Fach Mathematik

Dissertationsthema

Automatic Generation of Volume Conductor Models of the Human Head for EEG Source Analysis

Inaugural-Dissertation
zur Erlangung des Doktorgrades
der Naturwissenschaften im Fachbereich
Mathematik und Informatik
der Mathematisch-Naturwissenschaftlichen Fakultät
der Westfälischen Wilhelms-Universität Münster

vorgelegt von
Benjamin Lanfer
aus Vreden
- 2014 -

Dekan:	Prof. Dr. Martin Stein
Erster Gutachter:	Priv.-Doz. Dr. Carsten H. Wolters
Zweiter Gutachter:	Prof. Dr. Martin Burger
Tag der mündlichen Prüfung:	
Tag der Promotion:	

Acknowledgment

Many people supported me in writing this thesis. Here, I'd like to acknowledge their support. I would like to thank

PD Dr. Carsten Wolters for years of invaluable advice, always helpful and sometimes lively discussions, his pleasant company, frequent dinner invitations to his home, and not at least for him being the supervisor for this thesis.

Dr. Michael Scherg for offering me the chance to write this thesis on the fascinating topic of EEG source analysis, and the endless insights into neuroscience.

Prof. Dr. Martin Burger for always being willing to provide rapid and helpful feedback on my work, and for reviewing this thesis.

Prof. Dr. Christo Pantev for giving me a "home" at the Institute for Biomagnetism and Biosignalanalysis in Münster.

Theo and Dr. Tobias Scherg for the ability to work in the very pleasant, familiar atmosphere of the BESA company.

My colleagues in the workgroup at the Institute for Biomagnetism and Biosignalanalysis, Ümit Ayden, Felix Lucka, Johannes Vorwerk, and Sven Wagner for stimulating discussions, for proof-reading the thesis, and for doing the even more tedious manual segmentation for the validation chapter.

Dr. Harald Kugel, Dr. Matthias Köchling, Dr. Wolfram Schwindt for providing some of the MRI data used in this work.

During the entire time of writing this thesis there has always been a welcoming home for me. For this and for their constant, manifold support, and for the many unforgettable moments I have to thank Sara and Ole Lanfer.

Last but not least, I would like to thank my parents who made it possible for me to go to university in the first place.

Zusammenfassung

Bei der Elektroenzephalographie (EEG) werden elektrische Potentialdifferenzen auf der Kopfoberfläche gemessen. Anhand dieser Daten können mit Hilfe der EEG-Quellenanalyse Rückschlüße auf die neu ro nalen Quellen der aufgezeichneten EEG-Signale gezogen werden. Diese Methode ist ein wichtiges Hilfsmittel für Hirnforscher und für Ärzte, die sich mit der Erforschung neurologischer Erkrankungen beschäftigen. Ein wesentlicher Bestandteil der Quellenanalyse ist die Lösung des so-genannten Vorwärts-Problems, das heißt, dem Problem die EEG Signale ausgehend von einer angenommenen Stromdichteverteilung zu simulieren.

Für eine genaue Lösung des Vorwärts-Problem wird eine Beschreibung der individuellen, elektrischen Eigenschaften des menschlichen Kopfes benötigt. Im Falle der quasi-statischen Gleichungen, die das Vorwärts-Problem beschreiben, bedeutet dies, dass ein Kopfmodell erzeugt werden muss, welches die Verteilung der elektrischen Leitfähigkeit wider spiegelt. Im Hinblick auf dieses Kopfmodell existieren einige offene Fragen von wissenschaftlichem Interesse.

Kopfmodelle werden in der Praxis fast immer erzeugt durch die Segmentierung des Kopfes in unterschiedliche Geweberegionen, denen dann zuvor gemessene, veröffentlichte Leitfähigkeitswerte für das jeweilige Gewebe zugewiesen werden. Zwei wichtige Aspekte in diesem Zusammenhang sind der notwendige Grad der Detailliertheit, das heißt, wieviele Gewebe unterschieden werden müssen, und die notwendige Genauigkeit der Kopfmodelle, das heißt, wie genau die einzelnen Geweberegionen segmentiert werden müssen.

Um diese Fragen zu beantworten, wurden im Rahmen dieser Arbeit zwei Simulationsstudien durchgeführt. Die erste Studie untersuchte den Einfluß einer Reihe von geometrischen Ungenauigkeiten bei der Modellierung des Schädels auf die EEG-Quellenanalyse. Die Ergebnisse dieser Studie erlaubten es Richtlinien für die Konstruktion von geeigneten Kopfmodellen zu formulieren. In einer zweiten Studie wurde der Effekt einer vereinfachten Modellierung der Liquor-Region betrachtet. Als Ergebnis konnte gezeigt werden, dass diese untersuchte, spezielle Vereinfachung in den meisten Anwendungen zulässig ist, und dass sie so die Anwendung von klassischen inversen Verfahren, wie dem Dipole Fit, in Kombination mit realistischen Kopfmodellen erlaubt.

Eine weitere wichtige und noch unbeantwortete Frage ist wie die Gewebe des menschlichen Kopfes auf effiziente Art und Weise und mit großer Genauigkeit segmentiert werden können zum Zwecke der Erzeugung von Kopfmodellen für die EEG-Quellenanalyse.

Diese Arbeit schlägt zwei Segmentierungsansätze vor, um diese Frage zu beantworten. Der erste Ansatz basiert auf der Multiphase Active Contours Methode von Chan und Vese. Er erlaubt die Segmentierung unter Einbeziehung von mehreren Bildgebungsmodalitäten und Randbedingungen in Bezug auf die Krümmung der segmentierten Grenzflächen. Der zweite Ansatz nutzt ein Markov Random Field-Modell, um das Vorwissen über die geschichtete Struktur der Gewebe des Kopfes zu berücksichtigen. Beide Segmentierungsansätze sind des weiteren in der Lage anatomisches Vorwissen auf Grundlage eines Atlases der Gewebewahrscheinlichkeiten einzubeziehen. Ein geeigneter Atlas wurde ebenfalls im Rahmen dieser Arbeit entwickelt. Die vi-

suelle Inspektion der Segmentierungsergebnisse, die anhand von Magnetresonanztomographien erzeugt wurden, offenbarte eine gute Genauigkeit für beide Ansätze.

Um die Genauigkeit der vorgeschlagenen Ansätze tiefer zu untersuchen wurden die Segmentierungsergebnisse mit Referenzsegmentierungen auf Basis von manuellen Markierungen und auf Basis von Computertomographien verglichen. Der Schwerpunkt bei dieser Studie wurde auf die Genauigkeit der Schädelsegmentierung gelegt.

Die hohe Präzision der beiden vorgeschlagenen Segmentierungsansätze insbesondere im Hinblick auf den Schädel konnte bestätigt werden. Die Validierung gegen manuelle Referenzergebnisse zeigte, dass der Ansatz unter Ausnutzung eines Markov Random Field-Modells und des Gewebeatlasses die besten Ergebnisse lieferte. Die Genauigkeit der Schädelsegmentierung wurde dabei im Vergleich zu einer weit verbreiteten, konkurrierenden Methode wesentlich verbessert.

Zusammengefasst konnten in dieser Arbeit relevante und zuvor offene Fragen im Hinblick auf den benötigten Detailgrad und die benötigte Genauigkeit von Kopfmodellen für die EEG-Quellenanalyse beantwortet werden. Des weiteren, wurden zwei Ansätze für die automatische Erzeugung von EEG-Kopfmodellen vorgeschlagen. Es wurde gezeigt, dass die von den Ansätzen erzeugten Modelle sich sehr gut für die Anwendung in der Quellenanalyse eignen. Die Verfügbarkeit einer effizienten Methode zur Erzeugung solcher Modelle ermöglicht die zunehmende Verwendung von realistischen, individuellen Kopfmodellen in vielen Anwendungen. Dies wiederum wird die Genauigkeit und Aussagekraft der Quellenanalyse verbessern und es Wissenschaftlern ermöglichen tieferen Einblick in die Arbeitsweise des menschlichen Gehirns zu gewinnen.

Abstract

Electroencephalography (EEG) source analysis is the process of infering the neural sources that produce voltage differences measured on the surface of the scalp. It is a valuable tool for neuroscientist and medical researchers. Integral part of EEG source analysis is the solution of the forward problem, that is, the problem of simulating the EEG signals given an assumed current density distribution.

To accurately solve the forward problem in human subjects an accurate description of the individual electrical properties of the human head is necessary. For the quasi-static equations governing the forward problem this means that a head model must be created describing the conductivity distribution. Some open questions exist with regard to these head models.

The models are in practice nearly always constructed by segmenting the head into different tissue regions and assigning previously measured and published conductivity values to each of the regions. Here, two aspects that are of great interest to many researchers are the required level of detail, that is, how many and which tissues have to be differentiated, and the required accuracy, that is, how accurately do the tissue regions have to be segmented.

To answer these questions two simulation studies were performed in this thesis. The first study investigated the effects of a wide variety of inaccuracies in the geometrical modeling of the skull on EEG source analysis. Results of this study allowed to conclude guidelines for the construction of suitable head models. In a second study we investigated the effect of modeling the cerebro-spinal fluid compartment in a simplified way. As a result we could show that a certain simplification is valid and, thus, allows for the combination of realistic head models with classical inverse methods like the moving dipole fit.

Another important open question is how to efficiently and accurately segment the tissues of the human head for the purpose of generating EEG head models.

This thesis proposes two segmentation approaches to answer this question. The first approach is based on the multiphase active contours method of Chan and Vese. It allows the segmentation taking into account multiple image modalities and constraints on the smoothness of the segmented interfaces. The second approach utilizes a Markov random field model to represent a-priori knowledge on the layered anatomy of the tissues of the human head. Both approaches are able to exploit anatomical a-priori information provided by a tissue probability atlas which was also developed as a part of this work. Visual inspection of the segmentation results obtained on real MRI data revealed a good accuracy for both approaches.

The performance of the proposed segmentation approaches was in more depth investigated by comparing the results to manual reference segmentations and to reference segmentations based on computed tomography images. During the investigations special focus was put on the accurate reconstruction of the skull geometry.

In these validation studies the good accuracy of the segmentation approaches was confirmed especially with regard to the skull. The validation against reference segmentations based on manual labelings showed that the proposed approach utilizing the Markov random field model

and exploiting the atlas-based a-priori information performed best. In comparison to a widely-used competing method skull segmentation accuracy was substantially improved.

In conclusion, this work was able to answer relevant open questions with regard to the required level of detail and the required accuracy of head models for EEG source analysis. Furthermore, two segmentation approaches for the automatic generation of head models were proposed. It was shown that the generated segmentations are suited as a basis for the generation of EEG head models. The availability of an efficient method for the generation of such head models will make it more feasible in many applications to use realistic, individual models. This in turn will facilitate EEG source analysis and help researchers gain further insight into the workings of the human brain.

Contents

Nomenclature

Abbreviations

$D(R,S)$	Dice coefficient for masks R and S
AFCM	Adaptive fuzzy c-means
AMG-CG method	Algebraic multi-grid preconditioned conjugate gradient method
BEM	Boundary element method
BG	Background
CFL condition	Courant-Friedrich-Lewy condition
CG method	Conjugate gradient method
CSF	Cerebrospinal fluid
CT	Computed (X-ray) tomography
DW-MRI	Diffusion-weighted MRI
EEG	Electroencephalography
EM	Expectation-maximization
FDM	Finite difference method
FE	Finite element
FEM	Finite element method
FID	Free induction decay
FMM	Fast marching method
FN	False negative
FP	False positive
FVM	Finite volume method
GM	Gray matter
GRF	Gibbs random field

ICM	Iterated conditional modes
ISA	Isolated skull approach
MAG	Magnitude error
MAP	Maximum a-posteriori
MEG	Magnetoencephalography
MRA	Magnetic resonance angiography
MRF	Markov random field
MRI	Magnetic resonance imaging; magnetic resonance image
PD	Proton density
RDM	Relative difference measure
RF	Radio-frequency
ROI	region of interest
SA	Simulated annealing
SCT	Subcutaneous tissue
SE	Spin-echo
SNR	Signal-to-noise ratio
tDCS	Transcranial direct current stimulation
TE	Echo time
TM	Test model
TMS	Transcranial magnetic stimulation
TR	Repetition time
WM	White matter

Symbols

$\boldsymbol{j}_p$	Primary current
$\boldsymbol{M}^i$	Moment of the i-th mathematical dipole
Γ	Surface of the volume conductor Ω
Ω	The volume conductor, i.e., solution domain for the forward problem

ϕ	Electric potential
σ	Electrical conductivity tensor
L	Partial differential operator
u	Electric potential
$V \subset H_0^1$	Linear space of solutions for the FEM minimization problem
$\boldsymbol{J}$	FEM source vector
$\boldsymbol{R}$	Restriction matrix (transfer matrix approach)
$\boldsymbol{T}^{\text{EEG}}$	EEG transfer matrix
$\boldsymbol{u}$	FEM potential vector
$\boldsymbol{u}^{\text{EEG}}$	Electric potentials at the EEG electrodes
ϕ^{∞}	Potential of a mathematical dipole in an infinite homogeneous medium (subtraction approach)
ϕ^{corr}	Correction potential (subtraction approach)
${}^{(e)}\phi_h$	Finite element ansatz function
${}^{(e)}\phi_k$	Finite element node based form functions
H_0^1	Sobolev space of order 1
J_k	Monopol sources of the blurred dipole model
K	Stiffness matrix of the finite element method
N	Number of finite element nodes
$V_h \subset V$	Subspace of finite element solutions
$\boldsymbol{B}_0$	External static magnetic field (MRI)
γ	Gyromagnetic ratio
ω_0	Larmor frequency
μ	Attenuation coefficient (Computed tomography)
ρ	Material density (Computed tomography)
HU	Hounsfield units (Computed tomography)
σ_{tot}	Total photon atomic cross section (Computed tomography)

