Framework Coupling Registration and Segmentation	76
3.3.2	The Registration Parameters \mathcal{R}	82
3.3.3	An Implementation of the Joint Registration and Segmentation Model	87
3.3.4	Validation	91
3.4	Conclusion	97
4	Modeling Shape Constraints	99
4.1	Shape Atlas	101
4.1.1	Shape Representations	101
4.1.2	Principal Component Analysis on Distance Maps	102
4.2	Coupling Shape Deformation with Brain MRI Segmentation	107
4.2.1	Unified Framework for Image Inhomogeneity Correction, Shape Modeling, and Segmentation	109
4.2.2	Estimating the Shape	112
4.2.3	Validation	115
4.3	Modelling Atlas Registration and Shape Deformation within a EM Segmentation Method	120
4.3.1	A Formulation of Shape Based Segmentation and Registration	121
4.3.2	A Non-Rigid Registration Model	124

4.3.3	Validation	126
4.4	Discussion	130
5	Hierarchical Segmentation	133
5.1	An Anatomically Guided EM Approach	135
5.1.1	Represent Anatomical Hierarchy in a Tree	136
5.1.2	Populate Information within the Tree	138
5.1.3	Integrate EM Approach into Hierarchical Segmentation	140
5.2	Incorporate Prior Information into the Tree	146
5.2.1	Develop Intensity Based Confidence Parameters	147
5.2.2	Regulate Influence of Spatial Priors	149
5.3	Experiments and Validation	152
5.3.1	Parcellating 31 Structures	153
5.3.2	Comparison of Four Different Segmentation Methods	156
5.3.3	Multiple Rater Experiment of the Superior Temporal Gyrus	158
5.3.4	Comparison of Joint Registration and Segmentation Approach Using the Hierarchical Framework	162
5.4	Discussion	165
6	Conclusion	167
6.1	Future Work	168

List of Figures

1-1	Different types of magnetic resonance images	18
1-2	Variations in manual segmentations	19
1-3	Image artifacts in magnetic resonance images	20
1-4	Segmentation of the thalamus	21
1-5	Examples of a probabilistic and a shape atlas	22
2-1	An iteration of an EM approach	39
2-2	Image inhomogeneities in magnetic resonance images	47
2-3	Labelmap derived from the posterior probability of two structures . .	51
3-1	Problems in outlining the thalamus	59
3-2	Spatial prior generation	61
3-3	Spatial priors of the major brain tissue	62
3-4	Difficulties in outlining sulci	63
3-5	Registration of the spatial atlas to magnetic resonance images	65
3-6	Segmentations of the right superior temporal gyrus	70
3-7	Overlap between the reference standard and the spatial prior	71
3-8	Similarities in the intensity pattern of white matter and thalamus . .	73
3-9	Segmentations of the thalamus with EM-Affine and EM-NonRigid . .	74
3-10	A suboptimal segmentation with joint registration and segmentation .	80
3-11	Impact of biased atlas on automatic segmentations	81
3-12	A hierarchical affine registration framework	83
3-13	Similarity measure of the registration parameters	86
3-14	An example of statistical inconsistency	88

3-15	Iteration sequence of joint segmentation and registration approach . . .	89
3-16	Examples of the experiment on synthetic images	92
3-17	Results of the experiment on synthetic images	93
3-18	Examples of segmentations performed by EM-Affine, EM-NonRigid, and EM-Simultaneous-Affine	94
3-19	Pipeline versus uniform EM approach	95
4-1	Difficulties in segmenting the caudate	100
4-2	Example of segmentation with corresponding distance map	101
4-3	Synthetic example of the PCA analysis	103
4-4	The first eigenvector of a joint PCA model	106
4-5	A distance map with corresponding logistic function	113
4-6	Automatic generated 3D models of subcortical structures	115
4-7	Comparison of EM approaches using spatial priors and shape atlases	116
4-8	3D segmentation models produced by EM-Shape	117
4-9	Segmentations of subcortical structures	118
4-10	Error plotted of non-shape versus shape constraint method	119
4-11	A non-rigid alignment approach composed by global rigid registration and local shape deformation	125
4-12	Accuracy of EM approaches for subcortical parcellation	127
4-13	A comparison between different EM approaches	128
4-14	Iteration sequence of EM-Simultaneous-Shape	129
4-15	Iteration sequence of EM-Simultaneous-Shape	131
5-1	Problems with spatial priors for multiple structures	134
5-2	Comparison between binary and multiple structure segmentation . . .	136
5-3	Anatomical structures ordered in a tree	137
5-4	Influence of the hierarchical model on the registration objective function	143
5-5	Hierarchical tree for 31 anatomical structures	154
5-6	Segmenting a brain into 31 different structures	155
5-7	Segmentation of a magnetic resonance image by four methods	156

5-8	Comparison between four different segmentation methods	159
5-9	Example segmentations of the multi rater experiment	160
5-10	Average DICE measure of the superior temporal gyrus	161
5-11	Comparison of automatic and manual segmentations using DICE . . .	162
5-12	Pipeline based EM algorithms versus hierarchical EM approach . . .	164

List of Tables

3.1	Comparison between EM implementation using affine and non-rigid registration	74
5.1	Impact on hierarchical framework on the EM approaches	163

Chapter 1

Introduction

In the last two decades, the field of medical image analysis has greatly influenced many areas in neuroscience. For example, with the advancement of functional imaging modalities, neuroscientists have been increasingly interested in methodologies that can identify cortical and subcortical structures in anatomical imaging modalities. Image analysis methods support this identification process by providing automated segmentation of structures from multimodal images. The surgical removal of brain tumors is another field influenced by medical image analysis. The surgeons sometimes use image-guided navigational systems, which identify the pathology in the multimodal images and track surgical probes with respect to the pathology.

Many neuroscience studies aim to find new disease related anatomical characteristics in order to increase the reliability of diagnosing the illness or improving the effectiveness of treatment methods against the disease. For example, a recent advancement in the study of schizophrenia suggests that the superior temporal gyrus, an anatomical structure partially responsible for auditory system, is altered by the disease [57, 95]. This might explain why some schizophrenic patients experience auditory hallucinations. These studies are often based on the analysis of Magnetic Resonance (MR) images. Figure 1-1 shows an example of three types of MR images. They visualize different characteristics of the brain, which simplifies the detection of anatomical structures that are altered by a disease. The goal of this thesis is to aid this type of analysis by developing principled statistical approaches that automati-




Figure 1-1: Examples of different types of magnetic resonance images that are often used by neuroscientist to better understand diseases. The T1-weighted image was produced through a 3D spoiled gradient-recalled acquisition sequence. The T2-weighted and proton density images were acquired through a double-echo spin-echo sequence resulting in 2D multi slices throughout the brain.

cally segment anatomical structures in MR images of different acquisition formats.

Neuroscientists categorize segmentations of anatomical structures by attributes such as shape, size, or location. These subject specific characterizations are collected for studies testing the hypothesis that the disease alters an anatomical structure. In the schizophrenia example, a scientist might determine the average size of the superior temporal gyrus in the controlled and diseased patient group to evaluate the hypothesis that the structure is enlarged in schizophrenic patients.

The reliability of a study often depends on the number and quality of segmented cases. Neuroscientists mostly rely on human experts specifically trained for this task to perform the segmentations. This process, however, is very labor intensive and increases the risks related to inter- and intra-observer reliability [50] as shown in examples of Figure 1-2. In this experiment, six experts were asked to segment the superior temporal gyrus in the image to the left. None of the resulting segmentations are equivalent.

The field of medical image analysis has developed a variety of automatic segmentation methods to increase the impact of neuroscience studies. Automated methods reduce the manual labor for outlining the structures and the results are generally reproducible. However, these methods generally do not achieve the quality of human



Figure 1-2: The image to the left shows the area around the superior temporal gyrus on an MR image. The other images are the corresponding segmentations of six medical experts. None of the six segmentations are equivalent.

experts so that neuroscientists continue to rely on manual segmentations.

The task of automatically segmenting medical images is challenging as the images are corrupted by several artifacts. Brain MR images, such as in Figure 1-3, include image inhomogeneities, noise, ambiguous intensity patterns with respect to inside and outside the brain, and partial volume effects. The partial voluming effect describes voxels whose intensities are influenced by multiple structures. Furthermore, regions inside and outside the brain may have ambiguous intensity patterns in MR images.

In the last decade the field of medical imaging has made tremendous progress in outlining the three main brain “tissue” classes - corticospinal fluid, white matter and gray matter - in MR images, e.g. [112, 103, 19, 65]. Most neuroscientists, however, are keenly interested in further parcellating these tissue classes into their substructures such as [87, 54].

The detection of the substructures is even more difficult than outlining the basic classes as many of them are defined by weakly visible boundaries. For example, the intensity pattern of the thalamus in MR images are very similar to the neighboring white matter (Figure 1-4.) Algorithms cannot rely on the MR images alone in order to distinguish these two structures. Instead, they are guided by prior information related to these anatomical structures. Prior information captures common attributes of anatomical structures in a population. For example, in the case of the thalamus one might measure the average volume of the thalamus relative to the ventricles within a healthy population group. Automatic segmentation methods could use this information in order to reduce the space of possible labelmaps of the thalamus as the ventricles are easily identifiable.




Figure 1-3: MR images are corrupted by various image artifacts. Those include noise, image inhomogeneities, and partial voluming. In addition, regions inside and outside the brain have ambiguous intensities.

1.1 Prior Information

Atlases capture prior information about anatomical structures. Their use and organization varies depending on the intended purpose [96, 97]. For example, if an algorithm outlines structures with weakly visible boundaries, the spatial distributions of the structures can simplify the task [34, 80], as the location and not the intensity pattern defines the boundary. Spatial distributions are defined by probabilistic atlases [67, 94] capturing the inter-subject variability in brain architecture (see Figure 1-5.) These types of atlases are often generated by analyzing the variability within manually segmented training data [75, 28].

Another type of probabilistic atlas is an intensity-based atlas, in which the combination of multiple scans of the same subject allows a high-resolution representation of the complex brain anatomy [48]. Intensity based atlases are also often used to capture the variability in intensity pattern across a large population. Such atlases can be stationary [112, 60, 108], which means that the prior is independent from the voxel location, or spatially varying [34] so that the prior depends on the voxel location.




Figure 1-4: (a) shows a MR image of the area around the thalamus, which is this image the structure below the dark ventricles. The corresponding manual segmentation of the thalamus is shown in (b). From the intensity pattern in (a) the boundary of the thalamus to the neighboring white matter is not clearly defined.

A very different category of prior information is represented by deformable atlases, which describe the mean shape as well as standard shape differences of anatomical structures. Defining the mean shape itself is a complex problem. For example, simple averaging of subjects generates a blurred image that is not particularly useful. Several approaches have been developed to generate an average model without losing the high initial image resolution associated with MR images [96, 10, 90, 6]. More recently, algorithms [70, 49, 40] use the same concept but calculate implicit correspondences, which are obtained using the voxel sum of squared differences as a voxel agreement measure.

A critical aspect in generating a shape atlas is the representation of an anatomical structure’s boundary. As discussed in [37], shapes can be represented by spherical harmonics [52], point-based models [15], skeleton or medial representations [79, 8], and landmark-based descriptors [7, 17]. The best representation strongly depends on the application. For example, we use signed distance maps [38, 60] for the shape representation of subcortical structures in Chapter 4 because the representation is relatively simple to integrate into our segmentation approach.

To capture the shape variability across a population, researchers in medical image analysis often use Principal Component Analysis [16]. Principal Component Analysis transforms the training data into a lower dimensional eigenvector space. The variabil-

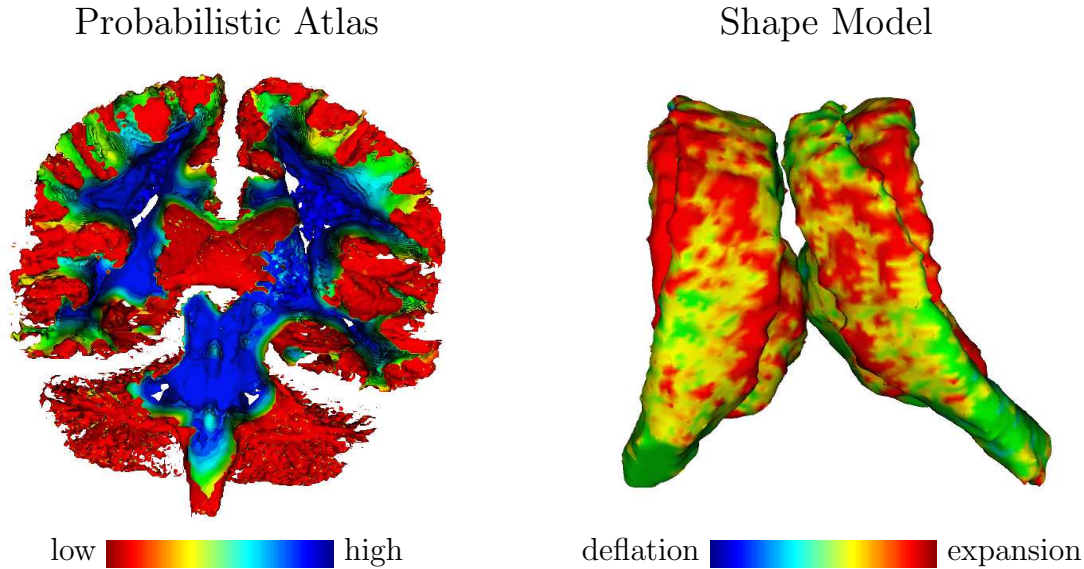


Figure 1-5: These are examples of prior information used by segmentation methods. The image to the left shows part of a probabilistic atlas of the white matter in three dimensions. Areas in red indicate low probability of the structure and blue define high probability regions. The image to the right is a shape model of the ventricles, caudate, and thalamus where red indicates areas of expansions and blue represents deflation with respect to the mean.

ity within the training data is then represented by the most dominant eigenvectors. The method is especially useful when the training data under-represents the variation within a population.

1.2 Registration

In order to incorporate prior information into the segmentation algorithm, the atlas must be aligned to the subject. This task is performed by registration methods, which compute the mapping between two data sets. Registration algorithms can be separated into three categories: rigid, affine, and more general non-rigid. Rigid and affine registration algorithms compute a single, global transform. For example, the affine registration method by Warfield et al.[107] defines the global registration parameters as the setting that produces the minimum entropy between the two data sets.

In contrast to affine and rigid registration methods, non-rigid registration algo-

gorithms produce correspondence fields to match individual voxels of both data sets. A voxel is defined as the smallest unit in a three-dimensional image, similar to a pixel in two dimensions. An interesting class of non-rigid registration methods are physically based elastic matching models, where the constraints are defined by the material's elastic properties [3, 21, 35, 26]. To reduce computational complexity, the constraints are often modeled by a finite element method [11, 30, 109]. Less computationally expensive algorithms are methods that model the matching from the source to the target by spline functions [84, 6, 22, 72, 28, 78]. The simplicity of spline functions is favorable for medical image analysis as they can be specifically tailored towards brain matching by including specific characteristics of the internal cortex [93].

For this thesis, we are specifically interested in non-rigid registration methods, which align MR images of the prior information to the MR images of a specific subject of interest. This task can be achieved by registration methods constrained by the intensity patterns of the source and target images. Based on the theory of optical flow, Thirion [92] and later Guimond [41] experimented with Maxwell Demons using attraction and diffusion attributes for areas within each image. As such, areas within the source images are assigned attraction attributes in correlation to areas in the target images. The algorithm now estimates the mapping between source and target by aligning areas that attract each other.

1.3 Segmentation

Given a range of possible atlas types and registration methods, we now give a brief overview of segmentation algorithms applied to medical images. A segmentation algorithm separates a medical image into different anatomical structures based on the image data and prior information about anatomical appearance and layout. In this thesis we specifically focus on methods that automatically segment MR images. As mentioned, the task is difficult as the MR images are corrupted by image inhomogeneities, noise, and partial volume effects.

1.3.1 Modeling Image Artifacts

Initially, the robustness of automatic segmentation methods was increased by using multiple input channels [105, 61]. In the last decade, explicit models of the image artifacts have been developed. For example, the implementations of [1, 23, 111, 29] focus on the detection and correction of image inhomogeneities. The importance of correctly identifying inhomogeneities for accurately segmenting the image into structures of interest is also emphasized by the EM approach of Wells et al. [112], which simultaneously solves both problems. We will give a detailed account of this method in Chapter 2.

The partial volume effect is another artifact in MR images. This artifact, for example, causes the small gaps between adjacent folds of the neocortical gray matter to be invisible on MR images. Several groups have addressed this issue by explicitly modeling the effect [32, 91] or by performing a topology correction on the anatomical structures [42, 86], as smoothing over small gaps generates anatomical abnormalities.

The third major artifact in MR images - noise - has been addressed by anisotropic filtering [36, 55]. These types of filters prefer intra-region smoothing rather than smoothing across boundaries. Such filters consider intensity data only and neglect anatomical information of the structures of interest. In contrast, Markov random fields can reduce the impact of noise by considering the intensity of the input image as well as modeling the neighborhood relationship of multiple anatomical structures. Various solutions within this framework have been proposed for medical imaging [44, 74, 33, 117, 65, 51, 103, 76].

1.3.2 Including Prior Information

Most of the previous methods rely solely on the image intensity information. They robustly identify anatomical structures like the three major brain tissue classes, whose boundaries are clearly visible in MR images. However, when faced with structures with weakly visible boundaries, such as the superior temporal gyrus, prior knowledge about attributes such as shape or location is necessary.

In light of this problem, template-driven segmentation methods complement the information extracted from patient-specific image data with generic anatomical templates. These approaches generate the segmentations by aligning previously segmented brains to the subject, such as in [12]. They often rely heavily on the accuracy of the registration, which is itself a challenging problem [45]. This concern was addressed by Rohlfing et al. [83] by producing a segmentation based on the alignment results of multiple registration techniques. In general, template-driven segmentation methods strongly depend on the atlas information. This drawback impacts the accuracy of the methods especially for the identification of pathologies, such as tumors, that are difficult to capture in an atlas.

Deformable Model Based Methods

Unlike template driven methods, deformable model based algorithms explicitly deform the prior information within the segmentation process. Examples of deformable model-based approaches include active contour methods [99, 60, 2, 46, 102, 113, 62, 85, 114], which evolve an initial shape of the tissue class until it fits the patient's specific situation. These methods typically only segment one tissue class or organ at a time, but extensions to multiple structures have been developed [116, 100]. A variation of this approach is the T-Snake approach, where tissue classes are represented by deformable organisms [68]. Another type of border growing approach is based on a fuzzy clustering technique [106, 101, 77]. These methods use dynamic programming, where the mapping between input and segmentation output is largely influenced by the input data itself. Inaccuracies in the methods' modeling approach are therefore difficult to resolve.

While these methods often focus on the integration of prior information, they only implicitly model the image artifacts of MR images mentioned in Section 1.3.1. This weakness may cause difficulties in initializing the algorithms to larger data sets as the parameter settings of the algorithms have to be adjusted to the individual noise level or intensity pattern of the subject.

Serial Methods

To overcome this problem, serial methods combine several different operations. For example, the segmentation pipeline by Collins et al. [13] first determines the image artifacts, then registers the atlas to the patient, extracts the surfaces from the aligned atlas, and finally uses these aligned surfaces to initialize an artificial neural network. Warfield [108] uses an optimal tissue class boundary estimator in a feature space consisting of multiple signal intensities and spatial context from an anatomical atlas. Fischl et al. [34] first filter out the image artifacts, then erase intensities connected with the skull, determine the gray matter/white matter border on these modified images, split the border into cortical and subcortical structures, apply a topological correction tool to the cortical gray matter, map the corrected boundary on a sphere to register the subject with an atlas, and then further parcellate the cortical surface into its substructures. Serial methods can achieve high quality segmentations but they are generally difficult to adjust to new segmentation problems or image acquisition sequences.

EM Based Methods

Unlike serial methods, unified approaches simultaneously determine the solution to different aspects of a given segmentation problem. Many of these methods first adjust their model to the segmentation problem and then segment the subjects into the structures of interest. As an example Ballester et al. [4] uses the result of Wells EM segmenter [112] to construct a 3D mesh, which detects the outer surface of the structure to be segmented. Grau [39] first aligns an atlas to a patient and then applies a watershed algorithm to the segmentation problem. Unified methods are generally adjustable to different segmentation environments with a minimal amount of effort. However, these methods often approximate the solution to simplified segmentation models, which reduces the robustness of the approach.

This thesis will focus on a unified framework inspired by the EM segmentation approach of Wells [112]. As mentioned above, the method simultaneously estimates

the image inhomogeneities and segments the image into structures of interest. The method has the advantage of being applicable to a wide variety of problems as the intensity correction within the model adjusts to the given data set. To increase the robustness of the method Van Leemput [103] first registers spatial priors to the subject of interest using an affine registration method and then guides the segmentation algorithm with the aligned priors. This process, however, cannot accurately parcellate brain tissues into their substructures due to inaccuracies in the the registration of the spatial priors to the subject. For example, Srivastava et al. [88] have shown that the atlas aligned with an affine registration method cannot properly capture the variability within the brain. These shortcomings increase the risk of systematic errors in the segmentation process, which generally lowers the accuracy of the methods. As we will show in this thesis, these shortcomings can be overcome by more flexible registration methods than affine alignment approaches.

In summary, the field of medical imaging has developed methods that are generally either not robust enough to determine cortical structures or that are targeted towards specific imaging acquisition protocols. This deficiency has restricted the use of automatic segmentation approaches in many neuroscience applications.

1.4 Objective and Contribution of the Thesis

The objective of this thesis is the development of statistical frameworks that are targeted towards automatically and robustly parcellating the major tissue classes into their substructures. The design of the models are guided by the overall philosophy that they should be applicable to a wide variety of problems with a minimal amount of training effort. We are specifically interested in robustly segmenting anatomical structures with weakly visible boundaries in MR images to increase the use of automatic segmentation methods in the field of neuroscience.

The approaches of this thesis rely on prior information to accurately segment these structures. While prior information simplifies the automatic detection of many anatomical structures, it might also negatively bias the results. For example, the

atlas might favor uncommon characteristics in an anatomical structure, which will mislead the segmentation approach. In order to create robust unified segmentation models, this thesis attempts to answer the following two questions:

- What type of prior information is useful for the automatic segmentation of anatomical structures?
- How should the information be introduced into the framework?

Our answers to these questions are embedded in the novel automatic segmentation models developed throughout this thesis.

The major contributions of this thesis with respect to unified segmentation methods are summarized as:

- combining registration of prior information with segmentation of MR images,
- modeling explicitly the boundary of anatomical structures through a shape atlas,
- defining a tree to represent the hierarchical relationship between structures.

We prove the robustness of the developed models by applying them to an EM approach originally proposed by Wells et al. [112]. We are aware, that a variety of extensions of Wells' EM approach already incorporate prior information [51, 103, 66]. However, none of these methods have been applied to the substructures of the major brain tissue classes. We show that our new methods reliably segment these substructures in multiple experiments. In these experiments, we compare the results of our new approach to segmentations produced by other methods.

The remainder of this section explains in further detail the major contributions of this thesis.

1.4.1 Combine Registration and Segmentation

When using prior information in a segmentation algorithm, the alignment of the atlas to the MR images is very important. Any incoherence between the two coordinate

systems reduces the effectiveness of the prior information in assisting the segmentation process. We therefore develop two very different alignment strategies. We first suggest a pipeline approach, which uses a non-rigid registration method to align the spatial priors to the subject and then applies an EM segmenter to the images. We demonstrate the accuracy of our algorithm by identifying in multiple subjects the thalamus and the superior temporal gyrus, both substructures of the gray matter with indistinct boundaries. The experiment, however, also demonstrates the extreme sensitivity of the implementation to the initial alignment errors of the registration method.

In order to increase the robustness of unified segmentation methods, we propose a novel statistical framework, which couples the registration of the atlas with the segmentation of the MR images. We determine a solution within this framework through an instance of an EM algorithm. The method simultaneously updates both registration and segmentation parameters at each iteration. We test this approach by segmenting 22 subjects into the thalamus and the caudate. We illustrate its superior performance by comparing the accuracy of the results to the outcome of other EM methods.

1.4.2 Model Shape Constraints in Statistical Framework

Spatial priors simplify the segmentation of many anatomical structures but they often fail to properly capture the variability of these structures. For example, the caudate is defined by two “horns” whose subject specific bending is difficult to represent with spatial priors. In addition, these priors are generally characterized by slowly varying values indicating the variability in location of the anatomical structure. The impact of these priors is therefore limited for the automatic segmentation of structures with weakly visible boundaries as the boundaries are neither identified by the intensity patterns nor the prior information. An alternative to spatial priors are shape models, which analyzes the boundary of a structure explicitly by comparing it to the expected shape and standard deformations of the anatomical structure.

For our approach, we choose a shape model based on signed distance maps. We

integrate the atlas into a unified statistical framework that tightly couples the shape analysis of a structure with the segmentation of MR images. Again, we determine the solution within this framework using the EM approach that simultaneously deforms the shape structure and segments the image. Unlike most methods in the field, the final segmentation is only indirectly influenced by the shape constraints. This is generally advantageous as the training data normally is too small to capture all shape variations within a population. Our model can therefore handle subject specific abnormalities that are not represented in the training data.

We further extend this shape based segmentation method with the registration of the shape atlas to the image space. The resulting algorithm simultaneously aligns the atlas to the image space, detects the shape of anatomical structure, segments the images into the structures of interests, and estimates the image inhomogeneities.

1.4.3 Represent Hierarchical Relationships in a Tree

As mentioned above, the effectiveness of prior information greatly depends on the anatomical structures to which it is applied. For example, some anatomical structures are defined by indistinct boundaries in which case prior information is essential for robustly outlining these structures. Other boundaries are clearly visible on MR images, so that the segmentation process might ignore them by relying too much on the prior information. To our knowledge, previous segmentation frameworks do not explicitly model these circumstances.

To overcome this deficiency we develop a hierarchical model, which defines anatomical dependencies between structures. These dependencies are expressed in a data tree, which guides the segmentation process. Guided by the data tree the algorithm partitions the segmentation problem into less difficult subproblems. At each subproblem, the prior model is redefined based on specific requirements of the subproblem as well as the attached subtree. Unlike other hierarchical approaches [65, 34, 74] the atlas space defines the dependency between anatomical structures so that the approach is adjustable to a variety of segmentation problems with a minimal amount of training effort. To show the method’s versatility and accuracy, we apply the approach to

MR images with different signal to noise ratios and we outline different anatomical structures.

1.5 Overview of Validation Techniques

Validating segmentation methods for medical images is very difficult due to the absence of ground-truth. Instead, the field of medical image analysis has developed phantoms [58] where the ground-truth is known or compare the results to a reference standard [110], or the approach is applied to multiple scans of the same subject [14]. As mentioned, the goal of this thesis is to increase the reliability of automatic segmentation methods so that they can be applied to a wider variety of neuroscience studies. Many of these studies currently rely on manual segmentations to find disease-related characteristics. Therefore, we compare the automatic generated results to manually segmentations. When multiple manual segmentations of the same subject are present we compare our method to a reference standard produced by STAPLE[110]. This reference standard represents the highest consensus between the manual segmentations. The remainder of this section focuses on metrics measuring the accuracy of the automatic segmentations with respect to the reference standard.

While the literature proposes a wide variety of metrics for this purposes [118] we apply the volume overlap measure DICE [63] and the Positive Predictive Value (PPV) [31]. DICE measures the volume overlap between two segmentations by normalizing the volume of intersection between the two segmentations with the sum of the volumes of both segmentations. The measure ranges from 0, indicating no similarity between two segmentations, to 1, indicating complete agreement.

Unlike DICE, PPV considers the foreground, outlined by the segmentation, as well as the background. It measures the agreement of both areas with respect to a reference standard by normalizing the sensitivity by both specificity and sensitivity. In the medical literature, “sensitivity” is defined as the proportion of the image correctly outlining the foreground, while “specificity” is the proportion of the image correctly outlining the background. PPV is zero if the foreground and background do not

overlap in both segmentations and one if they are equivalent.

In summary, we show the accuracy of the approaches developed throughout this thesis by comparing the resulting automatic segmentations to a reference standard. When multiple segmentations of the same subject are present we use STAPLE to generate a reference standard. Otherwise, the reference standard is represented by the only manual segmentation present. The agreement between automatic segmentation and reference standard is measured by different matrices that deviate in the aspects they capture.

1.6 Organization of the Thesis

Chapter 2 introduces the EM algorithm by deriving the method from a lower bounding viewpoint. In this light, EM is interpreted as a lower bound optimization algorithm, which tightens the lower bound to an objective function at each iteration. We then rederive and compare two EM approaches targeted towards medical imaging segmentation [112, 103].

The following three chapters analyze the use of different prior information in an EM approach. Chapter 3 discusses a variety of alignment methods that register an atlas to the image coordinate system. Chapter 4 focuses on the integration of shape constraints into the EM framework. Chapter 5 develops an anatomically guided segmentation approach that incorporates allows the EM approaches of the previous chapters. The resulting algorithm simultaneously registers the atlas to the image space, detects the shape of an anatomical structure, models the image inhomogeneities in the MR images, and is explicitly guided by the hierarchical relationship of anatomical structures. The final chapter, Chapter 6, summarizes the results and discusses future research directions resulting from this thesis.

Chapter 2

The Expectation-Maximization Algorithm

“The difficulties associated with the use of the EM algorithm may have been the cause for its considerable popularity” Stoica and Selen [89]

The automatic segmentation of medical images separates an image into different anatomical structures based on image and training data. This is a challenging problem because the images are corrupted by image artifacts such as image inhomogeneities and noise. The accuracy of the segmentation depends on properly modeling the artifacts as well as correctly incorporating the training data. In this thesis, we address this issue as a parameter estimation problem with respect to an incomplete data model.

In general, problems are put in the context of an incomplete data model if no accurate analytical solution can be determined. In this section, the incomplete data model describes the complicated relationship between the parameters representing the image inhomogeneities and the observed medical image data. This relationship is greatly simplified if the underlying ground-truth segmentation is known. Throughout this thesis, however, we are interested in problems where the ground-truth segmentation is unknown. Instead, we determine the solution to the corresponding estimation problem using an instance of the Expectation-Maximization (EM) algorithm [25].

At each iteration, the algorithm calculates the most likely parameter setting of the image inhomogeneity based on the current segmentation of the observed image and then updates the segmentation based on the current parameter setting.

This chapter first gives a general overview of the EM algorithm to motivate its use for the automatic segmentation of medical images. Next, we adopt the algorithm to a multiple clustering example that closely relates to our original segmentation problem. We then derive two different instances of the EM algorithm specifically target towards medical image segmentation. The first approach, originally proposed by Van Leemput et al.[103], solves a maximum likelihood problem. The second approach, suggested by Wells [112], determines the solution to a maximum *a posteriori* probability estimation problem.

2.1 Background

This section revisits the most important terminology necessary to understand the EM algorithm in detail. We assume the reader is familiar with basic concepts in statistics, which are also reviewed in [5].

As mentioned, the EM-algorithm estimates the optimal parameter setting within an incomplete data model. The following definition further formalizes this concept:

Definition 1 (Incomplete Data Model)

*A data model is incomplete if the relationship between the **observed data** \mathbf{O} and the **parameter space** Ψ is too complicated to be described in a statistical framework. In this case, the relationship between \mathbf{O} and Ψ is represented by **incomplete log likelihood** $L(\Psi) \triangleq \log p(\mathbf{O}|\Psi)$. This relationship is simplified by adding the **hidden data** \mathbf{H} to the model. The relationship of \mathbf{O} , \mathbf{H} , and Ψ is now defined by the **complete log likelihood** $L_c(\Psi) \triangleq \log p(\mathbf{O}, \mathbf{H}|\Psi)$.*

Example:

The segmentation of medical images can be modeled as an incomplete data problem, where the medical images are the observed data, the hidden data is the true segmentation, and the parameter space captures the image inhomogeneities. Many other

modeling approaches for the same problem are possible.

The EM algorithm is generally applied to incomplete data problems that can be solved iteratively by calculating the maximum to a closely related but simpler lower bound. This lower bound is derived from the expected value over the hidden data with respect to the current parameter estimate and the observed data. We will define the expected value explicitly as its notation varies in the literature.

Definition 2 (Expected Value)

The expected value with respect to the distribution $p(z)$ of the random vector z is defined by $E_z(f(z)) \triangleq \int_z p(z) \cdot f(z) dz$ where $f(\cdot)$ is a function defined in the random vector space of z . In discrete space, the expected value is defined as

$E_z(f(z)) \triangleq \sum_z P(z) \cdot f(z)$. Furthermore, the expected value with respect to the distribution of z conditioned on random vector or parameter y is defined as

$$E_{z|y}(f(z)) \triangleq \int_z p(z|y) \cdot f(z).$$

Example:

The expected value of the log likelihood with respect to the hidden variables is defined as

$$E_{H|O,\Psi}(L_c(\Psi)) = \int_H L_c(\Psi) \cdot p(x|O, \Psi) dx = \int_H \log(p(O, x|\Psi)) \cdot p(x|O, \Psi) dx$$

The definition of the simpler lower bound is based on Jensen’s Inequality with respect to the expected value of the hidden data. We now show the proof of Jensen’s Inequality for the expected value of concave functions. However, the inequality is applicable to a larger class of functions.

Definition 3 (Concave)

Let $f : (a, b) \rightarrow \mathbb{R}$ be a scalar function, for which the derivative f' exists in the open interval (a,b) . f is concave iff

$$\forall x, y \in (a, b) : f(y) \leq f(x) + f'(x)(y - x)$$

Lemma 4 (Jensen's Inequality)

For any concave function $h : \mathbb{R} \rightarrow \mathbb{R}$ over the random vector x , the following inequality holds: $E_x(h(x)) \leq h(E_x(x))$.

Proof:

The proof simply re-applies the definition of concavity to expected values. If we define $\tilde{x} \triangleq E(x)$ and $\tilde{y} \triangleq x$ then an upper bound for the concave function $h(\cdot)$ at \tilde{y} is defined according to Definition 3 as

$$h(\tilde{y}) \leq h(\tilde{x}) + h'(\tilde{x}) \cdot (\tilde{y} - \tilde{x}).$$

\tilde{x} and \tilde{y} have the same expected value, since $E(\tilde{x}) = E(E(x)) = \tilde{x}$ and $E(\tilde{y}) = E(x) = \tilde{x}$. Taking the expected value on both sides of the inequality results in the Jensen inequality

$$\begin{aligned} E(h(x)) = E[h(\tilde{y})] &\leq E[h(\tilde{x}) + h'(\tilde{x}) \cdot (\tilde{y} - \tilde{x})] \\ &= h(\tilde{x}) + h'(\tilde{x}) \cdot (E(\tilde{y}) - \tilde{x}) = h(E(x)) \blacksquare \end{aligned}$$

We end this overview of terminology with a notation simplifying the understanding of the mathematical derivations in this thesis.

Definition 5 (Extended Explanation of Derivation)

In the remainder of this thesis ' $\stackrel{x}{=}$ ' refers to footnote x for further explanation.

2.2 Deriving the EM Algorithm

This section is an introduction to the EM algorithm. First, it describes the EM algorithm as a lower bound optimization method. Then, it proves the convergence properties of the method. We end the section with a summary of this topic.

2.2.1 EM: A Lower Bound Optimization Method

Over the years, many different interpretations of the EM algorithm have been presented in the literature, e.g. [43, 71, 73, 89, 69]. This section views the algorithms as a lower bound optimization technique originally described by [64, 24]. We choose this interpretation as it nicely displays the convergence properties of the method. From the lower bound viewpoint, the EM algorithm estimates the solution to the maximum likelihood problem

$$\hat{\Psi} = \arg \max_{\Psi} L(\Psi), \quad (2.1)$$

by defining a lower bound for the incomplete log likelihood $L(\cdot)$ as displayed in Figure 2-1. In the Expectation-Step (E-Step), the method defines this lower bound based on the current estimate Ψ' of the optimal solution $\hat{\Psi}$. The Maximization-Step (M-Step) calculates the maximum of the lower bound to update the estimate Ψ' . The algorithm iterates between the E-Step and M-Step until the series of estimates Ψ' converges.

The method is in general applied to problems where no analytical solution to Equation (2.1) can be found. However, we can restate the problem by introducing the hidden data H and defining the parameter estimation as

$$\hat{\Psi} = \arg \max_{\Psi} \log \left(\sum_H p(O, H|\Psi) \right). \quad (2.2)$$

Note, that the complete likelihood $p(O, H|\Psi)$ of Equation (2.2) is the marginal of the incomplete likelihood $p(O|\Psi)$ of Equation (2.1).

Instead of solving Equation (2.2), the EM algorithm finds a maximum to the simpler objective function

$$Q(\Psi, \Psi') \triangleq E_{H|O, \Psi'} \left(\log \frac{p(O, H|\Psi)}{p(H|O, \Psi')} \right) \text{ with } p(H|O, \Psi') > 0 \quad (2.3)$$

where Ψ' is an approximation of the optimal solution $\hat{\Psi}$. As we show later, the objective function $Q(\cdot, \Psi')$ has two nice properties with respect to the original problem

defined in Equation (2.1). First, it defines a lower bound on the incomplete log likelihood $L(\cdot)$ touching $L(\cdot)$ at least in one point. Second, the maximum of the objective function $Q(\cdot, \Psi')$ is an improvement over the initial estimate Ψ' with respect to maximum $\widehat{\Psi}$. Based on these properties of the objective function, the EM method improves its estimate Ψ' by repeating the following two steps:

Definition 6 (Expectation-Maximization Algorithm)

Expectation Step (E-Step): Calculate the lower bound $Q(\cdot, \Psi')$ of $L(\cdot)$ based on the current estimate Ψ' of the optimal parameter setting $\widehat{\Psi}$:

$$Q(\Psi, \Psi') \triangleq E_{H|O, \Psi'} \left(\log \frac{p(O, H|\Psi)}{p(H|O, \Psi')} \right)$$

Maximization Step (M-Step): Update the estimate Ψ' by the maximum of the current lower bound

$$\Psi' \leftarrow \arg \max_{\Psi} Q(\Psi, \Psi').$$

The EM algorithm iterates between the E- and M-Steps until the lower bound function converges to a local extrema, which is guaranteed by the EM framework if the iteration sequence has an upper bound [69].

Before deriving important properties of the lower bound we note that the objective function $Q(\cdot, \Psi')$ of the EM algorithm is generally defined as

$$Q(\Psi, \Psi') = E_{H|O, \Psi'} (\log p(O, H|\Psi)) = E_{H|O, \Psi'} (L_c(\Psi)), \quad (2.4)$$

which is the expected value of the complete log likelihood. The lower bound defined in Equation (2.3) and the objective function of Equation (2.4) have the same maximum, because the denominator $p(H|O, \Psi')$ in Equation (2.3) is constant with respect to Ψ . Thus, both objective functions define the same algorithm.

2.2.2 Properties of the Lower Bound

This section focuses on the convergence properties of the algorithm. In particular, we will show that the EM algorithm improves the approximation Ψ' at each iteration unless $Q(\Psi', \Psi')$ is the maximum of the bound $Q(\cdot, \Psi')$ (see also Figure 2-1.) In other words, at each iteration the EM-algorithm tightens the bound with respect to the log likelihood $L(\cdot)$ unless the current estimate defines a local extremum.

In comparison, gradient descent methods find a local optimum by first approximating the objective function with a polynomial function. As an example, Figure 2-1 shows a linear approximation at point Ψ' . Gradient descent methods take a step along this polynomial function. The proper step size is unknown because of the uncertainty of the quality of the approximated function. Gradient descent methods therefore do not guarantee an improvement at each iteration.




Figure 2-1: The lower bound defined by $Q(\cdot, \Psi')$ touches the incomplete likelihood at least at the approximation Ψ' . The new guess is the maximum of the current lower bound, which is at least as good as the previous estimation. Unlike gradient descent methods, the EM algorithm improves its approximation at each iteration, unless the current approximate already is a local maximum. The graph was originally motivated by [71].