$\boldsymbol{\mu}_c$	Mean intensities in region c
$\boldsymbol{\Sigma}_c$	Covariance matrix of the image intensities in region c
$\boldsymbol{y}$	Vector of potentially multimodal observations at all grid sites $i \in \mathcal{S}$
$\boldsymbol{y}_i$	Multimodal observation, i.e., MR intensity, at grid site i
$\Omega \subset \mathbb{R}^3$	Image domain
$C(x^*, x)$	Cost function for Bayesian risk
i	Single grid site index
j, k, l	Grid site indices
M	Binary mask
n^j, n^k, n^l	Number of sites along each of the three principal grid axes
$P(x)$	A-priori probability mass function for the labeling x
$R(x^*)$	Bayesian risk
$u_0(\boldsymbol{r})$	Continous image data
x	Labels at all grid sites $i \in \mathcal{S}$; labeled image
X, Y, Λ	Random fields
x_i	Label at grid site i
y_i	Monomodal observation, i.e., MR intensity, at grid site i
$\mathcal{L} = \{0, \ldots, m-1\}$	Set of possible label values
$\mathcal{P}(\Omega)$	Partition of the image domain Ω
$\mathcal{S}$	Set of grid sites
$\boldsymbol{U}_i^*$	Update matrix term for level set function i
χ_i	Characteristic function for region i
Δt	Time step for active contours update equation
ϵ	Smoothness parameter for the active contours segmentation
ν	Inflation parameter for active contour segmentation
ψ_i	Level set function i
$H(y)$, $\boldsymbol{H}(\boldsymbol{y})$	Heaviside function and vector valued Heaviside function

$c \in C$	Clique (Gibbs random field)
C_1, C_2	Single site and pairwise cliques
T	Temperature parameter (Gibbs random field)
$t(\boldsymbol{y}, x)$	Sufficient statistic
$U(x)$	Gibbs energy function
$V_c(x)$	Clique potential (Gibbs random field)
Z	Normalization constant (Gibbs random field)
C	Set of all possible cliques with sites from S
$\mathcal{L} = \{0, \ldots, m-1\}$	Set of labels
$\mathcal{N}$	Neighborhood system
$\mathcal{N}_i$	Neighborhood of site i
λ_1, λ_2	Regularization parameters for the AFCM algorithm
μ_k	Centroid of the k-th cluster
∇	Gradient operator
c	Number of clusters
D_r	Finite difference operator with respect to the r-th coordinate
m_i	Value of the multiplier field at site i
p_i^k	Tissue probability weighting factor for the modified AFCM algorithm
u_i^k	AFCM membership function for site i and cluster k
$d^{\text{scalp}}, d^{\text{WM}}$	Distances of current voxel to scalp and WM reference surface
$h_j^l(d^{\text{scalp}}, d^{\text{WM}})$	Histogram count for label l at vertex j of the projection surface
$P_j^l(d^{\text{scalp}}, d^{\text{WM}})$	Probability for label l at vertex j of the projection surface
w_j	Weighting factor for vertex j of the projection surface

1 Introduction

1.1 Motivation of This Work

Neuroscience is the scientific study of all aspects of the nervous system and the brain. It is a very active area of research with a lot of ongoing work concerned with, for example, how the brain processes information, how it makes decisions, how neurological pathologies arise, and how diseases of the brain can be treated.

Electroencephalography (EEG) [13] is a useful tool for this research. EEG measures the scalp potentials which originate from electric currents in the brain. These currents are directly related to neuronal activity. Already by inspecting the EEG some knowledge can be gained on the behavior of the neuronal circuits of the brain. Even more insight can be obtained by trying to infer the neuronal sources of the EEG from the measured signal. This is what EEG source analysis does [90].

A central problem in EEG source analysis — the forward problem — is to estimate the scalp potentials for an assumed current density distribution in the brain. For an accurate solution of the forward problem an individual, realistic model of the electric properties of the subject's head must be constructed, and the partial differential equation describing the relation between the electric neuronal currents and the scalp potentials must be solved for this model [37, 162, 135].

The construction of the individual, realistic volume conductor model is a critical point. There are open questions regarding how accurate and detailed the models have to be so that an accurate forward solution is obtained. It was, for example, until now not studied how accurately the geometry of the skull has to be modeled, or to which extent the cerebrospinal fluid (CSF) has to be incorporated into the model. These questions will be answered in Chapters 4 and 5 of this thesis.

In the application of EEG source analysis volume conductor models are commonly generated by segmenting the different tissues of the head from magnetic resonance images (MRIs) of the subject. Previously measured and published conductivity values are then assigned to the tissue regions to complete the model. An open question is how individual head models that are accurate and that have a sufficient level of detail can be generated efficiently from MR image data. Manual segmentation is too labor intensive to be feasible. Thus, an automatic method for the segmentation of the relevant tissues is needed. In the literature some methods [67, 46, 115, 147] were suggested to solve the problem of head tissue segmentation but these methods either rely strongly on manually chosen parameter values, or they are too inaccurate, or they are computationally too expensive. Therefore, in Chapters 7 and 8 we propose two automatic methods for the segmentation of volume conductor models from MRI data that overcome the disadvantages of previous methods.

The first segmentation approach is based on the multiphase active contours segmentation of Vese and Chan [143] that is known to be robust against noise, and that was successfully em-

ployed in other applications [5]. The second proposed method for the automatic generation of head tissue segmentations utilizes a Markov random field (MRF) model to encode our a-priori knowledge on the layered structure of the head tissues.

A probabilistic tissue atlas based on local tissue probability distributions depending on the distances to two reference surfaces is employed to incorporate additional a-priori knowledge on the anatomy of the human head (Chap. 9).

For both proposed segmentation approaches results on experimental data are presented. The visual inspection of these results already gives some indication on the performance of the methods. Nevertheless, to demonstrate how accurate the methods really are a more thorough validation has to be performed. For this reason, we perform two validation studies in which we compare the results obtained by our approaches to reference segmentations based either on manual segmentations, or on computed tomography (CT) data. In this way, we finally prove that our methods yield accurate segmentation results which are a suitable basis for the construction of volume conductor models (Chap. 10).

1.2 Contributions

The contributions of this work can be grouped into three parts. The most important contribution is the development, presentation, and validation of two new methods for the automatic segmentation of volume conductor models of the human head from MRI data which can be used for EEG source analysis. The generated models might also be suitable for simulating brain stimulation techniques like transcranial direct current stimulation (tDCS) and transcranial magnetic stimulation (TMS).

Most notably, a segmentation approach was presented that incorporates a-priori knowledge on the anatomy of the head via a MRF model. Experiments on real data have shown that this method delivers accurate segmentations of all tissue regions relevant for EEG source analysis.

Furthermore, a probabilistic tissue atlas was developed which encodes a-priori information on the head anatomy in the form of local tissue probability distributions depending on the distances to scalp and WM reference surfaces. Accuracy and robustness of the proposed segmentation approaches could be further improved using this atlas.

The proposed segmentation methods were thoroughly validated by comparison to manual segmentations and to a reference segmentation based on a CT image. It was proven that our approaches are able to accurately segment the skull, and other tissues relevant for EEG source analysis, and that they improved upon a previously published, widely-used method.

The second major contribution of this thesis are two simulation studies that answer important questions regarding the accuracy and level of detail that are required for a volume conductor model that is to be used in EEG source analysis. The first study investigates the effects of skull geometry inaccuracies on the accuracy of the forward and inverse problem. From the results of the study criteria for the accuracy of the skull segmentation are derived.

The second study is concerned with the question how detailed the volume conductor model must be with regard to the CSF. In particular, a simplification of the CSF is investigated that is necessary for the combination of realistic head models and the dipole fit inverse method. It is found that the simplification is acceptable, thus, demonstrating a possible way to use the dipole

fit inverse method together with realistic EEG head models. The two studies have been published in peer-reviewed journals [75, 77] by the author of this thesis as corresponding author.

The presentation of an alternative pipeline for the extraction of anatomical information from MRI data for the purpose of EEG source analysis is the third major contribution of this thesis. Most notably a novel modification to the classic adaptive fuzzy c-means (AFCM) algorithm of Pham and Prince [105] was proposed, and it was shown that the modified algorithm is more accurate and more robust than the original AFCM algorithm. In addition, alternative methods for the reconstruction of a scalp surface and a cortex-based source space from MRI data are described.

Besides the previously mentioned two studies [75, 77] the author of this thesis during his Ph.D. work also contributed either as first or second author to two other publications [76, 39].

1.3 Organization of This Work

The contents of this thesis are organized as follows. Chapter 2 introduces the reader to EEG source analysis which is the context for the segmentation problem discussed in this thesis. The EEG forward problem and the role of realistic volume conductor models are described. Furthermore, previous publications on the generation of EEG head models are reviewed.

In Chapter 3 the segmentation problem is introduced. The basics of MR and CT imaging are explained to define our input data. A short description of the anatomy of the human head further specifies the segmentation problem. Next, the theoretical concepts underlying the segmentation approaches proposed in Chapters 7 and 8 are explained.

Chapters 4 and 5 present two simulation studies that were conducted to answer some open questions with regard to the required accuracy and level of detail of volume conductor models for EEG source analysis.

In Chapters 7 and 8 two different segmentation approaches are proposed that are intended to solve the segmentation problem detailed in Chapter 3. The methods are described in detail, experiments demonstrating their effectiveness are presented, and the performance of the methods is discussed.

Both approaches can be enhanced by incorporating a-priori knowledge provided by the probabilistic tissue atlas described in Chapter 9. It is explained how the atlas was generated and how the a-priori information can be combined with the proposed segmentation approaches.

Chapter 10 presents two validation studies that were performed to prove the efficacy of the proposed segmentation methods. Results of the validation are shown and discussed.

Finally, Chapter 11 summarizes the results of this thesis and gives an outlook on further related work that might be done in the future.

2 Individual, Realistic Volume Conductor Models for EEG Source Analysis

2.1 EEG Source Analysis

Neuroscience, that is, studying how the human brain processes information and how neurological diseases affect the human nervous system is a very wide and active area of research. A lot can be learned about the functioning of the brain by directly or indirectly observing the activity of the brain's neurons. EEG is a widely-used method to observe neuronal activity at a high temporal resolution. It measures electric potential differences on the surface of the scalp which originate from the electrical activity of neuronal cells, in particular, from the excitatory or inhibitory post-synaptic potentials [124]. A signal can be recorded in the EEG if at least a certain number of neurons (10000 to 50000) are synchronously active [93].

As the EEG signals are directly related to the activity of the neuronal cells it is possible to infer the neuronal sources from the EEG signal. This is the aim of EEG source analysis [90, 21]. In other terms, EEG source analysis tries to find the current density distribution in the brain which best explains the measured EEG signal. From the reconstructed current density distribution it is then possible to draw conclusions which neuronal sources are active and how they interact with each other.

Presurgical epilepsy diagnosis is a practical application of EEG source analysis. Epilepsy is a disease where neurons in some brain area exhibit highly synchronous, pathological activity patterns. The increased activity might spread to other brain areas and cause seizures. In epilepsy patients where the treatment with anti-epileptic drugs is unsuccessful resection of the brain area from which the pathological activity originates might render the patient seizure free. A prerequisite for surgical resection is the accurate localization of the pathological brain area. Commonly used diagnostic procedures for localizing the resection area are video monitoring, MRI, and single photon emission computed tomography (SPECT). Additional information for the localization of the origin of the seizures can also be gained by measuring the pathological activity patterns using EEG. EEG source analysis can then be used to reconstruct the location of the brain area from which the activity originates. Brodbeck et al. [23], for example, have shown that the location of the irritative zone can be predicted with a sensitivity of 84% and a specificity of 88% by analyzing high-density EEG recordings from patients with epilepsy.

Due to the high temporal resolution on the order of milliseconds EEG is also a valuable tool for observing the rapid and complex processing of informations in the brain. Auditory processing and the mismatch negativity are one example (see, e.g., [101, 35, 102, 64], or [96] for a review on the mismatch negativity). In the classical mismatch negativity paradigm a series of two different tones is presented to the subject where one tone occurs much more frequently than the other tone. Whenever the less frequent tone is presented to the subject a characteristic response can be observed in the EEG. The characteristic EEG response is called the mismatch

negativity. Source analysis can contribute to the localization of the cortical origin of this EEG component, and the high temporal resolution of the EEG allows it to investigate the fast interactions between different neuronal sources.

2.2 The EEG Forward Problem

In source analysis, two important problems must be solved: the inverse problem and the forward problem. The problem of finding the current density distribution which (in some sense) best explains the measured EEG signal is called the inverse problem. A variety of different methods were proposed to solve the inverse problem (see, e.g., [90, 21] for reviews).

All of the inverse methods have in common that they need to simulate the EEG for a given current density distribution. The simulation of the scalp EEG potentials from an assumed current density distribution is called the forward problem. An introduction to the forward problem and its solution is given here.

The relationship between the current density distribution and the electric fields and potentials is described by Maxwell's equations of electrodynamics. The general set of Maxwell's equations also governs temporal phenomena like magnetic induction or displacement currents. When considering the typical electrical activity of the neurons in the brain and the electrical properties of the human head, however, these temporal phenomena are small and can be ignored as it was shown by Plonsey and Heppner [106]. Ignoring the time-dependencies in Maxwell's equations leads to the quasi-static approximation which is the basis for the EEG forward problem.

$$\nabla \cdot (\sigma \nabla \phi) = \nabla \cdot \boldsymbol{j}_p \tag{2.1}$$

$$(\sigma \nabla \phi) \cdot \boldsymbol{n} = 0 \quad \text{on} \quad \Gamma \tag{2.2}$$

Here, σ, ϕ, and $\boldsymbol{j}_p$ denote the electrical conductivity of the head tissue, the electric potential and the primary currents, respectively. The second equation is the homogeneous Neumann boundary condition on the surface Γ of the volume conductor.

The current density $\boldsymbol{j}_p$ is commonly expressed as a superposition of mathematical dipoles.

$$\boldsymbol{j}_p(\boldsymbol{r}) = \sum_i^n \boldsymbol{M}^i \delta(\boldsymbol{r} - \boldsymbol{r}_0^i) \tag{2.3}$$

This formula describes the superposition of n mathematical dipoles at locations $\boldsymbol{r}_0^i$ and with moments $\boldsymbol{M}^i$.

Analytical solutions to the quasi-static Maxwell equations can be formulated for some particular volume conductor geometries only. A solution for a multi-layered spheroidal volume conductor in form of a series expansion was proposed by de Munck [40]. Berg and Scherg [12] describe another solution approximating the potentials on the surface of a four-layered spherical volume conductor as the superposition of the potentials of multiple sources in a homogeneous sphere. The advantage of these analytical approaches is that they can be evaluated very efficiently. This is useful when a large number of forward simulations has to be performed as it is usually the case in EEG source analysis. For this reason, the geometry of a subject's head is still

in some applications approximated as a multi-layered sphere so that these analytical solutions can be used to solve the EEG forward problem.

The rather crude approximation of the head geometry by multilayer sphere models introduces errors into source analysis as it was shown, for example, by Yvert et al. [162]. To avoid these errors the forward problem has to be solved in realistic volume conductor models. To do this, numerical approaches have to be employed.

One of the numerical approaches used is the Boundary Element Method (BEM) [56, 89, 41, 74, 136]. In the BEM the domain of the volume conductor Ω is represented by disjunct homogeneous compartments Ω_i with interfaces Γ_i. Integral equations for the potentials over the interfaces can be derived. The Γ_i are then discretized, most commonly using triangular meshes, and the potentials across the interfaces are represented using either linear or constant basis functions. The resulting system of equations can be solved to obtain the potentials at the scalp surface which solves the EEG forward problem.

The most current BEM approaches, the symmetric BEM or the Galerkin BEM using the isolated skull approach (ISA), both achieve very high accuracies in three compartment spherical models [74, 136]. A problem of all present BEM approaches is the fact that the computational demand increases quadratically with the number of unknowns. In four layer models where the CSF compartment is included the BEM integral equations also have to be solved on the cortex surface, that is, the interface between the brain and the CSF compartment. Due to the complex geometry of the cortex surface this adds on the order of at least several thousand unknowns to the BEM system of equations. The computational demand of solving the BEM in four compartment models is, thus, substantially higher and might even be considered too high for practical purposes. In the work of Kybic et al. [74], for example, six hours computation time on a 30 PC cluster were necessary for the simulation of a single source in a four compartment model using the symmetric BEM approach. The Galerkin BEM with ISA was reported [2] to be able to compute the lead field for approximately 7500 sources in a four compartment model in less than 3h which can be considered acceptable for most applications. The BEM mesh used for the simulations had around 20000 nodes which only allowed for a coarse approximation of the cortex surface. Due to the high computational demand for computing the BEM in high-resolution four compartment models, in practice commonly three compartment models representing the tissues scalp, skull, and brain are used. In these models the presence of the CSF is ignored. Another disadvantage of the BEM is its inability to model anisotropic conductivities. Using the BEM it is, thus, not possible to account for the anisotropic conductivity of the WM fiber bundles. Furthermore, model generation for the BEM is not trivial. The triangular surfaces describing the tissue interfaces must not intersect themselves or each other, each surface must be closed, and to achieve high numerical accuracies the meshing size must be adjusted to the distance between two interfaces. In practice, these conditions are hard to filfill.

The Finite Difference Method (FDM) [54], the Finite Volume Method (FVM) [32], and the Finite Element Method (FEM) [16, 26, 158] are used besides the BEM to solve the EEG forward problem in realistic head models.

The FEM has some beneficial properties because of which it was our method of choice for the studies conducted for this thesis. Background on the FEM is presented in the next section.

2.3 FEM Based Solution of the EEG Forward Problem

The FEM is a variational method widely used in many scientific and engineering applications. It is able to numerically solve a variety of partial differential equations, and it is especially suited for the highly accurate solution of elliptic partial differential equations, including the Poisson type equation 2.1, as it is specifically tailored to the corresponding variational formulation [47].

The FEM is used in this thesis to compute the EEG scalp potentials for the simulation studies presented in Chapters 4 and 5. Therefore, the basics of the FEM based solution of the EEG forward problem will be presented here mainly following the book of Braess [20].

In the context of the forward problem we have to deal with an elliptic differential equation that can be expressed using a differential operator L.

$$Lu = f \tag{2.4}$$

$$Lu \coloneqq -\sum_{i,k=1}^{n} \partial_i \left(a_{ik}\partial_k u\right) \tag{2.5}$$

$$f \coloneqq \nabla \boldsymbol{j}_p \tag{2.6}$$

Here, $u = \phi$ is our potential as defined in Equation 2.1, and $(a_{ik}) = \boldsymbol{\sigma}$. The differential operator L also defines a bilinear form $a(u, v)$.

$$a(u, v) \coloneqq \int_\Omega \left[\sum_{i,k} a_{ik}\partial_i u \partial_k v\right] \mathrm{d}x \tag{2.7}$$

A solution to the differential Equation 2.4 can be found by solving the variational problem

$$J(u) \coloneqq \frac{1}{2}\, a(u, u) - \, < l, u > \rightarrow \min! \tag{2.8}$$

$$\langle l, u\rangle \coloneqq \int_\Omega f\, u \,\mathrm{d}x \tag{2.9}$$

The characterization theorem [20, Thm. 2.2] states that the solution u minimizes $J(u)$ if and only if the following condition holds:

$$a(u, v) = \langle l, v\rangle \qquad \forall v \in V \tag{2.10}$$

V is here the linear space of solutions.

The conductivity $\boldsymbol{\sigma}$ jumps at the interface between two tissues. For this reason a classical solution $u \in C^2(\Omega) \cap C^0(\delta\Omega)$ for Equation 2.4 does not exist. A weak solution u in the Sobolev space $H_0^1(\Omega)$, however, does exist. The linear space of solutions is then also a subspace of $H_0^1(\Omega)$.

The idea of the FEM is to find a solution u_h in the finite element (FE) subspace V_h that satisfies Equation 2.10.

$$a(u_h, v) = \langle l, v\rangle \qquad \forall v \in V_h \tag{2.11}$$

The FE subspace V_h is made up from local polynomial basis functions. The solution domain, that is, in our application the head domain, is discretized into smaller subdomains. On these subdomains ansatz functions ${}^{(e)}\phi_h$ are defined. In our work trilinear ansatz functions defined on hexahedral subdomains are used (Eq. 2.12). The ansatz functions are zero everywhere outside of the subdomain.

$$ {}^{(e)}\phi_h = c_1 + c_2 x + c_3 y + c_4 z + c_5 xy + c_6 xz + c_7 yz \tag{2.12} $$

From the ansatz functions defined on the subdomains we can turn to node based form functions ${}^{(e)}\psi_k$.

$$ {}^{(e)}\phi_h = \sum_{k=1}^{p} {}^{(e)}u_k \, {}^{(e)}\psi_k \tag{2.13} $$

The u_k are called node variables, and p is the number of node based form functions that contribute to the ansatz function. The form functions ψ_k only have local support. They are equal to 1 at node k and 0 at all other nodes.

We can express our solution u_h to Equation 2.11 from the subset V_h as a linear combination of N form functions.

$$ u_h(x) = \sum_{i=1}^{N} u_i \, \psi_i(x) \tag{2.14} $$

For the derivation of the FEM matrix equation the Ritz-Galerkin approach is applied: we substitute the approximation of the potential from Equation 2.14 for the u_h in Equation 2.11. Equation 2.11 must be fulfilled for all functions $v \in V_h$. As $a(u, v)$ is a bilinear form and as V_h is a linear space it is sufficient that a solution fulfills the equation for the functions ψ_i which form a basis of V_h.

$$ \sum_{k=1}^{N} a(\psi_i, \psi_k) u_k = \langle l, \psi_i \rangle \tag{2.15} $$

Equation 2.15 can also be written in the form of a matrix equation:

$$ \boldsymbol{K}\boldsymbol{u} = \boldsymbol{J} \tag{2.16} $$

$$ K_{ik} = a(\psi_i, \psi_k) \tag{2.17} $$

$$ J_i = \langle l, \psi_i \rangle \tag{2.18} $$

$\boldsymbol{K}$ is called the stiffness matrix, $\boldsymbol{u}$ is the potential vector containing the potential values at the nodes of the FE mesh, and $\boldsymbol{J}$ is the source vector.

Different approaches exist to discretize the mathematical dipole source and to compute the source vector $\boldsymbol{J}$. For the FEM computations in this thesis the *blurred dipole* model [26] is used which is also called Venant approach. In the blurred dipole model the mathematical dipole (Eq. 2.3) is approximated by monopole current sources and sinks J_k distributed at nodes of the finite elmement mesh in the vicinity of the source. The monopole values are chosen in such a way that the resulting moment of the current density distribution approximates the dipole moment $\boldsymbol{M}^i$.

A different way of embedding the source singularity is the subtraction approach [158, 47, 146]. The idea of this approach is to split up the potential into a potential ϕ^{∞} which is due to a mathematical dipole in an infinite homogeneous medium, and a potential ϕ^{corr} which corrects for the heterogeneity of the volume conductor. ϕ^{∞} can be computed analytically. The correction potential ϕ^{corr} is obtained by solving a Poisson equation with inhomogeneous Neumann boundary conditions. Using the FEM this leads to a matrix equation of similar structure as that in Equation 2.16.

To compute the potential values at each FE node caused by a dipole source represented by the source vector $\boldsymbol{J}$ the FEM matrix equation 2.16 has to be solved with respect to the potential vector $\boldsymbol{u}$. A direct inversion of the stiffness matrix $\boldsymbol{K}$ which has the dimensions $N \times N$ is not possible for typical problem sizes where the number of FE nodes N is on the order of at least several hundred thousands. For this reason, iterative solvers are employed. The most widely used solver is the conjugate gradient (CG) method which is combined with a variety of different preconditioning schemes. The CG method with an algebraic multi-grid preconditioner (AMG-CG), for example, was shown to be very efficient in computing the solution of the FEM matrix equation [80].

Although efficient FEM solvers are available computation time is still an important factor when applying the FEM to EEG source analysis. When solving the inverse problem of EEG the potentials at the EEG electrodes have to be simulated for hundreds or thousands of different dipole sources. Directly solving the matrix equation for each source would take a too long time for any practical application. To accelerate the FEM computations for the solution of the inverse problem the transfer matrix approach was proposed [150, 157]. The transfer matrix approach allows to very efficiently compute the potentials at a subset of the FE nodes for a large number of different sources.

In EEG source analysis our interest is in the potentials at the EEG electrodes. When assuming that each electrode can be identified with an FE node then the electrode potential vector $\boldsymbol{u}^{\text{EEG}}$ can be obtained by applying the restriction matrix $\boldsymbol{R}$ to $\boldsymbol{u}$. The restriction matrix $\boldsymbol{R}$ is a $(s-1) \times N$ matrix where s is the number of electrodes. The element R_{ij} is 1 if electrode i is identified with the FE node j. Otherwise, the element is 0. Formally inverting the FEM matrix equation and applying $\boldsymbol{R}$ then leads to Equation 2.19 which defines the EEG transfer matrix $\boldsymbol{T}^{\text{EEG}}$, and which describes how it can be used to compute the potentials at the EEG electrodes for any given source.

$$\boldsymbol{u}^{\text{EEG}} = \underbrace{\boldsymbol{R}\boldsymbol{K}^{-1}}_{=:\boldsymbol{T}^{\text{EEG}}} \boldsymbol{J} \tag{2.19}$$

From the definition of the transfer matrix Equation 2.20 can be derived by multiplying $\boldsymbol{K}$ from the right and by transposing left- and right-hand side.

$$\boldsymbol{K}\left(\boldsymbol{T}^{\text{EEG}}\right)^{\mathsf{T}} = \boldsymbol{R}^{\mathsf{T}} \tag{2.20}$$

The transfer matrix $\boldsymbol{T}^{\text{EEG}}$ is computed row-by-row by solving Equation 2.20 $(s-1)$ times while using the columns of $\boldsymbol{R}^{\mathsf{T}}$ as right-hand sides. After having obtained the transfer matrix a forward computation can be performed by multiplying the sparse source vector $\boldsymbol{J}$ to the transfer matrix. Using the blurred dipole model with its sparse source vectors $\boldsymbol{J}$ this can be done very efficiently

requiring only few multiplication operations. In this way, the EEG forward computation is sped up whenever the number of sources is larger than the number of sensors. In the application of EEG source analysis this is virtually always the case.