To prove the convergence property of the EM algorithm we first show that $Q(\cdot, \Psi')$

is a lower bound of $L(\cdot)$. In the second step, we prove that $Q(\cdot, \Psi')$ touches $L(\cdot)$ at least at one point, which is the parameter setting Ψ' . From these two properties it follows that the log likelihood $L(\Psi'')$ of the maximum Ψ'' of $Q(\cdot, \Psi')$ is greater than or equal to $L(\Psi')$ because

$$L(\Psi'') \stackrel{1}{\geq} Q(\Psi'', \Psi') \stackrel{2}{\geq} Q(\Psi', \Psi') \stackrel{3}{=} L(\Psi')$$

The next two corollaries prove the previously mentioned properties:

Corollary 7 ($Q(\cdot, \Psi)$ is a Lower Bound of $L(\cdot)$)

$$\forall \Psi', \Psi : L(\Psi) \geq Q(\Psi, \Psi')$$

Proof:

The proof marginalizes over the hidden data H and applies Jensen's Inequality

$$\begin{aligned} L(\Psi) &= \log p(O|\Psi) \\ &\stackrel{4}{=} \log \left(\int p(H, O|\Psi) dH \right) \\ &= \log \left(\int \frac{p(H, O|\Psi)}{p(H|O, \Psi')} p(H|O, \Psi') dH \right) \\ &\stackrel{5}{=} \log \left(E_{H|O, \Psi'} \left[\frac{p(H, O|\Psi)}{p(H|O, \Psi')} \right] \right) \\ &\stackrel{6}{\geq} E_{H|O, \Psi'} \left(\log \frac{p(H, O|\Psi)}{p(H|O, \Psi')} \right) = Q(\Psi, \Psi') \blacksquare \end{aligned}$$

Corollary 7 implies that the maximum of the lower bound determined in the M-Step is equal or suboptimal to the solution of Equation (2.1). The next corollary proves that

¹ $Q(\cdot, \Psi')$ is a lower bound of $L(\cdot)$

² Ψ'' is the maximum of $Q(\cdot, \Psi')$

³ $Q(\cdot, \Psi') = L(\cdot)$ at point Ψ'

⁴Marginalizing over H

⁵Apply Definition 2 of the expected value

⁶Apply Jensen Inequality defined in Lemma 4

at least one point of the lower bound touches the log likelihood $L(\cdot)$; an important quality of this bound.

Corollary 8 ($Q(\cdot, \Psi')$ touches $L(\cdot)$)

$$\forall \Psi' \exists \Psi : Q(\Psi, \Psi') = L(\Psi)$$

Proof:

The proof simply uses Bayes' rule

$$\begin{aligned} Q(\Psi, \Psi') &\stackrel{7}{=} -E_{H|O, \Psi'} \left(\log \frac{p(H|O, \Psi) \cdot p(O|\Psi)}{p(H|O, \Psi')} \right) \\ &= -E_{H|O, \Psi'} \left(\log \frac{p(H|O, \Psi)}{p(H|O, \Psi')} \right) + \log p(O|\Psi) \\ \Rightarrow Q(\Psi', \Psi') &= -E_{H|O, \Psi'} \left(\log \frac{p(H|O, \Psi')}{p(H|O, \Psi')} \right) + \log p(O|\Psi') \\ &\stackrel{8}{=} \log p(O|\Psi') = L(\Psi') \blacksquare \end{aligned}$$

The previous corollary is very important for the understanding of the EM algorithm as it proves that the EM algorithm converges to a local extrema of the objective function $L(\cdot)$. In order to understand the connection between the corollary and the convergence behavior of the algorithm, let $\hat{\Psi}$ be the point to which algorithm converges. By the definition of EM formulation, $\hat{\Psi}$ has to be the global maximum of $Q(\cdot, \hat{\Psi})$ so that the derivative $\frac{\partial Q(\Psi, \hat{\Psi})}{\partial \Psi} \Big|_{\Psi=\hat{\Psi}}$ is zero. Based on the previous proof we know that $Q(\cdot, \hat{\Psi})$ touches $L(\cdot)$ in point $\hat{\Psi}$, which implies that the derivative $\frac{\partial Q(\Psi, \hat{\Psi})}{\partial \Psi} \Big|_{\Psi=\hat{\Psi}}$ is equivalent to $\frac{\partial L(\Psi)}{\partial \Psi} \Big|_{\Psi=\hat{\Psi}}$. We conclude that $\hat{\Psi}$ is a local extremum of the original objective function $L(\cdot)$.

⁷Bayes Rule: $p(A, B|C) = p(A|B, C) \cdot p(B|C)$

⁸ $\log \frac{p(H|O, \Psi')}{p(H|O, \Psi')} = 0$

2.2.3 Summary

In summary, the EM algorithm finds the maximum likelihood solution to the incomplete data problem by transforming it into a problem of finding the optimal lower bound in the complete data model. As we show in Section 2.2.2, the method improves its bound at each iteration unless the parameter setting of the bound already defines a local maxima.

Note, that the right choice of the hidden data H is critical for the adoption of the EM algorithm to a specific problem. This task is in general not trivial as the hidden data should simplify the initial maximization problem so that the maximum of the expected value with respect to H can be found easily. Motivated by this challenge, the next section will derive the EM algorithm for a class of problems, for which the implementation is quite simple.

2.3 Examples of EM-Approaches

This section first derives a solution for a generic example closely related to the problem of segmenting medical images. We will then apply the concept to two different models of the segmentation of MR images.

2.3.1 Applying EM to a Multiple Clustering Example

The complete model is defined by n independent samples $O = \{o_1, \dots, o_n\}$. Each sample o_i is drawn from a group of l different clusters. The assignment of sample space O to the clusters is unknown. In this model, the unknown assignment is represented by the independent hidden parameters $H = (h_1, \dots, h_n)$. Each hidden data h_x is an indicator random variable, so that $h_x \in \{e_1, \dots, e_l\}$. The vector e_j is zero at every position except j , where its value is one. As an example, if $h_x = e_j$, then sample x is assigned to cluster j . An important aspect of the model is the relationship between the parameter space Ψ , the hidden data H , and the observed data O . If we assume independence of the parameter space with respect to the hidden data then

the relationship is encoded in the complete likelihood

$$P(H, O|\Psi) = P(H|\Psi) \cdot P(O|H, \Psi) \stackrel{9}{=} \prod_{x=1}^n P(h_x) \cdot P(o_x|h_x, \Psi)$$

We define the optimal parameter setting Ψ with respect to this model as the solution to a maximum likelihood problem with the objective function of Equation (2.4). We therefore focus first on finding a simpler expression for the complete data log likelihood $L_c(\cdot)$, which is defined as

$$L_c(\Psi) = \log P(H, O|\Psi) = \sum_{x=1}^n \log P(h_x) + \sum_{x=1}^n \log P(o_x|h_x, \Psi) \quad (2.5)$$

Equation (2.5) can be expressed in terms of vector functions, because h_x is an indicator random variable. To do so, we define the vector over the spatial distribution of each cluster

$$V_x \triangleq (P(h_x = e_j))_{j=1, \dots, l} \quad (2.6)$$

and the vector function over the likelihood of the parameter space

$$U_x(\Psi) \triangleq (P(o_x|h_x = e_j, \Psi))_{j=1, \dots, l}. \quad (2.7)$$

The complete log likelihood of Equation (2.5) is defined as

$$L_c(\Psi) = \sum_{x=1}^n h_x^t \log V_x + \sum_{x=1}^n h_x^t \log U_x(\Psi).$$

⁹The group of samples O are drawn independently. In addition, we assume that $P(O|H, \Psi) = \prod_{x=1}^n P(o_x|h_x, \Psi)$.

Based on Equation (2.4) the objective function of the complete model is

$$\begin{aligned}
Q(\Psi, \Psi') &= E_{H|O, \Psi'}(L_c(\Psi)) = \sum_{x=1}^n E_{h_x|o_x, \Psi'}(h_x^t \log V_x) + E_{h_x|o_x, \Psi'}(h_x^t \log U_x(\Psi)) \\
&= \sum_{x=1}^n E_{h_x|o_x, \Psi'}(h_x^t) \cdot \log V_x + \sum_{x=1}^n E_{h_x|o_x, \Psi'}(h_x^t) \cdot \log U_x(\Psi).
\end{aligned} \tag{2.8}$$

To simplify the objective function we replace the expected value by vector \mathcal{W}_x , which is defined as

$$\begin{aligned}
\mathcal{W}_x &\triangleq E_{h_x|o_x, \Psi'}(h_x) \\
&= (0 \cdot P(h_x \neq e_j|o_x, \Psi') + 1 \cdot P(h_x = e_j|o_x, \Psi'))_{j=1, \dots, l} \\
&= (P(h_x = e_j|o_x, \Psi'))_{j=1, \dots, l} \\
&\stackrel{10}{=} \left(\frac{P(o_x|h_x = e_j, \Psi')P(h_x = e_j|\Psi')}{P(O|\Psi')} \right)_{j=1, \dots, l}
\end{aligned} \tag{2.9}$$

\mathcal{W}_x defines the cluster specific posterior probability of the hidden data h_x . In a machine learning context, $\mathcal{W}_x(a) \triangleq E_{h_x|o_x, \Psi'}(h_x(a))$ captures the weight that the sample o_x is assigned to the cluster a . In the remainder of this thesis we therefore call \mathcal{W} the *weights*. Based on the previous independence assumptions, the weights \mathcal{W} transform into

$$\begin{aligned}
\mathcal{W}_x &= \left(\frac{P(o_x|h_x = e_j, \Psi')P(h_x = e_j)}{P(O|\Psi')} \right)_{j=1, \dots, l} \\
&\stackrel{11}{=} \left(\frac{P(o_x|h_x = e_j, \Psi') \cdot P(h_x = e_j)}{\sum_{k=1}^l P(o_x|h_x = e_k, \Psi') \cdot P(h_x = e_k|\Psi')} \right)_{j=1, \dots, l} \\
&\stackrel{12}{=} \left(\frac{P(o_x|h_x = e_j, \Psi') \cdot P(h_x = e_j)}{\sum_{k=1}^l P(o_x|h_x = e_k, \Psi') \cdot P(h_x = e_k)} \right)_{j=1, \dots, l} \\
&= \left(\frac{U_x(\Psi', j) \cdot V_x(j)}{U_x(\Psi')^t V_x} \right)_{j=1, \dots, l}
\end{aligned} \tag{2.10}$$

Substituting the expected value in Equation (2.8) with the weights \mathcal{W}_x changes the

¹⁰Bayes Rule

objective function to

$$Q(\Psi, \Psi') = \sum_{x=1}^n \mathcal{W}_x^t \log V_x + \sum_{i=1}^n \mathcal{W}_x^t \log U_x(\Psi) \quad (2.11)$$

As mentioned, at each iteration the EM algorithm finds the parameter setting Ψ'' , which maximizes the objective function $Q(\cdot, \Psi')$ given the current estimates Ψ' . In order for Ψ'' to be a maximum, the derivative of the objective function at point $\Psi = \Psi''$ has to be zero:

$$\frac{\partial Q(\Psi, \Psi')}{\partial \Psi} = \frac{\partial}{\partial \Psi} \sum_{x=1}^n \mathcal{W}_x^t \log U_x(\Psi) = \sum_{x=1}^n \mathcal{W}_x^t \frac{\partial \log U_x(\Psi)}{\partial \Psi} = 0. \quad (2.12)$$

For some underlying models, Equation (2.12) can be turned into a system of linear equations. The exact form of the system of linear equations depends on the definition of $U_x(\Psi)$. In general, if the complete-data probability belongs to the class of the regular exponential family, the solution to the system is easily determined [69].

An example of such a model is one where the conditional probability has the Gaussian distribution $P(o_x | h_x = e_j, \Psi) \sim \mathcal{N}(\mu_x(e_j, \Psi), \Upsilon_j)$. The mean $\mu_x(\cdot, \cdot)$ depends on the complete data model, and the variance Υ_j on the hidden data $h_x = e_j$. If $\mu_x(e_j, \Psi)$ is differentiable with respect to Ψ then Equation (2.12) changes for this example to

$$\sum_{x=1}^n \mathcal{W}_x^t \left((o_x - \mu_x(e_j, \Psi))^t \Upsilon_j^{-1} \frac{\partial \mu_x(e_j, \Psi)}{\partial \Psi} \right)_{j=1, \dots, l} = 0 \quad (2.13)$$

and the EM algorithm is defined as

E-Step: Calculate the Weights \mathcal{W} based on Equation (2.10).

M-Step: Update the approximation Ψ' by the solution of Equation (2.13) with respect to Ψ .

Note, that the above definition of the EM algorithm never explicitly determines the lower bound or objective function $Q(\cdot, \cdot)$ of Equation (2.8). However, the function

¹¹Bayes Rule

¹²Independence assumption between H and Ψ

is fully defined by the weights \mathcal{W} as all other parameters are fixed. As EM is only interested in the maximum of the objective function an explicit definition of $Q(\cdot, \cdot)$ is unnecessary.

In the next two subsections, we apply the previous approach to the segmentation problem of MR images. This results in two slightly different methods as each subsection makes different modeling assumptions.

2.3.2 Image Inhomogeneity Correction of MR Images as Maximum Likelihood Problem

The model derived in this section is closely related to the work by Van Leemput [103]. The approach simultaneously approximates the image inhomogeneities and generates soft assignment of the voxels to the structure of interests, which was originally suggested by Wells [112]. When the algorithm converges, the soft assignments are turned into labelmaps, which segment the anatomical regions on the MR images.

To apply the method derived in the last section to the segmentation problem of MR images, the incomplete data model has to be defined. Similar to the example for Definition 1 of Section 2.1 the observed data O represents the log intensities of the MR images, where o_x is the log intensity of voxel x . The remainder of this thesis will only consider log intensities of MR images and will simply refer to them as *intensities*.

The parameter space Ψ captures the image inhomogeneities in the image. The image inhomogeneities are caused by the radio frequency coil of the MR scanner and are slowly varying (Figure 2-2 (b).) In Van Leemput's model, the slowly varying values are represented by the spline function

$$\nu_x(\Psi) \triangleq \sum_{i=1}^k f_i(o_x) \cdot \Psi(i) \quad (2.14)$$

At each voxel x , the image inhomogeneity is defined by a linear combination of the polynomial functions $f_j(\cdot)$ and the weight $\Psi(i)$. The dimension of Ψ is equal to the number of smoothing functions f_i .

The optimal parameter setting $\hat{\Psi}$ of the image inhomogeneities with respect to the image data O is defined in Van Leemput’s model by the most likely setting of the incomplete log likelihood

$$\hat{\Psi} \triangleq \arg \max_{\Psi} L(\Psi) = \arg \max_{\Psi} \log P(O|\Psi) \quad (2.15)$$

Equation (2.15) defines an incomplete problem as the relationship between observed data O and parameter space Ψ generally too complex to allow an analytical solution.




Figure 2-2: Image (a) shows an MR image corrupted by image inhomogeneities, noise, partial voluming and other image artifacts. The image inhomogeneities of (a) are shown in (b). Unlike noise, image inhomogeneities are characterized by a slowly varying values within the brain. Image (c) is as image inhomogeneity corrected MR image of (a).

Finding the solution to Equation (2.15) is simplified if we add the labelmap of the image to the model. The labelmap assigns each voxel x to an anatomical region inside the brain. Combining the intensity pattern of each structure of interest with the labelmap of the unknown true segmentation defines the ideal MR image without image artifacts. The pattern of image inhomogeneities, which is a multiplicative field to the MR images, is then defined by the difference between the log intensities of the real and ideal MR images. This definition oversimplifies the problem as MR images are corrupted by noise, partial voluming effects and other image artifacts. Nevertheless, Van Leemput indirectly accounts for these artifacts as Equation (2.14) enforces the smoothness constraints of the inhomogeneities.

From now on, our complete model is defined by the MR images O , image inhomo-

genieties Ψ , and the hidden data H , which capture the labelmap of the image. The indicator random variable h_x represent the assignment of voxel x to an anatomical structure, such as gray matter or corticospinal fluid. To keep the concept general, we refer to an anatomical structure by a number instead of explicitly mentioning the name. For example, if we would like to segment the image O into l anatomical structures and $h_x = e_a$ then voxel x is assigned to the anatomical structure a , where $a \in \{1, \dots, l\}$.

Having defined the parameters and data, we focus on the relationship between these components of the complete model. In our framework, their relationship is characterized by the following maximum likelihood problem:

$$\hat{\Psi} \triangleq \arg \max_{\Psi} L_c(\Psi) = \arg \max_{\Psi} \log P(O, H | \Psi)$$

where the complete log likelihood $L_c(\Psi)$ is equivalent to Equation (2.5).

According to Section 2.3.1 the solution to the previous maximum likelihood estimation problem can be computed by an instance of the EM algorithm. Vector V_x of Equation (2.6) is defined by the structure specific spatial prior $P(h_x = e_a)$. This is the probability of structure a being present at voxel x . The vector U_x of Equation (2.7) represents the likelihood of the image inhomogeneity, which is defined by $P(o_x | h_x = e_a, \Psi) \sim \mathcal{N}(\mu_a + \nu_x(\Psi), \Upsilon_a)$. The mean μ_a and variance Υ_a define the intensity pattern of the ideal MR images within the region of the anatomical structure a . These parameters characterize the Gaussian distribution, which best captures the structure specific intensity patterns in the training data.

In the last section, the E-Step calculates the weights \mathcal{W}_x , which captures the posterior probability $P(h_x = e_a | o_x, \Psi')$. Based on the current estimate of the incomplete model defined by O and Ψ' , $\mathcal{W}_x(j)$ represents the posterior probability that structure j is present at voxel x . Thus, the E-Step defines the soft assignment of structure j to voxel x .

To derive the M-Step, we must specify Equation (2.13) to turn it into a system of linear equations. The mean of the distribution is composed of the image inhomogeneity

geneties of Equation (2.14) and the structure specific means, so it has the form

$$\mu_x(e_a, \Psi) \triangleq \mu_a + \nu_x(\Psi) = \mu_a + \sum_{i=1}^k f_i(o_x) \cdot \Psi(i).$$

The derivative of the mean is $\frac{\partial \mu_x(e_a, \Psi)}{\partial \Psi(i)} = f_i(o_x)$. If we define

$f(o_x) \triangleq (f_i(o_x))_{i=1, \dots, k}$, then Equation (2.13) turns into system of linear equations of the form

$$\begin{aligned} 0 &= \sum_{x=1}^n \mathcal{W}_x^t [(o_x - \mu_a - f(o_x)^t \cdot \Psi)^t \Upsilon_a^{-1} f_i(o_x)]_{a=1, \dots, l} \\ &= \sum_{x=1}^n \left(\mathcal{W}_x^t [(o_x - \mu_a)^t \Upsilon_a^{-1} - \Psi^t f(o_x) \Upsilon_a^{-1}]_{a=1, \dots, l} \right) f_i(o_x) \end{aligned}$$

with $i = 1, \dots, k$.

Because Υ_a is symmetric, the equation can be rearranged, so that for $i = 1, \dots, k$:

$$\sum_{x=1}^n f_i(o_x)^t \left(\sum_{a=1}^l \mathcal{W}_x(a) \Upsilon_a^{-1} \right) f(o_x)^t \cdot \Psi = \sum_{x=1}^n f_i(o_x)^t \left(\sum_{a=1}^l \mathcal{W}_x(a) \Upsilon_a^{-1} (o_x - \mu_a) \right).$$

If we define the mean covariance at voxel x as $\bar{\Upsilon}(x) \triangleq \left(\sum_{a=1}^l \mathcal{W}_x(a) \Upsilon_a^{-1} \right)$ then the above equation is for $i = 1, \dots, k$ equivalent to

$$\sum_{x=1}^n f_i(o_x)^t \bar{\Upsilon}(x)^{-1} f(o_x)^t \Psi = \sum_{x=1}^n f_i(o_x)^t \bar{\Upsilon}(x)^{-1} \left(o_x - \bar{\Upsilon}(x) \left(\sum_{a=1}^l \mathcal{W}_x(a) \Upsilon_a^{-1} \mu_a \right) \right). \quad (2.16)$$

To further simplify Equation (2.16) we define the derivative of the mean μ over the entire image space as

$$M \triangleq \left(\frac{\partial \mu_x(e_1, \Psi)}{\partial \Psi(i)} \right)_{x=1, \dots, n; i=1, \dots, k} \quad \text{with} \quad \frac{\partial \mu_x(e_1, \Psi)}{\partial \Psi(i)} = f_i(o_x), \quad (2.17)$$

the mean covariance over the entire image space as

$$\bar{\Upsilon} \triangleq \text{diag}(\bar{\Upsilon}(x)), \quad (2.18)$$

and the mean residual as

$$\bar{R} \triangleq \left(o_x - \bar{\Upsilon}(x) \left(\sum_{a=1}^l \mathcal{W}_x(a) \Upsilon_a^{-1} \mu_a \right) \right)_{x=1, \dots, n}. \quad (2.19)$$

Equation (2.16) is equivalent to

$$M^t \bar{\Upsilon}^{-1} M \Psi = M^t \bar{\Upsilon}^{-1} \bar{R}$$

so that the solution for Ψ is defined as

$$\Psi = \left(M^t \bar{\Upsilon}^{-1} M \right)^{-1} M^t \bar{\Upsilon}^{-1} \bar{R}.^{14} \quad (2.20)$$

In summary, this section developed an EM-approach specific to the problem of segmenting MR images. We first defined the complete model and then applied the model to the generic EM approach of Section 2.3.1. This instance of an EM algorithm is now of the form:

Expectation Step (E-Step): Calculate the weights \mathcal{W}_x according to Equation (2.10), which are a soft assignment of voxel x to the structures of interest based on the MR images O and the current setting of the image inhomogeneities Ψ .

Maximization Step (M-Step): Calculate the new parameter setting

$$\Psi \leftarrow \left(M^t \bar{\Upsilon}^{-1} M \right)^{-1} M^t \bar{\Upsilon}^{-1} \bar{R}$$

¹³The diagonal matrix over G_x is defined as $diag(G_x) \triangleq (g_{ij})_{i=1, \dots, n; j=1, \dots, n}$ with

$$g_{ij} \triangleq \begin{cases} G_i & , \text{ if } i = j \\ 0 & , \text{ otherwise} \end{cases}$$

¹⁴For simplicity, the previous definition of M , $\bar{\Upsilon}$, and \bar{R} are given for the case where O captures the log intensities of a single input channel. If, however, the observed data O represents the log intensities of m input channels then the column of M and \bar{R} are of dimension $n \cdot m$ as each voxel is defined by an entry of dimension m . $\bar{\Upsilon}$ is a band matrix of dimension $n \cdot m \times n \cdot m$ with the covariance matrixes $\bar{\Upsilon}_x$ along the diagonal. These changes, however do not impact Equation (2.20).

based on the matrix M defined by the image inhomogeneity smoothing f_i in Equation (2.17), the mean inverse covariance matrix $\overline{\Upsilon}^{-1}$ defined by the weights \mathcal{W} in Equation (2.18), and the mean residual \bar{R} defined by the weights \mathcal{W} in Equation (2.19).

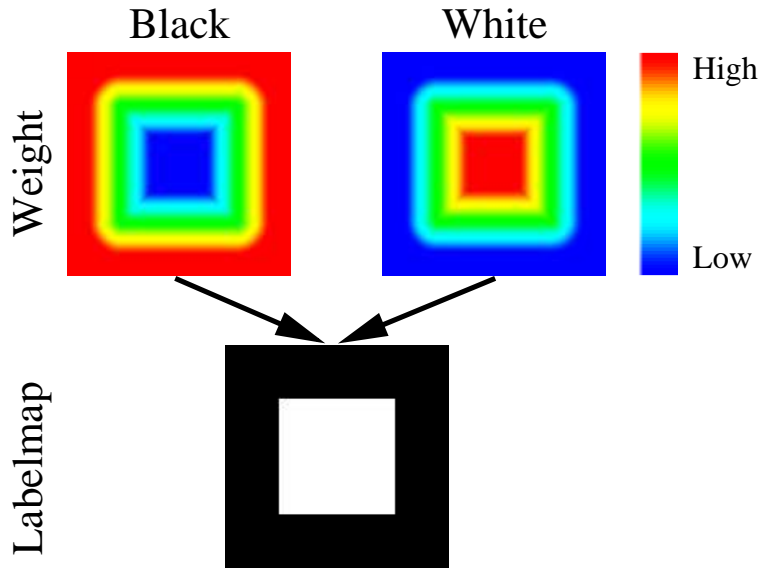


Figure 2-3: The labelmap in the second row is based on the weights of each structure in the first row. The label of each voxel in the labelmap corresponds to the structure (black and white) with the largest weight at that location. Areas in red represent high values of the weights and blue shows low values.

When the algorithm converges, we turn the final weights \mathcal{W}_x from a soft to a hard assignment \mathcal{L} :

$$\mathcal{L}(x) \triangleq \arg \max_a \mathcal{W}_x(a) \quad (2.21)$$

In other words, a voxel x is assigned to the structure j with the largest posterior probability defined by $\mathcal{W}_x(j)$. For example, Figure 2-3 shows the labelmap resulting from the synthetic weights \mathcal{W} of 'Black' and 'White'.

2.3.3 Image Inhomogeneity Correction of MR Images as a Maximum *a Posteriori* Probability Estimation Problem

In this section, a maximum *a posteriori* probability estimation problem defines the relationship between the image inhomogeneity, the labelmap, and the intensity patterns of the images[112]. Like Van Leemput’s complete data model, the hidden data H defines the labelmap, the observed data O represents the MR images, and Ψ captures the parameter space. Unlike the example of Section 2.3.2, the parameter space Ψ is constrained by the pre-defined Gaussian Distribution $P(\Psi) \sim \mathcal{N}(0, \Upsilon_{\mathcal{B}})$, where $\Upsilon_{\mathcal{B}}$ is the variance of the image inhomogeneities. In addition, the dimension of Ψ is n , the number of voxels in the image. The maximum *a posteriori* probability estimation problem is defined as

$$\hat{\Psi} = \arg \max_{\Psi} P(\Psi|O) = \arg \max_{\Psi} \log P(\Psi|O). \quad (2.22)$$

This section first reiterates the derivations of Section 2.3.1 to apply the EM algorithm to the current problem. As in Section 2.3.2, we then derive the M-Step specific to our current problem.

The adaptation of the approach of Section 2.3.1 starts with a careful analysis of the posterior log probability of Equation (2.22). Applying Bayes rule, we interpret the posterior log probability as

$$\log P(\Psi|O) = \log \left(\frac{P(O|\Psi) \cdot P(\Psi)}{P(O)} \right) = \log L(\Psi) + \log P(\Psi) - \log P(O). \quad (2.23)$$

In Section 2.3.1, the objective function $Q(\cdot, \cdot)$ of Equation (2.4) characterizes the instance of the EM algorithm. Based on the posterior probability of Equation (2.23) and the fact that $\log P(O)$ does not depend on Ψ , the objective function Q_{MAP} for the current model is defined as

$$Q_{MAP}(\Psi, \Psi') = E_{H|O, \Psi'}(\log P(\Psi|H, O)) = Q(\Psi, \Psi') + \log P(\Psi).$$

$Q_{MAP}(\Psi, \Psi')$ defines estimation problem of

$$\hat{\Psi} = \arg \max_{\Psi} P(\Psi, O),$$

which has the same solution as Equation (2.22).

With respect to the version of Section 2.3.2, this objective function is extended by the term $\log P(\Psi)$. The additional log prior over the parameter space transforms Ψ into a random variable, which is characterized by the prior $P(\Psi)$. The convergence properties of the algorithm with respect to $\log P(\Psi, O)$ do not change. By using the result of Corollary 7, we first show that $Q_{MAP}(\cdot, \Psi')$ is a lower bound of $\log P(\Psi, O)$

$$\begin{aligned} Q_{MAP}(\Psi, \Psi') &= Q(\Psi, \Psi') + \log P(\Psi) \stackrel{15}{\leq} L(\Psi) + \log P(\Psi) = \log P(O|\Psi) + \log P(\Psi) \\ &= \log P(\Psi, O). \end{aligned}$$

Similarly, we prove that $Q_{MAP}(\cdot, \Psi')$ touches $\log P(\Psi, O)$ at least in one point using Corollary 8:

$$Q_{MAP}(\Psi', \Psi') = Q(\Psi', \Psi') + \log P(\Psi') \stackrel{16}{=} L(\Psi') + \log P(\Psi') = \log P(\Psi', O).$$

In the remainder of this section we derive an EM approach that is very similar to the one of Section 2.3.2. To do so we apply Equation (2.11) to the new objective function

$$Q_{MAP}(\Psi, \Psi') = \sum_{x=1}^n \mathcal{W}_x^t \log V_x + \sum_{x=1}^n \mathcal{W}_x^t \log U_x(\Psi) + \log P(\Psi)$$

¹⁵Result of Corollary 7

¹⁶Result of Corollary 8

so that the zero-condition for the maximum is defined as

$$\begin{aligned} \frac{\partial Q_{MAP}(\Psi, \Psi')}{\partial \Psi} &= \frac{\partial}{\partial \Psi} \sum_{x=1}^n \mathcal{W}_x^t \log U_x(\Psi) + \log P(\Psi) \\ &= \sum_{x=1}^n \mathcal{W}_x^t \frac{\partial \log U_x(\Psi)}{\partial \Psi} + \frac{\partial \log P(\Psi)}{\partial \Psi} = 0. \end{aligned}$$

Similar to Equation (2.13) the M-Step is characterized by

$$\sum_{x=1}^n \mathcal{W}_x^t \left((o_x - \mu_x(e_a, \Psi))^t \Upsilon_a^{-1} \frac{\partial \mu_x(e_a, \Psi)}{\partial \Psi} \right)_{a=1, \dots, l} + \frac{\partial \log P(\Psi)}{\partial \Psi} = 0 \quad (2.24)$$

This completes the revision of Section 2.3.1.

Similar to Section 2.3.2, we turn Equation (2.24) into a system of linear equations using the definition of the weights \mathcal{W} , the vector V , and vector U of Section 2.3.2. In contrast to Section 2.3.2, the parameter space Ψ directly represents the image inhomogeneities. The smoothness constraints of the image inhomogeneities are encoded in the prior of the parameter space $P(\Psi) \sim \mathcal{N}(0, \Upsilon_B)$. Thus, the variance Υ_B defines the degree of smoothness as any large deviations from the zero mean are penalized.

Different from Section 2.3.2, the distribution of conditional intensity probability is defined by [112] as $P(o_x | h_x = e_a, \Psi) \sim \mathcal{N}(\mu_a + \Psi(x), \Upsilon_a)$. The mean function changes to $\mu_x(e_a, \Psi) \triangleq \mu_a + \Psi(x)$, so that its derivative turns into

$$\frac{\partial \mu_x(e_a, \Psi)}{\partial \Psi(y)} = \begin{cases} 1 & , y = x \\ 0 & , \text{otherwise} \end{cases} \quad (2.25)$$

In addition to the derivative of the mean function, Equation (2.24) requires the derivative of the log prior with respect to voxel x , which is defined as

$$\frac{\partial \log P(\Psi)}{\partial \Psi(x)} = -\Psi^t \Upsilon_B^{-1} e_x.$$

The system of linear equations Equation (2.24) simplifies for voxel $x = 1, \dots, n$ to

$$\mathcal{W}_x^t \left((o_x - \mu_a - \Psi(x))^t \Upsilon_a^{-1} \right)_{a=1, \dots, l} - \Psi^t \Upsilon_B^{-1} e_x = 0$$

The symmetry of Υ_a and Υ_B allows us to further resolve the system of linear equations for voxel $x = 1, \dots, n$ to

$$\left[\sum_{a=1}^k \mathcal{W}_x(a) \Upsilon_a^{-1} (o_x - \mu_a) \right] - \left[\sum_{a=1}^k \mathcal{W}_x(a) \Upsilon_a^{-1} \right] \Psi(x) - e_x^t \Upsilon_B^{-1} \Psi = 0 \quad (2.26)$$

If we define the mean inverse covariance $\bar{\Upsilon}^{-1}$ as in Equation (2.18) and the mean residual \bar{R} as in Equation (2.19), then Equation (2.26) turns into the following linear equation system

$$\bar{\Upsilon}^{-1} \cdot \bar{R} - (\bar{\Upsilon}^{-1} + \Upsilon_B^{-1}) \Psi = 0$$

so that Ψ is defined as

$$\Psi = (\bar{\Upsilon}^{-1} + \Upsilon_B^{-1})^{-1} \bar{\Upsilon}^{-1} \cdot \bar{R} \quad (2.27)$$

Equation (2.27) is very similar to the results of Section 2.3.2 presented in Equation (2.20). If we extend Equation (2.27) with the derivative matrix M of Equation (2.17), which is the identity matrix according to Equation (2.25), the solution to the parameter space is defined as

$$\Psi = \left(M^t \bar{\Upsilon}^{-1} M + \Upsilon_B^{-1} \right)^{-1} M^t \bar{\Upsilon}^{-1} \bar{R}$$

This solution only differs from Equation (2.20) by the inverse covariance of the image inhomogeneity Υ_B^{-1} . In this section Υ_B^{-1} defines the smoothness constraints of the image inhomogeneities instead of the spline function ν of Equation (2.14) in the last section.

In summary, this section defined the problem of segmenting MR images as a maximum *a posteriori* probability estimation problem. We adjusted the instance of the EM algorithm of Section 2.3.1 to our new problem. We then compared the approach of this section to Section 2.3.2. The comparison highlighted the dependency between

the modeling the image inhomogeneities and resulting system of linear equations.

2.4 Discussion

This chapter described the EM algorithm as a lower boundary optimization technique. The algorithm is in general superior to gradient-based methods for incomplete data problems where the introduction of hidden data greatly simplifies their solution. Choosing the correct hidden data, which simplifies the problem, is often quite challenging. However, Section 2.3.1 gives an example for a class of models, for which the EM algorithm is quite simple to implement. We then apply the results of Section 2.3.1 to two popular modeling approaches in medical imaging. The model by Van Leemput [103] describes the segmentation of MR images as a maximum likelihood problem whereas the model by Wells [112] defines it as a maximum *a posteriori* probability estimation problem. For the first time in the medical imaging community, we are able to show the similarity of the two approaches.

The model of Van Leemput described in Section 2.3.2 slightly differs from its original publication [103] as the parameter space is extended by the structure specific mean and variance. Due to the definition of the parameter space in Section 2.3.1, the image inhomogeneity and the structure specific mean statistically depend on each other. Thus, the resulting system of linear equations has generally no unique and closed form solution. To overcome this problem, Van Leemput maximizes the objective function Q in two steps. First, Q is maximized with respect to the structure specific mean and variance. Next, the algorithm optimizes the parameters representing the image inhomogeneities. The resulting algorithm is an implementation of the General-Expectation-Maximization algorithm, which determines an improved instead of optimal parameter setting at each iteration. The implementation with the extended parameter space can be applied to a larger class of image acquisition sequences compared to the approach of Section 2.3.2. The risk of getting stuck in local maxima due to the definition of the structure dependent intensity pattern is reduced. On the downside, the algorithm converges more slowly.

Wells' final EM approach [112] also differs from the derivation in Section 2.3.3 as the covariance matrix defining the prior of the image inhomogeneities Ψ is difficult to compute. Wells [112] simplifies the model by assuming that the distribution of the image inhomogeneities is stationary. As a result, the prior is replaced by a computationally more efficient low pass filter.

If the structure specific distribution is voxel-wise independent, the methods of Section 2.3.2 and Section 2.3.3 fail to consider the noise in MR images. This causes both methods to produce fragmented segmentations [104, 51]. This issue has been addressed by modeling the structure specific posterior probability as a Markov random field [27]. Markov random fields can smooth over the noise, which violates the spatial neighborhood relationships of anatomical structures. Excluding the binary case, determining the exact behavior of the Markov random field is an intractable problem. A solution to the Markov random field is approximated using simulated annealing [44, 74], iterated conditional modes [33, 117, 44], Newton descent methods [65], or mean field approximation [51, 104].

Another possible solution to reduce the fragmentation in the segmentations is the use of non-stationary spatial priors as done by Van Leemput [104]. The spatial priors clearly define the inside of a structure, which favors uniform labeling in this region. In addition, the method can also handle anatomical structures with weakly visible boundaries.

As outlined by [104, 33, 80], the use of spatially varying priors in combination with Markov random fields can smooth over subtle differences in images such as the thin sulci located between the folds of the outer cortex. Spatial priors of an anatomical structures are normally characterized by high values inside the structure, which enforce uniform labeling in this region and therefore reduces fragmentation. Only along the boundaries of structures, where spatial priors are not certain about the labeling, Markov random fields reduce the fragmentation. However, most methods in medical image segmentation related to Markov random fields [51, 103, 44] are defined by stationary pair potentials, which assign high probabilities to uniform labeling independent of the location. These methods do not distinguish between the inside

and the boundary of a structure, which causes them to neglect subtle differences in the image.

Instead, Fischl [34] advocates the use of non-stationary pair potentials. However, non-stationary pair potentials are hard to generate because manual segmentations of all structures of interest must be present in each training subject. In addition, this framework does not guarantee a segmentation without the unwanted fragmentation.

We end this discussion with the approach by Marroquin et al. [65]. This publication suggests that his gradient projection Newtonian descent method is superior to EM approaches like the one of Van Leemput [103]. Unfortunately, the paper is not convincing as it only proves the efficiency of the approach but fails to properly compare it to other EM approaches in the field.

In conclusion, we base the remainder of this thesis on the method by Wells [112]. As suggested by Van Leemput [103], we extend the approach by using non-stationary priors, which allow us to segment structures with similar intensity patterns. We exclude Markov random fields, as the reduction of fragmentation can be achieved with simpler filters. We also will not estimate the structure dependent intensity patterns within the EM framework as suggested by Van Leemput [103]. Using Wells' approach, we were able to adjust the method to different image acquisition protocols through a histogram-normalization filter [9].

Chapter 3

Integration of Spatial Prior into EM Framework

Many neuroscience studies rely on accurate segmentations of cortical and subcortical gray matter. In MR images, most of these anatomical structures have weakly visible boundaries because they are surrounded by structures with similar intensity patterns. For example, the thalamus has a very similar intensity pattern to the neighboring white matter so that the boundary between them is mostly invisible on MR images (Figure 3-1).




Figure 3-1: The MR image (a) shows the region around the thalamus. The corresponding segmentation (b) identifies the right thalamus as pink and the left thalamus as blue. In (a) the boundary of the thalamus is not clearly defined because it has a very similar intensity pattern to the neighboring white matter. The image in (c) shows the spatial priors of both anatomical structures with red indicating higher prior probability of the right thalamus and blue of the left thalamus. Spatial priors capture the spatial variation of an anatomical structure across a population. Many algorithms use them for segmenting structures with weakly visible boundaries.

In the last decade, the medical imaging community has been advocating the use of spatial priors to overcome this problem (Figure 3-1.) Spatial priors define the probability of the presence of an anatomical structure with respect to each individual voxel in the image space. This prior is often captured by an atlas, which defines the probability in accordance with a specific population.

This chapter carefully analyzes the integration of spatial priors into the EM framework of Section 2.3.3. A detailed discussion of this topic is important as the spatial priors not only guide the EM approach in the segmentation of anatomical structures with weakly visible boundaries but also increase the risk of systematic biases in the resulting segmentation.

In this chapter we first discuss a method that generates a spatial atlas from a set of manual segmentations. Next, we describe two different approaches for the integration of the atlas into the EM framework. The first approach uses a non-rigid registration method to align the atlas to the patient and then segments the MR images into anatomical structures. The second algorithm simultaneously registers the atlas to the MR image and segments the MR image into anatomical structures. To our knowledge, this is the first published EM approach to have these capabilities. Each method is validated on various segmentation tasks to show its strengths and weaknesses.

3.1 Defining Spatial Priors of Anatomical Structure

This section describes a method that transforms a set of manual segmentations into an atlas capturing the spatial probability of a specific structure. It is not obvious how to produce these types of atlases, as they should properly represent spatial variations across the training set and not favor characteristics present in a subgroup of subjects.