2.4 Individual, Realistic Volume Conductor Models

EEG source reconstruction strongly profits from using individual models that realistically describe the volume conduction properties inside of the subject's head. Several previously published studies prove this statement. In an early publication Cuffin [37] reconstructed dipolar sources that had been implanted at known locations into the brains of three subjects using individual, realistically shaped models and simple spherical approximations. By comparing the known locations of the implanted sources and the reconstructed source positions Cuffin was able to assess the accuracy of the different head models. His study showed that localization errors were lower in the realistically shaped model especially for high signal-to-noise ratio (SNR). Yvert et al. [162] demonstrated in their simulation study that spherical models introduced localization errors of 4 to 6 mm in the upper part of the head, and 15 to 25 mm in the lower part of the head as compared to realistically shaped BEM models. Accuracy improvements by using realistic head models can be observed for virtually all different inverse methods. Steinsträter et al. [135], for example, investigated in simulations the differences between beamformer results obtained in realistic FEM models and in spherical models. For the latter the authors found large localization errors of up to 37 mm as compared to the realistic model.

A prerequisite for using individual, realistic models to solve the EEG forward problem is, of course, the existence of a description of the volume conduction properties of the human head. In particular, for the solution of the quasi-static approximation to Maxwell's equations the head's conductivity distribution must be known. Although there are some methods [126, 127] trying to directly measure the electrical conductivity distributions of biological tissues in general this is not possible with sufficient accuracy. Therefore, individual realistic volume conductor models are constructed by segmenting the subject's head. The segmentation divides the head into disjunct regions where each region is assumed to consist of a different tissue type. Individual conductivity distributions are then derived by assigning previously measured and published conductivity values to the segmented tissue regions.

Important aspects in the context of volume conductor modeling are the required level of detail, and the required geometrical accuracy of the head models. The required level of detail, that is, how many and which different tissues must be included in the realistic individual volume conductor model, and therefore segmented from the MRI, is still a topic of on-going research.

Routinely three-compartment volume conductors differentiating the tissues scalp, skull and brain are used in EEG source analysis (e.g., [79, 137, 148]). In these models the brain compartment also encompasses the liquor filled spaces between the cortex and the inner bone surface, in the sulci of the cortex and the ventricular system. This means, CSF is not differentiated from other brain tissues.

Ramon et al. [110] found that the simulated scalp potentials differ substantially depending on if the CSF is incorporated into the model or if it is neglected. In their simulation study they compare the forward solutions obtained with realistically shaped volume conductor models at

varying levels of detail. In another study Wendel et al. [153] used a reciprocal approach to show the influence of CSF in realistic head models. In simulations they send a current from one source electrode to a second sink electrode. Both electrodes were placed on the scalp. The authors then compared the induced current density distributions to investigate the model differences. Significant differences were observed between the models with and without CSF layer.

We were able to replicate these results in one of our own studies (Chap. 5). We also demonstrated the effects of neglecting CSF on the inverse problem in single source scenarios.

Another interesting aspect is the modeling of the skull. In a recent study [39], we investigated different ways of modeling the internal structure of the human skull. Large forward errors and large single dipole localization errors of up to 20 mm were found when not differentiating between compact and cancellous bone. We, therefore, conclude that the skull should be modeled as consisting of compact and cancellous bone if these two tissues can be reliably identified in the volume conductor segmentation.

Further detail might be added to a head model by differentiating between gray matter (GM) and white matter (WM). Ramon et al. [110] also studied this aspect. When not distinguishing between GM and WM localization errors well below 1 mm were observed. These errors are small enough to state that GM and WM do not have to be differentiated in individual realistic volume conductor models.

Literature can also be found on the influence on EEG source analysis of modeling WM using anisotropic conductivities. Wolters et al. [156] observed in a simulation study using a high-resolution FE mesh that the introduction of WM anisotropy causes the directions of the volume currents to become more parallel to the WM fibre tracts. In the simulation study of Güllmar et al. [53] mostly only small localization errors below 2.6 mm (50% percentile) were found when ignoring the WM anisotropy. Although, the observed localization errors were small ignoring the anisotropy can substantially influence the orientation of the reconstructed sources [52]. This might be relevant as the orientation of a source also contains information on the source's location [121]. Unfortunately, modeling the anisotropic conductivity of the WM is only possible if information on the individual anisotropic conductivity distribution is available. This information is almost exclusively derived from diffusion-weighted MRI (DW-MRI) data [140]. Some problems hinder the use of WM anisotropy in EEG source analysis. First, DW-MRI data is rarely available in routine clinical or neuropsychological research applications. Secondly, preprocessing the data poses some challenges. DW-MRI data is acquired using MR sequences which can introduce strong distortions due to susceptibility artifacts. These distortions must be corrected to obtain a faithful geometric representation of the WM anisotropy which is only possible using advanced correction procedures [119]. For this reasons WM is currently modeled as an isotropic compartment in nearly all applications.

Besides the level of detail, the second important aspect is the accuracy of the geometric description of the volume conductor model. In this context, Cuffin [36] investigated how small local thickness variations of skull and scalp affect source localization from EEG data. The author concluded that these local variations caused localization errors much smaller than 10 mm. A systematic evaluation of the source reconstruction errors caused by small changes in the positions of the tissue boundaries was performed by von Ellenrieder et al. [145]. Using first-order perturbation analysis they found that a perturbation of the tissue boundaries with a standard deviation of 1 mm resulted in a standard deviation of also around 1 mm for the reconstructed

source positions. This means, that the localization errors due to the perturbations of the tissue interfaces were very small. Previous studies leave some open questions especially with regard to the required accuracy of the skull segmentation. For this reason, we performed another study specifically investigating how inaccuracies in the skull geometry affect EEG source analysis (Chap. 4). The results of our study in general agree very well with the results of the two studies described above. For local errors in the skull geometry (small skull holes, local skull thickness mis-specifications, ...) only small errors were observed for the solutions of the EEG forward and inverse problems. Only when drastically altering the skull geometry across a wide region substantial errors occurred.

Taking into account all these findings from the literature it can be concluded that at least the tissues scalp, skull, CSF and brain must be incorporated into the volume conductor model. A segmentation approach must, thus, be able to distinguish these four tissues. A segmentation approach that is able to also segment cancellous and compact bone would be a benefit. Regarding the required accuracy of the segmentation it can be stated after reviewing the previous publications that small, local segmentation errors on the scale of around 1 mm are acceptable for the construction of EEG head models.

3 Head Tissue Segmentation for the Construction of Volume Conductor Models

3.1 Anatomy of Head and Skull

Focusing on the skull, the human head can roughly be separated into three regions: intracranial tissue, extracranial tissue and the skull itself.

The skull comprises all bony structures in the head. Anatomically it can be sub-divided into the neurocranium and the viscerocranium (Fig. 3.1). The bones of the viscerocranium, or the facial skeleton, make up the basis for the face. Neurocranium refers to the upper part of the skull, that encompasses the brain. The sources, which generate the EEG and magnetoencephalography (MEG) signals are situated inside of the neurocranium, and the sensors are also placed mainly above the upper part of the head. Therefore, the main focus of this work is on modeling the neurocranium.

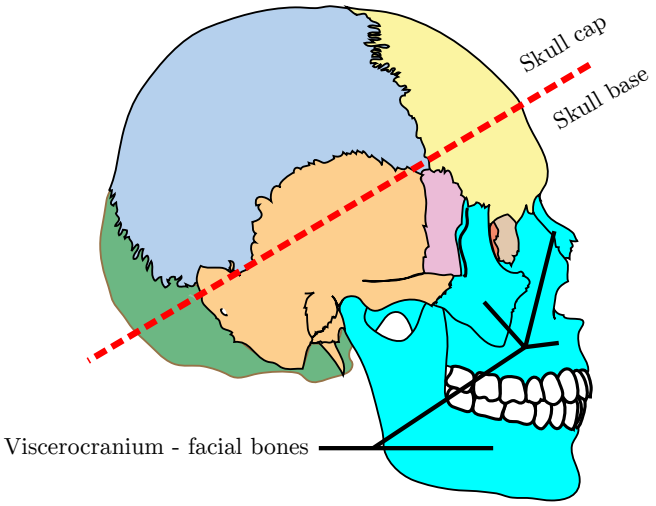

Figure 3.1: The basic anatomy of the skull showing the separation of the skull into the facial bones and the bones of the neurocranium. Also shown is the approximate location of the diving plane between skull base and skull cap. Figure derived from *Human skull side simplified (bones)*, M. R. Villareal, 2007, retrieved from *http://commons.wikimedia.org/wiki/File:Human_skull_side_simplified_(bones).svg*.

The neurocranium, again, can be divided into the calvaria, also called skull cap, and the skull base. There is no exact definition of the dividing plane between skull cap and skull base. The plane approximately passes through the center of the frontal and occipital bones, and passes

along the lower boundary of the parietal bones [123] (Fig. 3.1). Geometric properties of the calvaria differ strongly from that of the skull base. The calvaria has a dome-like shape with a thickness varying smoothly between around 1 and 12 mm [78]. It is thickest in occipital and thinnest in temporal areas.

The geometry of the skull base is much more complicated than that of the skull cap. Seen from inside the skull base can be divided into three regions of different indentation: the fossa anterior, the middle fossa and the fossa posterior. Furthermore the skull base exhibits more than a dozen holes through which arteries and nerves pass [123]. The largest of these holes is the foramen magnum in the fossa posterior through which the brain stem passes. Air filled cavities — the sinuses — are embedded into the bones of the skull base. The thickness of the skull base varies a lot. It can be thinner than a millimeter in the region above the orbitae, where the eyes are located, and it is much thicker, where sinus cavities are present.

In most regions the internal structure of the skull bones is not homogeneous. The bones reveal a layered internal structure with a spongy layer encompassed by two layers of compact bone — the inner and outer tables (Fig. 3.2). The spongy layer, called diploe, contains bone marrow, which is soft and fatty tissue. The typical thickness of the compact bone layers is around 1 to 2 mm [4], while the thickness of the diploe layer varies a lot. In regions where the skull is very thin there might be no spongious bone tissue whereas in thick regions of the skull the spongiosa can make up the larger part of the bone. In the vicinity of the suture lines there is no spongy bone tissue [123].

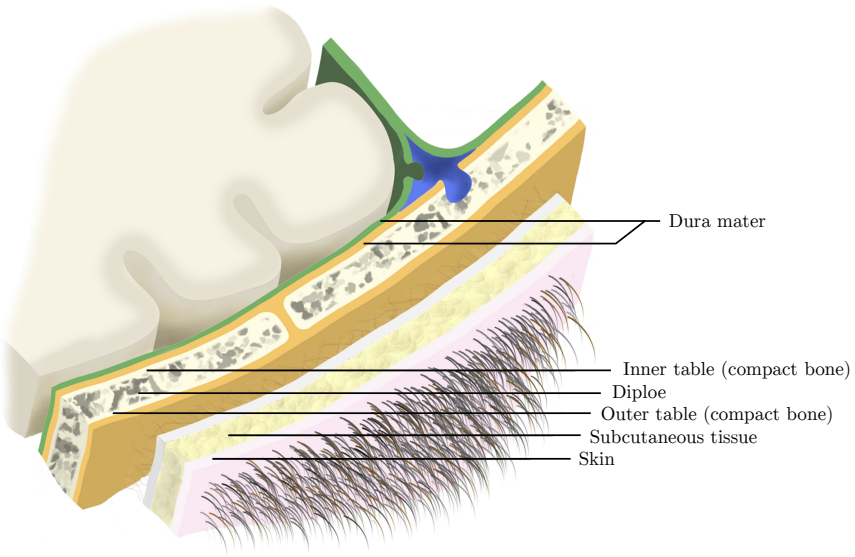

Figure 3.2: Anatomy of the skull, extra- and intracranial tissues. Figure derived from *Layers of the scalp and meninges*, F. Gaillard, 2010, retrieved from *http://commons.wikimedia.org/wiki/File:Layers_of_the_scalp_and_meninges.png*.

Directly adjacent to the inner skull boundary is the dura mater which is a membrane consisting of fibrous, connective tissue. Another type of membrane closely adheres to the surface of the

brain even dipping into the sulci. The space between the two membranes is mostly filled with CSF. The extent of the CSF filled spaces between the brain surface and the inner bone varies strongly between individuals.

The brain is accommodated in the intracranial cavity. Several brain structures can be differentiated anatomically, for example, the cerebrum, the cerebellum and the brain stem. Cerebellum and cerebrum are made up of two different kinds of tissues: the GM, consisting mainly of neurons, and the WM consisting mainly of nerve fibers which are connecting the different brain structures and regions.