To generate the atlas, we first select one case from the training set as shown in Figure 3-2. This case, which we call the generic case, defines the coordinate sys-

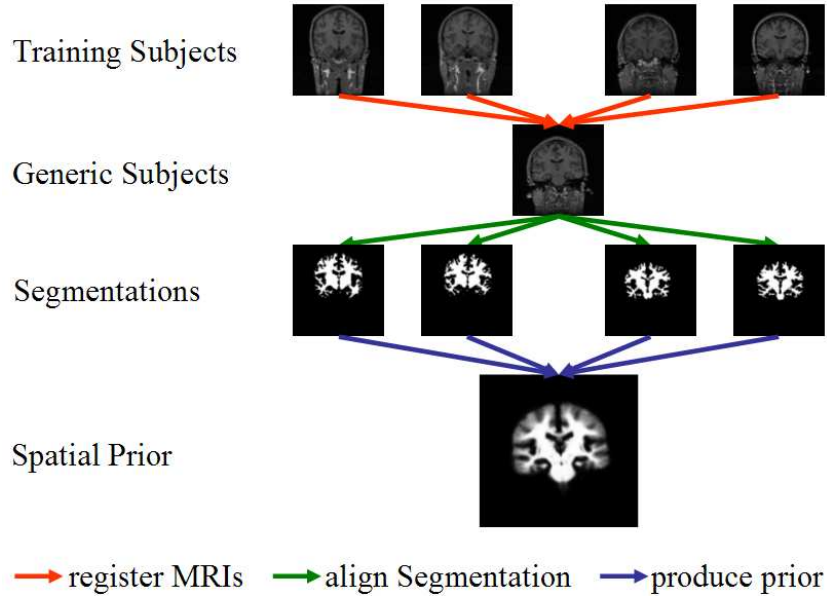


Figure 3-2: The graph visualizes the approach, which produces spatial priors of anatomical structures. First, the MR images of the training set are aligned to a generic subject. Second, the resulting deformation is used to align the corresponding segmentations of the structures of interest to the coordinate system of the generic subject. Finally, the spatial priors capture the ratio between the number of subjects where a specific structure is present at a voxel and the total number of training cases.

tem of the atlas. The choice for the generic case has to be carefully considered as any uncommon characteristics of the generic case might be favored by the atlas and consequently misguide our segmentation method.

The atlas is computed in three steps. First, the training data are aligned to the atlas coordinate system. We do so by individually registering each training case to the generic case. The choice of registration approach is important. The method has to project the anatomy of the specific training subject into the coordinate system of the atlas without significantly altering the shape of the anatomical structure. One choice for this task is affine registration methods, which naturally enforce this constraint as they only scale, displace, and rotate the source to the target. In the remainder of this thesis we will always refer to affine alignment methods being constrained to these three types of deformations.

The second step applies the subject-specific deformation map to the corresponding segmentation of the anatomical structures. As a result, the subject as well as its

segmentation are aligned to the generic case. The final step determines the spatial prior of the anatomical structure at each voxel by the ratio between the entire training set and those subjects that label the voxel with the anatomical structure.

In our specific experiment, we use a data set of 80 subjects where a few structures of interest are manually outlined in each subject. The subjects are registered to a generic case using the affine registration algorithm by Warfield [107]. The method is based on a minimum entropy metric of the labelmaps and has proven to be quite robust [47].




Figure 3-3: Examples of atlases defining the spatial prior of white matter (a), gray matter (b), and corticospinal fluid (c). Blue indicates high and red low probability of a structure being present at that voxel. We refer to the text for more detail.

After each segmentation is aligned to the coordinate system of the generic case, the structure specific spatial prior is computed. At each voxel x , the probability of x being labeled as the anatomical structure of interest a is defined by the ratio between the number of subjects where a is present at voxel x , and the total number of training cases. For example, if we use a training set of 80 volumes to generate the spatial prior of white matter and 40 volumes assign a specific voxel x to white matter then the spatial prior of white matter at voxel x is 50%.

Examples from our set of atlases are shown in Figure 3-3. Blue represents high and red, as well as black, low probability of the anatomical structure being present at that location. All three atlases are generally characterized by slowly varying values indicating the spatial variability of the exact boundary of the structure across a population. In difference to the other two anatomical structure, large areas of the




Figure 3-4: The image to the left shows a part of an MR image with blue highlighting the thin sulci filled with corticospinal fluid. Due to the great variance of the thin sulci the spatial priors under-represent the structure. This is also shown in the right image, where red indicates low spatial probability of corticospinal fluid being present.

spatial prior of white matter (a) are shown in blue because most of the white matter is defined by a relative large body that does not greatly vary spatially across patients. Most of the gray matter, however, is defined by the relatively thin folds of the outer cortex (indicated in green) that greatly vary across a population. Only the subcortical structures are fixated, which are indicated by blue in spatial priors. A similar phenomena is observed in the corticospinal fluid (Figure 3-3(c)), where most of the area close to the skull is shown in red as the structure is defined by very thin and spatially varying sulci located between the folds of the outer cortex. Figure 3-4 focuses on a small part of the corticospinal fluid with blue lines outlining the thin sulci in the MR image and corresponding spatial prior. In Figure 3-3 (c) only the ventricles are shown in blue as they do not greatly vary across patients.

In summary, we generated an atlas capturing the spatial distribution of anatomical structures by first aligning 80 data sets to one generic case and then measuring the correspondence between manual segmentations across the set of training cases. The remainder of this chapter discusses different EM approaches utilizing the spatial atlas.

3.2 Separate Framework for Registration and Segmentation

This section integrates the atlas generated in Section 3.1 into the EM framework of Section 2.3.3 by first registering the atlas to the patient and then segmenting the patient MR into the anatomical structures of interest. We initially focus on the implementation itself by discussing two different registration methods and revisiting the approach of Section 2.3.3. Afterwards, we test the approach by segmenting different substructures of the gray matter and comparing the results to manual segmentations. Most of the work presented in this section originally was published by us in [81]. At that time, we presented for the first time in medical image analysis an approach that couples a non-rigid registration method with an EM model for the segmentation of subcortical structures.

3.2.1 A Sequential Approach

As mentioned, the algorithm developed in this section is defined in two steps. First, the atlas of each structure is aligned to the subject. In the second step, the EM algorithm of Section 2.3.3 segments the subject into the structures of interest. We now review these two steps in detail.

1. Step: Registration of Atlas

The alignment of the atlas to the patient is similar to the registration of the test cases to the atlas coordinate system in Section 3.1. As shown in Figure 3-5, we first register the MR images of the generic case to the MR images of the patient. This process results in a correspondence field, which maps points from the coordinate system of the atlas defined by the generic case to the patient coordinate system. The alignment of the structure specific spatial priors to the patient is achieved by applying the previously generated correspondence field to the spatial priors themselves.

To highlight the dependency between the accuracy of the atlas registration and ro-




Figure 3-5: The figure shows the registration of the atlas to the patient. The MR images of the generic subject are initially registered to the MR images of the patient. The resulting correspondence field is re-applied to the atlases to align the spatial priors with the patient.

bustness of the segmentation approach we develop two different implementation. The first implementation (EM-Affine) uses the affine registration method of Warfield et al. [107], which we previously used for the alignment of the training cases to the atlas coordinate system. The registration method is useful for generating an atlas because it preserves the shape of each anatomical structure. However, for the alignment of the atlas to the subject we need a method that deforms each anatomical structure of the atlas space to its shape in the subject. For this task, affine registration methods are generally too constrained [88]. Instead of an affine method, the second implementation (EM-NonRigid) utilizes the non-rigid registration approach of Guimond et al. [41].

Originally proposed by Thirion [92], the method initially performs an intensity mapping to transform the intensity histogram of the source to that of the target image. The algorithm then determines the displacement between source and target using the demons algorithm (see also Section 1.2). Applied to our atlas, the registration approach essentially turns the spatial priors of a population into a patient specific one. As we will show later in Section 3.2.2, the approach robustly registers cortical structures, such as the superior temporal gyrus, but has too many degrees of freedom for subcortical structures, such as the thalamus.

2. Step: EM Segmentation

After aligning the atlas to the MR images, we segment the images into the anatomical structures of interest. For this purpose, we use the instance of the EM algorithm defined in Section 2.3.3. While Section 2.3.3 focuses on the similarity between Wells' and van Leemput's approach, this section describes the EM approach with respect to structure specific parameters.

We revisit the derivation of Section 2.3.3 by substituting the rather general notation for the complete data model of Chapter 2 with variables more easily associated with terms in medical imaging. The observed log image data \mathcal{I} replaces O , the labelmap \mathcal{T} replaces H , and the image inhomogeneities or bias \mathcal{B} represents the parameter space Ψ . As in Chapter 2, ' a ' represents an anatomical structure of interest and ' x ' defines a voxel in the image.

Within the segmentation model, the intensity distribution of each anatomical structure a is defined by a Gaussian distribution $\mathcal{N}(\mu_a, \Upsilon_a)$. While Gaussians do not properly represent the distribution of large anatomical structures, like gray matter, this representation simplifies the calculations in the M-Step as discussed in Section 2.3. Each anatomical structure a is also characterized by the aligned spatial prior $f_a(\cdot)$, which are defined by the atlas. The aligned spatial prior defines the prior probability of a voxel x being assigned to specific anatomical structure a , which is represented in our statistical framework as

$$P(\mathcal{T}_x = e_a) \triangleq f_a(x). \quad (3.1)$$

Another important relationship in our model is expressed by the likelihood $P(\mathcal{I}_x | \mathcal{B}_x, \mathcal{T}_x = e_a)$, which is the probability of the log intensity \mathcal{I}_x at voxel x conditioned on image inhomogeneities \mathcal{B}_x and the anatomical structure a . As mentioned in Section 2.3.3, the likelihood $P(\mathcal{I}_x | \mathcal{B}_x, \mathcal{T}_x = e_a)$ is defined by the Gaussian distribution $\mathcal{N}(\mu_a + \mathcal{B}_x, \Upsilon_a)$, which is equivalent to

$$P(\mathcal{I}_x | \mathcal{B}_x, \mathcal{T}_x = e_a) \triangleq \frac{1}{\sqrt{(2\pi)^n |\Upsilon_a|}} e^{-\frac{1}{2}(\mathcal{I}_x - \mathcal{B}_x - \mu_a)^t \Upsilon_a^{-1} (\mathcal{I}_x - \mathcal{B}_x - \mu_a)} \quad (3.2)$$

After defining the structure dependent probabilities of the segmentation model, we now reiterate the EM approach of Section 2.3.3.

Expectation Step (E-Step)

The E-Step calculates the weights \mathcal{W} , which is the structure specific posterior probability with respect to the current estimate of the image inhomogeneity \mathcal{B}' and the image data \mathcal{I} . According to Equation (2.10) the weights \mathcal{W} are defined as

$$\mathcal{W}_x(a) \triangleq \frac{P(\mathcal{I}_x | \mathcal{B}'_x, \mathcal{T}_x = e_a) \cdot P(\mathcal{T}_x = e_a)}{\sum_{a'} P(\mathcal{I}_x | \mathcal{B}'_x, \mathcal{T}_x = e_{a'}) \cdot P(\mathcal{T}_x = e_{a'})}, \quad (3.3)$$

which is the soft assignment of a voxel x to the anatomical structure a based on the current model. If we apply Equation (3.2) and Equation (3.1) to Equation (3.3), and define the normalization factor

$$Z_x \triangleq \sum_{a'} |\Upsilon_{a'}|^{-\frac{1}{2}} \cdot e^{-\frac{1}{2}(\mathcal{I}_x - \mathcal{B}'_x - \mu_{a'})^t \Upsilon_{a'}^{-1} (\mathcal{I}_x - \mathcal{B}'_x - \mu_{a'})} \cdot f_{a'}(x)$$

then the weights at voxel x for the anatomical structure a become

$$\mathcal{W}_x(a) = \frac{1}{Z_x} |\Upsilon_a|^{-\frac{1}{2}} \cdot e^{-\frac{1}{2}(\mathcal{I}_x - \mathcal{B}'_x - \mu_a)^t \Upsilon_a^{-1} (\mathcal{I}_x - \mathcal{B}'_x - \mu_a)} \cdot f_a(x). \quad (3.4)$$

In summary, the E-Step updates the weights \mathcal{W}_x based on the current estimate of the inhomogeneity \mathcal{B}' as all other parameters in Equation (3.4) are fixed.

Maximization Step (M-Step)

The M-Step updates the approximation of the image inhomogeneities \mathcal{B}' based on the results of the E-Step. As mentioned in Section 2.4, modeling the inhomogeneities as a Gaussian distribution and determining \mathcal{B} based on Equation (2.27) is computationally relative expensive. Instead, we achieve good results with the approach suggested by Wells [112]. The “inverse variance” $(\overline{\Upsilon}^{-1} + \Upsilon_{\mathcal{B}}^{-1})^{-1}$ in Equation (2.27) is approximated by a simple low pass filter \mathcal{F} , which enforces the smoothness constraint of the image inhomogeneities. The approximation of the image inhomogeneity \mathcal{B}' is updated by

the product between the \mathcal{J} and a weighted residual \tilde{R} :

$$\mathcal{B}' \leftarrow \mathcal{J} \cdot \tilde{R} \quad (3.5)$$

where the residual is defined as $\tilde{R}(x) \triangleq \sum_a \mathcal{W}_x(a) \Upsilon_a^{-1}(\mathcal{I}_x - \mu_a)$.

Algorithm 1: EM SEGMENTATION WITH SPATIAL PRIORS()

procedure REGISTRATION(*ATLAS*)

define correspondence field $C \leftarrow$ align MRIs of *ATLAS* to image \mathcal{I}

define the structure specific spatial prior of the subject:

for all structures $a \in \{1, \dots, N\}$

define $f_a \leftarrow$ apply C to the specific spatial prior of structure “ a ”
 defined by *ATLAS*

return (f_1, \dots, f_N)

procedure SEGMENTATION(f_1, \dots, f_N)

repeat

E-Step: Calculate $\mathcal{W}_x(a) \leftarrow \frac{1}{Z} P(\mathcal{I}_x | \mathcal{T}_x(a) = 1, \mathcal{B}'_x) \cdot f_a(x)$,
 where $P(\mathcal{I}_x | \mathcal{T}_x(a) = 1, \mathcal{B}'_x)$ is defined as in Equation (3.2)

M-Step: Update $\mathcal{B}' \leftarrow \mathcal{J} \cdot \sum_a \mathcal{W}_x(a) \Upsilon_a^{-1}(\mathcal{I}_x - \mu_a)$

 where \mathcal{J} is a low pass filter

 and (μ_a, Υ_a) are structure specific intensity parameters

until \mathcal{B}' converges

define labelmap: $\mathcal{T}_x \leftarrow \arg \max_a \mathcal{W}_x(a)$

main

define aligned spatial priors $f \leftarrow$ REGISTRATION(*ATLAS*)

define labelmap $\mathcal{T} \leftarrow$ SEGMENTATION(f)

The EM approach iterates between the E- and M-Step until the algorithm con-

verges to a local maximum (see [25].) To determine the final labelmap \mathcal{L} according to Equation (2.21), we assign each voxel x to the structure a with the largest weight $\mathcal{W}_x(a)$.

This completes the review of the EM approach, which first aligns the atlas to the subject and then segments the MR images of the subject into anatomical structures. We derived two different implementations in order to assess the accuracy of our approach with respect to the alignment method. The first implementation, EM-Affine, uses an affine registration approach and the second approach, EM-NonRigid, uses a non-rigid method. An abstract representation of both implementations is given in Algorithm 1.

3.2.2 Validation

We focus on two different experiments, which analyze the accuracy of each implementation. In each experiment, the automatic segmentation results are compared to manual segmentations, which substitute for the missing ground truth. The first experiment uses four test cases on which the superior temporal gyrus is outlined by five different experts. The superior temporal gyrus (Figure 3-6) defines a complex shape whose boundary is often difficult to define because it is adjacent to structures with similar intensity patterns. The second experiment uses 22 test cases on which the thalamus is manually outlined by a single expert. Compared to the superior temporal gyrus, the thalamus has a much simpler shape (see Figure 3-8.) Unlike the superior temporal gyrus, the thalamus is difficult to register because it is mostly outlined by weakly visible boundaries on MR images.

A Multiple Rater Experiment of the Superior Temporal Gyrus

The multiple rater experiment is based on four test cases generated by different brain MRIs (Gradient-echo T1-weighted: 256x256x124, resolution of 0.9375mm x 0.9375mm x 1.5 mm; T2-weighted: 256x256x62, resolution of 0.9375mm x 0.9375mm x 3.0 mm.) For the experiment, both implementations as well as five experts outline the right

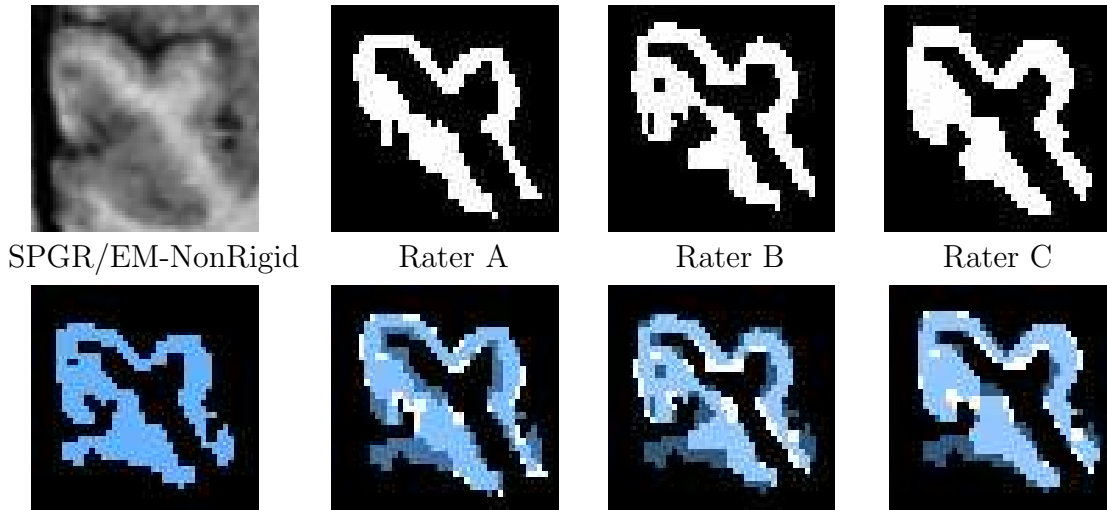


Figure 3-6: Comparing manual and automatic segmentation of the right superior temporal gyrus. In comparison to the different manual segmentations, blue indicates areas of overestimation of EM-NonRigid, light blue represents agreement between the two segmentations, and white shows areas underestimated by EM-NonRigid. While all three manual segmentations differ, they agree that the area outlined by EM-NonRigid to the left is not part of the structure.

and left superior temporal gyri. The experts differ in their level of expertise.

The superior temporal gyrus was chosen because its intensity pattern does not clearly indicate the boundary of the structure. For example, in Figure 3-6 the structures are segmented by three different experts and EM-NonRigid. All four segmentations differ and EM-NonRigid seems to have mislabeled the “arm” to the left of the structure because of the superior temporal gyrus’ similar intensity pattern to the neighboring cortical structures.

The superior temporal gyrus is also ideal for testing the strengths and weaknesses of automatic methods because of its relatively large variation across patients and its small volumetric size. Its large variations are reflected in the slowly varying spatial priors. The spatial priors do not clearly indicate the boundary of the structure, which reduces their efficiency in guiding the automatic segmentation of this structure. As the superior temporal gyrus only defines about 1% of the brain volume [87], volumetric comparison studies of this structure are very sensitive to segmentations with outliers.

To compare the segmentations, three slices around the centroid of the structure are chosen. We would prefer to conduct the experiment for a larger set of slices but it

is quite time consuming to manually outline this structure. The five experts manually outlined the superior temporal gyrus on all three slices. The resulting segmentations are then used by STAPLE to generate a reference standard. STAPLE is an instance of the EM algorithm, which estimates the reference standard in the E-Step and rates each segmentation in the M-Step (see Section 1.5.) Afterwards, we compare the automatic segmentations to the reference standard using the volume overlap measure DICE of Section 1.5. The metric is defined by the volume of the intersection between automatic segmentation and reference standard normalized by the volumes of the segmentation and the reference standard. The measure is zero if two regions do not overlap and one if they are equivalent.




Figure 3-7: The images compare the reference standards (in blue and red) to the spatial prior (the white areas in the background indicate high likelihood) of EM-Affine and EM-NonRigid. The global affine registered spatial priors of EM-Affine show large displacement to the reference standards. The areas, however, favored by the non-rigidly aligned spatial priors of EM-NonRigid correspond to the reference standard of the anatomical structure.

The accuracy of the two methods greatly differs for this task. EM-Affine only receives an average score of $61 \pm 5.3\%$ (mean \pm standard error) for the left and $65.4 \pm 3.4\%$ (mean \pm standard error) for the right superior temporal gyrus where EM-NonRigid receives an average score of $79.3 \pm 2.0\%$ (mean \pm standard error) and $80.7 \pm 0.7\%$ (mean \pm standard error) accordingly. The very poor performance of EM-Affine is caused by initial alignment errors of the affine registration method. These alignment errors are also visualized in Figure 3-7, which compares the reference standards based on the manual segmentations to the aligned spatial priors. Blue defines the reference standard for the right and red for the left superior temporal gyrus; the white areas indicate spatial priors with high likelihoods.

Clearly visible in these images is the displacement between the reference standard and the spatial priors used by EM-Affine. Consequently, EM-Affine is misguided by the spatial priors, which explains the relative low score of the method. The folds of the cortex are characterized by large spatial variations across a population. The affine registration approach of EM-Affine is apparently too constrained to capture these variations. In contrast, the spatial priors of EM-NonRigid are accurately aligned to the reference standards. The non-rigid registration method of EM-NonRigid performs well as the clearly visible folds properly constrain the alignment process. Unlike the affine approach, EM-NonRigid has enough degrees of freedom to capture the corresponds between the folds in the atlas space to the image space.

In conclusion, this experiment shows how sensitive atlas based EM methods are with respect to errors in the spatial priors. EM-Affine performs poorly here, because the affine registration method is too constrained to properly align the atlas to the patient. EM-NonRigid, however, achieves high quality segmentations of the superior temporal gyrus, because the non-rigid registration method accurately defines the displacement between the atlas and image space.

A Large Experiment Based on the Thalamus

For this experiment, both implementations segment 22 test cases into white matter, gray matter, corticospinal fluid and further parcellate gray matter into right and left thalamus. To simplify the task, both methods first skull strip the images so that only the brain itself is visible on the MR images. Afterwards, the methods segment the brain region into the anatomical structures of interest. Chapter 5 describes in detail how to perform such a task in an EM framework.

After all cases are segmented, we compare the automatic segmentations of the thalamus to manual ones. The experiment focuses on the thalamus because most of its boundary is weakly visible on MR images so that it is a challenging structure for registration and segmentation. The weakly visible boundary is caused by the large amount of fiber tracts passing through the thalamus, which create a very similar intensity pattern to the neighboring white matter (Figure 3-8.) EM without spatial




Figure 3-8: The labelmap to the right shows the thalamus in purple and red. The white matter is segmented in white. As can be observed in the MR image to the left, the boundary between white matter and thalamus is only weakly visible as both structures are defined by very similar intensity patterns in this area. Segmentation and registration methods therefore have difficulties in properly analyzing this region.

priors fails and EM with spatial priors heavily relies on these priors. Thus, registration errors greatly influence the segmentation quality of our implementations.

To measure the quality of the automatically generated results, we compare them to the manual segmentations using DICE, the volume overlap measure described in Section 1.5. Unlike the last experiment, each subject is only segmented by one human expert. The experiment can only compare both methods with each other; it cannot indicate how well the methods perform in relation to human experts.

In general, EM-NonRigid performed worse than EM-Affine with an average DICE measure of $82.4 \pm 1.2\%$ (mean \pm standard error.) As visible on Figure 3-9, the EM-NonRigid method greatly underestimates the structure because the intensity-based registration method is too unreliable for structures with smooth boundaries. On average, we found a volume decrease of 10% when comparing the automatic to the manual segmentations.

EM-Affine is more reliable, with an average score of $87.3 \pm 1.2\%$ (mean \pm standard error.) Unlike the non-rigid approach, the alignment method of EM-Affine defines the mapping between atlas and image space globally so that the mapping is not greatly influenced by the weakly visible boundaries between the thalamus and white matter. However, errors in the initial alignment influence the quality of the results as the segmentation approach assigns voxels to the wrong structure based on the misaligned




Figure 3-9: The images to the left show the manual segmentation of the thalamus in 3D and outlined in black on the slice of the MR image. The other two columns feature the segmentation of the two implementations EM-Affine and EM-NonRigid. EM-NonRigid underestimates the size of the thalamus as the non-rigid registration approach cannot properly register this region due to the similar intensity pattern between thalamus and white matter. EM-Affine performed much better but a mis-registration along the vertical direction is visible in the images.

spatial priors. For example, the segmentation shown in Figure 3-9 shows an offset along the vertical direction of the image that is caused by the misalignment between the generic case to the image space.

Method	Superior Temporal Gyrus	Thalamus
EM-Affine	too constrained	good
EM-NonRigid	good	too flexible

Table 3.1: The table summarizes the results of the two experiments of this section. EM-Affine performs poorly on the superior temporal gyrus but does well for the thalamus. The opposite is true for the EM-NonRigid. We refer to the text for more detail.

In summary, this section integrated a spatial atlas into an EM framework by first registering the atlas to the MR images and then segmenting the subject into the structures of interest. We generated two different implementations of the pipeline approach by using an affine and a non-rigid registration method. To determine the accuracy of each approach we applied them to two experiments whose results are summarized in Table 3.1. While the experiments highlighted the necessity of spatial priors for the automatic detection of cortical and subcortical structures, they also

brought to the surface the problems associated with our pipeline approach. In order to segment these anatomical structures, both implementations rely on an accurate alignment of the atlas to the subject. While EM-NonRigid is reliable for the superior temporal gyrus, it is less accurate for the segmentation of the thalamus, because the non-rigid registration method greatly underestimates this area. The opposite is true for EM-Affine. However, the quality of the superior temporal gyrus segmentation is negatively influenced by the overconstrained affine registration method.

3.3 Integrated Registration Framework

In the previous section we encountered problems with our application due to the sequential execution of atlas registration and image segmentation. Active contour methods, such as [59, 100], have addressed this issue by joining these two steps within one framework. They integrate atlas registration into their model by extending the definition of the evolving shape to include its pose.

Motivated by their success, we now propose a novel EM framework joining atlas registration and image segmentation. Unlike Bayesian frameworks developed for active contour methods, our approach considers the anatomical structure associated with each voxel within image. In addition, we explicitly model the image inhomogeneities of MR images to segment large data sets without manual intervention.

In Section 3.3.1, we develop a Bayesian framework targeted towards atlas registration, inhomogeneity correction, and segmentation of anatomical structures. We use the EM approach to solve the resulting estimation problem. The method uses an initial imperfect estimation to converge to a good approximation for these problems.

In Section 3.3.2, we apply the concept to a hierarchical registration framework, which models global- and structure-dependent deformations. With the exception of the background, each anatomical structure is associated with an individual set of affine registration parameters. The mapping of the background to the MR images is modeled as a non-rigid deformation to further increase the accuracy of our approach.

In Section 3.3.4 we perform two different experiments. We first test different

variations of our approach on synthetic images. Afterwards, we repeat the validation experiment of the thalamus of Section 3.2.2 to compare the accuracy of our new approach to the pipeline approaches of the last section.

3.3.1 Derivation of the EM Framework Coupling Registration and Segmentation

The accuracy of our EM approach outlining structures with indistinct boundaries on MR images significantly depends on properly modeling the image inhomogeneities as well as correctly registering the atlas to the subject. To model the image inhomogeneities, Chapter 2 points out the complicated relationship between the image inhomogeneities \mathcal{B} and the MR images \mathcal{I} . In this section, we extend their relationship by the registration parameters \mathcal{R} . \mathcal{R} defines the mapping of the structure specific spatial priors to the MR images. The spatial priors do not contain any information about the inhomogeneities \mathcal{B} so that without additional assumptions it is very difficult to extract the inhomogeneities \mathcal{B} and the registration parameters \mathcal{R} from the MR images \mathcal{I} . However, this problem is greatly simplified when the solution is determined for the incomplete model within an EM framework.

The incomplete model is defined by observed data \mathcal{I} and parameter space $(\mathcal{B}, \mathcal{R})$. Their relationship is captured by the maximum *a posteriori* probability estimation problem

$$(\hat{\mathcal{B}}, \hat{\mathcal{R}}) = \arg \max_{\mathcal{B}, \mathcal{R}} \log P(\mathcal{B}, \mathcal{R} | \mathcal{I}) \quad (3.6)$$

In general, this framework results in a system of equations for which there is no analytical solution.

Similar to Chapter 2, we simplify the problem by adding the labelmap \mathcal{T} to the model. Comparing the images generated from \mathcal{T} to the observed MR images \mathcal{I} , we determine the solution to Equation (3.6). We restate Equation (3.6) by marginalizing with respect to all possible labelmaps \mathcal{T}

$$(\hat{\mathcal{B}}, \hat{\mathcal{R}}) = \arg \max_{\mathcal{B}, \mathcal{R}} \log \left(\sum_{\mathcal{T}} P(\mathcal{B}, \mathcal{R}, \mathcal{T} | \mathcal{I}) \right). \quad (3.7)$$

Next, we incorporate $P(\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}')$, where $(\mathcal{B}', \mathcal{R}')$ are estimates of $(\hat{\mathcal{B}}, \hat{\mathcal{R}})$, into Equation (3.7)

$$\begin{aligned} (\hat{\mathcal{B}}, \hat{\mathcal{R}}) &= \arg \max_{\mathcal{B}, \mathcal{R}} \log \left(\sum_{\mathcal{T}} \frac{P(\mathcal{B}, \mathcal{R}, \mathcal{T}|\mathcal{I})P(\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}')}{P(\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}')} \right) \\ &= \arg \max_{\mathcal{B}, \mathcal{R}} \log E_{\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}'} \left(\frac{P(\mathcal{B}, \mathcal{R}, \mathcal{T}|\mathcal{I})}{P(\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}')} \right). \end{aligned} \quad (3.8)$$

The purpose of these operations is to put Equation (3.8) into a form such that we can exploit the following bound derived via Jensen's Inequality from Corollary 7

$$\log E_{\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}'} \left(\frac{P(\mathcal{B}, \mathcal{R}, \mathcal{T}|\mathcal{I})}{P(\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}')} \right) \geq E_{\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}'} \left(\log \frac{P(\mathcal{B}, \mathcal{R}, \mathcal{T}|\mathcal{I})}{P(\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}')} \right). \quad (3.9)$$

The right side of Equation (3.9), which defines a lower bound on the objective function $\log E_{\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}'} \left(\frac{P(\mathcal{B}, \mathcal{R}, \mathcal{T}|\mathcal{I})}{P(\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}')} \right)$, is more easily maximized using an EM approach. The Expectation Step (E-Step) computes the objective function

$$Q[(\mathcal{B}, \mathcal{R}); (\mathcal{B}', \mathcal{R}')] \triangleq E_{\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}'} \left(\log \frac{P(\mathcal{B}, \mathcal{R}, \mathcal{T}|\mathcal{I})}{P(\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}')} \right). \quad (3.10)$$

The Maximization Step (M-Step) updates the approximations \mathcal{B}' and \mathcal{R}' with the maximum of the lower bound

$$(\mathcal{B}', \mathcal{R}') \leftarrow \arg \max_{\mathcal{B}, \mathcal{R}} Q[(\mathcal{B}, \mathcal{R}); (\mathcal{B}', \mathcal{R}')]. \quad (3.11)$$

Finding the solution to Equation (3.11) is complicated as we simultaneously estimate the independent inhomogeneities \mathcal{B} and registration parameters \mathcal{R} . We will make certain assumptions to simplify the problem. Before we do so, we rewrite

Equation (3.11) by applying Bayes' rule

$$\begin{aligned}
(\mathcal{B}', \mathcal{R}') &\stackrel{1}{\leftarrow} \arg \max_{\mathcal{B}, \mathcal{R}} E_{\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}'} (\log P(\mathcal{B}, \mathcal{R}, \mathcal{T}|\mathcal{I}) - \log P(\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}')) \\
&\stackrel{2}{=} \arg \max_{\mathcal{B}, \mathcal{R}} E_{\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}'} (\log P(\mathcal{B}, \mathcal{R}, \mathcal{T}|\mathcal{I})) \\
&\stackrel{3}{=} \arg \max_{\mathcal{B}, \mathcal{R}} E_{\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}'} (\log P(\mathcal{B}, \mathcal{R}|\mathcal{T}, \mathcal{I}) + \log P(\mathcal{T}|\mathcal{I})) \\
&\stackrel{4}{=} \arg \max_{\mathcal{B}, \mathcal{R}} E_{\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}'} (\log P(\mathcal{B}, \mathcal{R}|\mathcal{T}, \mathcal{I})) \\
&\stackrel{5}{=} \arg \max_{\mathcal{B}, \mathcal{R}} E_{\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}'} (\log P(\mathcal{I}|\mathcal{T}, \mathcal{B}, \mathcal{R}) + \log P(\mathcal{B}, \mathcal{R}|\mathcal{T}) - \log P(\mathcal{I}|\mathcal{T})) \\
&\stackrel{6}{=} \arg \max_{\mathcal{B}, \mathcal{R}} E_{\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}'} (\log P(\mathcal{I}|\mathcal{T}, \mathcal{B}, \mathcal{R}) + \log P(\mathcal{R}|\mathcal{T}, \mathcal{B}) + \log P(\mathcal{B}|\mathcal{T}))
\end{aligned} \tag{3.12}$$

As mentioned, finding a solution to the equation is simplified by making certain independence assumptions. The first factor in Equation (3.12) is the likelihood $P(\mathcal{I}|\mathcal{T}, \mathcal{B}, \mathcal{R})$. In the last section, the likelihood $P(\mathcal{I}_x|\mathcal{T}_x = e_a, \mathcal{B}_x)$ is defined by the Gaussian distribution $\mathcal{N}(\mu_a + \mathcal{B}(x), \Upsilon_a)$ of Equation (3.2). In this framework, the structure dependent parameters μ_a and Υ_a are spatially independent so that the parameters are not influenced by the registration parameters \mathcal{R} . This leads us to the assumption that \mathcal{I} is independent of \mathcal{R} conditioned on \mathcal{T} and \mathcal{B} . Based on the same reasoning, we also assume that the registration parameters \mathcal{R} are independent of \mathcal{B} conditioned on \mathcal{T} . The last factor in Equation (3.12) is the conditional inhomogeneity probability $P(\mathcal{B}|\mathcal{T})$ for which we assume independence between \mathcal{B} and \mathcal{T} as the image inhomogeneities are caused by the radio frequency coil of the scanner [108]. Based on three independence assumptions, Equation (3.12) simplifies to

$$(\mathcal{B}', \mathcal{R}') \leftarrow \arg \max_{\mathcal{B}, \mathcal{R}} E_{\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}'} (\log P(\mathcal{I}|\mathcal{T}, \mathcal{B}) + \log P(\mathcal{R}|\mathcal{T}) + \log P(\mathcal{B})) \tag{3.13}$$

We demonstrate the utility of the independence assumptions in the experiment of

¹Replace $Q[\cdot; \cdot]$ with its definition of Equation (3.10)

² $P(\mathcal{T}|\mathcal{I}, \mathcal{B}', \mathcal{R}')$ does not depend on \mathcal{B} and \mathcal{R}

³Apply Bayes' rule

⁴ $P(\mathcal{T}|\mathcal{I})$ does not depend on \mathcal{B} and \mathcal{R}

⁵Apply Bayes' rule

⁶ $P(\mathcal{I}|\mathcal{T})$ does not depend on \mathcal{B} and \mathcal{R} and apply Bayes' rule on $P(\mathcal{B}, \mathcal{R}|\mathcal{T})$

Section 3.3.4 when we segment 22 subjects into subcortical structures.

In its current form, Equation (3.13) is characterized by an expected value, which sums over all possible settings of \mathcal{T} . To further simplify the problem, we assume spatial independence of \mathcal{T} and independence between \mathcal{R} and \mathcal{B} on the basis of the previous discussion. Equation (3.13), which characterizes the EM approach, is restated as the sum over all voxels x of the expected value with respect to \mathcal{T}_x

$$\begin{aligned}
(\mathcal{B}', \mathcal{R}') &\leftarrow \arg \max_{\mathcal{B}, \mathcal{R}} \sum_x E_{\mathcal{T}_x | \mathcal{I}, \mathcal{B}', \mathcal{R}'} (\log P(\mathcal{I} | \mathcal{T}, \mathcal{B}) + \log P(\mathcal{R} | \mathcal{T}) + \log P(\mathcal{B})) \\
&\stackrel{7}{=} \arg \max_{\mathcal{B}, \mathcal{R}} \sum_x \sum_a E_{\mathcal{T}_x | \mathcal{I}, \mathcal{B}', \mathcal{R}'} (\mathcal{T}_x(a)) [\log P(\mathcal{I}_x | \mathcal{T}_x = e_a, \mathcal{B}_x) \\
&\quad + \log P(\mathcal{R} | \mathcal{T}_x = e_a) + \log P(\mathcal{B})] \quad (3.14) \\
&= \arg \max_{\mathcal{B}, \mathcal{R}} \sum_x \sum_a \mathcal{W}_x(a) [\log P(\mathcal{I}_x | \mathcal{T}_x = e_a, \mathcal{B}_x) \\
&\quad + \log P(\mathcal{R} | \mathcal{T}_x = e_a) + \log P(\mathcal{B})].
\end{aligned}$$

The weights, which are calculated in the E-Step instead of Equation (3.10), capture the posterior probability of the structure a being present at voxel x and are now defined as

$$\mathcal{W}_x(a) \triangleq P(\mathcal{T}_x(a) | \mathcal{I}, \mathcal{B}', \mathcal{R}') = \frac{P(\mathcal{I}_x | \mathcal{T}_x(a) = 1, \mathcal{B}'_x) \cdot P(\mathcal{T}_x(a) = 1 | \mathcal{R}')}{P(\mathcal{I}_x | \mathcal{B}'_x, \mathcal{R}')} \quad (3.15)$$

As each factor in Equation (3.14) only depends on either \mathcal{R} or \mathcal{B} , the M-Step updates the approximations of the inhomogeneities and the registration parameters separately

$$\mathcal{R}' \leftarrow \arg \max_{\mathcal{R}} \sum_x \sum_{\alpha} \mathcal{W}_x(a) \cdot \log P(\mathcal{T}_x = e_a | \mathcal{R}) + \log P(\mathcal{R}) \quad (3.16)$$

$$\mathcal{B}' \leftarrow \arg \max_{\mathcal{B}} \sum_x \sum_{\alpha} \mathcal{W}_x(a) \cdot \log P(\mathcal{I} | \mathcal{T}_x = e_a, \mathcal{B}) + \log P(\mathcal{B}) \quad (3.17)$$

Before we focus on the modelling of the registration parameter \mathcal{R} we must point

⁷At this point, we assume that the two conditional probabilities $P(\mathcal{I} | \mathcal{T}_x, \mathcal{B})$ and $P(\mathcal{R} | \mathcal{T})$ are defined by the product of the corresponding conditional probabilities over all the voxels in the image space. For the likelihood $P(\mathcal{I} | \mathcal{T}_x, \mathcal{B}) = \prod_x P(\mathcal{I}_x | \mathcal{T}_x, \mathcal{B}_x)$ this assumption was previously stated in [112] and for the conditional registration probability $P(\mathcal{R} | \mathcal{T}) = \prod_x P(\mathcal{R} | \mathcal{T}_x)$ we describe a distribution in Section 3.3.2 that fulfills this constraint.




Figure 3-10: For this MR image our novel approach achieved a suboptimal result as it failed to differentiate the neck from the brain (outlined in red.) See text for detail discussion.

out the method’s sensitivity towards uncommon characteristics favored by the atlas; a minor drawback of this formulation. The aligned atlas is represented in Equation (3.16) and Equation (3.15) by $P(\mathcal{T}_x = e_a | \mathcal{R})$. To match the atlas with the segmentation problem, the algorithm compensates for any biases through intensity correction and atlas realignment. For example, if the atlas does not properly capture the brain intensity distribution, the algorithm might identify the neck as part of the brain such as in Figure 3-10. The inhomogeneity correction will then normalize the intensity pattern of the neck to the ones inside the brain and the registration might scale the atlas of the brain to cover both brain and neck. This causes the algorithm to converge to a suboptimal solution but a re-calibration of the atlas to the intensity pattern of the MRI protocol can overcome this problem.

Figure 3-11 shows another example of the influence of a biased atlas on the segmentation results. The first row in Figure 3-11 presents a biased spatial prior of white matter and the resulting average error in the thalamus over 22 test cases. In comparison to the normal spatial prior in the second row, the biased spatial prior strongly favors white matter in the area of the left thalamus. This unevenness between the right and left brain hemisphere causes an offset in the horizontal direction of the

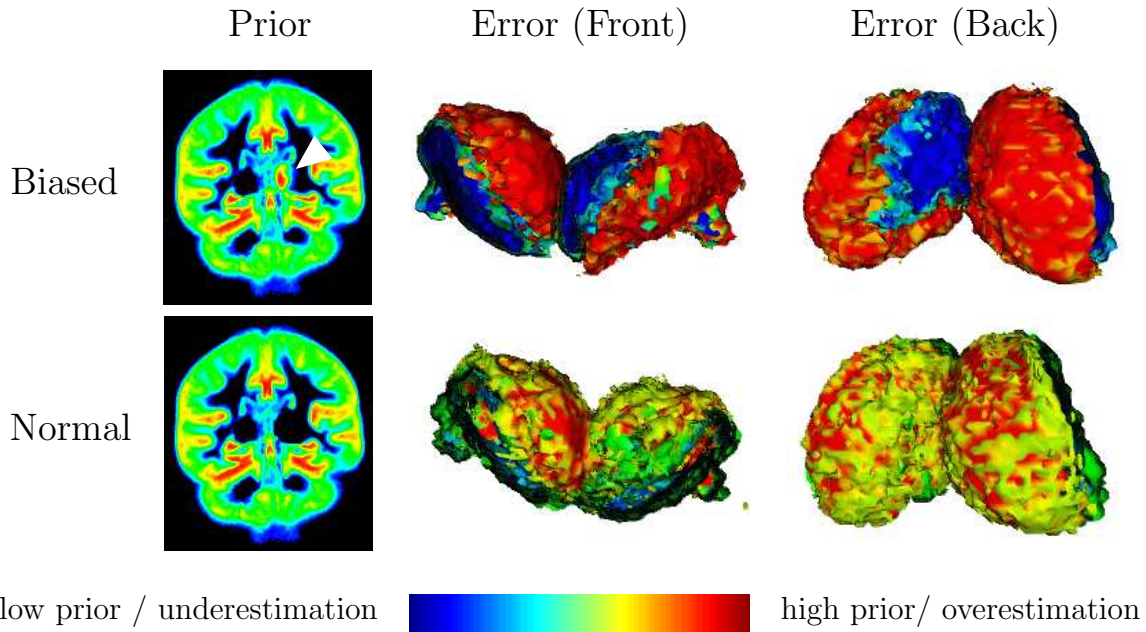


Figure 3-11: This example highlights the effect of a biased spatial prior on the overall performance of the segmentation method. The first row shows a biased spatial prior of the white matter (biased area is indicated by white arrow head) and the resulting average error in the thalamus over 22 test cases. In comparison to the normal spatial prior in the second row, the biased spatial prior over-represents white matter in the area of the left thalamus. This unevenness between the right and left brain hemisphere causes an offset in the horizontal direction in the automatic segmentation of the thalamus. While the average error of the normal segmentation occurs almost symmetrically with respect to both brain hemispheres, the average error of the biased distribution shows an underrepresentation (shown in blue) on one side of both thalamus and an over-representation (shown in red) to the other side.

resulting thalamus segmentations. This offset is also shown in the average error of the thalamus to the left of the corrupted spatial prior, where blue indicates an under- and red an over-representation by the automatic segmentations in comparison to the manual ones. While the average error of the normal segmentation, shown in the second row, occurs almost symmetrically with respect to both brain hemispheres, the average error of the biased distribution shows an underrepresentation on one side and an over-representation to the other side of right and left thalamus.

In summary, this section defines the problem of accurately mapping the atlas to the image space as a maximum *a posteriori* probability estimation problem. Finding the solution to this problem is generally too difficult so that we iteratively solve the

simpler lower bound optimization problem of Equation (3.9) in an instance of an EM algorithm. Even computing the solution to the optimization problem is generally too complicated as we simultaneously have to maximize the underlying objective with respect to the image inhomogeneity \mathcal{B} and the registration parameters \mathcal{R} . We therefore split Equation (3.9) into two separate estimation problems for \mathcal{B} and \mathcal{R} by making certain simplifying independence assumptions. The resulting EM approach is composed by an E-Step determining the \mathcal{W} of Equation (3.15) and the M-Step, which solves the estimation problems of Equation (3.16) and Equation (3.17). While Section 3.2 discusses in detail the solution to Equation (3.17), the next section focuses on finding the optimal registration parameter setting of Equation (3.16).

3.3.2 The Registration Parameters \mathcal{R}

To solve the estimation problem of Equation (3.16), we first define the registration parameters \mathcal{R} and the conditional log probability $\log P(\mathcal{T}|\mathcal{R})$. We model the parameters \mathcal{R} as a hierarchical registration framework, which distinguishes between global- and structure-dependent deformations. We then apply the registration framework to the estimation problem and determine its solution with another optimization algorithm.

The hierarchical registration parameters \mathcal{R} capture the correspondence between atlas, brain, and structures within the brain. The mapping of the atlas to the image space is performed by an interpolation function $r(\mathcal{R}, x)$, which maps voxel x into the coordinate system defined by \mathcal{R} . $r(\mathcal{R}, \cdot)$ can be rigid, affine, or more general non-rigid. As a rigid registration can be interpreted as a special case of an affine transformation, we focus our discussion on affine and non-rigid registration methods.

The parameters \mathcal{R} can be structure-dependent or -independent. The structure-independent parameters capture the correspondence between atlas and image space without knowledge of the underlying brain anatomy. Based on the discussion of Section 3.2, an affine interpolation model is too constrained to capture the characteristics of individual brain structures. The alternative is a non-rigid framework, which often has problems aligning structures with weakly visible boundaries. Instead, we favor an approach that registers each anatomical structure based on the specific requirements

of that structure.

Structure-dependent registration parameters treat the relationship between the atlas and image space for each structure independently. Since most of the misalignment is structure-independent, e.g the head of the subject is not aligned with the atlas space, we expect small differences between structure-dependent parameters of different structures. However, forcing the parameters to be similar across structures is difficult to express in Equation (3.16) so that we need a framework that consists of structure-dependent and -independent parameters.

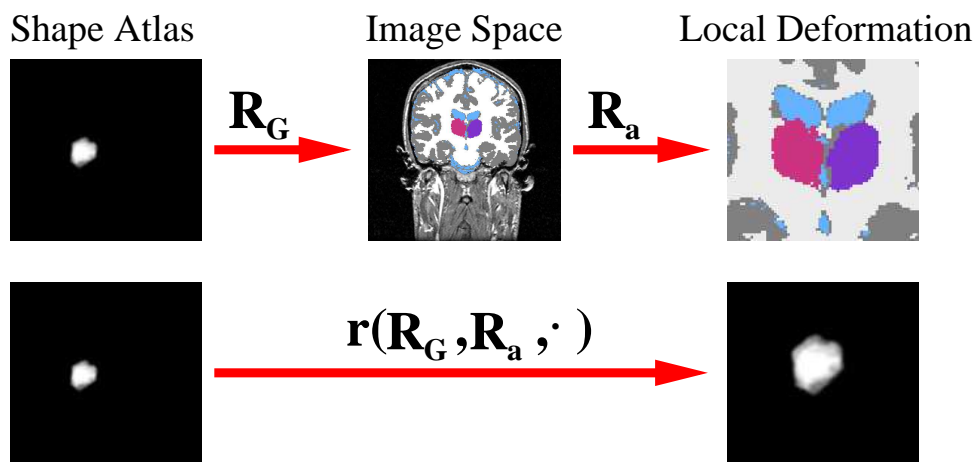


Figure 3-12: The graph represents the hierarchical affine registration framework developed throughout this section. In this approach, the registration parameter \mathcal{R}_G captures the global and \mathcal{R}_a the structure specific alignment of the spatial prior to the location of the structure in the image space.

We model dependency across structures with a hierarchical registration framework $\mathcal{R} = (\mathcal{R}_G, \mathcal{R}_C)$ as shown in Figure 3-12. \mathcal{R}_G are the global registration parameters, which describe the non-structure dependent deformations between atlas and image. The structure dependent parameters $\mathcal{R}_C \triangleq (\mathcal{R}_1, \dots, \mathcal{R}_N)$ are defined in relation to \mathcal{R}_G and capture the residual structure-specific deformations that are not adequately explained by \mathcal{R}_G . In the remainder of this thesis, we refer to \mathcal{R}_a , the a^{th} entry of \mathcal{R}_C , as the registration parameters specific to structure a with $a \in \{1, \dots, N\}$.

The similarity between structure specific parameters is encoded in Equation (3.16) through the prior probability $P(\mathcal{R})$. Section 3.3.3 defines \mathcal{R} explicitly and then model $P(\mathcal{R})$ with a simple Gaussian distribution. The distributions enforce tight bounds on

\mathcal{R}_C and weak constraints on \mathcal{R}_G .

To further specify the model, we define the conditional probability $P(\mathcal{I}_x = e_a | \mathcal{R})$. While in Section 3.2, f_a captures the spatial distribution of structure a in the image space, f_a is now defined in atlas space. Instead $f_a(r(\mathcal{R}_G, \mathcal{R}_a, \cdot))$ defines the spatial distribution in the structure specific coordinate system of the patient. We can model the conditional structure probability:

$$P(\mathcal{I}_x = e_a | \mathcal{R}) \triangleq \frac{f_a(r(\mathcal{R}_G, \mathcal{R}_a, x))}{\sum_{a'} f_{a'}(r(\mathcal{R}_G, \mathcal{R}_{a'}, x))}, \quad (3.18)$$

which is the conditional probability of the presence of an individual structure normalized by the spatial prior probability of all structures. The normalization across all structures is necessary as the coordinate system of each structure is characterized by the structure dependent registration parameters \mathcal{R}_a . Unlike global affine registration methods, this results in structure dependent coordinate systems that are not aligned with each other. In other words, multiple voxels in the atlas space can be mapped to one location in the image space so that the hierarchical mapping $(r(\mathcal{R}_G, \mathcal{R}_1, \cdot), \dots, r(\mathcal{R}_G, \mathcal{R}_N, \cdot))$ does not define a homeomorphism.

Substituting Equation (3.18) into Equation (3.16) changes the maximum *a posteriori* probability estimation problem to

$$\begin{aligned} \mathcal{R}' &\leftarrow \arg \max_{\mathcal{R}} \sum_x \left(\sum_a \mathcal{W}_x(a) \cdot \left(\log f_a[r(\mathcal{R}_G, \mathcal{R}_a, x)] \right. \right. \\ &\quad \left. \left. - \log \sum_{a'} f_{a'}[r(\mathcal{R}_G, \mathcal{R}_{a'}, x)] \right) \right) + \log P(\mathcal{R}) \\ &= \arg \max_{\mathcal{R}} \sum_x \left(\sum_a \mathcal{W}_x(a) \cdot \log f_a[r(\mathcal{R}_G, \mathcal{R}_a, x)] \right) \\ &\quad - \log \left(\sum_a f_a[r(\mathcal{R}_G, \mathcal{R}_a, x)] \right) + \log P(\mathcal{R}) \end{aligned} \quad (3.19)$$

Finding a closed form solution to Equation (3.19) is generally difficult. Instead, we express Equation (3.19) through a objective function $Q(\cdot)$ for which the maximum is found in Section 3.3.3. We replace $\log f_a(r(\mathcal{R}_G, \mathcal{R}_a, x))$ by $\log(f_a(r(\mathcal{R}_G, \mathcal{R}_a, x)) + \epsilon)$ with $\epsilon > 0$, so that the estimation problem is defined for $f_a(r(\mathcal{R}_G, \mathcal{R}_a, x)) = 0$. If we

define $Q(\cdot)$ as

$$Q(\mathcal{R}) \triangleq \sum_x \left[\sum_a \mathcal{W}_x(a) \cdot \log(f_a[r(\mathcal{R}_G, \mathcal{R}_a, x)] + \epsilon) - \log \left(\sum_{a'} f_{a'}[r(\mathcal{R}_G, \mathcal{R}_{a'}, x)] + \epsilon \right) \right] + \log P(\mathcal{R}) \quad (3.20)$$

then Equation (3.19), which represents the M-Step, changes to

$$\mathcal{R}' \leftarrow \arg \max_{\mathcal{R}} Q(\mathcal{R}). \quad (3.21)$$

To get a better understanding of the objective function $Q(\cdot)$, let $y_a \triangleq r(\mathcal{R}_G, \mathcal{R}_a, x)$ be the coordinate of voxel x in the atlas space of structure a . If voxel x is clearly assigned to structure a' then $\mathcal{W}_x(a') = 1$ and $f_{a'}(y_{a'}) = \sum_a f_a(y_a)$ so that

$$\begin{aligned} & \sum_a [\mathcal{W}_x(a) \log(f_a(y_a) + \epsilon)] - \log(f_{a'}(y_{a'}) + \epsilon) \\ &= \log(f_{a'}(y_{a'}) + \epsilon) - \log(f_{a'}(y_{a'}) + \epsilon) = 0. \end{aligned}$$

In this case, the value of $Q(\cdot)$ is not influenced by x as the sum over all structures at this voxel is zero. In other words, $Q(\cdot)$ is determined by all voxels that are not clearly assigned to one structure.

Figure 3-13 shows a practical example of how the misalignment between atlas and image space impacts the analysis of the EM approach. The intensity of the images corresponds to the value of the underlying function with white indicating high and black low values. In this example, we especially focus on the region around the ventricles, whose location in the MR image in (a) is pointed out by the blue arrow head and in the spatial atlas in (b) by the red arrow head. The misalignment between (a) and (b) impacts the weights \mathcal{W} of the brain in (c) as it detects the ventricles twice - once according to the location of the image and once according to the spatial prior. The initial objective function $Q(\cdot)$ defined by the M-Step in (d) detects this ambiguity as it measures the disagreement between spatial prior of the brain $f_a(\cdot)$ in (b) and the weights $\mathcal{W}_a(\cdot)$ in (c). The objective function shows large disagreement in the area around the blue arrow head as (b) assigns this region to the brain and (c) does not.




Figure 3-13: We applied the EM approach to the MR image (a) and the misaligned spatial atlas in (b). The result of the analysis are shown in (c), which is the weights \mathcal{W} of the E-Step, and (d), which is the initial objective function $Q(\cdot)$ of the M-Step. The impact of the misalignment between atlas and image space is especially apparent in the region of the ventricles, whose location are pointed out in the image space by the blue arrow head and in the atlas space by the red arrow head. We refer to the text for more detail.

In order to resolve this disagreement the EM approach aligns the atlas to the image space as we will observe later in the example of Figure 3-15.

In summary, this section developed an EM approach based on the newly defined registration parameters \mathcal{R} . The E-Step calculates the weights \mathcal{W} at voxel x of Equation (3.15) based on the aligned spatial priors $f_a(r(\mathcal{R}'_G, \mathcal{R}'_a, \cdot))$, intensity \mathcal{I} , and image inhomogeneities \mathcal{B}' :

$$\begin{aligned}
 \mathcal{W}_x(a) &= \frac{P(\mathcal{I}_x | \mathcal{I}_x(a) = 1, \mathcal{B}'_x) P(\mathcal{I}_x(a) = 1 | \mathcal{R}')}{P(\mathcal{I}_x | \mathcal{B}'_x, \mathcal{R}')} \\
 &= \frac{P(\mathcal{I}_x | \mathcal{I}_x(a) = 1, \mathcal{B}'_x) \cdot P(\mathcal{I}_x(a) = 1 | \mathcal{R}')}{\sum_{a'} P(\mathcal{I}_x | \mathcal{I}_x(a') = 1, \mathcal{B}'_x) \cdot P(\mathcal{I}_x(a') = 1 | \mathcal{R}')} \\
 &= \frac{|\Upsilon_a|^{-0.5} \cdot e^{-\frac{1}{2}(\mathcal{I}_x - \mathcal{B}_x - \mu_a)^T \cdot \Upsilon_a^{-1} \cdot (\mathcal{I}_x - \mathcal{B}_x - \mu_a)} \cdot f_a(r(\mathcal{R}'_G, \mathcal{R}'_a, x))}{\sum_{a'} |\Upsilon_{a'}|^{-0.5} \cdot e^{-\frac{1}{2}(\mathcal{I}_x - \mathcal{B}_x - \mu_{a'})^T \cdot \Upsilon_{a'}^{-1} \cdot (\mathcal{I}_x - \mathcal{B}_x - \mu_{a'})} \cdot f_{a'}(r(\mathcal{R}'_G, \mathcal{R}'_{a'}, x))}
 \end{aligned}$$

The M-Step updates the approximation of the inhomogeneities \mathcal{B}' and registration parameters \mathcal{R}' based on the current weights \mathcal{W} . As in Equation (3.5), the inhomogeneity \mathcal{B} is approximated by the product between the simple low pass filter \mathcal{J} and the weighted residual \tilde{R} :

$$\mathcal{B}' \leftarrow \mathcal{J} \cdot \tilde{R} = \mathcal{J} \sum_a \mathcal{W}_x(a) \Upsilon_a^{-1} (\mathcal{I}_x - \mu_a)$$

\mathcal{R}' is updated by the maximum of objective functions $Q(\cdot)$ as defined in Equation (3.21)

$$\begin{aligned} \mathcal{R}' &\leftarrow \arg \max_{\mathcal{R}} Q(\mathcal{R}) \\ &= \arg \max_{\mathcal{R}} \sum_x \left[\sum_a \mathcal{W}_x(a) \cdot \log(f_a[r(\mathcal{R}_G, \mathcal{R}_a, x)] + \epsilon) \right. \\ &\quad \left. - \log \left(\sum_a f_a[r(\mathcal{R}_G, \mathcal{R}_a, x)] + \epsilon \right) \right] + \log P(\mathcal{R}) \end{aligned}$$

The solution to this update rule is approximated by a maximization algorithm that we discuss in detail in the next section.

3.3.3 An Implementation of the Joint Registration and Segmentation Model

This section describes an implementation of the previously derived EM approach. We give an example for the interpolation function $r(\cdot, \cdot, \cdot)$, the corresponding registration parameters \mathcal{R} , a probability density function $P(\mathcal{R})$, and a maximization algorithm to solve the maximum *a posteriori* probability estimation problem defined in Equation (3.21).

The interpolation function $r(\cdot, \cdot, \cdot)$ of Equation (3.21) can represent a variety of mapping approaches. For simplicity, we choose an restricted affine interpolation function so that the parameters for $\mathcal{R}_z = (\vec{t}_z^t, \vec{r}_z^t, \vec{s}_z^t)^t$ with $z \in \{G, 1, \dots, N\}$ define displacement \vec{t}_z , rotation \vec{r}_z , and scaling \vec{s}_z . The mapping is defined by the function

$$r(\cdot, \cdot, \cdot) : \mathbb{R}^{9 \times 9 \times 3} \rightarrow \mathbb{R}^3, (\mathcal{R}_G, \mathcal{R}_a, x) \rightarrow A_{\mathcal{R}_G} \cdot A_{\mathcal{R}_a} \cdot (x^t, 1)^t$$

where $A_{\mathcal{R}_z}$ is an affine transformation matrix based on the parameter setting \mathcal{R}_z .

This framework makes no assumptions about the correspondence between the atlas and the image space by uniformly distributing the global registration parameter \mathcal{R}_G . As opposed to \mathcal{R}_G , the structure or class specific parameters $\mathcal{R}_C \triangleq (\mathcal{R}_1, \dots, \mathcal{R}_N)$ describe the residual of structure specific deformations that are not well explained by \mathcal{R}_G . In general, our model should penalize large deviations of \mathcal{R}_C from the expected mean, which is approximated by the average structure-specific registration parame-




Figure 3-14: This is an synthetic example, where the structures of the foreground - GRAY and WHITE - are of different scale in the image but not in the atlas space. To secure statistical consistency within our affine warping framework, the atlas of the background, which is black in the image, is implicitly defined by the foreground.

ters of the training data. We enforce this penalty by modeling the probability density function of \mathcal{R}_C as a Gaussian distribution $N(\mu_{\mathcal{R}_C}, \Upsilon_{\mathcal{R}_C})$ with structure dependent mean $\mu_{\mathcal{R}_C}$ and variance $\Upsilon_{\mathcal{R}_C}$ based on the mapping parameters of the training data. We choose a Gaussian distribution as small variance $\Upsilon_{\mathcal{R}_C}$ discourages large deformations from the mean $\mu_{\mathcal{R}_C}$. In addition, Gaussian distributions simplify the calculations in the M-Step as discussed in Section 2.3.

A problem with this modeling approach is shown in Figure 3-14. WHITE and GRAY are defined by spatial distributions of equivalent size but their scale in the image space differs. The black structure, the background (BG), is the opposite of the foreground composed by WHITE and GRAY. The affine registration parameters of BG are too constrained to cope with the enlarged WHITE and shrunken GRAY object. To solve this problem and therefore increase the statistical consistency of the model, the spatial distribution of BG is determined implicitly. If we define the structure-specific coordinates in the atlas space as $y_a \triangleq A_{\mathcal{R}_G} \cdot A_{\mathcal{R}_a} \cdot (x^t, 1)^t$ then the spatial distribution of BG is

$$f_{BG}(x) \triangleq \begin{cases} 1 - \sum_{a \neq BG} f_a(y_a) & , \text{ if } \sum_{a \neq BG} f_a(y_a) < 1 \\ 0 & , \text{ otherwise} \end{cases} .$$

The impact of implicitly defined spatial probability on the robustness of our method will be further investigated in Section 3.3.4, where we perform a experiment on the

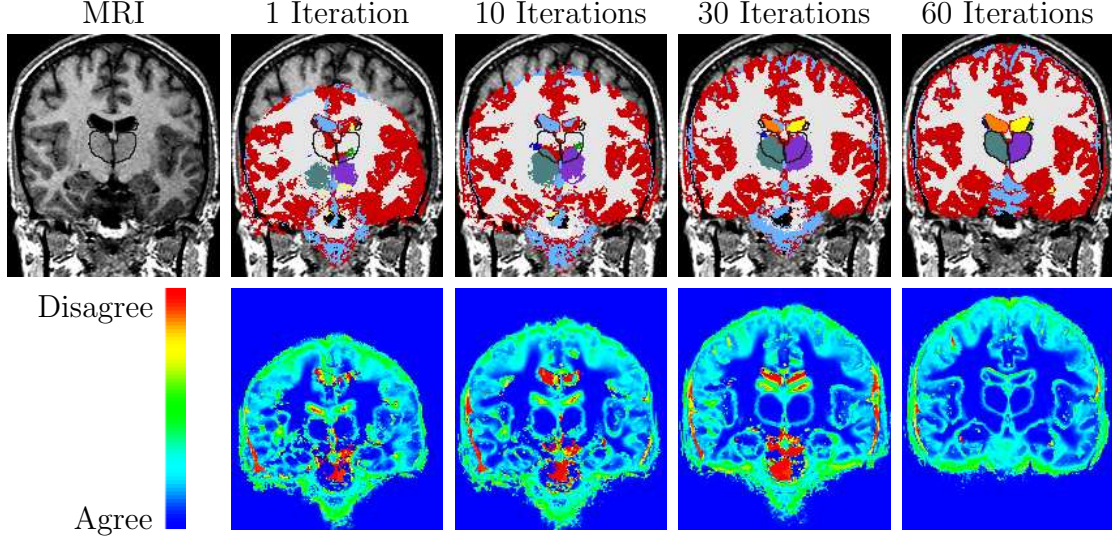


Figure 3-15: The first row shows a slice of an MR volume with black lines outlining the thalamus and the caudate. The segmentations to the right are the segmentation results of our approach after 1, 10, 30, and 60 iterations. The images in the row below are the corresponding objective function $Q(\cdot)$. Red indicates large and blue small disagreement measured by the objective function values. Initially, only the ventricles and corticospinal fluid close to the neck are correctly outlined, which is also indicated in bright red and yellow in the image of the objective function. In the final iteration, the atlas is aligned to the image space so that disagreement is mostly due to the smooth boundaries of the spatial priors.

synthetic images of Figure 3-14.

Based on the previous modeling assumptions, the objective function $Q(\cdot)$ of Equation (3.20) is defined as

$$\begin{aligned}
 Q(\mathcal{R}) \triangleq & \sum_x \left(\sum_a \mathcal{W}_x(a) \cdot \log [f_a(A_{\mathcal{R}_G} \cdot A_{\mathcal{R}_a} \cdot (x^t, 1)^t) + \epsilon] \right. \\
 & \left. - \log \left[\sum_a f_a(A_{\mathcal{R}_G} \cdot A_{\mathcal{R}_a} \cdot (x^t, 1)^t) + \epsilon \right] \right) \\
 & - \frac{1}{2} (\mathcal{R}_C - \mu_{\mathcal{R}_C})^t \Upsilon_{\mathcal{R}_C}^{-1} (\mathcal{R}_C - \mu_{\mathcal{R}_C})
 \end{aligned} \tag{3.22}$$

The example of Figure 3-15 illustrates the dependency between the accuracies of both registration and segmentation. The first row shows a slice of a MR volume with black lines outlining the thalamus and the caudate. The segmentations to the right are the result of our algorithm after 1, 10, 30, and 60 iterations. To show the strength of our model we initialized the approach with an unusual position of the

atlas space. This also explains why the method needs 60 iterations instead of the normal 20 iterations until convergence.

The second row shows objective function $Q(\cdot)$ corresponding to the segmentations above. Red indicates high and blue low disagreement between the atlas space and the weights of the M-Step. Initially, the area around the ventricles and corticospinal fluid close to the neck are highlighted in red, as they are correctly segmented regions that disagree with their location within the atlas space. As the method progresses, the overall accuracy of the registration as well as segmentation increases. The disagreement captured by the objective function in the final iteration is mostly caused by the smooth boundaries of the spatial priors.

The remainder of this section focuses on how to determine a solution for the *a posteriori* probability estimation problem defined by $Q(\cdot)$. First, we decouple the search for \mathcal{R}_G and \mathcal{R}_C , as dependencies between these two parameter settings can cause instability. We then estimate the solution to these problems with the help of a maximization algorithm that finds the optimal solution without the derivative of $Q(\cdot)$. Such methods include the Downhill Simplex algorithm and Powell’s method [82], whose performance is tested with respect to the accuracy of our segmentation approach in Section 3.3.4.

The Downhill Simplex method is initialized by a simplex, which is a geometrical figure consisting of $N + 1$ points or vertices in N dimensions. The edges of the figure connect all vertices directly with each other. At each iteration, the Downhill Simplex method moves the simplex towards the extremum of the function by re-scaling the simplex or reflecting the simplex along one of its edges. As the method gets easily stuck in local extrema, the algorithm is often repeatedly applied to the same problem by initializing the algorithm with slightly varying simplices to the best known solution. The method is generally easy to implement but computationally not efficient.

An often faster method is Powell’s algorithm, which turns the N -dimensional optimization problem into N separate line maximizations. At each iteration, the method first determines the line maxima along each of the N orthogonal directions from the current maximum. The new maximum is then defined by the combination of these N

line maxima. Like the Downhill Simplex method, the algorithm is easy to use. We complete this section with the pseudo code below, which describes the integration of Powell’s method into our approach.

Algorithm 2: SEGMENTATION AND REGISTRATION()

repeat

E-Step: Update soft assignment of anatomical structures

$$\mathcal{W}_x(a) \leftarrow \frac{1}{Z} P(\mathcal{I}_x | \mathcal{I}_x(a) = 1, \mathcal{B}'_x) \cdot f_a(r(\mathcal{R}'_G, \mathcal{R}'_a, x))$$

with \mathcal{R}'_a being the structure specific entry of a in vector \mathcal{R}'_C

M-Step: Update parameter space

$$\mathcal{B}' \leftarrow \mathcal{J} \cdot \sum_a \mathcal{W}_x(a) \Upsilon_a^{-1}(\mathcal{I}_x - \mu_a)$$

$$\mathcal{R}'_G \leftarrow \text{Result of Powell's method with } Q((\cdot, \mathcal{R}'_C))$$

$$\mathcal{R}'_C \leftarrow \text{Result of Powell's method with } Q((\mathcal{R}'_G, \cdot))$$

until \mathcal{B}' and $(\mathcal{R}'_G, \mathcal{R}'_C)$ converge

define labelmap: $\mathcal{T}_x \leftarrow \arg \max_a \mathcal{W}_x(a)$

3.3.4 Validation

This section evaluates the previously derived approach in two steps. The first experiment applies variations of the approach discussed in Section 3.3.3 to synthetic images. The second experiment uses a 22 brain MRI set to evaluate different mapping strategies of the atlas to the image space within an EM segmentation framework.

Experiment on Synthetic Images

To illustrate the reliability of our approach, we apply variations of the implementation discussed in Section 3.3.3 to synthetic images similar to Figure 3-14. The images are composed of three structures: WHITE, GRAY, and BLACK. Unlike in




Figure 3-16: The synthetic images are examples of our experiment, which tests the accuracy of various implementations from Section 3.3.3. In this experiment, only the scale of object WHITE changes. The atlas space of Figure 3-14 remains constant and corresponds to the image with scale 1.0.

Figure 3-14, the structures WHITE and GRAY have the same intensity pattern in the synthetic images (see Figure 3-16.) All structure specific parameters stay fixed with the exception of WHITE in the image space.

The accuracy of the various implementations is tested with respect to the scaling of structure WHITE in the synthetic image. The structure is scaled within a range of 0.1 to 2.5 of its original size for which the atlas space in Figure 3-14 was constructed. For each test image, the automatically generated segmentations are compared to the ground truth by using the volume overlap measure DICE of Section 1.5.

The first experiment shown in Figure 3-17 (a) compares the reliability of our approach using Powell to one using the Downhill Simplex method. Powell’s method outperforms the Downhill Simplex method, when comparing their robustness with respect to the scaling of structure WHITE. The Downhill Simplex method fails on images where WHITE and the non-aligned atlas (Figure 3-14) do not overlap. Powell’s method is unreliable for scaling parameters greater than 2.4 for which WHITE almost disappears from the synthetic image (see Figure 3-16.)

The experiment in Figure 3-17 (b) compares the robustness of global- to structure-specific registration parameters. As expected, structure specific registration parameters are superior because they can better capture the scale differences between WHITE and GRAY.

The graph in Figure 3-17 (c) shows the performance of an implicitly- and an explicitly-defined spatial distribution of BLACK structure. The superior implicit

spatial distribution is the inverse of the spatial distribution of the aligned foreground, which is defined by WHITE and GRAY. As mentioned, the explicit spatial distribution increases the risk of statistical inconsistency within our model. This also complicates the search for the solution of the maximum *a posteriori* probability estimation problem of Equation (3.21), which greatly reduces the reliability of our implementation.




Figure 3-17: The graphs show the robustness of different implementations proposed in Section 3.3.3. The robustness of the method is determined by the quality of the automatic generated result with respect to the scaling of structure WHITE (Figure 3-16.) The graph in (a) shows that Powell’s method is more robust than Simplex approach for our registration problem. The graph in (b) indicates that our approach should use structure specific registration parameters as it is more robust than just using global registration parameters. (c) displays the impact of the explicitly defined spatial background prior on the robustness of our approach. In general, using an implicitly defined spatial prior for the background increases the robustness of our method.

In summary, the most robust approach uses a hierarchical registration framework, an implicitly defined spatial background, and determines the solution to the maximum *a posteriori* probability estimation problem using Powell’s method.

Comparative Experiment on 22 Test Cases

This section repeats the experiment of Section 3.2.2 but compares the accuracy of the pipeline approaches to the novel method of this section. The two pipeline approaches, EM-Affine and EM-NonRigid of Section 3.2, first map the atlas to the subject to initialize the EM approach. In contrast, our new algorithm, EM-Simultaneous-Affine, solves the registration and segmentation problem simultaneously.




Figure 3-18: The images show the segmentation results of four different methods on the thalamus (red /purple) and the caudate (green/blue.) We refer to the text for further detail.

The three methods segment 22 test cases into the three brain tissue classes and further parcellate gray matter into the subcortical structures of the thalamus and the caudate. In order to determine the accuracy of the automatic segmentations with respect to the two subcortical structures we compare them to manual segmentations, which we view as ground truth. The overlap between automatic and manual segmentation is measured using volume overlap measure DICE of Section 1.5. Using DICE we can compare the accuracy of the automatic methods but we cannot make any conclusion about the accuracy of the automatic methods with respect to manual segmentations. We do so later in Chapter 5.

This experiment focuses on the thalamus and caudate as they are challenging structures for registration and segmentation. Purely intensity based segmentation methods, such as EM without spatial priors, cannot outline these structures because part of the boundary is invisible on MRI (Figure 3-18.) Consequently, EM relies heavily on spatial priors making it sensitive towards any misalignment between atlas and image space. In addition, the registration of the spatial priors to the subcortical area is challenging. As we observed in Section 3.2.2, the intensity based non-rigid registration of EM-NonRigid has difficulties mapping the priors to the thalamus because of the structure’s similar intensity pattern to the neighboring white matter. The

caudate, on the other hand, is defined by two “horns” (Figure 3-18,) whose subject specific bending is difficult to capture with an affine registration model such as used by EM-Affine and EM-Simultaneous-Affine. In conclusion, automatic segmentation and registration methods have difficulties in correctly identifying both anatomical structures so that this experiment highlights the strength and weaknesses of each implementation.




Figure 3-19: The graph shows the average DICE score over 22 cases for the pipeline based approaches EM-Affine and EM-NonRigid (EM-NRigid), and the new method EM-Simultaneous-Affine (EM-Sim-Af), which simultaneously performs registration and segmentation. Out of the three methods, EM-Simultaneous-Affine is the only approach that performs well for the thalamus and the caudate.

The results of the experiment are summarized in the Figure 3-19. The graph shows the mean and standard error of the DICE measure for the three algorithms in the 22 cases. For the thalamus, EM-NonRigid (EM-NRigid) performed worst ($82.4 \pm 1.2\%$ - average DICE score \pm standard error) because the intensity based registration method is too unreliable for structures with smooth boundaries. As mentioned, the method often overestimates white matter and underestimated the thalamus in this region, which can be also observed by the segmentation in Figure 3-18 (EM-NonRigid). EM-Affine performs much better than EM-NonRigid ($87.3 \pm 1.2\%$) but the method is sensitive towards initial misalignments. In Figure 3-18 (EM-Affine), this sensitivity causes a vertical offset in the segmentation generated by

EM-Affine. For the thalamus, both methods are outperformed by EM-Simultaneous-Affine (EM-Simultaneous-Affine: $88.7 \pm 0.4\%$). Our new approach performed much better than the other two methods on cases where the deformation between atlas and image space was difficult. Unlike the other two methods it can overcome initial misalignment issues and therefore increase the accuracy of the segmentation.

The observations made with respect to the thalamus cannot be generalized to all anatomical structures in the brain. For example, the caudate is more accurately segmented by EM-NonRigid ($86.0 \pm 1.0\%$) than EM-Simultaneous-Affine ($85.0 \pm 0.7\%$.) Unlike the thalamus, the caudate has a different intensity profile than white matter. Only the relatively small portion of the boundary neighboring the putamen, another subcortical structure, is invisible on MRI. Thus, the intensity based registration method of EM-NonRigid correctly registers the spatial priors to this region and allows the EM segmenter to perform well.

As mentioned, the affine registration methods of EM-Affine ($83.2 \pm 1.7\%$) and EM-Simultaneous-Affine are too constrained to capture the bending of the horn-shaped caudate. This causes both methods to misclassify voxels especially at the tip of the structure, which is also visible in the 3D models in Figure 3-18 and explains the lower score. However, EM-Simultaneous-Affine performs much better than EM-Affine.

We believe that EM-Simultaneous-Affine outperforms EM-Affine on both anatomical structures because modeling the dependency between segmentation and registration further constrains our approach. The constraints reduce the space of possible solutions, which simplifies the search for the optimal solution. In addition, the spatial priors of the structures are directly mapped to the segmentation model. In contrast, EM-Affine aligns an MR image in the atlas space to the image of the patient. It uses the resulting deformation map to align the spatial priors, which inherently increases the risk of systematic biases in the model.

We end the discussion by further investigating the previously made observation with respect to the ventricles. As pointed out in Section 3.3.2, the ventricles are very important in guiding the joint registration and segmentation method. In order to validate this statement, we repeat the previous experiment but ignore the ventricles

in the registration model. This change in our model reduces the average DICE score for the two structures by 5.2% and the standard error increases by factor of 2.5. We therefore conclude that the accuracy of our new approach greatly depends on anatomical structures that are easily identifiable in the segmentation problem.

In summary, this section developed a statistical framework combining atlas registration, segmentation, and inhomogeneity estimation of MR images. We implemented this framework using a hierarchical affine mapping approach for explicit anatomical structures in combination with an implicit spatial distribution for the background. Our approach was validated by automatically segmenting 22 sets of MR images into the thalamus and caudate. We then compared our results to other EM methods as well as manual segmentations. For the caudate our novel approach is not as reliable as EM-NonRigid as the underlying affine registration model cannot properly capture the bending of the horns of the caudate. However, our approach was the only automatic method performing well on both anatomical structures. It achieved the lowest standard error, implying a higher predictive power than the other two methods. In addition, the method outperforms the pipeline approaches for the thalamus. The improvement is primarily due to the joint registration-segmentation estimation problem.

3.4 Conclusion

In summary, spatial priors increase the capability of our EM segmentation method. On the downside, they also increase the risk of systematic biases in our framework. This chapter showed the importance of carefully analyzing the influence of spatial priors in automatic segmentation methods. It discussed two very different philosophies for the integration of a spatial prior into the segmentation process of MR images.

In Section 3.2, we presented an approach that coupled a non-rigid registration method with an EM approach by first aligning the spatial prior to the MR images and then segmenting the MR images into structures of interest. While the approach enables the segmentation tool developed in Section 2.3.3 to segment cortical and

subcortical structures, it is generally too unreliable for neuroscience studies.

The second EM approach developed in Section 3.3 directly couples segmentation and registration. The approach achieves accurate results on the thalamus and the caudate unlike the previously mentioned EM approach. The improved accuracy is caused by joining registration and segmentation into a unified Bayesian framework.

Chapter 4

Modeling Shape Constraints

The segmentation methods of Chapter 3 define anatomical structures by their intensity patterns and spatial distributions. As mentioned, spatial priors do not properly capture the anatomical variability of many structures in the brain. For example, Figure 4-1 shows the average segmentation error of EM-Affine in the experiment of Section 3.3.4. The average error indicates a bias by EM-Affine to overestimate the caudate’s superior (in red) and underestimate its inferior surface (in blue.) This bias is due to the subject specific spatial prior, which does not properly capture the bending of the horns.

Another drawback of spatial priors of this sort is their ”soft” boundaries representing the large spatial variability of a structure within a population (Figure 3-3.) Spatial priors are mostly needed when the boundary between two structures is not clearly defined as in the example of the thalamus in Figure 3-1. In this scenario, neither the slowly varying spatial prior nor the intensity pattern clearly define the boundary, which reduces the accuracy of the EM approach.