On the outside the skull is lined by several membranes belonging to the scalp. In some regions of the skull, for example, in temporal or occipital regions, muscle tissue attaches to the skull. Next, a layer of subcutaneous tissue (SCT) follows. Finally, at the surface of the skin the dermis and epidermis are found.

3.2 Appearance of the Head in Different Imaging Modalities

This section intends to describe how the human head appears in typical volumetric medical images. First, an introduction into the basic principles of MRI will be given. Afterwards, artifacts commonly occurring in MR images which might influence the skull segmentation are discussed and the appearance of the different head tissues in MR images is described. Finally, CT imaging will be introduced as a reference modality for the segmentation of the human skull.

3.2.1 Basic Principles of Magnetic Resonance Imaging

In this subsection the basic principles of MRI will be presented. The generation of the MR signal and the image contrast in different MRI modalities will be explained. Furthermore, the spatial encoding, which is the basis for recording three-dimensional MR images, will be described. Enough detail will be given in the explanations, that it can be understood how the human head appears in the different MRI modalities and why some relevant imaging artifacts occur. The description in this section in parts follows Ax et al. [7], Petersen et al. [104] and Reiser [113].

MRI is based on the manipulation and measurement of the magnetization of hydrogen nuclei in the measured sample. The spin is a fundamental quantum mechanical property of elementary particles and atomic nuclei. A magnetic moment is linked to the spin. Although the spin is a quantum mechanical property in an external magnetic field the proton with its magnetic moment behaves like any other classical particle with a magnetic moment. Therefore we can use a classical model for our explanations of MRI.

When the sample containing protons is placed in a static magnetic field $\boldsymbol{B}_0$ which is oriented in z-direction by convention then two things happen. First, the magnetic moments of the protons start to precess around the direction of $\boldsymbol{B}_0$. And second, with time a net magnetization along the direction of the external magnetic field builds up as the proton spins partially align with the external magnetic field. The protons precess around the direction of $\boldsymbol{B}_0$ with the Larmor frequency ω_0:

$$\omega_0 = \gamma \cdot B_0$$

In this equation γ is the gyromagnetic ratio, which does have a constant value for protons, and B_0 is the strength of the static magnetic field.

Energy can be transfered to the system of precessing spins by exciting it with electro-magnetic waves at the Larmor frequency. The electro-magnetic waves are then in resonance with the frequency at which the spins precess. By exciting the spins with waves which are oscillating in a transverse plane to the direction of $\boldsymbol{B}_0$ the net magnetization of all spins is rotated away from the direction of the static magnetic field. In this way a net transverse magnetization builds up. This is analogous to applying a torque to a rotating top under the influence of a gravitational force. It has to be noted, that the exciting electro-magnetic waves only act on the spins contributing to the longitudinal magnetization. Furthermore, it is important that the spins which are rotated away from the direction of $\boldsymbol{B}_0$ are all precessing in phase after the excitation. How far the magnetization is rotated away from the direction of the $\boldsymbol{B}_0$ field is determined by how long the spins are exposed to the exciting waves and how strong these are. Electro-magnetic pulses which rotate the magnetization completely into the transverse plane are called 90° pulses while pulses rotating the magnetization into the opposite direction of the $\boldsymbol{B}_0$ field are called 180° pulses. In general, the exciting pulses are called radio-frequency (RF) pulses because the Larmor frequency for typical static magnetic fields is in the range of radio waves.

The net transverse magnetization of the spin system precessing around $\boldsymbol{B}_0$ constitutes an oscillating magnetic field and, therefore, electro-magnetic radiation at the frequency of their precession, that is, the Larmor frequency, is emitted. After excitation the emitted, observable signal quickly decays. This exponentially decaying signal oscillating at the Larmor frequency is called *Free Induction Decay* (FID) signal. During MR measurements the FID signal is measured using sensitive receiver coils.

The exponential decay of the FID is due to two basic relaxation mechanisms — spin-lattice and spin-spin-relaxation. When the sample is no longer exposed to the exciting RF pulse then the excited spins contributing to the transverse magnetization will with time return to parallel alignment with the static magnetic field. At first the transverse magnetization will be large and therefore the emitted FID signal is also strong. At a later point in time the transverse magnetization will have decreased and the FID signal will also decrease until the longitudinal magnetization is fully restored. This mechanism is called longitudinal or spin-lattice relaxation and its time constant is called T1. The time constant depends on the Larmor frequency — and therefore the B_0 field —, on the presence of large molecules, and the presence of paramagnetic ions and molecules among other factors.

An MR signal which is weighted by the T1 time constant of the sample can be measured by choosing a short repetition time (TR) between successive RF pulses. After an RF pulse it takes a certain time, specified by the time constant T1, until the excited spins have rotated back to parallel alignment with the external magnetic field and the longitudinal magnetization of the probe is restored (Fig. 3.3). As the RF pulse only acts on the spins contributing to the longitudinal magnetization the net transverse magnetization is smaller when the longitudinal magnetization was not fully restored at the time of excitation. When exciting with short TR then only in materials with small values of T1 the longitudinal magnetization will have time to restore before the next pulse, a larger net transverse magnetization will build up, and the generated signal will be stronger. On the other hand, in samples with long T1 times there is not enough time before the next RF pulse for the longitudinal magnetization to recover, and during

the successive excitation a smaller transverse magnetization will build up which in turn will lead to a FID signal with smaller amplitude. In this way at short TR materials with shorter T1 times will generate stronger signals than materials with longer T1 times.

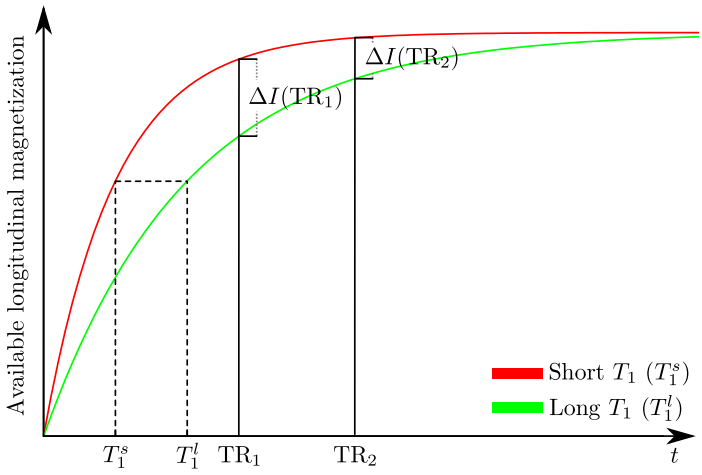

Figure 3.3: Longitudinal magnetization for tissues with short and long T1 times.

During the experiment we are only able to pick up a signal oscillating with the Larmor frequency because a significant number of excited spins are in phase after excitation, and the emitted electro-magnetic waves are then in phase, too. Due to local inhomogeneities of the magnetic field caused by neighboring spins, or by imperfections in the external magnetic fields the spins in the sample precess with slightly differing Larmor frequencies. With time the spins will therefore go out of phase so that the emitted electro-magnetic waves will also be out of phase and cancel each other out. In this way the measurable signal decays exponentially with time constant T2*. Using the spin-echo (SE) sequence of RF pulses, as it is explained below, the decay caused by inhomogeneities in the external magnetic fields is reversed and the decay of the measured signal is determined only by the time constant T2. The time constant T2 depends on the temperature and on tissue characteristics among other factors but not on inhomogeneities of the external magnetic fields.

A T2-weighted signal can be measured employing the SE pulse sequence (Fig. 3.4). In this sequence first an initial 90° RF pulse is applied to rotate the magnetization into the *x*-*y*-plane. Immediately after the RF pulse all spins are precessing in phase in the *x*-*y*-plane but with time dephasing will occur and the spins will go out of phase. Spins with higher Larmor frequency will soon be ahead of spins with lower Larmor frequency. After a certain time a second RF pulse is applied which rotates the spins by 180°. In this way the order of the precessing spins is reversed. The spins precessing slowest are now ahead and the spins precessing faster are behind. As the field inhomogeneities are still present the spins precessing at higher Larmor frequency catch up to the slower spins after exactly twice the time between the first and the second RF pulse. At this point in time the dephasing caused by the imperfections of the external magnetic fields, which are the same before and after the 180° pulse, is reversed and an echo of the initial FID signal can

be recorded. The dephasing caused by the interaction of the magnetic moments of neighboring spins are not reversed due to their random nature. Therefore the maximum amplitudes of the initial signal after the 90° pulse and the echo decay with time constant T2. The time at which the echo appears is called the *echo time* (TE).

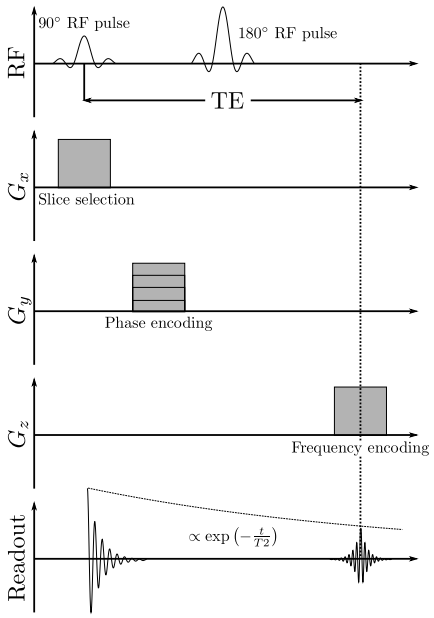

Figure 3.4: Sequence diagram for the SE sequence. Also shown is the timing of the spatial encoding gradients.

The protons which are involved in the generation of the MR signal, for example, during MR imaging on a human subject, are not isolated but they exist embedded in larger molecules. In these molecules the static external magnetic field $\boldsymbol{B}_0$ is shielded by the electrons accompanying the electrically neutral molecules so that the effective static magnetic field the protons see depends on the electron distribution of the molecule. This also affects the Larmor frequency with which the spins of the protons oscillate.

The MR signal measured during MR imaging on human subjects originates nearly exclusively from protons belonging to either water molecules or methylene groups [113]. Methylene groups occur as a part of fatty acids. A high proton density and for this reason also a strong signal is therefore found for tissues which contain a lot of water or fat molecules. In some situations it might be advantageous to suppress the signal from the protons embedded in the fat molecules and only record the signal originating from the water protons. Fat suppression can be achieved by using, for example, the following two techniques. The first technique is fat suppression by saturation [42]. In this method a 90° saturation pulse tuned to the Larmor frequency of the protons embedded in lipids and a spoiler gradient are applied before the actual excitation RF pulse. Due to the saturation pulse no longitudinal magnetization of the lipid spins remains which could be excited when the actual excitation RF pulse is applied. The spoiler gradient causes a fast dephasing of the spins excited by the saturation pulse so that the lipid spins do not emit

an observable signal. The second technique that shall be mentioned here is the fat suppression by selective water excitation [139, 59]. Specially tailored RF pulses are used in this method to excite solely the spins of protons embedded in water molecules. Those in lipids are not excited and, thus, do not contribute to the MR signal.

In MRI one is interested not in the MR signal generated by the sample as a whole but one is interested in measuring the MR signal from small volume elements so that a two- or three-dimensional image can be constructed. To achieve this the spatial location of the volume element, or voxel, must be encoded in the MR signal. The basis for the spatial encoding of the MR signal is the dependence of the Larmor frequency on the static magnetic field. The three mechanisms for encoding the location of a voxel in the MR signal are *slice selection*, *phase encoding* and *frequency encoding*.

For the slice selection a magnetic field gradient along, for example, the x-axis is switched on in addition to the $\boldsymbol{B}_0$ field. The Larmor frequencies of the spins in the sample are now larger where the gradient field is stronger and smaller where the gradient field is weaker. When exciting the sample with an RF pulse with small bandwidth then the Larmor frequency of spins only in a transverse slice of the probe are in resonance with the RF pulse and thus only the spins in this slice will be excited. The FID signal after the excitation will then also originate only from the selected slice.

When using phase encoding the location of the volume element is encoded in the phase of the FID signal. Before reading out the MR signal emitted from the excited spins an additional magnetic field gradient, for example, along the y-direction is switched on for a short time. During this time the spins at different positions along the y-axis experience different magnetic field strengths and therefore precess with different Larmor frequencies. This introduces phase shifts in the precession of spins along the y-axis. Measurements of the MR signal are repeated with phase encoding gradients of varying slope. When we plot the signal emitted from a volume element at a fixed y-coordinate against the slope of the phase encoding gradient, we obtain a curve oscillating with a certain frequency. This frequency is called phase-change ratio. Voxels further away from the center of the slice, that is, with larger y-coordinates, will see larger variations in the magnetic field and therefore these voxels will experience larger phase shifts. In this way, the phase-change ratio directly depends on the y-coordinate of the volume element. The MR signal we measure is then a superposition of signals with different phase-change ratios.

Frequency encoding utilizes a magnetic field gradient, for example, along the z-direction during the readout of the MR signal. Due to the magnetic field gradient spins along the z-axis now precess at different Larmor frequencies while we measure the signal they emit. Thus, the measured signal becomes a superposition of signals oscillating at different Larmor frequencies.