Deformable models offer an alternative type of prior information as they capture the shape and permissible modes of variation within a population. In contrast to simple space-conditional label probabilities, deformable models capture the structure specific boundary conditions such as the smoothness of its boundary or the length of the boundary in relation to others.

In this chapter, we improve the discriminatory power of our method by adding a




Figure 4-1: This image shows the average segmentation error of EM-Affine segmenting the caudate in 22 cases. The blur in the image represents the spatial variability of the error function. The average error indicates a general overestimation of the caudate’s outer (in red) and underestimation (blue) of its inner surface by EM-Affine. This bias is due to the subject specific spatial prior.

deformable model to our segmentation approach. Inspired by the work of Tsai et al. [100], the boundary of the entire caudate and thalamus will be guided by the clearly defined border of the neighboring ventricles. This integration of shape constraints into the EM approach is divided into several steps. In Section 4.1, we describe a shape atlases based on the principal component analysis on signed distance maps. In Section 4.2, we discuss the integration of the how the atlas into the EM approach of Section 3.2. Motivated by the success of the joint registration and segmentation approach described in Section 3.3, we derive an EM method in Section 4.3, which simultaneously registers the shape atlas to the patient MR images, estimates the image inhomogeneities, and segments the MR images into the anatomical structures of interest. Each approach is validated by comparing it to other automatic segmentation methods as well as to manual segmentations. Finally, in Section 4.4 we review the results of this chapter.




Figure 4-2: The signed distance map in (b) corresponds to the labelmap in (a). The value of each voxel in (b) is determined by distance of the voxel to the boundary between black and white in (a). To distinguish voxels inside the white object from those outside, the value in the distance map of voxels outside the object is negated.

4.1 Shape Atlas

This section describes an approach to generate a shape atlas with respect to the representation of signed distance maps. This atlas captures the shape variability over the training data, which consists of manual segmentations. Section 4.1.1 briefly discusses the representation of signed distance maps before Section 4.1.2 details the construction of the shape atlas.

4.1.1 Shape Representations

It is challenging to characterize the shape differences and similarities for a specific anatomical structure over an entire population. One can choose from a variety of shape descriptors, such as skeleton based representations [79, 37] or distance transforms [38]. The proper choice greatly depends on the application as each method differs in terms of degree of manual interaction, mathematical properties, and sensitivity to noise. A detailed review of the current developments within the field of shape descriptors is given by Bouix [8]. This section will briefly explain the concepts behind distance transforms, as this is our choice of shape representation for the EM approaches developed later in this chapter.

EM segmenters, like the one presented in Chapter 3, belong to the class of voxel-

based classification methods, which consider the anatomical structure associated with each voxel within a Bayesian framework. In order to incorporate shape constraints into this Bayesian framework, we favor shape representations that operate in a similar manner to our EM approach by characterizing each voxel in the image space rather than by using dynamic meshes. One such representation is that of signed distance maps [20] (see Figure 4-2,) which we refer to as distance maps in the remainder of this thesis. The value of each voxel in a distance map is determined by its distance to an object's boundary and its relative location with respect to this object. If a voxel is inside the object, its value in the distance map is the distance to the closest voxel along the boundary of the object. However, if the voxel is outside the object its value is the negated minimum distance to the boundary.

Based on this definition, signed distance maps are composed of three types of voxels. The first type are voxels with value zero. They represent the boundary of the object. Voxels with positive values compose the second type, which embody the inside of the object. The third type are voxels with negative values that define the outside of the object. In other words, distance transforms represent the boundary of an object as a zero-level set function.

With respect to our segmentation algorithm this representation has the advantage of not requiring manual interaction in order to represent the shape of anatomical structures. In addition, we will show in Section 4.2 that signed distance maps implicitly define spatial probabilities. The integration of the shape representation in an EM approach therefore closely relates to the experience of the previous chapter where we extended our Bayesian framework with spatial priors.

4.1.2 Principal Component Analysis on Distance Maps

This section discusses a method that transforms segmentations of brain structures into a shape atlas based on signed distance maps [100]. An important task in this process is finding a training set large enough to represent the shape variations of an anatomical structure across a population. As this is a typical problem in medical imaging, we make the simplifying assumption that the variations within the shapes

of brain structures are Gaussian in nature. With this assumption, we can generate the shape atlas using Principal Component Analysis (PCA) [16].

PCA is based on the assumption that its input data are samples from a Gaussian distribution, whose variations are defined by a small subspace. The method defines these modes of variation by transforming the training data into a lower dimensional eigenvector space. The significance of an eigenvector capturing common variations within the training data is defined by its corresponding eigenvalue. The sum of eigenvalues defines the total energy of the PCA model and the set of eigenvalues composes the diagonal covariance matrix of a Gaussian distribution that defines the probability of a shape within this space.

The resulting shape atlas is defined by the mean shape, the eigenvectors, and the eigenvalues. The specific shape of a subject is represented by the atlas as the weighted sum of the eigenvectors to the mean shape. We will describe this representation in detail later in this section.




Figure 4-3: The graph shows the relation between the labelmaps in the training data sets and the resulting mean and eigenvector of the PCA approach. In general, red indicates positive and blue negative values. To determine the shape variations across the training data, the labelmaps are first converted into the signed distance maps, the shape representation of choice for our EM approach. The signed distance maps are then the input of the PCA, which determines the mean and the eigenvector. See the text for further detail.

Figure 4-3 shows an example of PCA applied to synthetic images. The training data consists of three images composed of differently sized rectangles. The corresponding distance maps are input to the PCA approach in order to determine the major shape variations across the training data. In the images of Figure 4-3, red indicates positive and blue negative values. Resulting from the PCA approach are the mean distance map and an eigenvector. As expected, the eigenvector captures the difference in size between the three different training images, as it only contains positive values.

In the remainder of this section, we apply the previously described concept to the manual segmentations of anatomical structures.

Preprocessing Training Data

Before applying PCA, we preprocess the training data, which consists of manual segmentations of the structures of interest. We exclude pose issues from the analysis by registering the data set to one central pose with the affine registration method developed by Warfield [107]. We performed this kind of alignment previously in Section 3.1, when we generated spatial priors from the segmentations. Next, we convert the labelmap of each data set $i = 1, \dots, m$ into a distance map, such as in the synthetic example of Figure 4-3. $\mathcal{D}_a^{(i)}$ defines the distance map of structure a for test case i . Each entry in $\mathcal{D}_a^{(i)}$ captures the distance of a voxel to the boundary (see Figure 4-2.)

The preprocessing of the training data concludes by calculating the mean shape $\bar{\mathcal{D}}_a$ and subtracting it from each training set. This step is necessary as PCA assumes that the training data set is a sample from a zero-mean Gaussian distribution. For non-linear spaces, like the one defined by distance maps, the mean is defined with respect to a distance function $d(\cdot, \cdot)$ as

$$\bar{\mathcal{D}}_a \triangleq \arg \min_{\hat{\mathcal{D}}} \sum_i d(\mathcal{D}_a^{(i)}, \hat{\mathcal{D}}). \quad (4.1)$$

$\bar{\mathcal{D}}_a$ corresponds to the shape with the minimum amount of distance to all training

subjects. However, finding the solution to Equation (4.1) is a problem, as current research in medical imaging shows [49, 70].

An alternative is the approximation of the mean $\bar{\mathcal{D}}_a$ of an anatomical structure a by the arithmetic mean

$$\bar{\mathcal{D}}_a \triangleq \frac{1}{n} \sum_i \mathcal{D}_a^{(i)} \quad (4.2)$$

and then project the approximated mean to its closest member on the manifold of distance maps. However, this approach would corrupt the PCA analysis as the first eigenmode would represent the difference between the mean of Equation (4.2) and its projection on the manifold. Instead, we simply assume that distance maps define a linear space and the mean is given by Equation (4.2); a reasonable assumption for small shape deformations. We complete the preprocessing of the training data by computing training example i , represented by

$$\tilde{\mathcal{D}}_a^{(i)} \triangleq \mathcal{D}_a^{(i)} - \bar{\mathcal{D}}_a,$$

which is the mean-corrected distance map of training example $i = 1, \dots, m$. In the remainder of this thesis $\tilde{\mathcal{D}}_a^{(i)}$ will be represented in the form of a vector so that we can apply it to the PCA method. Each entry in the vector represents the mean corrected distance map value of a specific voxel.

Definition of the Atlas

Similar to Tsai et al. [98] we apply PCA to vectors of distance maps across structures $\tilde{\mathcal{D}}^{(i)} \triangleq (\tilde{\mathcal{D}}_1^{(i)t}, \dots, \tilde{\mathcal{D}}_N^{(i)t})^t$ with $i = 1, \dots, m$. This analysis captures the covariance of structures, which is defined by the matrix of eigenvectors $U = (U^{(1)}, \dots, U^{(m-1)})$ and eigenvalues Λ . To reduce computational cost, U and Λ are only defined by the first K eigenvectors and eigenvalues, where K is selected so as to represent 99% of the eigenvalues' energy.

Figure 4-4 shows the outcome of PCA when applied to the ventricles, the caudate, and the thalamus. The images in the figure are different perspectives of the first eigenmode with red indicating an expansion and blue a compression along the

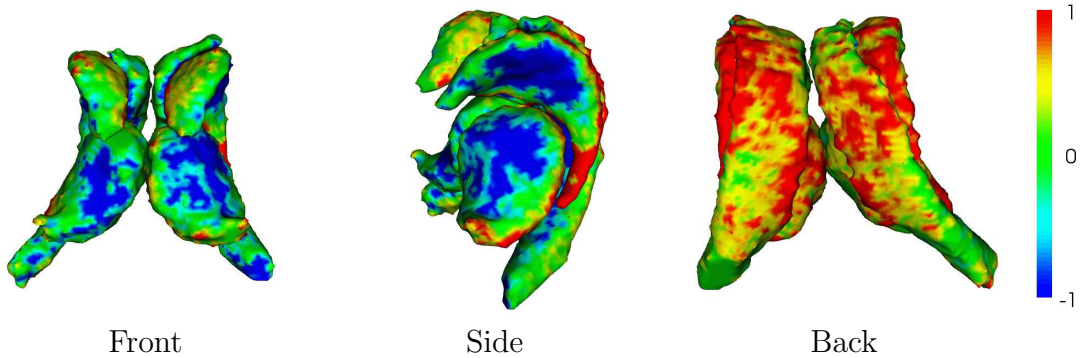


Figure 4-4: The images show different perspectives of the first eigenmode of the joint PCA model of the caudate, thalamus, and ventricles. Red indicates an expansion and blue a compression along the boundary of each structure relative to the mean shape.

boundary of the mean. For this example, we used 22 mean corrected distance map vectors $\tilde{\mathcal{D}}^{(i)}$. Each vector $\tilde{\mathcal{D}}^{(i)}$ consists of three distance maps representing the ventricles, the caudate, and the thalamus of the specific test case i . As mentioned, PCA analyzes the combined shape variations across the three structures so that resulting eigenvectors $U^{(1)}, \dots, U^{(5)}$ describe the covariance of the three structures within the training data set. We choose to define the shape atlas by the first five of the 21 eigenvectors as they represent most of the variability within the training data.

Later in Section 4.2, we develop an EM approach that uses this structure-related dependency to further constrain the space of possible solutions. In the case of the three anatomical structures, the EM approach is guided by the ventricles, an easily identifiable anatomical structure, in order to automatically segment the thalamus and caudate, two structures with weakly visible boundaries.

Now that we have defined a shape atlas, we can describe the structure specific shapes by the atlas-related expansion coefficient or shape parameter \mathcal{S} instead of the rather large shape vector \mathcal{D}_a . The relationship between \mathcal{S} and the structure specific distance map \mathcal{D}_a is defined as

$$\mathcal{D}_a \approx \bar{\mathcal{D}}_a + U_a \cdot \mathcal{S} \quad (4.3)$$

where U_a are the entries in eigenvector U that are associated with the anatomical structure a . Note, that the single shape vector \mathcal{S} with K components represents the

shape of all structures by defining the weighted sum of the eigenvectors to the mean shape.

With respect to the example of Figure 4-4, \mathcal{S} represents the shape across all three anatomical structures. To relate the parameter \mathcal{S} to the specific shape of one of the three structures we extract the entries from the five eigenvectors that represent the anatomical structure. We then apply Equation (4.3) to this shape model in order to generate the structure specific distance map.

As mentioned in the introduction of this section, an additional feature of PCA is the association of the following prior probability to the shape parameter \mathcal{S} :

$$P(\mathcal{S}) = \frac{1}{\sqrt{(2\pi)^K |\hat{\Lambda}|}} \exp\left(-\frac{1}{2} \mathcal{S}^t \hat{\Lambda}^{-1} \mathcal{S}\right) \quad (4.4)$$

where $\hat{\Lambda}$ is a diagonal matrix composed by the first K eigenvalues Λ of the PCA analysis. The prior $P(\mathcal{S})$ captures the probability of the shape of all structures and is grounded in the assumption that our training data are samples from a Gaussian distribution.

In summary, this section discussed a shape atlas based on distance maps. We generated the atlas using the PCA method to address the lack of training data. In the remaining two sections, we integrate this atlas into the EM approach of Section 2.3.3 and Section 3.3.

4.2 Coupling Shape Deformation with Brain MRI Segmentation

It is challenging to define an EM model that properly represents the relationship between the shape parameters \mathcal{S} , the labelmap \mathcal{T} , the image inhomogeneities \mathcal{B} , and the MR images \mathcal{I} . Specifically the interaction between the shape parameters and the labelmap is complicated, as the shape parameters \mathcal{S} capture global characteristics of structures and the hidden labelmap \mathcal{T} individually defines the value of each voxel.

However, we faced similar problems in Section 3.3 when we extended the EM model with the registration parameters. As a result, we closely relate the integration process of the shape constraints to the derivation of Section 3.3.1. As in Section 3.3.1, we will first develop a Bayesian framework and then determine an EM model, which calculates a solution within this framework.

A variety of active contour methods have advocated the use of PCA models on signed distance maps for the accurate segmentation anatomical structures. For example in [100, 60], active contour methods represent the shape by a level set function in an higher dimensional space. Yang et al. [115] further improved this line of research by developing a level set approach not restricted to the low-dimensional PCA space. Inspired by the work of Yang, the main contribution of this section is that while we make use of an implicit low-dimensional PCA-of-distance-maps representation of shape variations, we additionally derive from that an explicit space-conditioned probability model by way of the logistic function. Unlike the previous mentioned level set approaches, the space-conditioned probability model is combined with image-coupling and other terms in a unified Bayesian framework. This Bayesian framework accurately identifies shapes not captured by the prior model.

In contrast to other EM approaches [112, 51, 103, 80], our method explicitly models the boundary via the shape model. Consequently, we achieve smooth segmentations without underestimating fine structures; a common problem in EM approaches [80].

To integrate the shape constraints into our method, we first derive the EM approach using the Kullback-Leibler divergence. Note, that Section 3.3.1 already discussed a derivation we could apply to the shape parameters \mathcal{S} . However, this section presents an alternative viewpoint of the EM algorithm. We then define in Section 4.2.2 the relationship between the hidden data \mathcal{T} and the shape space \mathcal{S} . The last Section validates the approach by comparing the automatic generated segmentations of the caudate and thalamus to the ones created by human experts.

4.2.1 Unified Framework for Image Inhomogeneity Correction, Shape Modeling, and Segmentation

We now incorporate a shape atlas based on signed distance maps into the EM approach of Section 2.3.3. This EM approach determines a solution within the complete data model characterized by the observed MR image data \mathcal{I} , the hidden data \mathcal{T} representing the labelmap, and the parameter space defined by the image inhomogeneities \mathcal{B} . The shape constraints \mathcal{S} can be added to the complete model as an extension of the hidden data or the parameter space. Similarly to the image inhomogeneities \mathcal{B} , \mathcal{S} can be extracted from the hidden data \mathcal{T} and vice versa. This dependency would define a challenging estimation problem if the corresponding hidden data is composed by the labelmap \mathcal{T} and the shape \mathcal{S} . This leads us to the conclusion that the shape \mathcal{S} should be added to the parameter space.

Without additional assumptions, it is a complex problem to extract the inhomogeneities \mathcal{B} and the shape parameters \mathcal{S} from the MR images \mathcal{I} . However, this problem is greatly simplified when formulated as an incomplete data problem and a solution is iteratively determined via EM. Within this framework, we define the following maximum *a posteriori* probability estimation problem:

$$(\hat{\mathcal{S}}, \hat{\mathcal{B}}) = \arg \max_{\mathcal{S}, \mathcal{B}} \log P(\mathcal{S}, \mathcal{B} | \mathcal{I}). \quad (4.5)$$

Similar to Equation (3.6), this results in a system of equations for which there is no analytical solution. To simplify the problem, we combine the observed image intensities \mathcal{I} with the unknown true segmentation \mathcal{T} . If \mathcal{T} is known, \mathcal{B} and \mathcal{S} are more easily estimated from \mathcal{I} . Using the Kullback-Leibler divergence $D[\cdot || \cdot]$ (or relative entropy) [56] we improve the estimates $(\mathcal{S}', \mathcal{B}')$ of $(\hat{\mathcal{S}}, \hat{\mathcal{B}})$ of Equation (4.5) though

$$(\mathcal{B}'', \mathcal{S}'') = \arg \max_{\mathcal{S}, \mathcal{B}} D[P(\mathcal{T} | \mathcal{I}, \mathcal{S}, \mathcal{B}) || P(\mathcal{T} | \mathcal{I}, \mathcal{S}', \mathcal{B}')] + \log P(\mathcal{S}, \mathcal{B} | \mathcal{I}). \quad (4.6)$$

This update rule guarantees that the new estimates $(\mathcal{B}'', \mathcal{S}'')$ are equal or better than

the old $(\mathcal{S}', \mathcal{B}')$ due to the properties of the Kullback-Leibler divergence

$$\begin{aligned}
\log P(\mathcal{S}'', \mathcal{B}''|\mathcal{I}) &\stackrel{1}{=} -D[P(\mathcal{T}|\mathcal{I}, \mathcal{S}'', \mathcal{B}'')\|P(\mathcal{T}|\mathcal{I}, \mathcal{S}'', \mathcal{B}'')] + \log P(\mathcal{S}'', \mathcal{B}''|\mathcal{I}) \\
&\stackrel{2}{\geq} -D[P(\mathcal{T}|\mathcal{I}, \mathcal{S}'', \mathcal{B}'')\|P(\mathcal{T}|\mathcal{I}, \mathcal{S}', \mathcal{B}')] + \log P(\mathcal{S}'', \mathcal{B}''|\mathcal{I}) \\
&\stackrel{3}{\geq} -D[P(\mathcal{T}|\mathcal{I}, \mathcal{S}', \mathcal{B}')\|P(\mathcal{T}|\mathcal{I}, \mathcal{S}', \mathcal{B}')] + \log P(\mathcal{S}', \mathcal{B}'|\mathcal{I}) \\
&\stackrel{4}{=} \log P(\mathcal{S}', \mathcal{B}'|\mathcal{I}).
\end{aligned}$$

To see how the above derivation relates to an instance of the EM algorithm we substitute the divergence with its definition of the expected value

$D[P(Y)\|P(Z)] \triangleq E_Z \left(\log \frac{P(Y)}{P(Z)} \right)$ and apply Bayes' rule, so that Equation (4.6) changes to

$$\begin{aligned}
(\mathcal{S}'', \mathcal{B}'') &= \arg \max_{\mathcal{S}, \mathcal{B}} E_{\mathcal{T}|\mathcal{I}, \mathcal{S}', \mathcal{B}'} \left(\log \frac{P(\mathcal{T}|\mathcal{I}, \mathcal{S}, \mathcal{B})}{P(\mathcal{T}|\mathcal{I}, \mathcal{S}', \mathcal{B}')} \right) + \log P(\mathcal{S}, \mathcal{B}|\mathcal{I}) \\
&= \arg \max_{\mathcal{S}, \mathcal{B}} E_{\mathcal{T}|\mathcal{I}, \mathcal{S}', \mathcal{B}'} \left(\log \frac{P(\mathcal{T}, \mathcal{S}, \mathcal{B}|\mathcal{I})}{P(\mathcal{T}|\mathcal{I}, \mathcal{S}', \mathcal{B}')} \right),
\end{aligned} \tag{4.7}$$

which is the M-Step of an EM approach as discussed in Section 3.3.1.

Even though finding the solution to the maximum *a posteriori* probability estimation problem of Equation (4.7) is simpler than to Equation (4.6), we still need to simultaneously update the shape parameters \mathcal{S} and the image inhomogeneities \mathcal{B} . To split Equation (4.7) into two separate maximum *a posteriori* probability estimation problems, we return to the derivations of Section 3.3.1 as Equation (3.11) is equivalent to Equation (4.7) when replacing the registration parameter \mathcal{R} with the shape parameter \mathcal{S} . Based on Equation (3.12), Equation (4.7) is equivalent to

$$(\mathcal{S}', \mathcal{B}') \leftarrow \arg \max_{\mathcal{B}, \mathcal{S}} E_{\mathcal{T}|\mathcal{I}, \mathcal{S}', \mathcal{B}'} (\log P(\mathcal{I}|\mathcal{T}, \mathcal{S}, \mathcal{B}) + \log P(\mathcal{S}|\mathcal{T}, \mathcal{B}) + \log P(\mathcal{B}|\mathcal{T})) \tag{4.8}$$

¹The Kullback-Leibler divergence is zero for two equivalent distributions: $D[P(y)\|P(y)] = 0$

²The Kullback-Leibler divergence is non-negative for any two distributions

³ $(\mathcal{S}'', \mathcal{B}'')$ maximizes Equation (4.6)

⁴The Kullback-Leibler divergence is zero for two equivalent distributions: $D[P(y)\|P(y)] = 0$

Similar to Section 3.3.1, the optimization procedure decomposes nicely as a consequence of the following independence assumptions: First, we assume independence of \mathcal{I} with respect to \mathcal{S} conditioned on \mathcal{T} and \mathcal{B} because knowing \mathcal{I} and \mathcal{B} is enough to produce \mathcal{T} . Next, we assume independence of \mathcal{S} with respect to \mathcal{B} conditioned on \mathcal{T} , as the image inhomogeneities are caused by the radio frequency coil of the scanner and the shape of a structure is defined by the subject itself. Finally, we assume independence of \mathcal{B} with respect to \mathcal{T} , and \mathcal{T} with respect to the voxel x . As a result Equation (4.8) simplifies to

$$(\mathcal{S}', \mathcal{B}') \leftarrow \arg \max_{\mathcal{S}, \mathcal{B}} \sum_x E_{\mathcal{T}_x | \mathcal{I}, \mathcal{S}', \mathcal{B}'} [\log P(\mathcal{I} | \mathcal{T}_x, \mathcal{B}) + \log P(\mathcal{S} | \mathcal{T}_x)] + \log P(\mathcal{B})$$

As the previous equation is equivalent to Equation (3.14) the EM approach is defined according to Section 3.3.1 by:

- **E-Step** calculates the weights \mathcal{W} :

$$\mathcal{W}_x(a) \triangleq E_{\mathcal{T}_x | \mathcal{I}, \mathcal{S}', \mathcal{B}'}(\mathcal{T}_x(a)) = \frac{P(\mathcal{I}_x | \mathcal{T}_x(a) = 1, \mathcal{B}'_x) \cdot P(\mathcal{T}_x(a) = 1 | \mathcal{S}')}{P(\mathcal{I}_x | \mathcal{B}'_x, \mathcal{S}')} \quad (4.9)$$

- **M-Step** updates the approximations of $(\mathcal{B}', \mathcal{S})$ by finding the solution to the following two maximum estimation problems

$$\mathcal{S}' \leftarrow \arg \max_{\mathcal{S}} \sum_x \sum_a \mathcal{W}_x(a) \cdot \log P(\mathcal{T}_x = e_a | \mathcal{S}) + \log P(\mathcal{S}) \quad (4.10)$$

$$\mathcal{B}' \leftarrow \arg \max_{\mathcal{B}} \sum_x \sum_a \mathcal{W}_x(a) \log P(\mathcal{I} | \mathcal{T}_x = e_a, \mathcal{B}_x) + \log P(\mathcal{B}) \quad (4.11)$$

Equation (4.11) defines the approximation of the image inhomogeneities \mathcal{B}' based on the current definition of the weights \mathcal{W} . The solution to this problem was derived in Section 2.3.3. Equation (4.10) defines the approximation of the shape parameters \mathcal{S} , which strongly depends on the modeling assumptions of the relationship between the labelmap \mathcal{T} and the shape parameters \mathcal{S} expressed by the conditional probability $P(\mathcal{T}_x = e_a | \mathcal{S})$. The next section will define a distribution for $P(\mathcal{T}_x = e_a | \mathcal{S})$ and approximate the solution to Equation (4.10).

4.2.2 Estimating the Shape

This section defines the distribution of the conditional labelmap probability $P(\mathcal{T}_x = e_a | \mathcal{S})$ as a logistic function on distance maps. The logistic function is favorable for our EM approach as it provides an implicit representation of the shape and an explicit space conditioned probability model. In addition, the resulting model captures a broader class of shapes than those methods that are restricted to the PCA model.

To determine the solution to Equation (4.10) we first define the relationship of the unknown true segmentation \mathcal{T} and the shape parameter \mathcal{S} captured by the conditional probability $P(\mathcal{T}_x = e_a | \mathcal{S})$. The task is not straightforward because unlike active contour methods such as [100, 60, 115], we also model the hidden labelmap \mathcal{T} and the image inhomogeneities \mathcal{B} explicitly. The shape parameters \mathcal{S} capture global characteristics of structures, while \mathcal{T} and \mathcal{B} characterize local properties. In addition, the relationship between \mathcal{S} and \mathcal{T} is generally too complicated for finding a closed form solution to Equation (4.10). We model the relationship between \mathcal{S} and \mathcal{T} by a non-Gaussian distribution, for which a solution to Equation (4.10) can be approximated.

In our model, this relationship is captured in the unknown labelmap probability $P(\mathcal{T}_x = e_a | \mathcal{S})$ conditioned on the shape parameter \mathcal{S} (which corresponds to a signed distance map.) Since the random variable \mathcal{T}_x is discrete, we can further simplify the definition of the conditional probability in terms of a generic shape function $\mathcal{A}(\cdot, \cdot)$ to

$$P(\mathcal{T}_x = e_a | \mathcal{S}) \equiv \frac{\mathcal{A}(a, \mathcal{D}_{\mathcal{S},a}(x))}{\sum_{a'} \mathcal{A}(a', \mathcal{D}_{\mathcal{S},a'}(x))}$$

Given the motivation above, a natural choice for this formulation is the logistic function

$$\mathcal{A}(a, v) \equiv \frac{1}{1 + e^{-c_a v}},$$

which maps the distance map to the range $[0,1]$. $\mathcal{A}(a, \mathcal{D}_{\mathcal{S},a}(\cdot))$ depends on c_a , which captures the certainty of the method with respect to the shape model. As shown

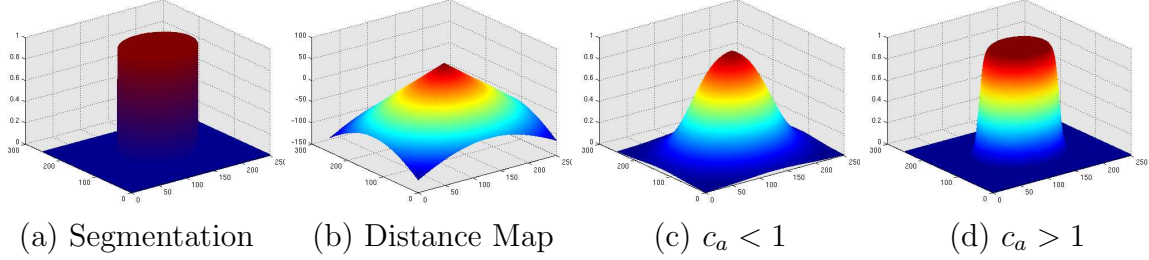


Figure 4-5: Image (a) is the segmentation of a circle whose distance map is shown in (b). Based on the distance map, two different logistic functions $\mathcal{A}(a, \cdot)$ are plotted in (c) and (d). The logistic function of (c) is defined by a large slope ($c_a < 1$) and the plot of (d) represents a logistic function with a steep one slope ($c_a > 1$.)

in the example of Figure 4-5, uncertainty about the shape model is represented by relatively small values of c_a . This results in a wide slope of the spatial distribution, which allows greater mobility of the boundary. Large c_a defines spatial priors with steep slopes, which tend to fixate the boundary of a structure. Currently, c_a is set manually but we would ultimately like to include c_a into the parameter space so that EM approach determines the value.

The probability of the labelmap conditioned on the shape is defined as

$$P(\mathcal{T}_x = e_a | \mathcal{S}) \equiv \frac{\mathcal{A}(a, \mathcal{D}_{\mathcal{S}, a}(x))}{\sum_{a'} \mathcal{A}(a', \mathcal{D}_{\mathcal{S}, a'}(x))} = \frac{\frac{1}{1+e^{-c_a \mathcal{D}_{\mathcal{S}, a}(x)}}}{\sum_{a'} \frac{1}{1+e^{-c_{a'} \mathcal{D}_{\mathcal{S}, a'}(x)}}}. \quad (4.12)$$

Equation (4.12) models the relationship between the labelmap \mathcal{T} , which is defined on the lattice of the image space, and the shape parameters \mathcal{S} , which are defined within the space spanned by the Eigenvector of the PCA analysis. In other words, the probability defines the relationship of an anatomical structure a at a voxel location x with respect to the global parameter \mathcal{S} capturing the shape constraints of the boundaries across structures.

To integrate this shape model into our EM approach, we substitute the conditional probability of the unknown labelmap \mathcal{T} with respect to the shape parameters \mathcal{S} of Equation (4.12) in the maximum *a posteriori* probability estimation problem of

Equation (4.10)

$$\begin{aligned}
\mathcal{S}' &\leftarrow \arg \max_{\mathcal{S}} \sum_x \sum_a \mathcal{W}_x(a) \cdot \log P(\mathcal{T}_x = e_a | \mathcal{S}) + \log P(\mathcal{S}) \\
&= \arg \max_{\mathcal{S}} \sum_x \left[\sum_a \mathcal{W}_x(a) \cdot \log (\mathcal{A}_a(\mathcal{D}_{\mathcal{S},a}(x))) - \log \left(\sum_{a'} \mathcal{A}_{a'}(\mathcal{D}_{\mathcal{S},a'}(x)) \right) \right] \\
&\quad - \frac{1}{2} \mathcal{S}^t \hat{\Lambda}^{-1} \mathcal{S} \\
&= \arg \min_{\mathcal{S}} \sum_x \left[\sum_a \mathcal{W}_x(a) \cdot \log (1 - e^{-c_a \mathcal{D}_{\mathcal{S},a}(x)}) + \log \left(\sum_{a'} \frac{1}{1 - e^{-c_{a'} \mathcal{D}_{\mathcal{S},a'}(x)}} \right) \right] \\
&\quad + \frac{1}{2} \mathcal{S}^t \hat{\Lambda}^{-1} \mathcal{S}
\end{aligned}$$

According to Section 4.2.1 the EM approach is defined as

E-Step Calculates the weights \mathcal{W} :

$$\mathcal{W}_x(a) = \mathcal{W}_x(a) = \frac{\frac{1}{\sqrt{|\Upsilon_a|}} e^{-\frac{1}{2}(\mathcal{I}_x - \mathcal{B}'_x - \mu_a)^t \Upsilon_a^{-1} (\mathcal{I}_x - \mathcal{B}'_x - \mu_a)} \frac{1}{1 - e^{-c_a \mathcal{D}_{\mathcal{S}',a}(x)}}}{\sum_{a'} \frac{1}{\sqrt{|\Upsilon_{a'}|}} e^{-\frac{1}{2}(\mathcal{I}_x - \mathcal{B}'_x - \mu_{a'})^t \Upsilon_{a'}^{-1} (\mathcal{I}_x - \mathcal{B}'_x - \mu_{a'})} \frac{1}{1 - e^{-c_{a'} \mathcal{D}_{\mathcal{S}',a'}(x)}}}$$

where (μ_a, Ψ_a) define the Gaussian intensity distribution of structure a .

M-Step Updates the parameters $(\mathcal{S}', \mathcal{B}')$ based on the weights \mathcal{W} .

The inhomogeneity \mathcal{B} is approximated by applying the simple low pass filter \mathcal{J} , represented by a large matrix, to the weighted residual (see Section 2.3.3):

$$\mathcal{B}' \leftarrow \mathcal{J} \cdot \sum_a \mathcal{W}_x(a) \Upsilon_a^{-1} (\mathcal{I}_x - \mu_a)$$

The shape parameters \mathcal{S}' are updated by

$$\begin{aligned}
\mathcal{S}' &\leftarrow \arg \min_{\mathcal{S}} \sum_x \left[\sum_a \mathcal{W}_x(a) \cdot \log (1 - e^{-c_a \mathcal{D}_{\mathcal{S},a}(x)}) + \log \left(\sum_{a'} \frac{1}{1 - e^{-c_{a'} \mathcal{D}_{\mathcal{S},a'}(x)}} \right) \right] \\
&\quad + \frac{1}{2} \mathcal{S}^t \hat{\Lambda}^{-1} \mathcal{S},
\end{aligned} \tag{4.13}$$

for which a solution is found by using Powell's method [82].




Figure 4-6: A 3D model generated by EM-Shape of the right (red) and left thalamus (purple,) right (green) and left caudate (blue,) and the right (pink) and left ventricle (turquoise.)

4.2.3 Validation

This section focuses on the impact of two different types of atlases on the robustness of the EM approach. We will compare the performance of EM-Affine - guided by spatial priors - to our new approach EM-Shape - constrained by a shape atlas. Both methods use the same initial affine alignment from the atlas to the image space to guarantee a fair comparison between the different types of atlases. As in Section 3.3.4, both methods segment 22 test cases into the three brain tissue classes - white matter, gray matter and corticospinal fluid. The ventricles (pink/turquoise) are extracted from the corticospinal fluid as shown in Figure 4-6. In addition, the parcellation of the gray matter into thalamus (red/purple) and the caudate (green/blue) is compared to the manual segmentations using the DICE measure.

We showed in Section 3.3.4 that EM-Affine can accurately segment the thalamus but has problems identifying the caudate. The initial affine registration is too constrained to properly adjust the spatial priors to the horn-shaped caudate. The goal of this experiment is to find out if the shape constrained EM-Shape can accurately segment both anatomical structures. This would indicate that the variations of anatomical structures are better captured in a shape atlas than by spatial priors.

Figure 4-7 shows the average DICE measures and standard error for the two methods with respect to the thalamus and caudate. For the thalamus both approaches




Figure 4-7: The graph shows the average DICE score and standard error over 22 cases for EM-Affine, which is guided by spatial priors, and the new method EM-Shape, which is constrained by a shape atlas. EM-Shape is superior to EM-Affine in segmenting the caudate because the variations of the caudate are better captured in the shape atlas than the spatial priors.

perform equally well ($87.3 \pm 1.2\%$; mean DICE score \pm standard error.) For the caudate, however, EM-Shape ($84.9 \pm 1.4\%$) is significantly more accurate than EM-Affine ($83.2 \pm 1.7\%$) as the shape variation of the horns of the caudate are better captured by the shape atlas than the spatial priors.

Generally, EM-Shape starts with a lower DICE score than EM-Affine as the initial shape model does not properly describe the patient specific structures. For example, Figure 4-8 shows the outcome of EM-Shape after every fifth iteration. Initially, the segmentation is noisy, which indicates discrepancy between the initial shape model defined by the mean shape and the patient specific shape. With each iteration, the arch of the caudate widens and the segmentations get smoother. After 20 iterations the method converges to a solution that generally outperforms EM-Affine.

As mentioned, it is difficult to determine the exact shape of a structures with weakly visible boundaries. The shape of the thalamus, for example, is oval with a hook attached to it (Figure 4-6 and Figure 4-8.) From the MR images, the size of the ellipse and the position of the hook are often not clearly defined. The top-left image of Figure 4-9 shows an example of such a scenario. The segmentations are the results




Figure 4-8: The 3D models are based on the segmentations generated by our new method through 20 iterations. The method is initialized with the mean shape of each structure. The noisy initial segmentation is an indication of the disagreement between the mean and the patient specific shape. As the algorithm proceeds the shape of the caudate and thalamus adjusts to the patient specific situation. After about 20 iterations the algorithm converges to a smoother segmentation.

of the two automatic segmentation methods where black indicates the outline of the human expert. In this example, EM-Affine underestimates the hook of the thalamus, which we found to be true throughout this experiment. EM-Shape can better cope with this problem as the shape model adds global constraints to the local analysis of the intensities. An example of a global constraint is the explicit definition of shape dependencies across anatomical structures. This causes the shape of the thalamus to be proportional to one of the easily segmentable ventricles. This impacts the accuracy of EM-Shape as it further constrains the space of possible segmentations.

The other structure of interest in this experiment is the caudate, which is defined by long, thin horns (Figure 4-8.) The structure is adjacent to the putamen, another subcortical structure with an identical intensity distribution. In the MR image of the middle column of Figure 4-9 the putamen is located on the outside of image. Neither the intensity pattern nor the spatial prior can properly separate these two structures, as indicated by the noisy segmentations of EM-Affine. The outliers visible in EM-Affine violate the shape constraints of EM-Shape as the boundary has to satisfy the conditions set by the ventricles and the thalamus.

For both structures, EM-Affine did not adequately segment the ends of the structure. In the right column of Figure 4-9 EM-Affine underestimates the tip of the caudate. The opposite is true for the thalamus where EM-Affine overestimates the

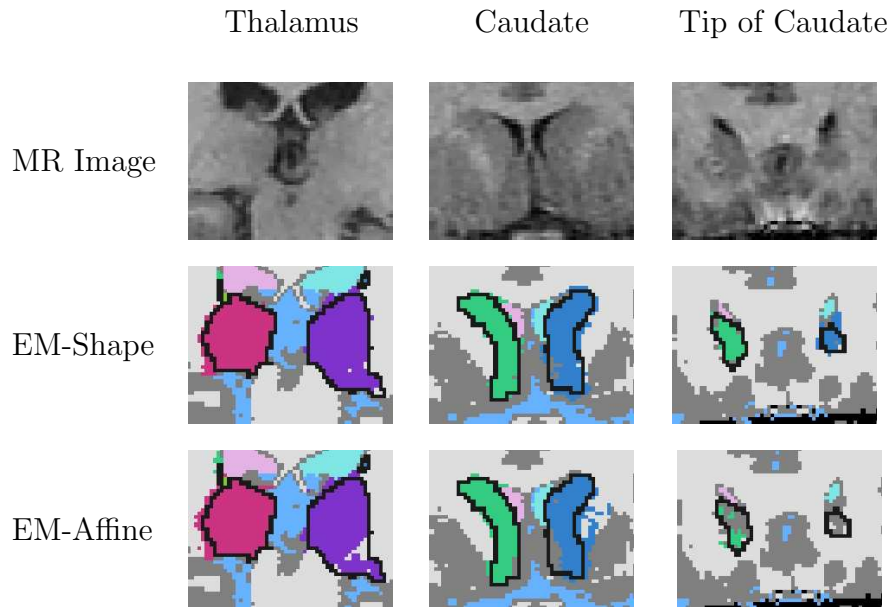


Figure 4-9: The figure is a collection of different subcortical regions. The black lines in the automatic segmentations are the thalamus or caudate outlined by the human expert. The left column shows a MR image with corresponding segmentations of the oval shaped body of the thalamus with attached hook. The middle column unveils part of the caudate, which is adjacent to the putamen, another subcortical structure with identical intensity distribution. The right column presents the tip of the caudate, which generally is underestimated by EM-Affine. In all three examples, EM-Affine performs worse than EM-Shape because the shape prior better adopts to the boundary of the structure.

ends. Again, spatial and intensity distributions do not allow discrimination between anatomical structures in this area.