To record signals from which a two-dimensional image can be reconstructed a slice selection gradient is applied during RF excitation, a phase encoding gradient of varying slope is switched on before reading out the signal, and a frequency encoding gradient is employed during readout. As a result we obtain a two-dimensional matrix with values for n sampling times during the readout phase and k different phase encoding gradients. It is also possible to directly measure three-dimensional images by applying phase encoding gradients in two directions while not applying the slice selection gradient to excite the whole sample. This way one ends up with a three-dimensional matrix with values for n sampling times and $k_1 \cdot k_2$ different phase encoding gradients. Both, the frequency and phase encoded signals, can be regarded as superpositions

of signals with different frequencies, and the frequencies of the contributing signals are directly related to the coordinates of the volume elements from which the contributing signals originated. The individual frequency components of the measured MR signal can be reconstructed using the Fourier transformation. Thus, from the two- or three-dimensional matrices of MR signals measured at different sampling times and for different phase encoding gradients the signals from each individual volume element can be reconstructed.

3.2.2 Artifacts in MR Imaging

Typical MR images measured in the context of EEG source analysis are often affected by artifacts which impair an accurate segmentation. Some of the most common and most relevant artifacts will be discussed here.

Ideally, the MR signal should only depend on factors such as the proton density (PD) and the T1 and T2 time constants of the imaged tissue. In reality, the signals are also affected by noise from a number of possible sources. Sources include random thermal variations in the investigated sample and electric noise from the measurement equipment [10].

It can be assumed that the noise distribution in the measured FID signal can be described by a Gaussian distribution. To produce an image the FID signal is recorded for many different combinations of encoding gradients. An inverse Fourier transform is then applied to these values. This yields a signal with a real and an imaginary part. As the inverse Fourier transform is a linear operation the real and imaginary parts of the signal each for themselves are still normally distributed. This changes, however, when the magnitude of the complex signal is computed which is commonly done. Computing the magnitude is a non-linear mapping and, thus, the Gaussian distribution of the noise is not preserved. Instead, the measured magnitude signal is now distributed according to a Rician distribution [51].

$$p_I(I) = \frac{I}{\sigma^2} \exp\left(-\frac{(I^2 + A^2)}{2\sigma^2}\right) I_0\left(\frac{I\,M}{\sigma^2}\right) \tag{3.1}$$

Here, I is the measured signal corrupted with noise, A is the ideal noise-free signal, σ is the standard deviation of the Gaussian noise in the real and imaginary parts of the signal, and I_0 is the zeroth order Bessel function. Comparing Rician and Gaussian distributions it can be observed that they differ from each other substantially for low SNR values A/σ. Already for a SNR around $A/\sigma \approx 3$, however, the Gaussian distribution begins to be a reasonable approximation of the Rician distribution. For images with a SNR $\gtrapprox 3$ the Gaussian distribution can be used to describe the relation between the true magnitude signal and the measured signal.

Another artifact which often impairs the accurate segmentation of the skull from an affected MR image is the chemical shift artifact. It's origin is therefore explained below following Weinreb et al. [149].

Not only the protons embedded in water molecules but also protons embedded in lipids contribute to the MR signal. The Larmor frequency of the protons in water is approximately 420 to 440 Hz higher than in lipids. This difference is caused by the different effective magnetic field strengths acting on the protons in the two types of molecules. The differene between the Larmor frequencies interferes with the spatial encoding using the slice selection and frequency encoding

gradients. When applying the frequency encoding gradient protons in lipid and water molecules, thus, emit MR signals at different frequencies. The signals originating from the same volume element are then not reconstructed at the same spatial location along the gradient axis. The fat in the MR image seems to be shifted along the frequency encoding direction with respect to the aqueous part of the sample. In MR images with frequency encoding direction from superior to inferior, for example, the water-fat-shift artifact can cause parts of the fatty SCT to occlude the outer skull boundary (Fig. 3.5a). This would be fatal for segmentation accuracy.

The third artifact that will be discussed here is the intensity inhomogeneity or shading artifact (Fig. 3.5b). Ideally, the recorded MR signals should only depend on the biological properties of the tissue sample. A homogeneous region should then result in a region of homogeneous intensity in the MR image. Unfortunately, for common MR images this is not the case. Instead, a smooth spatial variation of the intensities of a homogeneous region can be observed. The inhomogeneity artifact can have several sources [133, 132, 131]. The excitation profile of the exciting RF coils can be non-uniform so that volume elements in some locations are excited stronger than volume elements in other locations. Similarly, a non-uniform sensitivity profile of the receiving coil can cause inhomogeneity of the measured intensities. Another possible source of intensity inhomogeneities are standing-wave and RF penetration effects caused by an interaction of the RF wave and the investigated tissue sample.

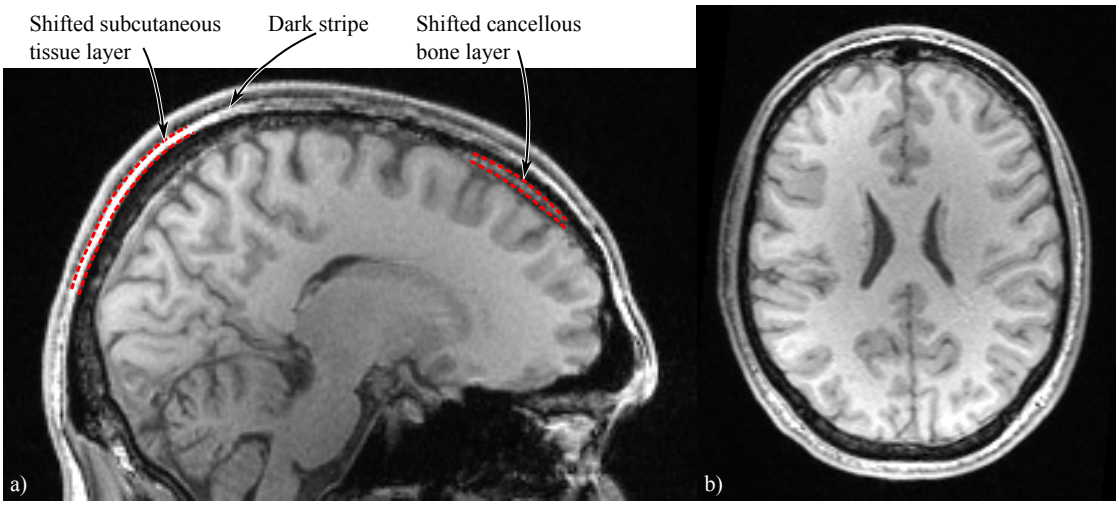

Figure 3.5: Illustration of two common MR artifacts. a) Water-fat-shift artifact. The signal from the fatty SCT and the cancellous bone layer is obviously shifted in inferior direction. A dark stripe can be seen where a part of the SCT is supposed to be located. b) Intensity inhomogeneity artifact. A clear inhomogeneity can be observed in the intensities of the WM especially in the occipital region;

3.2.3 Appearance of the Head in T1- and T2-weighted Magnetic Resonance Images

In this section, the appearance of the relevant head tissues in T1- and T2-weighted images will be described. Figure 3.6 presents slices of two typical T1-weighted MRIs and a T2-weighted

MRI for this purpose. In addition, a histogram of T1 gray values and a joint histogram of T1 and T2 gray values for each of the head tissues in a single MR data set are shown in Figures 3.7a and 3.7b, respectively.

Compact bone tissue is always displayed with low signal intensities in MRIs regardless of the used weighting. This is due to the low proton density in compact bone which is one of the main factors influencing the strength of the MR signal. The porous cancellous bone layer is partly composed of fatty tissue [112]. Therefore, its appearance is determined essentially by the appearance of that tissue. In T1- and T2-weighted images acquired without fat suppression the cancellous bone layer will be visible with medium to high intensities. It has to be noted that due to the chemical-shift artifact it might not be imaged at the correct location in the MRI but shifted by a few voxels. In images acquired with fat suppression sequences the cancellous bone layer is displayed with only low intensity and is nearly not distinguishable from the surrounding compact bone. In common T1-weighted images GM is displayed with medium intensity while WM appears brighter. With T2-weighting WM appears slightly darker than GM. The contrast between these two tissues is smaller in T2-weighted images than in T1-weighted images. Dura mater is displayed in both, T1-weighted and T2-weighted images, with comparable, medium intensities. This is different for the CSF. In T2-weighted images the CSF as all other liquids mainly consisting of water produces a strong signal [112] while it does produce nearly no signal in T1-weighted images. Skin and muscle tissue generates T1- and T2-weighted MR signals with medium intensities. The appearance of the fatty SCT again depends on the pulse sequence with which the MR was acquired. In sequences where fat suppression techniques are employed the SCT will be dark while without fat suppression it will generate a bright signal.

3.2.4 Basic Principles of X-Ray Computed Tomography

CT images only play a minor role in the studies presented in this work. This is because in most EEG source analysis applications CT images are not available for volume conductor segmentation due to the missing indication to expose the subject to the potentially harmful ionizing radiation of a CT measurement. A data set including a CT image will, however, be used for the generation of reference skull segmentations. Basic principles of CT, and the appearance of the human head in CT images are, thus, discussed in this section to explain why skull segmentations derived from CT images can play the role of reference segmentations.

CT imaging is based on measuring the attenuation of x-ray beams in the imaged sample. When a monochromatic beam propagates through a sample the intensity of the beam decreases as a function of the propagated distance and the attenuation characteristics of the sample. The relationship between the initial intensity $I(0)$ and the attenuated x-ray intensity after the beam has propagated a length of s can be given as [28]

$$I(s) = I(0) \cdot \mathrm{e}^{\int_0^s \mu(\eta)\,\mathrm{d}\eta}. \tag{3.2}$$

$\mu(\eta)$ is the attenuation coefficient at a distance η along the beam propagation path in the sample.

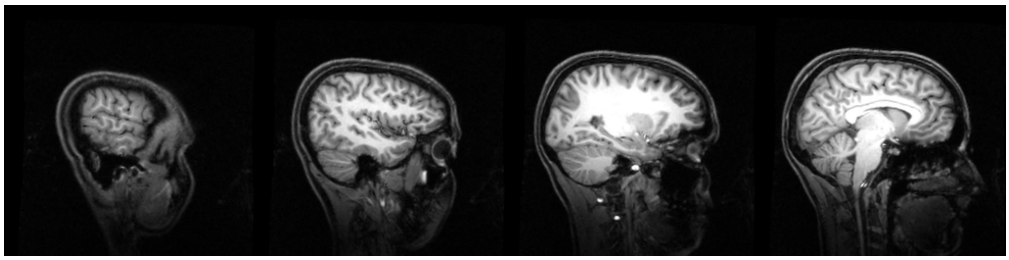

(a) Montage of sagittal slices of an exemplary T1 image with suppression of the fat signal achieved by using selective water excitation. The image was recorded using a 3D TFE sequence (TR = 9.32 ms, TE = 4.45 ms).

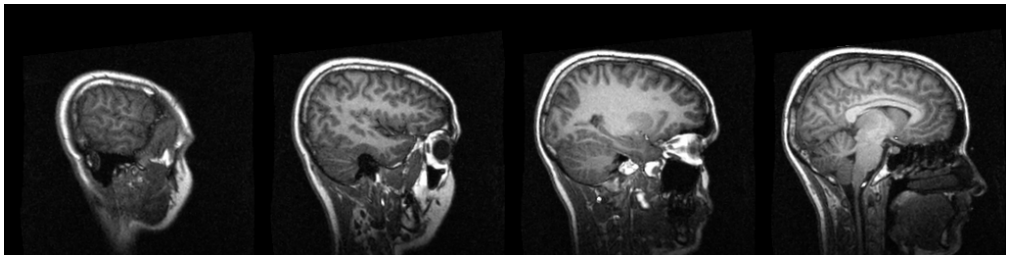

(b) Montage of sagittal slices of an exemplary T1 image with minimal water-fat-shift artifact. The image was recorded using a 3D TFE sequence (TR = 4.05 ms, TE = 1.81 ms, Bandwidth = 776.8 Hz).

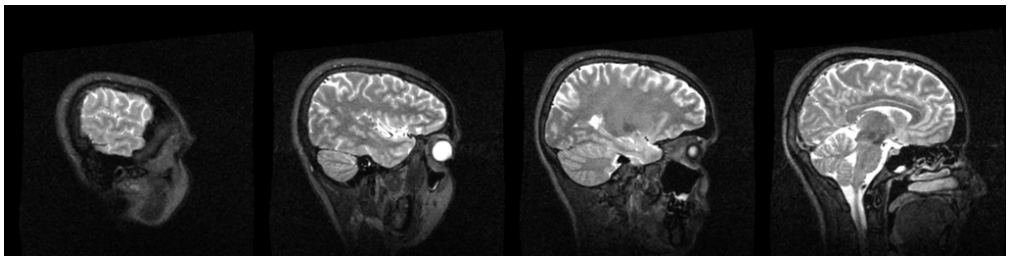

(c) Montage of sagittal slices of an exemplary T2 image. The image was recorded using a 2D TSE sequence (TR = 3656 ms, TE = 40 ms).

Figure 3.6: Sagittal slices of exemplary T1 and T2 images. All images were recorded with a resolution of $1.17 \times 1.17 \times 1.17\,\mathrm{mm}^3$.

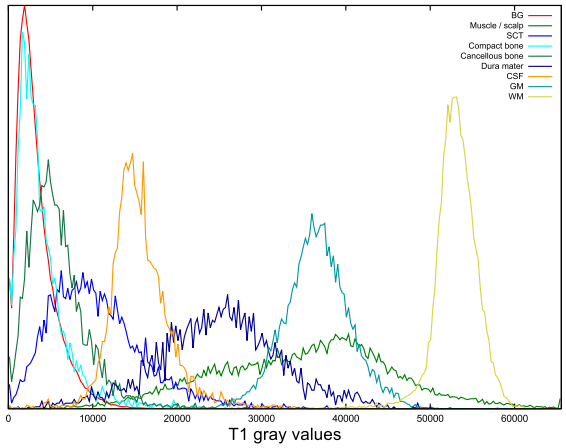

(a) Histogram of T1-weighted gray values for each of the head tissues. Tissues were labeled manually in a single T1-weighted image. The image was recorded using a 3D TFE sequence (TR = 9.32 ms, TE = 4.45 ms) with water-selective excitation.

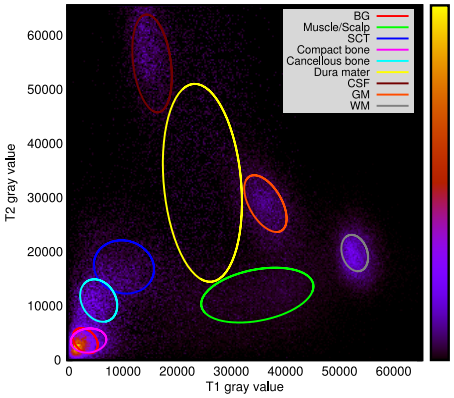

(b) Joint histogram of T1- and T2-weighted gray values for each of the head tissues. Tissues were labeled in a single data set containing a T1- and T2-weighted MRI. Ellipses depict the estimated covariance matrices for the gray values of each tissue type. The T1 image is identical to the one used in Fig. 3.7a. The T2 image was recorded using a 2D SE sequence (TR = 3429 ms, TE = 40 ms) with fat saturation.

Figure 3.7: Histograms of T1- and T2-weighted image intensities extracted from exemplary MRIs. BG: Background; CSF: Cerebro spinal fluid; SCT: Subcutaneous tissue; GM: Gray matter; WM: White matter.

From 3.2 the definition of the projection integral in 3.3 can be derived.

$$p(s) = -\ln\left(\frac{I(s)}{I(0)}\right) = \int_0^s \mu(\eta)\,\mathrm{d}\eta \tag{3.3}$$

For a fixed energy of the x-ray beam the attenuation coefficient is defined only by a few physical properties of the sample. The attenuation coefficient can be computed as [28]

$$\mu = \frac{\rho N_A}{A}\,\sigma_{\text{tot}} \tag{3.4}$$

From this relation it is obvious that the attenuation coefficient only depends on the density ρ, the atomic number A and the total photon atomic cross section σ_{tot} of the sample. N_A is the Avogadro constant.

As already mentioned above the attenuation coefficient also depends on the energy of the x-ray beam. To remove this dependency the attenuation is expressed in Hounsfield units by relating the attenuation of the sample to the attenuation of water. The relation between the actual attenuation and the attenuation in Hounsfield units (HU) is as follows:

$$\mu' = \left(\frac{\mu - \mu_{\text{water}}}{\mu_{\text{water}}}\right) \cdot 1000\,\text{HU} \tag{3.5}$$

In a CT measurement the attenuation of x-ray beams propagating along different paths through the sample is measured. In the original CT experiment [65], for example, an x-ray source emitting a needle-like beam is moved on a circular path around the sample and also tangentially to the circular path so that the x-ray beams propagate along different paths through the sample. On the opposite side of the x-ray source an x-ray sensor measures the intensity of the attenuated beam. This situation is depicted in Figure 3.8a.

The attenuation of the x-ray beam for a single path depends on the attenuation coefficients of all volume elements along the path. If we measure the attenuation of x-ray beams propagating through the sample along a multitude of paths then it is possible to reconstruct the attenuation coefficients of individual volume elements, $\mu(\boldsymbol{x})$, inside of the sample. This is the aim of the CT.

Fundamental work on this problem was presented in 1917 by the mathematician Johann Radon (reprinted in [109]). In it, Radon introduces an integral transform of a two-dimensional function $f(x, y)$ — the Radon transform. New coordinates (ξ, η) are introduced in a coordinate system rotated by γ with respect to the Cartesian coordinate system and the function f is integrated along a line with constant ξ. An illustration of the new coordinate system is shown in Figure 3.8b. This can be expressed as the convolution of the function f with the line function $\delta(L)$.

$$f(x, y) \overset{\mathcal{R}_2}{\circ\!\!-\!\!\bullet} f * \delta(L) = p_\gamma(\xi) = \iint f(x, y) \cdot \delta(L)\,\mathrm{d}x\,\mathrm{d}y \tag{3.6}$$

The two-dimensional Radon transform $\mathcal{R}_2$ can directly be related to the original CT experiment from Hounsfield [65]. It corresponds to the projection integral p_s from Equation 3.3 measured for an x-ray source rotated by γ around the object and translated tangentially by ξ. What is

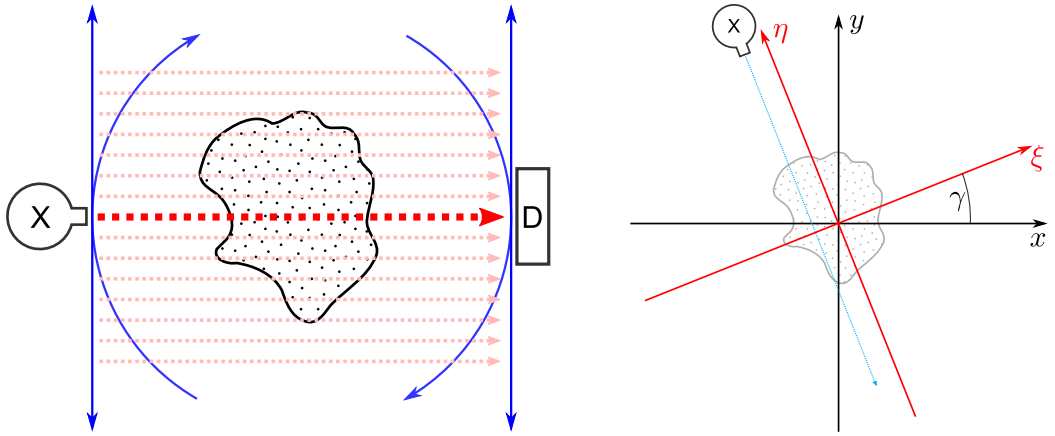

(a) Illustration of the CT measurement setup from [65]. *X* marks the x-ray source, *D* marks the detector.

(b) The Radon coordinate system in which the projection integrals are measured. *X* denotes the x-ray source.

Figure 3.8: The original CT measurement setup and the Radon coordinate system.

done in the CT measurement of Hounsfield is, thus, a sampling of the Radon transform of the attenuation coefficients $\mu(x, y)$ at discrete values of γ and ξ, and the inverse Radon transform will allow us to reconstruct the $\mu(x, y)$ values from the measured projection integrals. In principle, the inverse Radon transformation can be computed directly using Fourier transformation and exploiting the Fourier slice theorem.

The direct inversion of the projection integrals is rarely used in practice due to problems with the sampling in the polar space of the coordinates γ and ξ, the so-called, regridding problem.

Due to the difficulties in the direct inversion of the Radon transform, today, in almost all CT systems another approach, namely the filtered backprojection, is employed to reconstruct the attenuation coefficients from the measured projection integrals [28].

The original experiment of Hounsfield measures a two-dimensional image of the attenuation coefficients in the sample. Some approaches exist for recording three-dimensional CT images. In the conventional approach multiple two-dimensional images are recorded at different z-coordinates and combined to a three-dimensional image volume. A disadvantage of the conventional approach is the comparatively long measurement time, because the x-ray sensor and detector rotating on a circular path around the subject must be halted while the subject is translated perpendicularly to the image plane. Furthermore, due to an imperfect sensitivity profile of the x-ray beam throughout the slice thickness stair-case like artifacts can be observed in the 3D reconstruction of the imaged object. Another widely used method is spiral CT imaging [70, 144]. In spiral CT imaging the x-ray source is steadily rotating around the imaged object while the imaged object is steadily translated perpendicular to the image plane. From the perspective of the imaged object the x-ray source, thus, moves in a spiral path around the object. A problem for image reconstruction arises from the fact that projection integrals for a single volume element are not recorded from the necessary whole γ range of 180°. This is solved by interpolating the projection integrals for the complete γ range at discrete values of z. From the interpolated

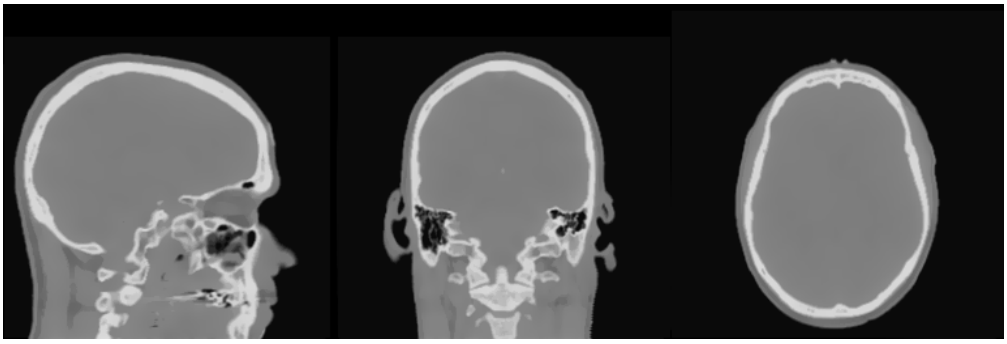

Figure 3.9: Sagittal (left), coronal (middle) and axial (right) slice of an exemplary CT image.

projection integrals slices of the three-dimensional image can then be reconstructed as in the two-dimensional case. Due to the continuous movement of the x-ray source measurement time can be greatly reduced. In addition, the stair-case like artifact in 3D structures are avoided in spiral CT imaging.

CT images typically exhibit a high contrast between soft tissue and bone structures when compared to the noise level. In addition, with current techniques high resolution images of the whole head can be acquired in reasonable short measurement times.

As can be seen in Figure (3.9) compact bone is represented by high intensities in CT images. This is due to the high density and, thus, higher attenuation of compact bone tissue. The less dense cancellous bone layer is for that same reason imaged with less intensity than compact bone. Soft tissues (e.g., brain, CSF, scalp, etc.) have a much lower attenuation and, therefore, appear much darker in CT images than cancellous and compact bone tissue. Thus, CT images exhibit a very good contrast between bone and surrounding soft tissues. Compared to MRI, CT also allows to differentiate between the dark air-filled sinuses and the surrounding bone.

A further benefit of CT imaging is that the resulting images do not exhibit any geometrical distortions as they are present in typical MRIs.

In summary, the good contrast between the skull bone and the surrounding soft tissues, the low noise, the available high-resolution images, and the faithful geometrical representation make CT a very well suited modality for the segmentation of the skull. Very accurate segmentations of the skull on the basis of a CT image can be achieved by using simple classification methods or even by thresholding of the image intensities.

3.3 Brief Survey on Existing Methods for Skull Segmentation From MRI

In the past decades a handful of methods have already been proposed for the segmentation of the head in the context of EEG source analysis. As before, the most critical aspect of the volume conductor segmentation is the delineation of the skull from MRI data. For this reason, many of the published methods focus on the segmentation of the skull from MRIs. We will briefly survey some of these methods.

One of the most widely used methods for generating skull and scalp segmentations is the method implemented in the *Brain Extraction Tool 2* (BET2) software [67]. BET2 segments the skull by first generating a smoothed, coarse mask of the brain surface [134]. For each of the vertices of the surface the MR intensity profile along a line perpendicular to that surface is inspected. Following some heuristic rules the intersection points of the line with the inner and outer skull boundary and the scalp surface are determined. Using a deformable surface approach, finally, the inner and outer skull, and the scalp surfaces are fitted to the boundary points found in the previous step.

Heinonen et al. [63] describe a semi-automatic skull segmentation approach in which they combine filtering, thresholding of the MRI intensities, and region growing.

In another publication by Dogdas et al. [46] thresholding of the MRI intensities, and morphological operations (closing, opening, selection of the biggest connected foreground component, ...) are combined to successively create masks for the brain, the scalp, and the inner and outer skull.

Thresholding and morphological operations are also used by Rifai et al. [115] to create an initial skull mask. This mask is then used to initialize a deformable surface model similar to the one used in the work of Malladi et al. [86]. The speed function of Rifai's model is designed in such a way that the surface propagation stops at the skull boundary taking also the partial volume effect into account.

Besides these approaches which mainly exploit the information contained in the MRI data and which only implicitly incorporate a-priori knowledge on the anatomy some approaches exist which strongly rely on anatomical a-priori models of the skull. Wang et al. [147], for example, build an active shape model [33] which captures the average geometry of the inner and outer skull boundaries and their principal variations. The model is then fitted to the individual MRI so that the skull boundaries are aligned with intensity gradients in the MR image.