To better understand the drawbacks of EM-Shape, Figure 4-10 shows the average error over the 22 test cases throughout the volume. The blurriness in the images represents the spatial variation within the average error. Blue indicates an underestimation and red an overestimation of the approach with respect to the manual segmentations.

With respect to EM-Shape, (Figure 4-10 (a)) shows an overestimation of the caudate's superior and underestimation of its inferior surface. (Figure 4-10 (c)) displays the highest underestimation in the region of the thalamus' hook. These artifacts are properly caused by the PCA based shape model. PCA defines the shape variations in a linear space characterized by quadratic cost function. The substantial bending of

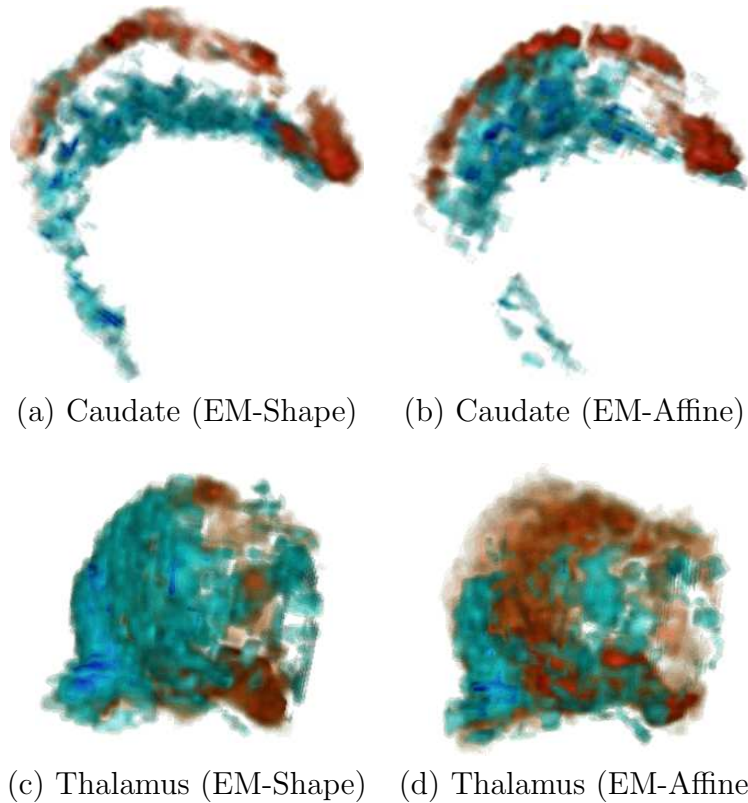


Figure 4-10: Average error of the EM-Shape and EM-Affine with respect to the caudate (a+b) and thalamus (c+d.) The red areas indicate an overestimation where the blue areas show an underestimation of the structures. The blurriness in the image represents the variations within the error. We refer to the text for a detailed discussion.

the caudate’s horns, however, are difficult to capture in a linear model. We therefore conclude that the PCA model is not perfect for representing anatomical shapes. The alternative prior of this experiment, spatial priors, causes even greater average under- and overestimation (Figure 4-10 (b) and (d).) Especially the caudate’s inferior surface is greatly underrepresented in the segmentations of EM-Affine (Figure 4-10 (b)) compared to the results of EM-Shape (Figure 4-10 (a).)

In summary, this section developed a statistical framework for the segmentation of anatomical structures in MR images. The approach is especially well suited for structures with weakly visible boundaries as it simultaneously estimates the image inhomogeneities, explicitly models the boundaries through a deformable shape model, and segments the MR images into anatomical structures. Our approach was validated by automatically segmenting 22 sets of MR images. We then compared our results to

a similar EM approach that does not use shape priors. In general, our new method performs significantly better. The improvement is due primarily to shape constraints on the boundary.

4.3 Modelling Atlas Registration and Shape Deformation within a EM Segmentation Method

The derivation of the EM framework in the previous Section was based on the assumption that the shape atlas is aligned to the MR images. In order to achieve this alignment, we first register the atlas to the image coordinate system and then use the EM approach to partition the MR images into anatomical structures. This two step process was originally introduced in Section 3.2.1, where we needed to align spatial priors to the image space.

We showed in Section 3.3 that we can increase the accuracy of our automatic segmentations if we perform the atlas registration within an EM segmenter. Motivated by these observations, Section 4.3.1 extends the estimation problem of the previous section so that it simultaneously registers the shape atlas to the patient, detects the shape of a specific anatomical structure, segments the image into the structures of interest, and performs an image inhomogeneity correction.

Similar to Section 3.3 we propose a hierarchical registration framework to capture the deformation between the shape atlas and the coordinate system of the image. In Section 3.3, this framework was composed by global and structure specific affine registration parameters. In this section, however, the global registration parameters will represent a rigid mapping as the PCA shape atlas of Section 4.1.1 is incompatible to scaling. The structure specific deformations are now captured by the shape parameters \mathcal{S} . This results in a registration model, which is related to the class of non-rigid registration algorithms incorporating structure specific deformation models.

Section 4.3.3 revisits the experiment of Section 4.2.3 segmenting 22 subjects into the three major brain tissue classes as well as the ventricles, the thalamus, and the cau-

date. We compare the accuracy of our approach with respect to other EM approaches by measuring the overlap to manual segmentations of the subcortical structures.

4.3.1 A Formulation of Shape Based Segmentation and Registration

The previous Section defined the parameter space by the shape model \mathcal{S} and the image inhomogeneities \mathcal{B} . This section extends this space with the registration parameters \mathcal{R} . The optimal setting $(\hat{\mathcal{B}}, \hat{\mathcal{S}}, \hat{\mathcal{R}})$ for these three instances is defined by the solution to the following maximum *a posteriori* probability estimation problem, where we are given the image data \mathcal{I} , and marginalize over all possible labelmaps \mathcal{T} :

$$(\hat{\mathcal{R}}, \hat{\mathcal{S}}, \hat{\mathcal{B}}) = \arg \max_{\mathcal{R}, \mathcal{S}, \mathcal{B}} \log \left(\sum_{\mathcal{T}'} P(\mathcal{R}, \mathcal{S}, \mathcal{B}, \mathcal{T}' | \mathcal{I}) \right)$$

As it is not clear how to find a solution to this problem, we simplify the search for the solution by deriving an EM framework. For this purpose, we determine a lower bound for the objective function via Jensen's Inequality

$$\begin{aligned} (\mathcal{R}', \mathcal{S}', \mathcal{B}') &\leftarrow \arg \max_{\mathcal{R}, \mathcal{S}, \mathcal{B}} E_{\mathcal{T} | \mathcal{I}, \mathcal{R}', \mathcal{S}', \mathcal{B}'} (\log P(\mathcal{R}, \mathcal{S}, \mathcal{B}, \mathcal{T} | \mathcal{I})) \\ &\stackrel{5}{=} \arg \max_{\mathcal{R}, \mathcal{S}, \mathcal{B}} E_{\mathcal{T} | \mathcal{I}, \mathcal{R}', \mathcal{S}', \mathcal{B}'} (\log P(\mathcal{R}, \mathcal{S}, \mathcal{B} | \mathcal{T}, \mathcal{I}) + \log(P(\mathcal{T} | \mathcal{I}))) \\ &\stackrel{6}{=} \arg \max_{\mathcal{R}, \mathcal{S}, \mathcal{B}} E_{\mathcal{T} | \mathcal{I}, \mathcal{R}', \mathcal{S}', \mathcal{B}'} (\log P(\mathcal{R}, \mathcal{S}, \mathcal{B} | \mathcal{T}, \mathcal{I})) \end{aligned}$$

Applying Bayes' rule similar to Equation (4.8), the previous equation expands to

$$\begin{aligned} (\mathcal{R}', \mathcal{S}', \mathcal{B}') &\leftarrow \arg \max_{\mathcal{R}, \mathcal{S}, \mathcal{B}} E_{\mathcal{T} | \mathcal{I}, \mathcal{R}', \mathcal{S}', \mathcal{B}'} (\log P(\mathcal{I} | \mathcal{T}, \mathcal{R}, \mathcal{S}, \mathcal{B}) \\ &\quad + \log P(\mathcal{R}, \mathcal{S} | \mathcal{T}, \mathcal{B}) + \log P(\mathcal{B} | \mathcal{T})) \end{aligned} \tag{4.14}$$

Equation (4.14) defines an estimation problem, which is still too complicated to be

⁵Bayes' rule

⁶ $P(\mathcal{T} | \mathcal{I})$ does not depend on \mathcal{R} , \mathcal{S} , or \mathcal{B} .

solved as the maximum over the entire parameter space has to be found. In order to simplify the problem, we first restate the independence assumptions of Section 4.2.1. As such, the image intensities \mathcal{I} are independent of the shape parameter \mathcal{S} conditioned on \mathcal{T}, \mathcal{R} , and \mathcal{B} . In addition, the image inhomogeneities \mathcal{B} are independent from the labelmap \mathcal{T} so that Equation (4.14) simplifies to

$$(\mathcal{R}', \mathcal{S}', \mathcal{B}') \leftarrow \arg \max_{\mathcal{R}, \mathcal{S}, \mathcal{B}} E_{\mathcal{T}|\mathcal{I}, \mathcal{R}', \mathcal{S}', \mathcal{B}'} (\log P(\mathcal{I}|\mathcal{T}, \mathcal{R}, \mathcal{B}) + \log P(\mathcal{R}, \mathcal{S}|\mathcal{T}, \mathcal{B})) + \log P(\mathcal{B})$$

The solution to this estimation problem strongly depends on the coordinate system in which it is computed. For instance, if we perform the calculations in the atlas space, then the registration parameters \mathcal{R} map the images \mathcal{I} into the atlas coordinate system. We could simply use the hierarchical affine registration approach of Section 3.3.2 to define the correspondence between atlas and image space. However, since \mathcal{B} is defined in the atlas space, the likelihood $P(\mathcal{I}|\mathcal{T}, \mathcal{R}, \mathcal{B})$ is not independent of \mathcal{R} . As such, a closed form solution no longer exists for the image inhomogeneities \mathcal{B} .

Instead, if we solve the estimation problem in the image space, then \mathcal{I} is independent of \mathcal{R} conditioned on \mathcal{T} and \mathcal{B} according to Section 3.3.1. The registration parameters \mathcal{R} now define the mapping from the atlas to the image space. We constrain \mathcal{R} to the class of rigid transformations as the distance maps defining the atlas are generally not invariant to scaling or local deformations. In addition, the first mode of variation of the shape atlas normally encodes implicitly the scaling of the structure. If we further assume independence of \mathcal{R} and \mathcal{S} to the image inhomogeneities \mathcal{B} conditioned on the labelmap \mathcal{T} , of the registration parameters \mathcal{R} with respect to the shape parameters \mathcal{S} , and voxel wise independence within the labelmap \mathcal{T} then the

problem simplifies to

$$\begin{aligned}
(\mathcal{R}', \mathcal{S}', \mathcal{B}') &\leftarrow \arg \max_{\mathcal{R}, \mathcal{S}, \mathcal{B}} E_{\mathcal{T}|\mathcal{I}, \mathcal{R}', \mathcal{S}', \mathcal{B}'} (\log P(\mathcal{I}|\mathcal{T}, \mathcal{B}) + \log P(\mathcal{R}, \mathcal{S}|\mathcal{T})) \\
&\quad + \log P(\mathcal{B}) \\
&\stackrel{7}{=} \arg \max_{\mathcal{R}, \mathcal{S}, \mathcal{B}} E_{\mathcal{T}|\mathcal{I}, \mathcal{R}', \mathcal{S}', \mathcal{B}'} [\log P(\mathcal{I}|\mathcal{T}, \mathcal{B}) + \log P(\mathcal{T}|\mathcal{R}, \mathcal{S})] \\
&\quad + \log P(\mathcal{R}, \mathcal{S}) + \log P(\mathcal{B}) \tag{4.15} \\
&\stackrel{8}{=} \arg \max_{\mathcal{R}, \mathcal{S}, \mathcal{B}} \sum_x E_{\mathcal{I}_x|\mathcal{I}, \mathcal{R}', \mathcal{S}', \mathcal{B}'} [\log P(\mathcal{I}|\mathcal{T}_x, \mathcal{B}) + \log P(\mathcal{T}_x|\mathcal{R}, \mathcal{S})] \\
&\quad + \log P(\mathcal{R}) + \log P(\mathcal{S}) + \log P(\mathcal{B}).
\end{aligned}$$

Note that the update rule comprises two separate maximum *a posteriori* probability estimation problems, one with respect to the parameter \mathcal{B} and the other with respect to the parameter $(\mathcal{R}, \mathcal{S})$. This separation greatly simplifies the problem.

The explicit formulation of the EM algorithm based on Equation (4.15) is characterized by the E-Step, which calculates the weights \mathcal{W} capturing the probability of structure a being assigned to x ,

$$\begin{aligned}
\mathcal{W}_x(a) &\triangleq E_{\mathcal{I}_x|\mathcal{I}, \mathcal{R}', \mathcal{S}', \mathcal{B}'}(\mathcal{T}_x(a)) = P(\mathcal{T}_x = e_a|\mathcal{I}, \mathcal{R}', \mathcal{S}', \mathcal{B}') \\
&\stackrel{9}{=} \frac{P(\mathcal{I}|\mathcal{T}_x = e_a, \mathcal{R}', \mathcal{S}', \mathcal{B}') \cdot P(\mathcal{T}_x = e_a|\mathcal{R}', \mathcal{S}', \mathcal{B}')}{P(\mathcal{I}|\mathcal{R}', \mathcal{S}', \mathcal{B}')} \\
&\stackrel{10}{=} \frac{P(\mathcal{I}|\mathcal{T}_x = e_a, \mathcal{B}') \cdot P(\mathcal{T}_x = e_a|\mathcal{R}', \mathcal{S}')}{P(\mathcal{I}|\mathcal{R}', \mathcal{S}', \mathcal{B}')} \\
&\stackrel{11}{=} \frac{P(\mathcal{I}|\mathcal{T}_x = e_a, \mathcal{B}') \cdot P(\mathcal{T}_x = e_a|\mathcal{R}', \mathcal{S}')}{\sum_{a'} P(\mathcal{I}|\mathcal{T}_x = e_{a'}, \mathcal{B}') \cdot P(\mathcal{T}_x = e_{a'}|\mathcal{R}', \mathcal{S}')}
\end{aligned}$$

⁷Apply Bayes' rule,; $\log P(\mathcal{R}, \mathcal{S})$ is independent from \mathcal{T} .

⁸Spatial independence of \mathcal{T} ; independence between \mathcal{S} and \mathcal{R} .

and the M-Step, which improves the estimates $(\mathcal{R}', \mathcal{S}', \mathcal{B}')$ of the parameter space,

$$\begin{aligned} (\mathcal{R}', \mathcal{S}') \leftarrow \arg \max_{\mathcal{R}, \mathcal{S}} \sum_x \sum_a \mathcal{W}_x(a) \cdot \log P(\mathcal{I}_x = e_a | \mathcal{R}, \mathcal{S}) \\ + \log P(\mathcal{R}) + \log P(\mathcal{S}) \end{aligned} \quad (4.16)$$

$$\mathcal{B}' \leftarrow \arg \max_{\mathcal{B}} \sum_x \sum_a \mathcal{W}_x(a) \cdot \log P(\mathcal{I} | \mathcal{I}_x = e_a, \mathcal{B}) + \log P(\mathcal{B}). \quad (4.17)$$

The solution to Equation (4.17) is solved according to Section 3.2. The remainder of this section focuses on the definition of the conditional probability $P(\mathcal{I}_x = e_a | \mathcal{R}, \mathcal{S})$ of Equation (4.16), which models the relationship of the labelmap with respect to the registration and shape parameters.

4.3.2 A Non-Rigid Registration Model

The hierarchical affine registration approach of Section 3.3.3 is not necessary anymore when using a shape atlas. Instead, we propose a rigid registration model, where the registration parameters \mathcal{R} represent the rigid global mapping between the atlas and the image space. The local deformations are captured by shape parameters \mathcal{S} . $(\mathcal{R}, \mathcal{S})$ therefore represents a non-rigid registration framework of the mean shape that is constrained by the structure specific shape model (see Figure 4-11.)

In order to properly characterize a rigid mapping between the atlas and image space, the registration parameters $\mathcal{R} = (\vec{t}^t, \vec{r}^t)^t$ are composed by $\vec{t} \in \mathbb{R}^3$ representing the displacement and $\vec{r} \in \mathbb{R}^3$ representing the angles of the rotations along the three axis. The interpolation function $r(\cdot, \cdot)$ captures the rigid mapping between two coordinate systems

$$r(\cdot, \cdot) : \mathbb{R}^{6 \times 3} \rightarrow \mathbb{R}^3, (\mathcal{R}, x) \rightarrow A_{\mathcal{R}} \cdot (x^t, 1)^t$$

where $A_{\mathcal{R}} \in \mathbb{R}^{4 \times 4}$ is the rigid transformation matrix based on the parameter \mathcal{R} .

⁹Apply Bayes' rule

¹⁰We extend the previously made independence assumption of $P(\mathcal{R}, \mathcal{S} | \mathcal{T}, \mathcal{B})$ with respect to \mathcal{B} to the conditional probability $P\mathcal{T} | \mathcal{R}, \mathcal{S}, \mathcal{B}$

¹¹Apply Bayes' rule




Figure 4-11: The graph shows the design of our non-rigid registration framework. In this approach, the registration parameter \mathcal{R} captures the global alignment of the shape based atlas to the image space. The local deformations are captured by the shape parameters \mathcal{S} .

We now apply this concept to the estimation problem of Equation (4.16). We focus on the definition of the condition probability $P(\mathcal{T}_x = e_a | \mathcal{R}, \mathcal{S})$ as it describes the relationship between shape parameters \mathcal{S} , registration parameters \mathcal{R} , and the labelmap \mathcal{T} . In Section 4.2, the conditional labelmap probability with respect to the shape $P(\mathcal{T}_x = e_a | \mathcal{S})$ was defined by the aligned shape atlas, which is represented by the distance map $\mathcal{D}_{\mathcal{S},a}$. In our current model $\mathcal{D}_{\mathcal{S},a}$ is given in the atlas space. As the interpolation function $r(\mathcal{R}, \cdot)$ maps a voxel from the image space to the corresponding coordinates in the atlas space, $\mathcal{D}_{\mathcal{S},a}(r(\mathcal{R}, \cdot))$ represents the aligned distance map in the image space. Based on Equation (4.12), we define the conditional probability of the labelmap with respect to the registration parameters \mathcal{R} and the shape model \mathcal{S} as

$$P(\mathcal{T}_x = e_a | \mathcal{S}, \mathcal{R}) = \frac{\frac{1}{1+e^{-c_a \mathcal{D}_{\mathcal{S},a}(r(\mathcal{R}, x))}}}{\sum_{a'} \frac{1}{1+e^{-c_{a'} \mathcal{D}_{\mathcal{S},a'}(r(\mathcal{R}, x))}}} \quad (4.18)$$

Note, that if \mathcal{R} is the identity matrix then Equation (4.18) defines the spatial distribution of the structure a in the atlas induced by the shape parameters \mathcal{S} .

With the help of Equation (4.13), we reformulate our original maximization problem of Equation (4.16) to define the following update rule of the M-Step

$$\begin{aligned}
(\mathcal{R}', \mathcal{S}') &\leftarrow \arg \max_{\mathcal{R}, \mathcal{S}} \sum_x \sum_a \mathcal{W}_x(a) \cdot \log P(\mathcal{T}_x = e_a | \mathcal{R}, \mathcal{S}) + \log P(\mathcal{S}) \\
&= \arg \max_{\mathcal{R}, \mathcal{S}} \sum_x \left[\sum_a \mathcal{W}_x(a) \cdot \log \left(1 - e^{-c_{a'} \mathcal{D}_{\mathcal{S},a'}(r(\mathcal{R}, x))} \right) \right. \\
&\quad \left. + \log \left(\sum_{a'} \frac{1}{1 - e^{-c_{a'} \mathcal{D}_{\mathcal{S},a'}(r(\mathcal{R}, x))}} \right) \right] + \frac{1}{2} \mathcal{S}^t \hat{\Lambda}^{-1} \mathcal{S}.
\end{aligned}$$

We dropped the prior of the registration parameters \mathcal{R} as it is a uniform distribution and therefore does not influence the outcome of the optimization problem.

We must note that the implementation used for our test results in the next section slightly differs from the proposed framework. First, we separated the search for \mathcal{R} and \mathcal{S} as it made our implementation five times faster. Second, we do not determine a closed form solution for the above equation but instead find an approximation for the estimation problem via Powell’s method [82]. Despite these changes, the principle idea behind the method has not been altered as the shape constrained method simultaneously segments the image and registers the atlas to the image space.

4.3.3 Validation

The experiment described in this section compares the new approach to other EM implementations of this thesis. The first two methods applied to this experiment are the pipeline based approaches EM-Affine and EM-NonRigid of Section 3.2.1, which initialize the EM segmenter by aligning spatial priors to the image space using affine and non rigid registration methods. The third implementation, EM-Simultaneous-Affine of Section 3.3, performs the alignment of the spatial atlas and the segmentation of the medical images simultaneously. The fourth algorithm, EM-Shape of Section 4.2, uses the same alignment strategy as EM-Affine but guides the segmentation via a shape model. Our new approach, EM-Simultaneous-Shape, uses the shape deformation concept of EM-Shape but also registers the shape model to the image space. As in Section 4.2.3 each method segments the 22 test cases into the major brain tissue classes and extracts the ventricles, the thalamus, and the caudate. We measure the accuracy of each method’s resulting automatic segmentations by comparing them to manual segmentations using the volume overlap measure DICE of Section 1.5.

The results of the DICE comparison are presented in Figure 4-12 (a). Our new approach EM-Simultaneous-Shape (EM-Sim-Shape) outperforms the other four methods on the thalamus and caudate. In specifics, the standard error of EM-Simultaneous-Shape is much smaller than EM-Shape because the global rigid alignment by EM-Simultaneous-Shape more accurately aligns the shape atlas to the image space than




Figure 4-12: Graph (a) shows the mean and standard error for the four methods segmenting the thalamus and caudate in 22 cases. Our new approach EM-Simultaneous-Shape (EM-Sim-Sh) achieves a higher accuracy than any of the other approaches with respect to automatically segmenting the thalamus and caudate. Graph (b) shows the average convergence behavior for the caudate of the four methods. EM-Simultaneous-Affine (EM-Sim-Affine) and EM-Simultaneous-Shape (EM-Sim-Shape) start with a bad initial score but converge to a solution, which is closer to the manual segmentations than EM-Affine and EM-Shape.

than initial global affine registration of EM-Shape. EM-Simultaneous-Shape outperforms EM-Simultaneous-Affine because the shape prior better captures the variability of these anatomical structures than the structure specific affine registration parameters of EM-Simultaneous-Affine. Based on the previous two observations it is therefore not surprising that the two pipeline based approaches EM-Affine and EM-NonRigid receive a lower score than EM-Simultaneous-Shape. These methods cannot recover from initial misalignment errors (similar to EM-Shape) and have to rely on the spatial prior (like EM-Simultaneous-Affine.)

For the remainder of this section we will exclude EM-NonRigid from the discussion as it is the only method relying on a more general non-linear global registration. For both structures, our joint registration and segmentation methods (EM-Integrated and EM-Hybrid) achieve a higher average mean score than the methods EM-Affine and EM-Shape, which sequentially perform this task. This increased performance is mainly caused by two test subjects, for which the affine registration approach (used by EM-Affine and EM-Shape) performs poorly. The first row in Figure 4-13 shows




Figure 4-13: The figure shows the manual and automatic segmentations for one case where the initial affine registration of EM-Affine and EM-Shape as well as the non-rigid registration of EM-NonRigid performed suboptimally. The black lines in the second row outline the manual segmentation for this slice. Comparing those with the automatic segmentations, an offset in the vertical direction of the segmentations produced by EM-Affine and EM-Shape is clearly visible. This offset is caused by poor initial registration of the atlas to the image space. EM-Nonrigid slightly overestimates the subcortical structures in this area as it cannot recover from any misalignment performed in the initial non rigid registration method.

the 3D segmentation results of all five methods in one of the two subjects. The second row shows an example slice of these segmentations indicating in black the boundary of the manual segmentation. An offset in the vertical direction is visible when comparing EM-Affine and EM-Shape to the manual segmentations. This offset is caused by the poor initial alignment.

Figure 4-14 shows the progress of EM-Simultaneous-Shape with respect to this specific case. The first row lists the segmentations starting with an initial bad guess and converging to a good approximation of the manual segmentation outlined in black. This observation is also captured in the objective function of the registration parameters (second row) and shape parameters (third row.) Red indicates large disagreement between atlas and image space. With each iteration, the number of voxels in red reduces, implying a better alignment of the atlas to the segmentation environment.

The graph of Figure 4-12(b) indicates that the previous observation holds for the 22 test cases. The graph lists the average DICE score at each iteration with respect to




Figure 4-14: This sequence shows the improvement of our new method EM-Simultaneous-Shape in the case of Figure 4-13. Initially, our method starts with a bad initial guess which is evident in the objective function of the registration parameters (second row) and the shape parameters (third row.) The color red in the objective functions indicates large disagreement between the atlas and the segmentation environment defined by the weights of the E-Step. At each iteration, the method improves its estimate of the alignment parameters which is also indicated by the reduction in red areas in the objective function. The final segmentation accurately identifies the structures of interest as it mostly agrees with the manual segmentation outlined in black.

caudate for four methods. In general, the two pipeline based approaches EM-Affine and EM-Shape start with a good initial guess but do not greatly improve the results. The opposite is true for EM-Simultaneous-Affine and EM-Simultaneous-Shape, which start with a poor initial segmentation to converge to a solution close to the manual ones. This increase in performance is due to our modelling approach based on the principle of least commitment between atlas registration, inhomogeneity correction, and image segmentation.

In summary, we developed a joint registration and segmentation framework, which was guided by the image inhomogeneity correction as well as the shape model. In order to compute a solution within this framework, we slightly simplified the model

and developed an EM approach to estimate its solution. We then compared our approach to other EM approaches, which do not simultaneously register and deform a shape model. The new approach achieves a higher accuracy than any of the other approaches in segmenting the thalamus as well as the caudate. In general, our new approach underlines the advantages of the EM framework as it turns an even a bad initialization within the complex segmentation model into a good solution.

4.4 Discussion

This chapter described the integration of shape constraints into an EM approach to increase the discriminatory power of the model with respect to the segmentation of brain structures with weakly visible boundaries. Unlike any other approaches in medical imaging, our framework simultaneously estimates the shape of multiple brain structures as well as the inhomogeneities of the image. To estimate the shape variations across a population, we use distance maps to describe the shape within the training set of manual segmentations. We first derive an approach defining the parameter space of the EM model with a shape atlas based on signed distance maps and the estimation of the image inhomogeneity. We then extend this concept to register the shape atlas to the MR images while segmenting the MR images into the structures of interest. Both implementations generate high quality segmentations of structures with weakly visible boundaries. Furthermore, we validate their accuracy by comparing them to other EM-like methods as well as manual segmentations.

The advantages of our approaches are numerous. First, our novel approaches are not restricted to variations defined by the shape atlas, which is normally under-trained. Instead, the algorithms define a boundary based on the intensity pattern, spatial distribution, and the shape constraints. Unlike many other shape based methods, our approaches may model patient specific abnormalities not captured in the shape atlas. Second, the generic design of our approaches is not customized towards a specific structure so that the automatic method can simultaneously segment brain tissue types as well as their substructures. Third, the approaches can be applied to a




Figure 4-15: The 3D models are the segmentation results of EM-Simultaneous-Shape after a certain number of iterations. They represent the thalamus (dark green/purple,) the caudate (blue/light green) , and the ventricles (orange/yellow.) The noisy initial segmentation is caused by the misalignment of the atlas to the image space and the incorrect deformation of the shape model. As the method progresses the quality of the segmentation improves. The experiment of Section 4.3.3 showed that the final segmentation is generally more accurate than the result of other EM approaches. The greater accuracy is due to the underlying uniform framework which simultaneously registers and deforms the shape atlas to the subject, determines the image inhomogeneities, and segments the images into structures of interest.

variety of imaging protocols as they simultaneously estimate image inhomogeneities as well as shape constraints. This results in a very robust segmentation approach that can recover from a bad initial guess to converge to a good solution as shown in Figure 4-15.

Chapter 5

Hierarchical Segmentation

The preceding two chapters focused on outlining structures with weakly visible boundaries. However, most segmentation problems also include boundaries that are clearly defined in MR images. For example in Figure 5-1, the algorithm has to determine the invisible boundary between two neighboring cortical structures and the visible boundary of the sulci. While spatial priors are essential for the determination of the invisible border, they reduce the accuracy of the outlining of the sulci. The smooth spatial priors under-represent sulci as discussed in detail in Section 3.1. This misrepresentation increases the likelihood of bridges or cavities across adjacent banks of sulci. Partial voluming in MR images, particularly with low-resolution T2-weighted data as shown in Figure 5-1, further amplifies this effect. One solution to this problem is a hierarchical approach, where the method divides the problem into much simpler subproblems.

Most hierarchical methods in the field [65, 34, 74] solve the problem through a sequence of separate operations. However, the dependency between the different approaches can make it difficult to adjust these methods to new segmentation scenarios. Instead, this chapter suggests a method where the division of the problem is decomposed into the segmentation of substructures having mutual constraints. These constraints are defined by the hierarchical relationship between anatomical structures and the influence of the prior information within the segmentation of these structures.

Section 5.1 focuses on the development of a segmentation algorithms guided by the




Figure 5-1: The graph shows the MR images, the corresponding automatic segmentation, and spatial priors of the region around the middle temporal gyrus. Based on T1- as well as T2-weighted data, the method relies on the spatial prior to detect the boundary between middle (blue) and superior temporal gyrus (red). The spatial priors, however, cause the algorithm to smooth over the clearly visible sulci between the two structures. This is due to the method’s inability to properly individualize the prior model to the segmentation problem.

hierarchical relationship of anatomical structures. The section starts with a method combining the class of segmentation methods related to EM with a data tree representing the anatomical relationships between structures. Afterwards, we alter an EM approach to determine the solution for the subproblems defined by the tree.

Section 5.2 expands the algorithm by adjusting the influence of the intensity and spatial prior information to each subtree of the tree. This regularization enables the method to accurately segment structures with weakly visible and clearly defined boundaries.

Section 5.3 tests the algorithm on a variety of experiments. The first experiment segments the brain into 32 structures. We then compare our new approach to other segmentation methods for the detection of the three major tissue classes of the brain. The third experiment is a repetition of the multi-rater experiment of Section 3.2.2 determining the accuracy of an EM approach with and without a hierarchical framework. For the last experiment, the joint registration and segmentation algorithms described in Section 3.3 and Section 4.3 are integrated into the hierarchical approach automatically segmenting the thalamus and caudate. The advantages of our new approach are discussed in Section 5.4.

5.1 An Anatomically Guided EM Approach

The section develops a hierarchical approach, which divides the segmentation problem into simpler subproblems. This division process is guided by a tree that represents the anatomical dependencies between structures. The solution to each subproblem is a labelmap computed by a segmentation approach. The final segmentation of the original problem is a combination of all the labelmaps of the corresponding subproblems.

The example of Figure 5-2 illustrates the dependency between the complexity of the problem and the accuracy of an EM approach. The smooth segmentation to the left was produced by an algorithm separating the left superior temporal gyrus from the rest of the image. The fragmented segmentation to the right was generated by the same approach but the algorithm outlines all the neighboring structures of the superior temporal gyrus, i.e. white matter, corticospinal fluid, and other cortical structures. The more complicated second problem generally results in less accurate segmentation by the EM approach.

In section 5.1.1 we develop a method guided by a tree, which represents the hierarchical relationship between anatomical structures. Section 5.1.2 discusses the population of prior information throughout the tree and Section 5.1.3 adopts the EM approaches of this thesis to the tree structure.




Figure 5-2: The graph shows the superior temporal gyrus and the corresponding segmentations. The image to the left is a “binary” segmentation, where the method separates the image into superior temporal gyrus versus non-superior temporal gyrus. The image to the right is produced by automatically segmenting the region into multiple structures, where the method separates the image into superior temporal gyrus, white matter, corticospinal fluid, and the other cortical structures. The lower accuracy of the multiple structure setting is caused by an increase in the statistical complexity of the segmentation problem.

5.1.1 Represent Anatomical Hierarchy in a Tree

We now introduce an algorithm, which is guided by the hierarchical relationship of anatomical structures of interest. This relationship is modelled as a tree as shown in Figure 5-3. In this tree X is a descendent of Y , when an anatomical structure X (e.g. white matter) is part of another structure Y (e.g. ICC.) The theory of trees is discussed in detail in Cormen et al. [18].

The segmentation algorithm proposed in this section uses the tree to recursively divide the initial segmentation problem into simpler subproblems (see also Algorithm 3.) It does so by first segmenting the region of interest associated with the root of the tree into the children of the root. Afterwards, a subproblem is defined for each child with a nonempty subtree. The subproblem associated with a child is composed by a subtree (TREE) rooted in the child and a region of interest (ROI). The region of interest is defined by those voxels assigned to the child in the current segmentation. This process of segmentation and dividing the problem into subproblems is repeated




Figure 5-3: The tree is an example of the hierarchical ordering between different anatomical structures. Here, background (BG) and intracranial capacity (ICC) define the tree’s top level while corticospinal fluid (CSF), gray matter (GM), and white matter (WM) compose the second level.

until the entire tree is traversed.

Algorithm 3: HIERARCHICAL SEGMENTATION($ROI, TREE$)

define labelmap $\mathcal{L} \leftarrow$ Segment ROI into structures of the $TREE$
for each child a of the root of the $TREE$ further segment ROI
 into the children of a
define $ROI' \leftarrow$ Each voxel with $\mathcal{L}_x = a$
define $TREE' \leftarrow$ The subtree rooted in a
if $ROI' \neq \emptyset$ **and** $TREE' \neq \emptyset$
then HIERARCHICAL SEGMENTATION($ROI', TREE'$)

For example, if the hierarchical approach defined in Algorithm 3 is guided by the tree of Figure 5-3 then the segmentation problem is divided into two subproblems. It first segments the region of interest composed by the entire MR image into the children background (BG) and intracranial capacity (ICC). The second problem is defined with respect to the intracranial capacity. The region of interest for this problem are all the voxels associated to intracranial capacity (in red) and the subtree is defined by the root ICC and the children corticospinal fluid (CSF), white matter(WM), and gray

matter(GM). After determining the solution to this problem the tree is traversed and the final labelmap is defined by the three children of the previous tree and the background.

5.1.2 Populate Information within the Tree

A class of segmentation methods, such as [34, 65, 103], can be integrated into Algorithm 3. The implementation of an EM segmenter into Algorithm 3 requires the population of the prior model to each of the subproblems. Doing so is not without risk, as the prior information might not be independent from each other. This might increase the amount of training amount necessary to properly define the model. For example, in the previous scenario (see Figure 5-3) we could define different atlases for the intracranial capacity and its substructures. This framework would increase the redundancy in our model as the substructures should contain the information that defines intracranial capacity. In addition, each task could be run separately by first separating the the image into background and intracranial capacity (Level 1) and then separating the intracranial capacity into its tissue classes (Level 2.) Essentially, two algorithms would have to be trained as is common in the field [65, 34]. This would limit the impact of the tree structure within the segmentation algorithm. We therefore further constrains Algorithm 3 to overcome this deficiency.

Instead of defining the atlas for each structure within the tree, the hierarchical segmentation approach propagates the information from the leafs to the internal nodes. In the previous scenario, the atlas of the ICC would be defined by the atlases of the corticospinal fluid, gray matter, and white matter. This concept is integrated into Algorithm 3 by defining two structure specific lists. The first list captures all children of an anatomical structure a and is defined as $a_{\text{SUB}} \triangleq \{a' | a \text{ is parent of } a'\}$. According to the example of Figure 5-3 the list of substructures for structure IMAGE is $\text{IMAGE}_{\text{SUB}} \triangleq \{\text{BG}, \text{ICC}\}$. The second list captures all the leaves in the subtree of

a structure, which is defined as

$$\begin{aligned} a_{\widehat{\text{SUB}}} &\triangleq \{a' | a'_{\text{SUB}} = \emptyset \wedge a' \text{ is on a descendent of } a\} \\ &= \{a' | a'_{\text{SUB}} = \emptyset \wedge (\exists a_1, \dots, a_k : a_1 = a \wedge a_k = a' \wedge a_{i+1} \in a_{i\text{SUB}} \text{ for } i < k)\} \end{aligned}$$

Note, that in the above definition $k-1$ corresponds to the depth of the leaf a' within the subtree of structure a . In the example of Figure 5-3, the list of all the leaves of the structure IMAGE is defined as $\text{IMAGE}_{\widehat{\text{SUB}}} \triangleq \{\text{BG}, \text{WM}, \text{GM}, \text{CSF}\}$.

The structures associated to the nodes of tree can be divided into two groups. The first group are all leaves ($a_{\text{SUB}} = \emptyset$.) They are characterized by the structure specific parameters of the atlas. The second group is composed by the internal nodes ($a_{\text{SUB}} \neq \emptyset$.) which are characterized by the atlas information of all their descendent in $a_{\widehat{\text{SUB}}}$.

Algorithm 4: HIERARCHICAL SEGMENTATION($ROI, TREE$)

define the two lists of substructures $\hat{a}_{\widehat{\text{SUB}}}$ and \hat{a}_{SUB} for root \hat{a} of $TREE$

define weights $\mathcal{W} \leftarrow$ Run EM on ROI and structures defined by $\hat{a}_{\widehat{\text{SUB}}}$

define labelmap $\mathcal{L} \leftarrow$ Assign voxels to structures of \hat{a}_{SUB} based on \mathcal{W}

for each structure a in \hat{a}_{SUB} further segments ROI into substructures

define $ROI' \leftarrow$ Each voxel x with $\mathcal{L}_x = a$

define $TREE' \leftarrow$ The subtree rooted in a

if $ROI' \neq \emptyset$ **and** $TREE' \neq \emptyset$

then HIERARCHICAL SEGMENTATION($ROI', TREE'$)

Algorithm 4 is a possible implementation of the previous described concept. As an extension of Algorithm 3, the approach recursively solves and divides the segmentation problem into subproblems until the entire tree is traversed. At each recursion, the method first defines the list of leaves $\hat{a}_{\widehat{\text{SUB}}}$ and the list of children \hat{a}_{SUB} based on the tree $TREE$ rooted in structure \hat{a} . Afterwards, an EM approach computes the weights \mathcal{W} of the leaves in $\hat{a}_{\widehat{\text{SUB}}}$ with respect to region of interest (ROI) and the

prior information of the leaves. It then turns the weights into a label map \mathcal{L} . These two steps will be discussed in more detail in the next section. Finally, it partitions the segmentation problem into the subproblems with respect to all internal nodes in a_{SUB} . This process is repeated for all the subproblems until the tree is traversed.

5.1.3 Integrate EM Approach into Hierarchical Segmentation

This section focus on the definition of the weights \mathcal{W} and the labelmap \mathcal{L} of Algorithm 4. As the weights are the results of an EM segmenter, we are specifically interested in adjusting the EM approaches of this thesis in order explicitly model the constraints defined by the tree $TREE$ and the region of interest ROI . For this purpose, we now derive a new instance of an EM algorithm that is easily customizable to any of the other EM segmenters of this thesis. We first describe a complete data model with respect the hierarchical concept of Section 5.1.2. Afterwards, we briefly derive an instance of an EM algorithm that determines a solution within this complete data model.

In Algorithm 4, the EM approach determines the weights \mathcal{W} with respect to the region of interest (ROI) and the structures given by $TREE$. In order to adjust our previous EM approaches to these new constraints, we define the method’s underlying complete data model with respect to ROI and $TREE$. The observed data is now composed of the image \mathcal{I} and the hierarchy specific data $\mathcal{H} \triangleq \{ROI, TREE\}$. The labelmap \mathcal{T} embodies the hidden data and Ψ represents the parameter space. To closely relate this model to the ones of the previous chapters, we are interested in finding the solution to the maximum *a posteriori* probability estimation problem of the complete data

$$\hat{\Psi} = \arg \max_{\Psi} \log \left(\sum_{\mathcal{T}} P(\Psi, \mathcal{T} | \mathcal{I}, \mathcal{H}) \right). \quad (5.1)$$

In general, Equation (5.1) describes a set of equations for which there is no analytical solution. This motivates the use of the EM algorithm that iteratively determines a solution for the estimation problem. We will now briefly describe the relationship

between Equation (5.1) and the resulting EM segmenter. More detailed derivations of similar EM frameworks have been presented in Chapter 2 to Chapter 4.

According to the derivations of Section 3.3.1, the corresponding EM approach to Equation (5.1) is characterized by the update rule

$$\begin{aligned} \Psi' &\leftarrow \arg \max_{\Psi} E_{\mathcal{T}|\mathcal{I},\mathcal{H},\Psi'} (\log P(\Psi, \mathcal{T}|\mathcal{I}, \mathcal{H})) \\ &\stackrel{1}{=} \arg \max_{\Psi} E_{\mathcal{T}|\mathcal{I},\mathcal{H},\Psi'} (\log P(\Psi|\mathcal{T}, \mathcal{I}, \mathcal{H}) + P(\mathcal{T}|\mathcal{I}, \mathcal{H})) \\ &\stackrel{2}{=} \arg \max_{\Psi} E_{\mathcal{T}|\mathcal{I},\mathcal{H},\Psi'} (\log P(\Psi|\mathcal{T}, \mathcal{I}, \mathcal{H})). \end{aligned}$$

As in the previous chapters, we assume voxel-wise independence of the labelmap \mathcal{T} so that the update rule changes to

$$\Psi' \leftarrow \arg \max_{\Psi} \sum_x E_{\mathcal{T}_x|\mathcal{I},\mathcal{H},\Psi'} (\log P(\Psi|\mathcal{T}_x, \mathcal{I}, \mathcal{H})). \quad (5.2)$$

Based on Equation (5.2), the E-Step calculates the weights \mathcal{W} , which capture the posterior probability of an anatomical structure a being present at a specific voxel x with respect to the current estimate of the parameter space Ψ' and the image \mathcal{I} :

$$\begin{aligned} \mathcal{W}_x(a) &\triangleq E_{\mathcal{T}_x|\mathcal{I},\mathcal{H},\Psi'}(\mathcal{T}_x(a)) = P(\mathcal{T}_x = e_a|\mathcal{I}, \mathcal{H}, \Psi') \\ &= \frac{P(\mathcal{I}|\mathcal{T}_x = e_a, \Psi', \mathcal{H}) \cdot P(\mathcal{T}_x = e_a|\Psi', \mathcal{H})}{P(\mathcal{I}|\Psi', \mathcal{H})} \end{aligned} \quad (5.3)$$

Note, that the conditional probability $P(\mathcal{T}_x = e_a|\Psi', \mathcal{H})$ was represented in the previous chapters by the structure specific atlases, which were explicitly defined by the leaves of the tree. The integration of this constraint in our current implementation is done in the E-Step by calculating the weights for all leaves $a \in \hat{a}_{\widehat{\text{SUB}}}$ of the root \hat{a} .

The M-Step determines an improved estimate Ψ' based on the weights \mathcal{W} and

¹Bayes' rule

² $P(\mathcal{T}|\mathcal{I}, \mathcal{H})$ does not depend on Ψ .

Equation (5.2), which is defined as

$$\Psi' \leftarrow \arg \max_{\Psi} \sum_x \sum_{a \in \hat{a}_{\text{SUB}}} \mathcal{W}_x(a) \cdot \log P(\Psi | \mathcal{I}_x = e_a, \mathcal{I}, \mathcal{H}). \quad (5.4)$$

Up to this point, the EM approach defined by Equation (5.3) (E-Step) and Equation (5.4) (M-Step) only takes limited advantage of the hierarchical relationships defined by TREE as it computes the weights for all leaves $a \in \hat{a}_{\text{SUB}}$. To change this, we now simplify the solution to the maximum *a posteriori* probability estimation problem in the M-Step by just considering the *first degree* children $\tilde{a} \in \hat{a}_{\text{SUB}}$ of the root \hat{a} in Equation (5.4). We do so by making the simplifying assumption that the posterior probability $P(\Psi | \mathcal{I}_x = e_a, \mathcal{I}, \mathcal{H})$ is the same for all structures $a \in \tilde{a}_{\text{SUB}}$ that have the same ancestor $\tilde{a} \in \hat{a}_{\text{SUB}}$. We then define the hierarchical weights $\tilde{\mathcal{W}}$ for all $\tilde{a} \in \hat{a}_{\text{SUB}}$ by the weights of all descendent of $a \in \tilde{a}_{\text{SUB}}$:

$$\tilde{\mathcal{W}}_x(\tilde{a}) \triangleq \begin{cases} \sum_{a \in \tilde{a}_{\text{SUB}}} \mathcal{W}_x(a') & , \tilde{a}_{\text{SUB}} \neq \emptyset \\ \mathcal{W}_x(a) & , \text{ otherwise.} \end{cases} \quad (5.5)$$

This allows us to reduce Equation (5.4) to

$$\Psi' \leftarrow \arg \max_{\Psi} \sum_x \sum_{\tilde{a} \in \hat{a}_{\text{SUB}}} \tilde{\mathcal{W}}_x(\tilde{a}) \cdot \log P(\Psi | \mathcal{I}_x = e_{\tilde{a}}, \mathcal{I}, \mathcal{H}). \quad (5.6)$$

The choice between the update rule defined by Equation (5.4) or Equation (5.6) solely depends on the parameter Ψ . In practice, we have found it useful to apply Equation (5.4) for updating the current estimates of the image inhomogeneity \mathcal{B}' , shape parameter \mathcal{S}' , and structure specific registration parameters \mathcal{R}'_C as the simplifying assumption incorporated into Equation (5.6) negatively impacts the accuracy of the method.

This observation is not true, however, for the global registration parameter \mathcal{R}_G , which is intensity and structure independent. In order to increase the computational performance of our algorithm, we update the registration parameter \mathcal{R}'_G by Equation (5.6). Even for our simple example of Figure 5-3 finding the solution to

The objective function therefore considers fewer voxels, which reduces the complexity of determining a solution for the estimation problem.

Before we complete the discussion of the EM approach in this section, we provide an additional interpretation of the hierarchical weight $\widetilde{\mathcal{W}}$ defined by Equation (5.5). We do so by substituting the weight $\mathcal{W}_x(a)$ of Equation (5.5) with its definition of Equation (5.3), which formalizes the hierarchical weight for a structure \tilde{a} with children $a \in \tilde{a}_{\widehat{\text{SUB}}}$ and normalization term Z_x as

$$\widetilde{\mathcal{W}}_x(a) = \frac{1}{Z_x} \sum_{a \in \tilde{a}_{\widehat{\text{SUB}}}} P(\mathcal{I}_x | \mathcal{T}_x = e_a, \Psi', \mathcal{H}) \cdot P(\mathcal{T}_x = e_a | \Psi', \mathcal{H}).$$

To relate the current EM approach to the solutions of the previous chapters, the likelihood $P(\mathcal{I}_x | \mathcal{T}_x = e_a, \Psi', \mathcal{H})$ is defined by a Gaussian distribution. For structures with children the soft assignment of structure $\widetilde{\mathcal{W}}_x(\tilde{a})$ is therefore a non-stationary Gaussian mixture model with respect to the intensities \mathcal{I} where the posterior probability $P(\mathcal{T}_x = e_a | \Psi', \mathcal{H})$ defines the relative weight within the model.

Equation (5.5) captures the notion that large structures with children, such as gray matter, are defined by spatially varying intensity patterns. In the previous chapters, we implicitly modeled this characteristic by the slowly varying image intensity correction or intensity inhomogeneities \mathcal{B} . With the hierarchical model, we can incorporate this constraint explicitly in the atlas space and thereby increase the realism of our model.

After integrating an EM method into the hierarchical approach, we focus on turning the hierarchical weights $\widetilde{\mathcal{W}}$ into the region specific segmentation \mathcal{L} . In order to adopt Algorithm 4 to the new definition of $\widetilde{\mathcal{W}}$, the segmentation of a specific ROI is now defined as

$$\mathcal{L}(x) \triangleq \arg \max_{\tilde{a} \in \tilde{a}_{\widehat{\text{SUB}}}} \widetilde{\mathcal{W}}_x(\tilde{a}).$$

Determining the labelmap \mathcal{L} , Algorithm 4 further divides \mathcal{L} by recursively applying itself to each structure \tilde{a} with non empty subtree $\tilde{a}_{\widehat{\text{SUB}}}$. The process is repeated until all leaves of the tree are reached.

In summary, this section integrated an EM approach into a hierarchical segmen-

tation model. A possible implementation of this hierarchical segmentation method is presented in Algorithm 5. The algorithm partitions an initial difficult problem into less complicated subproblems. The deviation is guided by a tree representing the hierarchical relationship between anatomical structures. Within the tree, prior information is populated from the leaves to the ancestors in order to reduce the training effort necessary to adjust the model to a different segmentation scenario. In the next section, we focus on further specifying Algorithm 5 in order to implementing constraints into the hierarchical framework that explicitly control the influence of the prior information on the segmentation results.

Algorithm 5: HIERARCHICAL SEGMENTATION(\mathcal{H})

define $TREE$ and ROI based on \mathcal{H}

define the two lists of substructures $\hat{a}_{\widehat{SUB}}$ and \hat{a}_{SUB} for root \hat{a} of $TREE$

repeat EM iteration for structures defined by $a \in \hat{a}_{\widehat{SUB}}$ and $\tilde{a} \in \hat{a}_{SUB}$

The E-Step: Calculate $\mathcal{W}_x(a) = \frac{1}{Z_x} P(\mathcal{I}_x | \Psi', \mathcal{T}_x = e_a, \mathcal{H}) \cdot P(\mathcal{T}_x = e_a | \Psi', \mathcal{H})$

$\widetilde{\mathcal{W}}_x(\tilde{a}) \leftarrow$ Define hierarchical weights according to Equation (5.5)

The M-Step: Update $\Psi' \leftarrow \arg \max_{\Psi} \sum_x \sum_{\tilde{a} \in \hat{a}_{SUB}} \widetilde{\mathcal{W}}_x(\tilde{a}) \cdot \log P(\Psi | \mathcal{T}_x = e_{\tilde{a}}, \mathcal{I}, \mathcal{H})$

until Ψ' converges in ROI

define labelmap $\mathcal{L} \leftarrow$ Assign voxels to structures of $\tilde{a} \in \hat{a}_{SUB}$ based on $\widetilde{\mathcal{W}}$

for each structure \tilde{a} in \hat{a}_{SUB} further segment ROI into substructures

define $ROI' \leftarrow$ Each voxel x with $\mathcal{L}_x = \tilde{a}$

define $TREE' \leftarrow$ The subtree rooted in \tilde{a}

if $ROI' \neq \emptyset$ **and** $TREE' \neq \emptyset$ **then**

define $\mathcal{H}' \leftarrow \{ROI', TREE'\}$

HIERARCHICAL SEGMENTATION(\mathcal{H}')

5.2 Incorporate Prior Information into the Tree

Algorithm 5 is a segmentation algorithm guided by the hierarchical relationship of anatomical structures. The approach reduces the computational complexity of segmentation problems by partitioning them into small tasks and explicitly defining the spatially varying intensity pattern of large structures. The solution for each subproblem is determined by an EM approach. However, the hierarchical approach does not define the influence of the prior information with respect to the subproblem so that it fails to properly address the issue of outlining structures with weakly visible and clearly defined boundaries. This section addresses this issue by introducing into the framework confidence parameters that control the importance of the observed data and prior information in the EM approach.

Before we develop this EM framework we make two observations with respect to weights of Equation (5.3), which are defined as

$$\mathcal{W}_x(a) = \frac{1}{Z_x} \cdot P(\mathcal{I}|\mathcal{T}_x = e_a, \Psi', \mathcal{H}) \cdot P(\mathcal{T}_x = e_a|\Psi', \mathcal{H}).$$

The first important component of the product is the conditional intensity probability $P(\mathcal{I}|\mathcal{T}_x = e_a, \Psi', \mathcal{H})$. In the implementations of the previous chapters the dependence of the image \mathcal{I} with respect to the entire parameter space was reduced to the image inhomogeneity \mathcal{B} . We now make the same independence assumptions for our hierarchical model so that the structure specific weights are defined as

$$\mathcal{W}_x(a) = \frac{1}{Z_x} \cdot P(\mathcal{I}|\mathcal{T}_x = e_a, \mathcal{B}', \mathcal{H}) \cdot P(\mathcal{T}_x = e_a|\Psi', \mathcal{H}). \quad (5.7)$$

The second observation concerns the conditional probability $P(\mathcal{T}_x = e_a|\Psi', \mathcal{H})$. The hierarchical independent approaches of this thesis describe $P(\mathcal{T}_x = e_a|\Psi')$ by a spatial prior function, which we name $\mathcal{F}(\Psi, a, \cdot)$. For example, in the approach of Section 4.3 $\mathcal{F}(\Psi, a, \cdot)$ represents the spatial prior of Equation (4.18) and for the EM approach of Section 3.2.1 $\mathcal{F}(\Psi, a, \cdot)$ simplifies to the spatial atlas distribution $f_x(a)$.

After having made these two observations, we focus the remainder of this section on the integration of confidence parameters into the EM method used by Algorithm 5.

5.2.1 Develop Intensity Based Confidence Parameters

The visibility of the structure specific boundaries often depends on the acquisition sequence of the MR images. Many medical imaging protocols acquire several channels of the same modality to properly detect these structures. However, the modalities are often different in terms of resolution, where lower resolutions increase the risk of partial volume effects.

For example, Figure 5-1 shows a T1- and T2-weighted image of the same cortical region. While the T2-weighted image provides good contrast for corticospinal fluid, the T1-weighted image clearly outlines white and gray matter. However, the T2-weighted image has lower spatial resolution compared to T1-weighted image so that the greater risk of partial voluming effect on T2-weighted images might negatively impact the accuracy of our approach. As this is common knowledge in the field, it has motivated researchers to generate two different methods, one for the detection of the outer boundary of the corticospinal fluid through brain stripping, and one for the detection of the boundary between gray and white matter [65]. We suggest instead to model both scenarios in the atlas space where the influence of each input channel or image modality depends on the segmentation task.

In order to regulate the influence of the input channels, we extend the definition of \mathcal{H} by c_i , the influence parameter of the input channel i with respect to the TREE. In order to incorporate c_i into our EM approach, we now modify the conditional intensity probability $P(\mathcal{I}_x|\mathcal{B}_x, \mathcal{T}_x = e_a, \mathcal{H})$ of Equation (5.7). According to the previous chapters $P(\mathcal{I}_x|\mathcal{B}_x, \mathcal{T}_x = e_a, \mathcal{H})$ is defined as a Gaussian distribution of the form

$$P(\mathcal{I}_x|\mathcal{T}_x = e_a, \mathcal{B}_x, \mathcal{H}) \triangleq \frac{1}{\sqrt{(2 \cdot \pi)^n |\Upsilon_a|}} e^{-\frac{1}{2}(\mathcal{I}_x - \mathcal{B}_x - \mu_a)^T \Upsilon_a^{-1} (\mathcal{I}_x - \mathcal{B}_x - \mu_a)}. \quad (5.8)$$

With respect to the example of Figure 5-1, the intensities of \mathcal{I}_x are defined by T1-

and T2-weighted images so that the previous equation represents a two dimensional Gaussian distribution.

To add the influence parameter c_i into this Gaussian model, the diagonal matrix $\gamma_{\mathcal{H}}$ is introduced. The entries in the matrix are defined as $\gamma_{\mathcal{H}}(i, i) = c_i^{-1}$ for $c_i > 0$ and $\gamma_{\mathcal{H}}(i, i) = 0$ otherwise. Applied to the example of Figure 5-3, we set $c_{T2} = 1$ for the Level 1 segmentation as T2-weighted images clearly show the boundary between background and brain. For the Level 2 segmentation, T2-weighted images can be neglected and thus $c_{T2} = 0$. If $c_{T1} = 1$ then the resulting diagonal matrix for the Level 2 segmentation is defined as

$$\gamma_{\mathcal{H}} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$$

The rank of $\gamma_{\mathcal{H}}$ is defined as n' , where n' can be smaller than the number of input channels. This allows us to selectively eliminate certain channels from taking influence in the segmentation process at this level of the tree. However, n' has to be greater than zero because otherwise the algorithm would totally ignore the MR images.

In the remainder of this section we define a variance for Equation (5.8) that takes into account the confidence parameters c_i . Simply multiplying $\gamma_{\mathcal{H}}$ with the structure specific variance Υ_a generally defines a matrix that is not regular, such as in the previous example. However, if we defined the transition matrix

$$\zeta_{\mathcal{H}}(i, j) \triangleq \begin{cases} 1 & , \text{ if } c_i > 0 \text{ and the number of } c_k > 0 \text{ for } k \leq i \text{ is equal to } j \\ 0 & , \text{ otherwise} \end{cases}$$

then we can define a regular matrix based on the product between $\zeta_{\mathcal{H}}$, $\gamma_{\mathcal{H}}$, and Υ_a . Note, that the matrix $\zeta_{\mathcal{H}}$ is composed by all columns in $\gamma_{\mathcal{H}}$ that do have entries greater than zero. In $\zeta_{\mathcal{H}}$, these entries are replaced by one. With respect to our previous

example of Figure 5-3, $\zeta_{\mathcal{H}}$ is defined for the Level 2 segmentation as

$$\zeta_{\mathcal{H}} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

If we now characterize for each structure $a \in \hat{a}_{\text{SUB}}$ the hierarchy specific variance by the regular matrix

$$\Upsilon'_{\mathcal{H},a} \triangleq \zeta_{\mathcal{H}}^T \gamma_{\mathcal{H}} \Upsilon_a \gamma_{\mathcal{H}} \zeta_{\mathcal{H}}$$

then the spatially conditional intensity probability is defined as

$$P(\mathcal{I}_x | \mathcal{T}_x = e_a, \mathcal{B}_x, \mathcal{H}) \triangleq \frac{1}{\sqrt{(2 \cdot \pi)^{n'} |\Upsilon'_{\mathcal{H},a}|}} e^{-\frac{1}{2}(\mathcal{I}_x - \mathcal{B}_x - \mu_a)^T \zeta_{\mathcal{H}} \cdot \Upsilon'_{\mathcal{H},a}^{-1} \zeta_{\mathcal{H}}^T (\mathcal{I}_x - \mathcal{B}_x - \mu_a)} \quad (5.9)$$

Note, that the above likelihood is not a real distribution if the confidence c_i is set to zero for at least one input i . However, setting c_i to zero is equivalent to ignoring input i , so that the conditional intensity probability $P(\mathcal{I}_x | \mathcal{T}_x = e_a, \mathcal{B}_x, \mathcal{H})$ is a distribution in the subspace of input channels with corresponding confidence parameter greater zero. In our previous example the matrix $\zeta_{\mathcal{H}} \Upsilon'_{\mathcal{H},a}^{-1} \zeta_{\mathcal{H}}^T$

$$\begin{aligned} &= \begin{pmatrix} 1 \\ 0 \end{pmatrix}^t \left[\begin{pmatrix} 1 \\ 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \Upsilon_a(1,1) & \Upsilon_a(1,2) \\ \Upsilon_a(2,1) & \Upsilon_a(2,2) \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix}^t \right]^{-1} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \\ &= \begin{pmatrix} 1 \\ 0 \end{pmatrix}^t \Upsilon_a(1,1)^{-1} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} \Upsilon_a(1,1)^{-1} & 0 \\ 0 & 0 \end{pmatrix}. \end{aligned}$$

In conclusion, we have extended the hierarchical framework by the *intensity based influence parameters* c , which controls the importance of the input channel in the resulting segmentation.

5.2.2 Regulate Influence of Spatial Priors

The second important component of the weights \mathcal{W} of Equation (5.7) is the conditional probability $P(\mathcal{T}_x = e_a | \Psi, \mathcal{H})$, which is currently defined to dependent on

the spatial probability $\mathcal{F}(\Psi, a, \cdot)$. Independent of the specific parameter model Ψ , we want to regulate the influence of the spatial priors based on the visibility of the boundary of a structure. We therefore extend the definition of $P(\mathcal{I}_x = e_a | \Psi, \mathcal{H})$ in a similar fashion to the conditional intensity probability of the previous section. We add the parameter $\lambda_{\mathcal{H}} \in [0, 1]$ to the hierarchical parameter definition \mathcal{H} , where $\lambda_{\mathcal{H}}$ defines the influence of the spatial prior in \mathcal{H} . The conditional probability for each structure $a \in \hat{a}_{\text{SUB}}$ are defined as

$$P(\mathcal{I}_x = e_a | \Psi, \mathcal{H}) \triangleq (1 - \lambda_{\mathcal{H}}) + \lambda_{\mathcal{H}} \cdot \mathcal{F}(\Psi, a, x). \quad (5.10)$$

In the example of Figure 5-3, the atlas information is of great importance at the first level when the method robustly outlines the relatively smooth boundary between ICC and BG. We therefore set $\lambda_{\mathcal{H}}$ to 1. At the second level, the atlas causes errors because the large variations of each of the three tissue classes are not fully captured in the atlas so that $\lambda_{\mathcal{H}}$ is set close to 0. In general, $\lambda_{\mathcal{H}}$ is defined by the confidence in the atlas information for a given task, which can be measured by the overlap between the atlas of the structure and training cases.

After defining the conditional intensity distribution and the spatial prior, we revisit the EM segmenter of Section 5.1.3. The E-Step, which calculates the weights \mathcal{W} for each structure $a \in \hat{a}_{\text{SUB}}$ defined by Equation (5.7), is modified with respect to Equation (5.9) and Equation (5.10) to

$$\mathcal{W}_x(a) \triangleq \frac{1}{Z_x} \frac{e^{-\frac{1}{2}(\mathcal{I}_x - \mathcal{B}_x - \mu_a)^T \zeta_{\mathcal{H}} \cdot \Upsilon_{\mathcal{H}, a}^{-1} \zeta_{\mathcal{H}}^T (\mathcal{I}_x - \mathcal{B}_x - \mu_a)}}{\sqrt{(2 \cdot \pi)^n |\Upsilon_{\mathcal{H}, a}|}} \cdot ((1 - \lambda_{\mathcal{H}}) + \lambda_{\mathcal{H}} \mathcal{F}(\Psi, a, x)) \quad (5.11)$$

where Z_x is the normalization term defined by the sum over the product between the conditional intensity probability and conditional spatial probabilities over all structures .

In Section 5.2.1, we altered the definition of the intensity distribution. The changes also impact the closed form solution for the image inhomogeneity calculations in the M-Step. According to Equation (5.9) and Equation (2.27) the estimates of the

inhomogeneities \mathcal{B}' are updated by

$$\mathcal{B}' \leftarrow \mathcal{J} \cdot \sum_{a \in \widehat{a}_{\text{SUB}}} \mathcal{W}_x(a) \zeta_{\mathcal{H}} \cdot \Upsilon_{\mathcal{H},a}^{\prime-1} \zeta_{\mathcal{H}}^T (\mathcal{I}_x - \mu_a).$$

We note that the hierarchical model can also be used to decouple the optimization problems within the M-Step. In the hierarchical registration framework of Section 3.3.3 as well as the joint shape and registration model in 4.3.2, we perform the global registration and the assessment of the local deformation within the same iteration. However, this combination of global and local deformation is not very practical since determining the local deformations is only sensible once the global registration has converged. With the hierarchical representation of this chapter, we overcome this problem by performing a global registration at the top level of the tree. Once converged, we fix the global registration and enable the estimation of the local deformation for each of the subtrees. The experiment in Section 5.3.4 will specifically address this issue by comparing the performance of the two previously mentioned EM approaches segmenting the caudate and thalamus with and without an hierarchical anatomical framework.

In summary, this section presented an algorithm guided by the hierarchical relationship captured in a tree. Each segmentation problem defined by the tree is solved by an EM approach. We altered the EM approaches of the previous chapter so that the outcome of the results would be influenced by the confidence parameters with respect to the prior information. These confidence parameters can be individually adjusted to each subproblem defined by the tree.

As a result, the hierarchical approach is less difficult to scale than the EM approaches of the previous sections because the overall complexity of the problem depends on the hierarchical model and not on the total number of structures involved. Furthermore, information is populated from the leaves to the internal nodes within the hierarchical segmentation approach. This reduces the risk of redundancy in the model and therefore decreases the training effort. In addition, the prior information is customized to each level of the segmentation through the influence parameters of

the input channels and spatial distributions. The following pseudo code summarizes the method.

Algorithm 6: HIERARCHICAL SEGMENTATION(\mathcal{H})

define *TREE* and *ROI* based on \mathcal{H}

define the two lists of substructures \widehat{a}_{SUB} and \hat{a}_{SUB} for root \hat{a} of *TREE*

for all structures a in \widehat{a}_{SUB}

Based on the confidence parameters:

define conditional spatial probability $P(\mathcal{I}_x = e_a | \Psi, \mathcal{H})$

define conditional intensity probability $P_{\hat{a}}(\mathcal{I}_x | \mathcal{I}_x = e_a, \mathcal{B}_x, \mathcal{H})$

repeat EM iteration for structures defined by $a \in \widehat{a}_{\text{SUB}}$ and $\tilde{a} \in \hat{a}_{\text{SUB}}$

The E-Step: $\mathcal{W}_x(a) \leftarrow$ Define weights of structure according to Equation (5.11)

$\widetilde{\mathcal{W}}_x(\tilde{a}) \leftarrow$ Define hierarchical weights according to Equation (5.5)

The M-Step: Update $\Psi' \leftarrow \arg \max_{\Psi} \sum_x \sum_{\tilde{a} \in \hat{a}_{\text{SUB}}} \widetilde{\mathcal{W}}_x(\tilde{a}) \cdot \log P(\Psi | \mathcal{I}_x = e_{\tilde{a}}, \mathcal{I}, \mathcal{H})$

until Ψ' converges in *ROI*

define labelmap $\mathcal{L} \leftarrow$ Assign voxels to structures of $\tilde{a} \in \hat{a}_{\text{SUB}}$ based on $\widetilde{\mathcal{W}}$

for each structure \tilde{a} in \hat{a}_{SUB} further segment *ROI* into substructures

define *ROI'* \leftarrow Each voxel x with $\mathcal{L}_x = \tilde{a}$

define *TREE'* \leftarrow The subtree rooted in \tilde{a}

if *ROI'* $\neq \emptyset$ **and** *TREE'* $\neq \emptyset$ **then**

define $\mathcal{H}' \leftarrow \{\textit{ROI}', \textit{TREE}'\}$

HIERARCHICAL SEGMENTATION(\mathcal{H}')

5.3 Experiments and Validation

This section discusses four different experiments that show the strengths and weaknesses of the new hierarchical segmentation approach. The first three experiments use an EM approach of Section 3.2.1 in combination with Algorithm 6. This approach first registers the spatial priors to the MR images using a non-rigid registration method. In

the second step, the method segments the images into the structures of interest based on Algorithm 6. We further extended the E-Step with the Mean-Field approximation suggested by Kapur [51] to smooth out noise in the segmentation results.

The first experiment tests the capabilities of the method by segmenting the brain into 31 structures. The next two experiments highlight the increase in accuracy of the hierarchical concept with respect methods not implemented into Algorithm 6. We first compare different automatic segmentation approaches by segmenting four brains into gray matter, white matter, and corticospinal fluid. We then experiment repeats the multi-rater validation experiment of the superior temporal gyrus of Section 3.2.2. The experiment compares our new hierarchical method EM-Hierarchy to our previous method EM-NonRigid (Section 3.2.1) as well as manual segmentations.

The final experiment compares the accuracy of the joint segmentation and registration method of Section 3.3 and Section 4.3 integrated into Algorithm 6. As in Section 4.3.3, the accuracy of the 22 segmentations of the subcortical structures are compared to manual segmentations, which we view as ground truth for this experiment. The experiment shows that the hierarchical framework further increases the accuracy of these two methods as we can decouple global registration and local deformation between atlas and image space.

5.3.1 Parcellating 31 Structures

The capabilities of the method are demonstrated by segmenting multiple channel brain MRIs (SPGR - General Electric specific acquisition protocol for Gradient-echo T1-weighted: 256x256x124, resolution of 0.9375mm x 0.9375mm x 1.5 mm, T2-weighted volumes: 256x256x62, resolution of 0.9375mm x 0.9375mm x 3.0 mm) into 31 structures (Figure 5-6.) The hierarchical model of Figure 5-5, which guides the segmentation process, is based on the constraints given to us by the differences in resolution of the MRIs.

Similar to the example discussed in Section 5.1.1, the image is segmented into ICC and Background by relying on atlas information, T2-weighted images, and SPGR images. We define the Background as Air and Skin to properly model the intensity




Figure 5-5: The graph shows a possible hierarchical tree for segmenting a brain MRI into 31 structures. On the first level, the case is separated into Background and ICC. Furthermore, Background is divided into air and skin, and ICC is segmented into white matter (WM), gray matter (GM) and corticospinal fluid (CSF). Corticospinal fluid is parcellated into left ambient cistern (LAC), cerebral aqueduct (CA), fourth ventricle (FOR), third ventricle (THI), left lateral ventricle (LLV), the right ventricle (RVE), and the remaining structures in the corticospinal fluid (CSF*). Gray matter is parcellated into left and right thalamus (LTHA, RTHA), left and right gray matter (LGM, RGM), and the remaining structures of the gray matter (GM*). LGM as well as RGM are further divided into amygdala (LAMY, RAMY), anterior insular cortex (LANT, RANT), hippocampus (LHIP, RHIP), inferior temporal gyrus (LITG, RITG), middle temporal gyrus (LMTG, RMTG), parahippocampus (LPAR, RPAR), posterior insula cortex (LPOS, RPOS), superior temporal gyrus (LSTG, RSTG), and temporal lobe (LTEM, RTEM). The segmentation guided by this hierarchical structure is displayed in Figure 5-6.

distributions of the two structures. The ICC is further separated into corticospinal fluid, gray matter, and white matter. For this problem, the approach ignores the lower resolution T2-weighted data due to the higher risk of partial voluming effect and poor contrast. The method relies on T1-weighted data and lowers the influence of the spatial prior to reduce over-smoothing of sulci.

Further parcellating gray matter and corticospinal fluid into their substructures, the spatial priors of these structures are of great importance due to the low discriminative power of the MRI images in those regions. While we can complete the parcellation of the corticospinal fluid in the next level, the gray matter parcellation is




Figure 5-6: The images show an automatic segmentation of brain MR images into 31 anatomical structures. The segmentation is guided by an anatomical tree, which partitions the segmentation problem into four levels (see Figure 5-5.) This partition reduces the complexity of the segmentation problem by dividing it into simpler subproblems.

further divided into several tasks. This division is partly motivated by the memory constraints of the PC as well as differences in anatomical structures.

While we do not choose the hierarchical model of Figure 5-5 arbitrarily, similar segmentation results as shown in Figure 5-6 can be achieved with very different anatomical trees. As we define the atlas by propagating the prior information from the children to the parent structures, hierarchical models are changed with a minimal amount of training effort. For example, if our implementation is motivated by speed and memory allocation, then we define a binary tree with no more than two structures per level.

While Figure 5-6 shows a robust parcellation of the brain, we currently do not have the capabilities to validate the accuracy of this segmentation. We neither have subjects manually outlined in all structures of interest (for a validation study) nor multiple scans from the same patient (for a reliability study.) However, we will later validate the approach for multiple subcortical and cortical structures using data produced by the same image acquisition sequence.




Figure 5-7: The input data are shown in (a) and (b). Images (c) and (d) capture the reference standard for the various structures. The EM segmentation (e) is characterized by its many outliers, which are normally reduced by atlas information (see (f - h).) The atlas based EM approach (f) and the watershed method (g) smooth over small gyri resulting in an underestimation of corticospinal fluid. This problem is reduced with the new approach (h).

5.3.2 Comparison of Four Different Segmentation Methods

This section compares the segmentations of the hierarchical approach to these of three other implementations. Each method outlines the corticospinal fluid, white matter and gray matter in four MR data sets of the brain (SPGR and T2: 256x256x124, resolution of 0.9375mm x 0.9375mm x 1.5 mm.) This experiment uses a different imaging acquisition sequence than the previous experiments, which underlines the flexibility of our approach.

The first implementation, EM-Only, is very similar to the EM approach of Section 2.3.3, but uses stationary priors and the mean field approximation by Kapur et al. [51]. The mean field approximation generally smooths over noise and generates homogeneous segmentations. The next three approaches further constrain their segmentation model with spatial priors, which are aligned to the patient by the non-rigid registration approach of Section 3.2.1. Like EM-Only, the second implementation, EM-NonRigid of Section 3.2.1, uses the mean field approximation by

Kapur [51]. The third approach, WS-Atlas, is the watershed implementation by Grau [39]. The method turns the image into a topographic relief map, which defines the “peaks” and “valleys” of the image, and then “floods” the relief map to determine the corresponding segmentation.

The fourth method, EM-Hierarchy, incorporates the hierarchical model. Based on EM-NonRigid, EM-Hierarchy is guided by the hierarchical framework defined in Figure 5-3. Unlike the experiment of Section 5.3.1, SPGR and T2-weighted MR images have the same resolution so that the influence of T2-weighted data is not reduced when segmenting brain tissue on the second level.

The quality of each segmentation is defined using the STAPLE method of Section 1.5. To evaluate a segmentation, a reference standard is first estimated using STAPLE on the output of all methods except the one under evaluation. This reference standard captures the highest consensus between the segmentations, which is in general not the ground truth. However, we use this method in the hope that the the highest consensus between the segmentations is close to the ground truth. The consensus of the selected segmentation to the reference standard is defined by DICE, the volume overlap measure of Section 1.5.

Figure 5-8 shows the result of our comparison. Out of the four methods, EM-Only always receives the lowest score. Outliers, normally removed by an atlas, reduce its accuracy (see Figure 5-7 (e).) EM-NonRigid does well on white matter and gray matter but performs poorly on corticospinal fluid because it generally overestimates gray matter and underestimates corticospinal fluid (Figure 5-7 (f).) This problem is caused by the spatial priors that are needed for separating the three tissue classes from the rest of the image. We refer to the beginning of this chapter for more detail.

WS-Atlas (Figure 5-7 (g)) generally does better than EM-NonRigid but is outperformed by EM-Hierarchy for all but one case (white matter in case 4.) EM-Hierarchy (Figure 5-7 (h)) does well on all three tissue classes due to the new hierarchical definition of spatial prior, spatially conditional intensity probability intensity and the influence parameters λ and the intensity confidence parameter c (see Section 5.2.) Unlike the other methods, EM-Hierarchy adjusts the influence of the prior information

with respect to segmentation problem. It is guided by spatial priors for the segmentation of the intracranial capacity but ignores the priors when further segmenting the intracranial capacity into the three major tissue classes.

The influence of the hierarchical model on the segmentation results is very apparent when comparing the segmentation of EM-NonRigid to EM-Hierarchy. EM-NonRigid segments the entire MRI image in one step. It has to compromise between robustly outlining the smooth borders of the inter cranial cavity, which is the boundary of the brain defined by corticospinal fluid, gray and white matter, or the rather convoluted boundary between the corticospinal fluid and gray matter. As the outside and the inside of the inter cranial cavity have similar intensity patterns, the method heavily relies on spatial priors. However, this causes the algorithm to smooth over small sulci (see also picture (f) in Figure 5-7.) In contrast, EM-Hierarchy can properly individualize the atlas information towards both boundaries. It relies on the spatial prior to detect the smooth boundaries of the inter cranial cavity and reduces the influence of the spatial priors for the detection of the boundary between corticospinal fluid and gray matter. In general, this modeling approach reduces the overestimation of cortical gray matter by better individualizing the prior information to the segmentation scenario.

5.3.3 Multiple Rater Experiment of the Superior Temporal Gyrus

This section repeats the multiple rater experiment of Section 3.2.2 to further highlight the accuracy of the hierarchical guided segmentation process. The experiment compares six expert segmentations of the superior temporal gyrus in four cases to the automatically generated results of EM-NonRigid and EM-Hierarchy (see Figure 5-9.) Again, we use the volume overlap measure DICE (Section 1.5) and the Positive Predictive Value (PPV) of Section 1.5. In contrast to the approaches of Section 3.2.2, our new hierarchical approach achieves an accuracy in automatically segmenting the superior temporal gyrus that is comparable to the manual segmentations.

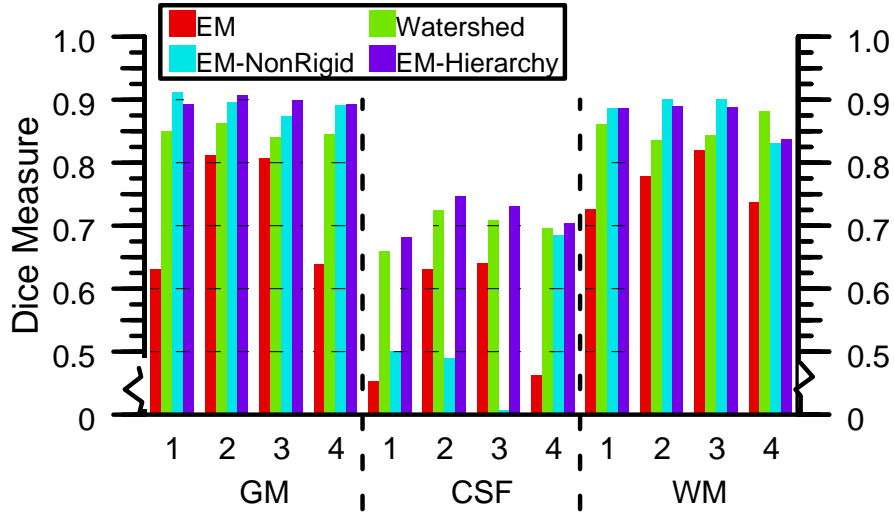


Figure 5-8: The graph shows the DICE score of the results of four different segmentation methods in comparison to a reference standard. Our new method EM-Hierarchy generally achieves the highest score for all three tissue classes because it can properly target the prior model to the segmentation problem.

Figure 5-10 displays the results of the DICE comparison. The brightness of a square indicates high correlation between the corresponding segmentations. Raters 1 - 6 represent medical experts, rater 7 is EM-NonRigid, and rater 8 is EM-Hierarchy. As mentioned, the medical experts differ in their levels of expertise with respect to segmenting the superior temporal gyrus. Rater 2 is the least experienced, which is also captured by our validation. The only other clear outlier is rater 7 (EM-NonRigid). While the method generally produces good results in comparison to EM-Affine of Section 3.2.1, it is not reliable enough to match the quality generated by manual segmentations. For the right superior temporal gyrus, however, the results (mean PPV score: 98.25%) developed by EM-Hierarchy (Rater 8) are within the variance of good manual segmentations (rater #1,#3 - # 6) (average mean PPV score: 0.98% variance within PPV score:1%). For the left superior temporal gyrus, the results of EM-Hierarchy (mean PPV Score: 96%) are within the range of all manual segmentations (average PPV score: 97.8% variance within PPV score: 2.3%). In general, EM-Hierarchy outperforms EM-NonRigid by 5 percentage points.

The DICE measure can also be correlated to the sample segmentations displayed in Figure 5-9. While no large differences can be observed among raters 1, 3-6 and




Figure 5-9: The images show segmentations of the right and left superior temporal gyrus. The upper row displays the original image with the segmentations of rater 1 - 6 and the segmentation of EM-NonRigid (EM-Non). The lower row overlays the result of the hierarchical approach (gray outline) with different segmentations (white outline.) No large visible differences can be found between the manual segmentations and the hierarchical segmentation. On the other hand, the segmentation of EM-NonRigid is very noisy when compared to the new approach.

EM-Hierarchy, rater 2 generally overestimates the structure and the results of EM-NonRigid are noisy and disconnected.

The observations of the DICE analysis are also confirmed by the PPV test. The first six entries in Figure 5-11 correspond to the experts, rater 7 is EM-NonRigid, and rater 8 is EM-Hierarchy. Rater 2 and EM-NonRigid clearly perform the worst. In all but one case, EM-Hierarchy performs similarly to rater 1, 3 - 6. The lower score for one case of the right superior temporal gyrus is caused by errors in the initial alignment process whose impact we already discussed in detail in Section 3.2.2. This exception also explains the lower mean score and higher standard error in the automatic segmentation of the right superior temporal gyrus in comparison to the left superior temporal gyrus. Even in this one case EM-Hierarchy scores higher than rater 2 and EM-NonRigid. If we compare the segmentation with the reference standard generated by STAPLE to the segmentation using the DICE measure, an overlap of




Figure 5-10: The graph shows the average DICE measure for the superior temporal gyrus over four cases. Number 1 - 6 correspond to manual expert segmentations, number 7 corresponds to EM-NonRigid, and number 8 represents EM-Hierarchy. The brightness of a square indicates high correlation between the corresponding segmentations. The only outliers for the right and left superior temporal gyrus are segmentation 2 and 7, which are characterized by the darker squares. Using DICE the new approach (#8) cannot be distinguished from the other good manual segmentations (#1, #3 - #6.)

95% is achieved, which is still an excellent rating.

The increased accuracy of EM-Hierarchy compared to EM-NonRigid is a result of the reduced statistical complexity of the hierarchical framework. In contrast to the last experiment, EM-NonRigid only focuses on the structures inside the inter-cranial cavity. Nevertheless, corticospinal fluid, white matter, and neighboring cortical substructures are competing for the region around the superior temporal gyrus. EM-NonRigid carefully compromises between the influence of the intensity information and the spatial prior. On the other hand, EM-Hierarchy is guided by the hierarchical structure presented in Figure 5-5. Based on Figure 5-5 it first outlines the border between gray matter, white matter, and corticospinal fluid mostly relying on the intensity constraints. When EM-Hierarchy later segments the superior temporal gyrus, the scenario is a lot simpler than for EM-NonRigid, because EM-Hierarchy only focuses on outlining the border between cortical structures, which is mainly guided by spatial constraints.

In summary, the hierarchical approach reduces the spatial complexity of segmen-

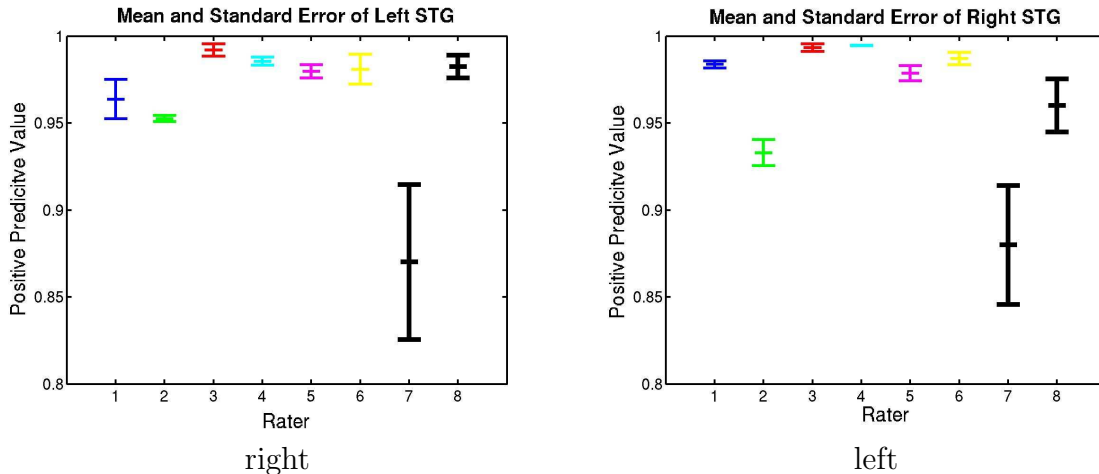


Figure 5-11: The results of STAPLE comparing 6 experts' (1-6) and 2 automatic segmentations (7 and 8). The algorithm presented in this chapter (8) is much more robust than the same approach without hierarchical structure (7). Additionally the results of our approach are close to the variance of the experts.

tation problems. The accuracy of the method is similar to manual segmentations at least for the superior temporal gyrus.

5.3.4 Comparison of Joint Registration and Segmentation Approach Using the Hierarchical Framework

We now analyze the impact on the accuracy of EM approach when decoupling the parameter estimation within the M-Step as suggested in Section 5.1.3. For this purpose, we repeat the experiment of the previous chapters segmenting 22 cases into the major brain tissue classes as well as the thalamus and the caudate. We then measure the accuracy of each automatic segmentation by comparing it to the corresponding manual segmentation using the volume overlap measure DICE of Section 1.5.

We first run the experiment on the two non-hierarchical methods EM-Simultaneous-Affine (Section 3.3.3) and EM-Simultaneous-Shape (Section 4.3.2,) which simultaneously determine the mapping from atlas to image space and segment the images into anatomical structures. The deformation model is encoded through global and local registration parameters. Determining the local deformation, however, is only practical when the optimal global registration is found. This constraint could

not be encoded in a principled manner by the models of Section 3.3.3 and Section 4.3.2. The hierarchical EM framework of this chapter can associate each registration parameter with the appropriate level of the tree.

To study the impact of the hierarchical model, we repeat the experiment of the 22 test cases. This time EM-Simultaneous-Affine and EM-Simultaneous-Shape are guided by the hierarchical relationship between anatomical structures. This relationship is captured in a tree of two levels. On the first level of the tree, the approaches determine the global mapping between the atlas and image space while simultaneously separating the image space into background and brain. On the second level of the tree, the methods fix the global registration parameters and perform an estimation for the local deformation as well as further parcellate the brain into white matter, gray matter, cortical spinal fluid, ventricles, caudate, and thalamus. We name the resulting methods EM-Hierarchy-Affine and EM-Hierarchy-Shape respectively.

Average DICE with standard error over 22 Cases

Method	Thalamus	Caudate
EM-Simultaneous-Affine	88.7 \pm 0.4%	85.0 \pm 0.7%
EM-Hierarchy-Affine	89.4 \pm 0.3%	86.1 \pm 0.6%
EM-Simultaneous-Shape	89.0 \pm 0.6%	86.2 \pm 0.6%
EM-Hierarchy-Shape	88.8 \pm 0.6%	86.6 \pm 0.7%

Table 5.1: The table shows the impact of the hierarchical framework on the EM approaches. For the integrated affine registration approach (EM-Simultaneous-Affine vs. EM-Hierarchy-Affine) the hierarchical framework increases the accuracy by an average DICE score of 0.9%. The greater accuracy is the result of separating the search of the optimal parameter setting with respect to global and local affine registration parameters. For the shape based approaches (EM-Simultaneous-Shape vs. EM-Hierarchy-Shape) the impact of the hierarchical model is insignificant as the global rigid registration and the local shape deformations are relatively independent.

The impact of the hierarchical approach is greater for EM-Hierarchy-Affine than EM-Hierarchy-Shape as shown in Table 5.1. In comparison to the non-hierarchical approach EM-Simultaneous-Affine, EM-Hierarchy-Affine improves the average DICE

score by 0.7% for the thalamus and 1.1% for the caudate. This improvement is due to the sequential optimization of the global and local affine registration parameters as defined by the hierarchical framework. EM-Simultaneous-Affine optimizes these two types of registration parameters simultaneously. This, however, creates instability within the EM approach as these two types of registration parameters are not independent.

The impact of the hierarchical framework on the shape based method EM-Hierarchy-Shape are rather negligible. Both shape based methods (EM-Simultaneous-Shape and EM-Hierarchy-Shape) perform well on both structures because the global rigid registration parameters and the structure-specific shape deformations describe two independent attributes within the segmentation model. Unlike the affine registration model, the simultaneous optimization of these two parameters does not destabilize the EM approach.




Figure 5-12: The graph compares the accuracy of the pipeline based approaches EM-Affine and EM-NonRigid (EM-NRigid) to our integrated methods EM-Hierarchy-Affine (EM-Hier-Af) and EM-Hierarchy-Shape (EM-Hier-Sh). For both structures our integrated methods perform better than the pipeline based approaches. See text for detailed discussion.

In Section 3.2.1, we developed the pipeline based approaches EM-Affine and EM-NonRigid. We again repeat this experiment to compare these two methods to the unified-hierarchical approaches EM-Hierarchy-Shape to EM-Hierarchy-Affine. Figure 5-12 shows the result of this comparison. For both anatomical structures, the

joint approaches achieve a higher average DICE score than the pipeline based methods. In addition, EM-Hierarchy-Affine significantly outperforms EM-Affine as the two intervals defined by the corresponding mean DICE score and the standard error do not overlap. The greater accuracy of our joint approaches is due to the seamless integration of the following three problems into an unified Bayesian framework:

- the registration of atlas information to the subject specific images,
- the segmentation of the images into anatomical structures, and
- the representation of the hierarchical relationships between anatomical structures in a tree.

5.4 Discussion

This chapter presented an anatomically guided segmentation model. The concept is based on a hierarchical data structure decoupling the complexity of the segmentation scenario from the number of structures to be segmented. To reduce the redundancy of the model, atlas information from the children is propagated to the parent structures. The resulting method adjusts to new segmentation problems efficiently as the dependency between structures is defined in the atlas space and not by the approach itself. In addition, the method properly tailors the prior information to each structure of interest, which reduces the risk of systematic biases as indicated by the experiments.

The hierarchical methodology presented in this section is a concept that is applicable to a class of segmentation methods such as [103, 33]. When applied to the different EM approaches the segmentation algorithm becomes more flexible and robustly handles complex segmentation scenarios. Parcellating a large number of structures can be handled by designing an appropriate tree, thus reducing the magnitude of memory and computing requirements. We expect similar results if the hierarchical concept is applied to other methods.

Chapter 6

Conclusion

This thesis developed an algorithm that automatically segments MR images into anatomical structures. The algorithm outlines structures with weakly visible boundaries as well as clearly defined perimeters. The main contribution of this thesis is the way we combine different types of prior information in a Bayesian framework, which captures image inhomogeneity correction, atlas registration, shape detection, and anatomical dependencies.

The success of the approach is due to the reduced space of possible solutions when compared to solving each problem individually. For example, the method detects misalignment between atlas and image space by comparing the prior information of the atlas to the current segmentation. The segmentation itself is guided by the shape atlas, whose deformations are constrained by the relationship between different anatomical structures. The relationship of different anatomical structures, however, is characterized by the aligned prior information. Capturing these dependencies results in a Bayesian framework that increases the accuracy of automatic segmentation methods in comparison to previous models in medical imaging.

To test the previous statement, we posed the optimal solution within this framework as a maximum *a posteriori* probability estimation problem. We find a solution for this estimation problem using an instance of the Expectation-Maximization algorithm. We then applied the algorithm to a variety of experiments ranging from the segmentation of the brain into the major tissue classes to the parcellation of the

tissue classes into the substructures with weakly visible boundaries. In general, these studies highlighted the superior performance of our novel approach in comparison to other Expectation-Maximization implementations. We conclude that our new way of integrating prior information into segmentation methods increases the robustness of certain automatic segmentation algorithms.

6.1 Future Work

The results of this thesis suggest possible extensions in several different directions. For example, this thesis implies that an increase in complexity of the parameter space can positively impact the robustness of the resulting segmentation method. We are interested in using more complete image inhomogeneity models (like in [29, 111]) as well as explicitly representing other image artifacts such as noise and partial voluming. The remainder of this section, however, focuses on new research opportunities resulting from the tree concept discussed in Chapter 5.

In this thesis, we used a tree to model the hierarchical relationship between different anatomical structures. The tree populates the current parameter estimates from the parent to the child classes and the prior information from the child to the parent classes. The algorithm traverses through the tree by segmenting the region of interest of the root into its subtrees and then repeats the process for the subtrees until it reaches the leafs of the tree. The algorithm never returns to the parent structures so that it cannot correct any inaccuracies at the higher levels of the tree. In order to overcome this drawback, we are interested in developing a graphical model based on the tree structure.

In the graphical model the nodes of the tree represent the anatomical structures of interest. This is a very different concept to the already existing graphical models in medical imaging, where each node represents a voxel in the image space [44, 74, 33, 117, 65, 51, 103, 76]. The random vectors associated with each node regularize the influence of the prior information onto the segmentation environment. The energy associated with the status of a node would be determined by the objective

function of the underlying EM implementation. As the edges in a graphical model represent the conditional probability of adjacent nodes, optimizing the tree would correspond to finding the most likely segmentation given the current observation and prior information.

The concept of graphical models is independent of the underlying segmentation approach so that the results of this research would be applicable to a variety of segmentation techniques. The new approach is integrated into the tree framework by substituting the EM implementation determining the energy of each node. This concept would not only increase the robustness of a wide range of already existing methods but also would reduce the efforts necessary to adjust these methods to different segmentation problems. For example, our current implementation requires manually defining the influence of the prior information at each level of the tree in the training phase. With the new approach, this task would be automatically performed.

As the tree determines the optimal parameter setting for each type of algorithms, the automatic segmentations of an approach are theoretically the best results achievable given the prior information and MR images. This framework would therefore present a validation platform for different segmentation philosophies, as the outcome would not be biased by the specific manual parameter setting of each algorithm.

To conclude, we developed a unified Bayesian framework in order to outline anatomical brain structures with weakly and clearly defined boundaries in MR images. The main contribution of this thesis is the way we integrate prior information into the unified Bayesian framework. The resulting EM approach simultaneously registers the atlas to the image space, detects the shape of an anatomical structure, and models the image inhomogeneities in the MR images. It uses the hierarchical relationship between anatomical structures to simplify the underlying optimization problem. We applied the method to a variety of experiments and showed that it outperforms other EM implementations on certain anatomical structures.

Bibliography

- [1] S. Aylward and J. Coggins. Spatially invariant classification of tissues in MR images. In *Conference on Visualization in Biomedical Computing*, 1994.
- [2] C. Baillard, P. Hellier, and C. Barillot. Segmentation of brain 3D MR images using level sets and dense registration. *Medical Image Analysis*, 5(3), 2001.
- [3] R. Bajcsy and S. Kovacic. Multiresolution elastic matching. *Computer Vision, Graphics and Image Processing*, 46:1–21, 1989.
- [4] M.A. Gonzalez Ballester, A. Zisserman, and M. Brady. Combined statistical and geometrical 3D segmentation and measurement of brain structures. In *IEEE Workshop on Biomedical Image Analysis*, pages 14–23, 1998.
- [5] Peter J. Bickel and Kjell A. Doksum. *Mathematical Statistics*. Prentice-Hall, Inc., 2nd edition, 2001.
- [6] F. Bookstein and W. Green. Edge information at landmarks in medical images. *Visualization in Biomedical Computing*, 242-252.
- [7] F.L. Bookstein. Landmark methods for forms without landmarks: morphometrics of group differences in outline shape. *Medical Image Analysis*, 1(3):225–243, 1996.
- [8] Sylvain Bouix. *Medial Surfaces*. PhD thesis, McGill University, 2003.
- [9] Sylvain Bouix, Lida Ungar, Chandlee C. Dickey, Robert W. McCarley, and Martha E. Shenton. Evaluating automatic brain tissue classifiers. In *Medical*

- Image Computing and Computer-Assisted Intervention*, volume 3216 of *Lecture Notes in Computer Science*, pages 1038–1039, Saint Malo, France, September 2004. Springer-Verlag.
- [10] L. Le Briquer and J.C. Gee. Design of a statistical model of brain shape. In *Information Processing in Medical Imaging Conference*, 1997.
- [11] M. Bro-Nielsen. Fast finite elements for surgery simulation. In *Studies in Health Technology and Informatics*, volume 39, pages 395–400, 1997.
- [12] D.L. Collins, C.J. Holmes, T.M. Peters, , and A.C. Evans. Automatic 3D model-based neuroanatomical segmentation. *Human Brain Mapping*, 3, 1995.
- [13] D.L. Collins, A.P. Zijdenbos, W.F.C. Barre, and A.C. Evans. Animal+insect: Improved cortical structure segmentation. *Information Processing in Medical Imaging*, 1613, 1999.
- [14] D.L. Collins, A.P. Zijdenbos, V. Kollokian, J. Sled, N.J. Kabani, C.J. Holmes, and A.C. Evans. Design and construction of a realistic digital brain phantom. *IEEE Transactions in Medical Imaging*, 17:463–468, 1998.
- [15] T.F. Cootes, G.J. Edwards, and C.J. Taylor. Active appearance model. In *Europeen Conference on Computer Vision (ECCV)*, volume 2, pages 484–498, 1998.
- [16] T.F. Cootes, A. Hill, C.J. Taylor, and J. Haslam. The use of active shape models for locating structures in medical imaging. *Imaging and Vision Computing*, 12(6):335–366, 1994.
- [17] T.F. Cootes, C.J. Taylor, D.H. Cooper, and J. Graham. Training models of shape from sets of examples. In *British Machine Vision Conference*, pages 9–18, 1992.
- [18] Thomas H. Cormen, Charles E. Leiserson, and Ronald L. Rivest. *Introduction to Algorithms*. MIT Press, 1998.

- [19] Anders M. Dale, Bruce Fischl, and Martin I. Sereno. Cortical surface-based analysis I. segmentation and surface reconstruction. *NeuroImaging*, 9:179–194, 1999.
- [20] P.-E. Danielsson. Euclidean distance mapping. *Computer Graphics and Image Processing*, 14, 1980.
- [21] C. Davatzikos. Spatial transformation and registration of brain images using elastically deformable models. *Computer Vision and Image Understanding*, 66(2):207–222, 1997.
- [22] M.H. Davis, A. Khotanzad, D.P. Flamig, and S.E. Harms. A physics-based coordinate transformation for 3-d image matching. *IEEE Transactions in Medical Imaging*, 16(3):317–328, 1997.
- [23] B.M. Dawant, E.J. Manders, and D.P. Lindstrom. Adaptive signal analysis and interpretation for real-time intelligent patient monitoring. *Methods of Information in Medicine*, 33(1):60–63, 1994.
- [24] Frank Dellaert. The expectation-maximization algorithm. Technical Report GIT-GVU-02-20, College of Computing GVU Center, Georgia Institute of Technology, 2002. <http://www.cc.gatech.edu/dellaert/em-paper.pdf>.
- [25] A. Dempster, N. Laird, and D. Rubin. Maximal likelihood form incomplete data via the EM algorithm. In *Proceedings of the Royal Statistical Society*, volume 39, pages 1–38, 1977.
- [26] J. Dengler, H. Bertsch, J.F. Desaga, and M. Schmidt. New trends of image analysis in the medical field. *Methods of Information in Medicine*, 27(2):53–57, 1988.
- [27] Ibrahim M Elfadel. *From Random Fields to Networks*. PhD thesis, Massachusetts Institute of Technology, 1993.

- [28] A.C. Evans, D.L. Collins, P. Neelin, D. MacDonald, M. Kamber, and T.S. Marrett. Three-dimensional correlative imaging: applications in human brain mapping. *Functional neuroimaging: technical foundations*, pages 145–162, 1994.
- [29] A. Fan, W.M. Wells, J.W. Fisher, M. Cetin, S. Haker, R. Mulkern, C. Tempny, and A. S. Willsky. A unified variational approach to denoising and bias correction in mr. In *Information Processing in Medical Imaging*, 2003.
- [30] M. Ferrant, A. Nabavi, B. Macq, F.A. Jolesz, R. Kikinis, and S.K. Warfield. Registration of 3-D intraoperative MR images of the brain using a finite-element biomechanical model. *IEEE Transactions in Medical Imaging*, 20(12):1384–1397, 2001.
- [31] A.H. Fielding and J.F. Bell. A review of methods for the assessment of prediction errors in conservation presence/absence models. *Environmental Conservation*, 24:38–49, 1997.
- [32] B. Fischl, A. Liu, and A.M. Dale. Automated manifold surgery: Constructing geometric accurate and topologically correct models of the human cerebral cortex. *IEEE Transactions in Medical Imaging*, 20(1), 2001.
- [33] B. Fischl, D.H. Salat, E. Busa, M. Albert, M. Dieterich, C. Haselgrove, A. van der Kouwe, R. Killiany, D. Kennedy, S. Klaveness, A. Montillo, N. Makris, B. Rosen, and A.M. Dale. Whole brain segmentation: Automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33, 2002.
- [34] B. Fischl, A. van der Kouwe, C. Destrieux, E. Halgren, F. Sgonne, D.H. Salat, E. Busa, L.J. Seidman, J. Goldstein, D. Kennedy, V. Caviness, N. Makris, B. Rosen, and A.M. Dale. Automatically parcellating the human cerebral cortex. *Cerebral Cortex*, 14:11–22, 2004.
- [35] J. Gee, C. Barillot, L.L. Briquer, D. Haynor, and R. Bajcsy. Matching structural images of the human brain using statistical and geometrical image features. In *Visualization in Biomedical Computing*, pages 191–204, 1994.

- [36] G. Gerig, O. Kübler, R. Kikinis, and F.A. Jolesz. Nonlinear anisotropic filtering of MRI data. *IEEE Transactions in Medical Imaging*, 11:221–232, June 1992.
- [37] G. Gerig and M. Styner. Statistical shape models for segmentation and structural analysis. In *IEEE International Symposium on Biomedical Imaging*, 2002.
- [38] Polina Golland. *Statistical Shape Analysis of Anatomical Structures*. PhD thesis, Massachusetts Institute of Technology, 2001.
- [39] V. Grau, M. Alcaniz A.J.U. Mewes, R. Kikinis, and S.K. Warfield. Improved watershed transform for medical image segmentation using prior information. *IEEE Transactions in Medical Imaging*, 23(4):447– 458, 2004.
- [40] A. Guimond, J. Meunier, and J.-P. Thirion. Average brain models: A convergence study. *Computer Vision and Image Understanding*, 77(2):192–210, 1999.
- [41] A. Guimond, A. Roche, N. Ayache, and J. Meunier. Three-dimensional multimodal brain warping using the demons algorithm and adaptive intensity corrections. *IEEE Transactions in Medical Imaging*, 20(1):58–69, January 2001.
- [42] X. Han, C. Xu, U. Braga-Neto, and J.L. Prince. Topology correction in brain cortex segmentation using a multiscale, graph-based algorithm. *IEEE Transactions in Medical Imaging*, 21(2):109–121, 2002.
- [43] H. Hartley. Maximum likelihood estimation from incomplete data. *Biometrics*, 14(2):174–194, 1958.
- [44] K. Held, E.R. Kops, B.J. Krause, W.M. Wells, R. Kikinis, and H.W. Muller-Gartner. Markov random field segmentation of brain MR images. *IEEE Transactions in Medical Imaging*, 16(6):878–86, 1997.
- [45] P. Hellier, C. Barillot, I. Corouge, B. Gibaud, G. Le Goualher, D.L. Collins, A. Evans, G. Malandain, N. Ayache, G.E. Christensen, and H.J. Johnson. Retrospective evaluation of inter-subject brain registration. *IEEE Transactions in Medical Imaging*, 22(9):1120–1130, 2003.

- [46] S. Ho, E. Bullitt, and G. Gerig. Level set evolution with region competition: Automatic 3-d segmentation of brain tumors. *IEEE International Conference on Pattern Recognition*, 2002.
- [47] Mark Holden, Derek L.G. Hill, E.R.E. Denton, J.M. Jarosz, Tim C.S. Cox, T. Rohlfing, Joanne Goodey, and David J. Hawkes. Voxel similarity measures for 3D serial MR brain image registration. *IEEE Transactions in Medical Imaging*, 19(2):94–102, 2000.
- [48] C.J. Holmes, R. Hoge, L. Collins, R. Woods, A.W. Toga, and A.C. Evans. Enhancement of mr images using registration for signal averaging. *Journal of Computer Assisted Tomography*, 22(2):324–333, 1998.
- [49] S. Joshi, B. Davis, M. Jomier, and G. Gerig. Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImaging*, 23:151–160, 2004.
- [50] T. Kapur, W.E.L. Grimson, W.M. Wells, and R. Kikinis. Segmentation of brain tissue from magnetic resonance images. *Medical Image Analysis*, 1(2):109–127, 1996.
- [51] Tina Kapur. *Model based three dimensional Medical Imaging Segmentation*. PhD thesis, Massachusetts Institute of Technology, 1999.
- [52] A. Kelemen, G. Szekely, and G. Gerig. Elastic model-based segmentation of 3-d neuroradiological data sets medical imaging. *IEEE Transactions in Medical Imaging*, 18:828 – 839, 1999.
- [53] Ron Kikinis, Martha E. Shenton, Guido Gering, John Martin, Mark Anderson, David Metcalf, Charles Guttman, Robert W. McCarley, William Lorensen, Harvey Line, and Ferenc A. Jolesz. Routine quantitative analysis of brain and cerebrospinal fluid spaces with MR imaging. *Magnetic Resonance Imaging*, 2(6):619–629, 1992.
- [54] K.R. Krishnan, W.M. McDonald, P.R. Escalona, P.M. Doraiswamy, C. Na, M.M. Husain, G.S. Figiel, O.B. Boyko, E.H. Ellinwood, and C.B. Nemeroff.

- Magnetic resonance imaging of the caudate nuclei in depression. Preliminary observations. *Archives of General Psychiatry*, 49:553–7, 1992.
- [55] K. Krissian. Flux-based anisotropic diffusion: Application to enhancement of 3D angiogram. *IEEE Transactions in Medical Imaging*, 22(11):1440–1442, 2002.
- [56] S. Kullback and R.A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22:79–86, 1951.
- [57] G.R. Kuperberg, M. Broome, P.K. McGuire, A.S. David, M. Eddy, F. Ozawa, D. Goff, W.C. West, S.C.R. Williams, A.J.W. van der Kouwe, D.H. Salat, A.M. Dale, and B. Fischl. Regionally localized thinning of the cerebral cortex in schizophrenia. *Archives of General Psychiatry*, 60:878–888, 2003.
- [58] R.K.-S. Kwan, A.C. Evans, and G.B. Pike. MRI simulation-based evaluation of image-processing and classification methods. *IEEE Transactions in Medical Imaging*, 18:1085–97, 1999.
- [59] M.E. Leventon, W.E.L. Grimson, and O.D. Faugeras. Statistical shape influence in geodesic active contours. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1316 – 1323, 2000.
- [60] Michael Emmanuel Leventon. *Statistical Models in Medical Image Analysis*. PhD thesis, Massachusetts Institute of Technology, 2000.
- [61] D.N. Levin, X.P. Hu, K.K. Tan, S. Galhotra, C.A. Pelizzari, G.T.Y. Chen, M.D. Cooper, J.F. Mullan, J. Hekmatpanah, and J.P. Spire. The brain: integrated threedimensional display of MR and PET images. *Radiology*, 172(3):783–9, 1989.
- [62] Liana M. Lorigo, Olivier D. Faugeras, W.Eric L.Grimson, Renaud Keriven, Ron Kikinis, Arya Nabavi, and Carl-Fredrik Westin. CURVES: Curve evolution for vessel segmentation. *Medical Image Analysis*, 5(3):195–206, 2001.

- [63] L.R.Dice. Measure of the amount of ecological association between species. *Ecology*, 26:297–302, 1945.
- [64] S.P. Luttrell. The partitioned mixture distribution: an adaptive Bayesian network for low-level image processing. In *IEEE Proceedings on Vision Image and Signal Processing*, 1994.
- [65] J.L. Marroquin, E.A. Santana, and S. Botello. Hidden markov measure field models for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25:1380–1387, 2003.
- [66] J.L. Marroquin, B.C. Vemuri, S. Botello, F. Calderon, and A. Fernandez-Bouzas. An accurate and efficient bayesian method for automatic segmentation of brain MRI. *IEEE Transactions in Medical Imaging*, 21:934–945, 2002.
- [67] J.C. Mazziotta, A.W. Toga, A. Evans, P. Fox, and J. Lancaster. A probabilistic atlas of the human brain: theory and rationale for its development. The international consortium for brain mapping (ICBM). *NeuroImaging*, 2(2):89–101, 1995.
- [68] T. McInerney and D. Terzopoulos. T-snakes: topology adaptive snakes. *Medical Image Analysis*, 4(2):73–91, 2000.
- [69] Geoffrey J. McLachlan and Thriyambakam Krishnan. *The EM Algorithm and Extensions*. John Wiley and Sons, Inc., 1997.
- [70] M.I. Miller. Computational anatomy: shape, growth, and atrophy comparison via diffeomorphisms. *NeuroImaging*, 23:19–33, 2004.
- [71] T.P. Minka. Expectation-maximization as lower bound maximization. Technical report, MIT, 1998.
- [72] M. Moshfeghi, S. Ranganath, and K. Nawyn. Three-dimensional elastic matching of volumes. *IEEE Transactions On Image Processing*, 3(2):128–138, 1994.

- [73] R. Neal and G. Hinton. A view of the EM algorithm that justifies incremental, sparse, and other variants. In M.I. Jordan, editor, *Learning in Graphical Models*, pages 355–368. Kluwer Academic Press., 1998.
- [74] Chahin Pachai, Yue Min Zhu, Charles R. G. Guttman, Ron Kikinis, Ferenc A. Jolesz, Gerard Gimenez, Jean-Claude Froment, Christian Confavreux, and Simon K. Warfield. Unsupervised and adaptive segmentation of multispectral 3D magnetic resonance images of human brain: a generic approach. In *Medical Image Computing and Computer-Assisted Intervention*, pages 1067–1074, 2001.
- [75] T. Paus, F. Tomaiuolo, N. Otaky, D. MacDonald, M. Petrides, J. Atlas, R. Morris, and A.C. Evans. Human cingulate and paracingulate sulci: pattern, variability, asymmetry, and probabilistic map. *Cereb Cortex*, 6(2):207–214, 1996.
- [76] D. Pham and J. Prince. A generalized EM algorithm for robust segmentation of magnetic resonance images. In *Conference on Information Sciences and Systems (CISS)*, pages 558–563, 1999.
- [77] D.L. Pham and J.L. Prince. Adaptive fuzzy segmentation of magnetic resonance images. *IEEE Transactions in Medical Imaging*, 18:737 – 752, 99.
- [78] S.M. Pizer, D.S. Fritsch, P.A. Yushkevich, V.E. Johnson, and E.L. Chaney. Segmentation, registration, and measurement of shape variation via image object shape. *IEEE Transactions in Medical Imaging*, 18(10):851–865, 1999.
- [79] Stephen M. Pizer, Guido Gerig, Sarang Joshi, and Stephen R. Aylward. Multi-scale medial shape-based analysis of image objects. In *Proceedings of the IEEE, Special Issue on: Emerging Medical Imaging Technology*, volume 91, pages 670 – 679, 2003.
- [80] K.M. Pohl, S. Bouix, R. Kikinis, and W.E.L. Grimson. Anatomical guided segmentation with non-stationary tissue class distributions in an expectation-maximization framework. In *IEEE International Symposium on Biomedical Imaging*, pages 81–84, 2004.

- [81] K.M. Pohl, W.M. Wells, A. Guimond, K. Kasai, M.E. Shenton, R. Kikinis, W.E.L. Grimson, and S.K. Warfield. Incorporating non-rigid registration into expectation maximization algorithm to segment MR images. In *Medical Image Computing and Computer-Assisted Intervention*, pages 564–572, 2002.
- [82] W.H. Press, B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling. *Numerical Recipes in C : The Art of Scientific Computing*. Cambridge University Press, 2 edition, 1992.
- [83] Torsten Rohlfing, Daniel B. Russakoff, C.R. Maurer, Robert Brandt, and Randolph Menzel. Performance-based multi-classifier decision fusion for atlas-based segmentation of biomedical images. In *IEEE International Symposium on Biomedical Imaging*, pages 404–407, 2004.
- [84] K. Rohr, H.S. Stiehl, R. Sprengel, T.M. Buzug, J. Weese, and M.H. Kuhn. Landmark-based elastic registration using approximating thin-plate splines. *IEEE Transactions in Medical Imaging*, 20(6):526–534, 2001.
- [85] M. Rousson, N. Paragios, and R. Deriche. Active shape models from a level set perspective. Technical Report 4984, Institut National de Recherche en Informatique et en Automatique, Sophia-Antipolis, 2003. <ftp://ftp.inria.fr/INRIA/publication/publi-pdf/RR/RR-4984.pdf>.
- [86] F. Ségonne, E. Grimson, and B. Fischl. Topological correction of subcortical segmentation. In *Medical Image Computing and Computer-Assisted Intervention*, pages 695–702, 2003.
- [87] M.E. Shenton, R. Kikinis, F.A. Jolesz, S.D. Pollak, M. LeMay, C.G. Wible, H. Hokama, J. Martin, D. Metcalf, M. Coleman, and R.W. McCarley. Left temporal lobe abnormalities in schizophrenia and thought disorder: A quantitative MRI study. *New England Journal of Medicine*, 327:604–612, 1992.
- [88] Siddarth Srivastava, Frederik Maes, Dirk Vandermeulen, Wim Van Paesschen, Patrick Dupont, and Paul Suetens. Effects of anatomical asymmetry in spatial

- priors on model-based segmentation of the brain mri: A validation study. In *Medical Image Computing and Computer-Assisted Intervention*, number 3216 in Lecture Notes in Computer Science, pages 327–334. Springer-Verlag, 2004.
- [89] P. Stoica and Y. Selen. Cyclic minimizers, majorization techniques, and the expectation-maximization algorithm: a refresher. *IEEE Signal Processing Magazine*, 21(1):112 – 114, 2004.
- [90] G. Subsol, J.P. Thirion, and N. Ayache. A scheme for automatically building three-dimensional morphometric anatomical atlases: application to a skull atlas. *Medical Image Analysis*, 2(1):37–60, 1998.
- [91] P. C. Teo, G. Sapiro, and B.A. Wandell. Creating connected representations of cortical gray matter for functional MRI visualization. *IEEE Transactions in Medical Imaging*, 16(6):852–863, 1997.
- [92] J.-P. Thirion. Image matching as a diffusion process: an analogy with Maxwell’s demons. *Medical Image Analysis*, 2(3):243–260, 1998.
- [93] P.M. Thompson, D. MacDonald, M.S. Mega, C.J. Holmes, A.C. Evans, and A.W. Toga. Detection and mapping of abnormal brain structure with a probabilistic atlas of cortical surfaces. *Journal of Computer Assisted Tomography*, 21(4):567–581, 1997.
- [94] P.M. Thompson, C. Schwartz, and A.W. Toga. High-resolution random mesh algorithms for creating a probabilistic 3D surface atlas of the human brain. *NeuroImaging*, 3(1):19–34, 1996.
- [95] P.M. Thompson, C. Vidal, J.N. Giedd, J. Blumenthal, P. Gochman, R. Nicolson, A.W. Toga, and J.L. Rapoport. Mapping adolescent brain change reveals dynamic profile of accelerated gray matter loss in childhood-onset schizophrenia. *NeuroImaging*, 13(6):1105, 2001.

- [96] P.M. Thompson, R.P. Woods, M.S. Mega, and A.W. Toga. Mathematical/computational challenges in creating deformable and probabilistic atlases of the human brain. *Human Brain Mapping*, 9(2):81–92, 2000.
- [97] A.W. Toga and P.M. Thompson. Maps of the brain. *The Anatomical Record*, 265(2):37–53, 2001.
- [98] A. Tsai, W.M. Wells, C. Tempany, E. Grimson, and A. Willsky. Coupled multi-shape model and mutual information for medical image segmentation. *Information Processing in Medical Imaging*, 2003.
- [99] A. Tsai, A. Yezzi, W.M. Wells, C. Tempany, D. Tucker, A. Fan, W. Grimson, and A. Willsky. Model-based curve evolution technique for image segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2001.
- [100] A. Tsai, A. Yezzi, W.M. Wells, C. Tempany, D. Tucker, A. Fan, W. Grimson, and A. Willsky. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Transactions in Medical Imaging*, 22(2):137 – 154, 2003.
- [101] J.K. Udupa, P.K. Saha, and R.A. Lotufo. Relative fuzzy connectedness and object definition: Theory, algorithms, and applications in image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(11):1485–1500, 2002.
- [102] G. Unal, D. Nain, G. Ben-Arous, N. Shimkin, A. Tannenbaum, and O. Zeitouni. Algorithms for stochastic approximations of curvature flows. In *IEEE International Conference on Image Processing*, 2003.
- [103] K. Van Leemput, F. Maes, D. Vanermeulen, and P. Suetens. Automated model-based bias field correction of MR images of the brain. *IEEE Transactions in Medical Imaging*, 18(10):885–895, 1999.
- [104] K. Van Leemput, F. Maes, D. Vanermeulen, and P. Suetens. Automated model-based tissue classification of MR images of the brain. *IEEE Transactions in Medical Imaging*, 18(10):897 – 908, 1999.

- [105] M.W. Vannier, R.L. Butterfield, D. Jordan, W.A. Murphy, R.G. Levitt, and M. Gado. Multispectral analysis of magnetic resonance images. *Radiology*, 154(1):221–4, 1985.
- [106] R.P. Velthuizen, L.O. Hall, and L.P. Clarke. Feature extraction for MRI segmentation. *Neuroimaging*, 9(2):85–90, 1999.
- [107] S.K. Warfield, J. Rexilius, P.S. Huppi, T.E. Inder, E.G. Miller, W.M. Wells, G.P. Zientara, F.A. Jolesz, and R. Kikinis. A binary entropy measure to assess nonrigid registration algorithm. In *Medical Image Computing and Computer-Assisted Intervention*, pages 266–274, October 2001.
- [108] S.K. Warfield, J. Rexilius, M. Kaus, F.A. Jolesz, and R. Kikinis. Adaptive, template moderated, spatial varying statistical classification. *Medical Image Analysis*, 4(1):43–55, 2000.
- [109] S.K. Warfield, K.H. Zou, and W.M. Wells. Validation of image segmentation and expert quality with an expectation-maximization algorithm. In *Medical Image Computing and Computer-Assisted Intervention*, 2002.
- [110] S.K. Warfield, K.H. Zou, and W.M. Wells. Simultaneous truth and performance level estimation (STAPLE): An algorithm for the validation of image segmentation. *IEEE Transactions in Medical Imaging*, 2004.
- [111] Neil Weisenfeld and Simon K. Warfield. Normalization of joint image-intensity statistics in MRI using the kullback-leibler divergence. In *IEEE International Symposium on Biomedical Imaging*, pages 101–104, 2004.
- [112] W.M. Wells, W.E.L. Grimson, R. Kikinis, and F.A. Jolesz. Adaptive segmentation of MRI data. *IEEE Transactions in Medical Imaging*, 15:429–442, 1996.
- [113] C. Westin, L. Lorigo, O. Faugeras, W. Grimson, S. Dawson, A. Norbash, and R. Kikinis. Segmentation by adaptive geodesic active contours. *Medical Image Computing and Computer-Assisted Intervention*, 2000.

- [114] Meihe Xu, P.M. Thompson, and A.W. Toga. An adaptive level set segmentation on a triangulated mesh. *IEEE Transactions in Medical Imaging*, 23(2):191–201, 2004.
- [115] Jing Yang, Lawrence H. Staib, and James S. Duncan. Neighbor-constrained segmentation with level set based 3D deformable models. *IEEE Transactions in Medical Imaging*, 23(8):940–948, 2004.
- [116] X. Zeng, L.H. Staib, R.T. Schultz, and J.S. Duncan. Segmentation and measurement of the cortex from 3D MR images using coupled surfaces propagation. *IEEE Transactions in Medical Imaging*, 18(10):927–937, 1999.
- [117] Y. Zhang, M. Brady, and S. Smith. Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Transactions in Medical Imaging*, 20:45–57, 2001.
- [118] A.P. Zijdenbos, B.M. Dawant, R.A. Margolin, and A.C. Palmer. Morphometric analysis of white matter lesions in MR images: Method and validation. *IEEE Transactions in Medical Imaging*, 13(4):716–724, 1994.