In the work of Bertelsen et al. [14] a-priori information on the anatomy of the skull is implicitly recorded in a set of reliable skull segmentations for different subjects, called atlases in [14]. The individual MRI which is to be segmented is first roughly aligned with a template MRI using an affine registration. Next, the subset of the atlases whose associated MRI is most similar to the registered individual MRI is determined. Each of the atlases from this subset is then non-linearly co-registered to the individual MRI. The labels of the multiple co-registered atlases are finally fused by majority vote, that is, a voxel of the individual MRI is labeled as skull if the voxel is also labeled as skull in the majority of the co-registered atlases.

More notable work on segmentation of the skull by fitting geometrical a-priori models of the skull was done by Lüthi et al. [84].

3.4 The Segmentation Problem

In this section the basic segmentation problem is defined. A discrete and a continuous definition of the segmentation problem is given.

This work is concerned with the segmentation of medical image data. In practice, this data is nearly always defined on a finite three-dimensional regular grid.

Definition 1 (Grid Sites) *The vertices of the finite regular grid can be indexed using the three*

grid indices (j, k, l)*. All together they constitute a set of sites* $\mathcal{S}$*.*

$$\mathcal{S} = \{(j, k, l) \mid 0 \leq j < n^j, 0 \leq k < n^k, 0 \leq l < n^l\} \tag{3.7}$$

n^j*,* n^k*, and* n^l *are the number of sites along each of the three principal axes of the grid.*

It is also possible to unambiguously identify each site of the finite grid using just a single site index i.

$$\mathcal{S} = \{i \mid 0 \leq i < n = n^j n^k n^l\} \tag{3.8}$$

Our medical image data, for example the MR images, consist of observations y_i for each of the sites $i \in \mathcal{S}$:

$$\boldsymbol{y} = \{y_1, \ldots, y_n\}, \quad y_i \in \mathcal{D} \tag{3.9}$$

The observed MR signal for each site can in principal take on values from a continuous domain. In practice, however, the MR signal is discretized with a certain bit rate of, for example, 8 or 16 bit as in most cases. An observation y_i can thus take on the values in $\mathcal{D} = \{0, \ldots, 255\}$ for 8 bit sampling rate, respectively, the values in $\mathcal{D} = \{0, \ldots, 65535\}$ for 16 bit sampling rate. If $m > 1$ imaging modalities, for example, T1- and T2-weighted MRIs, are available then we have at each site a vector of observations $\boldsymbol{y}_i \in \mathcal{D}^m$. In the following, except where noted otherwise, we will assume that we are dealing with multi-modal image data.

In the discrete case, our segmentation problem now is to find a labeling of the given image data which is in some sense optimal. For each site i it shall be determined to which tissue region it belongs. In other words, we want to assign a label x_i to each voxel to obtain a labeled image x:

$$x = \{x_1, \ldots, x_n\} \tag{3.10}$$

Labels x_i are in a discrete finite domain $\mathcal{L} = \{0, \ldots, m-1\}$. The number of different labels m depends on how many different tissues we are able to differentiate and how many tissues we need to differentiate for our application. The number of regions m we have to differentiate for the application of EEG source analysis is discussed in depth in Section 2.4.

In the following we will also make use of binary masks to define, for example, a single tissue region. The binary masks $M = \{m_1, \ldots, m_n\}$ consist of binary values $m_i \in \{0, 1\}$ for each of the sites $i \in \mathcal{S}$.

In the continuous case the image data is assumed to be a function $u_0(\boldsymbol{r})$ defined on a subset of the three-dimensional space $\Omega \subset \mathbb{R}^3$. We are furthermore assuming that the image function can take on positive, continuous values.

$$u_0(\boldsymbol{r}) : \Omega \to \mathbb{R}_+ \tag{3.11}$$

According to the image model of Mumford and Schah [92] the image u_0 is homogeneous on disjunct subsets $\Omega_{0,k} \subset \Omega$, that means, $u_{0,k} \in H^1(\Omega_{0,k})$ [143]. Our task is now to find a segmentation of the image domain Ω.

Definition 2 *Let* Ω *be the image domain on which* u_0 *is defined. A segmentation of the image*

domain is then a partition $\mathcal{P}(\Omega)$ *of* Ω *into regions* Ω_k *such that*

$$\Omega_i \cap \Omega_j = \emptyset, \text{ if } i \neq j \qquad \textit{Disjunct regions} \tag{3.12}$$

$$\bigcup_k \Omega_k = \Omega \qquad \textit{No vacuum} \tag{3.13}$$

The continuous segmentation problem is then the problem of finding an optimal partition $\mathcal{P}(\Omega)$. What an optimal partition is has to be defined for each specific segmentation problem.

3.5 Bayesian Image Analysis

Bayesian image analysis is a useful framework in which a-priori information from different sources can be combined in a very natural way. For this reason, we are using the Bayesian framework to derive objective functions for the two segmentation approaches proposed in this work. The objective functions complete the specification of the segmentation problem by defining what an optimal segmentation is. In this section we will describe the basics of Bayesian image analysis. For the sake of clarity the description will use only the discrete formulation of the segmentation problem. All definitions can be stated in an analogous way for the continuous case.

In the Bayesian framework each observation, that is the image data at site i, is regarded to be a random variable which takes on random values according to some probability distribution. The set of observations at all sites $i \in \mathcal{S}$ then constitutes a random field $Y = \{Y_1, \ldots, Y_n\}$ with the actually measured, possibly multimodal gray values $\boldsymbol{y} = \{\boldsymbol{y}_1, \ldots, \boldsymbol{y}_n\}$ being a realization of the random field Y.

In the same sense the label at site i is interpreted as a random variable X_i and the entirety of random variables at all sites of the image grid constitutes the random field $X = \{X_1, \ldots, X_n\}$. The actual value of the random variable X_i is written as x_i and $x = \{x_1, \ldots, x_n\}$ denotes a certain configuration or labeling. Using this notation the probability that a certain configuration of labels occurs is then equal to the probability of the event $X = x$ which is short for the joint event $(X_1 = x_1, \ldots, X_n = x_n)$.

Let us first assume that our model does not contain additional parameters. The relation between the observations and the labels in the Bayesian formulation of the segmentation problem is then defined by the conditional probability density $l(\boldsymbol{y} \mid x)$ to observe the fixed image data $\boldsymbol{y}$ assuming that a certain labeling x is given. This probability density is also called the likelihood function of x. The likelihood encodes our model of how the image data (i.e., MRI images) is generated by an assumed labeling x.

Bayesian image analysis furthermore offers the possibility to declare a probability mass function $P(x)$ which encodes a priori information on the labeling x. By specifying a certain a-priori probability $P(x)$ a model of how the labeling is generated can be incorporated into the segmentation problem. In our application the a-priori probability can be used to encode a model of the anatomy of the human head. Incorporating such a prior can strongly aid segmentation performance and will be used extensively in the two proposed segmentation approaches.

A central aspect of the Bayesian framework is the introduction of the a-posteriori probability via Bayes' theorem.

Theorem 1 (Bayes' Theorem)

$$P(x \mid \boldsymbol{y}) = \frac{l(\boldsymbol{y} \mid x)\, P(x)}{p(\boldsymbol{y})} \tag{3.14}$$

The a-posteriori probability $P(x \mid \boldsymbol{y})$ is proportional to the product of the likelihood and the prior on x.

The proportionality constant $p(\boldsymbol{y})^{-1}$ only depends on our fixed data $\boldsymbol{y}$ and is, thus, ignored in the following. Equation 3.14 describes a way in which a prior can easily be incorporated into the Bayesian formulation of the segmentation problem. By evaluating the a-posteriori probability we are able to obtain the probability of the configuration x given the observed data $\boldsymbol{y}$.

In general, we can postulate that a labeling x^* is the optimal labeling if it minimizes the Bayesian risk.

Definition 3 (Bayesian Risk) *The Bayesian risk for a configuration x^* is defined as*

$$R(x^*) = \sum_{x \in X} C(x^*, x)\, P(x \mid \boldsymbol{y}) \tag{3.15}$$

$C(x^*, x)$ denotes a cost function specifying the costs of the configuration x assuming that x^* is the true configuration.

Our segmentation problem can, thus, be formulated as finding the x^* which minimizes the Bayesian risk given the fixed observations $\boldsymbol{y}$ for some cost function C, and for a choice of likelihood functions and a-priori probabilities.

A special choice for the cost function is

$$C(x^*, x) = \delta(x^*, x) = \begin{cases} 0 & \text{if } x^* = x \\ 1 & \text{else} \end{cases}$$

For this case the risk in Equation 3.15 simplifies to $R(x^*) = 1 - P(x^* \mid \boldsymbol{y})$.

Definition 4 (Maximum A-posteriori Estimation) *Let $C(x^*, x) = \delta(x^*, x)$ be the cost function for the Bayesian risk. The optimal solution to the segmentation problem can then be found by determining the configuration x^* which maximizes the a-posteriori probability.*

$$x^* = \arg\max_x P(x \mid \boldsymbol{y}) \tag{3.16}$$

This maximum a-posteriori probability (MAP) estimation is a common approach in medical image analysis (cf., e.g., [49, 166] and many more).

One advantage of the MAP estimation is that we do not have to evaluate the a-posteriori probability across the complete solution space to compute the sum in Equation 3.15 which might not be feasible. Instead, we can employ a more efficient optimization scheme to find the optimal labeling.

3.6 The Likelihood Function

The likelihood function describes the probability with which the image data $\boldsymbol{y}$ is observed given the labeling x.

In MR imaging the observed intensity depends on several factors. First, it depends on tissue characteristics, for example, T1 and T2 relaxation times, and proton density. Second, random noise from various sources influences the observed MR signal. For a given tissue the observed MR intensity can be described using a certain probability distribution. As a result of the MR imaging process the noise is distributed according to a Rician distribution. At reasonably high SNRs of around 3 the Rician distribution can be approximated by a Gaussian distribution (cf. Sec. 3.2.2). This work is concerned with medical MRI data. For this data we are assuming a sufficiently high SNR so that the Gaussian approximation is feasible.

For a single site i inside of region c the probability density distribution for observing the intensity $\boldsymbol{y}_i$ is, thus, modeled by a multivariate Gaussian distribution.

$$l(\boldsymbol{y}_i \,|\, x_i, \lambda_c) = \frac{1}{\sqrt{(2\pi)^k |\boldsymbol{\Sigma}_c|}} \exp\left\{-\frac{1}{2}(\boldsymbol{y}_i - \boldsymbol{\mu}_c)^\intercal \boldsymbol{\Sigma}_c^{-1}(\boldsymbol{y}_i - \boldsymbol{\mu}_c)\right\} \tag{3.17}$$

$|\boldsymbol{\Sigma}_c|$ here denotes the determinant of the covariance matrix. The multivariate Gaussian distributions in each of the regions $c = 1, \ldots, m$ depends on the parameters

$$\lambda = \{\lambda_1, \ldots, \lambda_m\} = \{\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1, \ldots, \boldsymbol{\mu}_m, \boldsymbol{\Sigma}_m\} \tag{3.18}$$

Here, $\boldsymbol{\mu}_c$ and $\boldsymbol{\Sigma}_c$ are the mean intensities and the covariance matrices, respectively, in the regions $c = 1, \ldots, m$.

For the sake of simplicity it is sometimes useful to assume that the intensities in the different image modalities are mutually independent. In this case, it is possible to factorize the multivariate Gaussian distribution into multiple univariate distributions with standard deviations σ_c^m for each region c and modality h

$$l(\boldsymbol{y}_i \,|\, x_i, \boldsymbol{\mu}_c, \boldsymbol{\sigma}_c) = \prod_{h=1}^{H} \frac{1}{\sqrt{2\pi}\sigma_c^h} \exp\left(-\frac{(y_i^h - \mu_c^h)^2}{2\sigma_c^{h\,2}}\right) \tag{3.19}$$

Under the assumption that the intensities at each voxel are mutually independent the likelihood for the complete image data $\boldsymbol{y}$ can be expressed as

$$l(\boldsymbol{y} \,|\, x, \lambda) = \prod_{c=1}^{m} \prod_{i \in \{i \,|\, i \in \mathcal{S}, x_i = c\}} l(\boldsymbol{y}_i \,|\, x_i, \lambda_c) \tag{3.20}$$

Instead of evaluating the likelihood directly often the logarithm of the likelihood is used. The commonly used multivariate Gaussian likelihood then takes on a conveniently simple form. This is one of the reasons why the intensity distribution is approximated using the Gaussian instead

of dealing with the more complicated Rician distribution.

$$\log l(\boldsymbol{y} \mid x, \lambda) = \sum_{c=1}^{C} \sum_{i \in \{i \mid i \in \mathcal{S}, x_i = c\}} \left\{ -\frac{1}{2} \log\left((2\pi)^k \left| \boldsymbol{\Sigma}_c \right| \right) - \frac{1}{2} \left(\boldsymbol{y}_i - \boldsymbol{\mu}_c \right)^\intercal \boldsymbol{\Sigma}_c^{-1} \left(\boldsymbol{y}_i - \boldsymbol{\mu}_c \right) \right\} \tag{3.21}$$

3.7 Parameter Estimation

Earlier on during the discussion of our segmentation problem we assumed that our model does not contain any additional parameters. However, the likelihood which contributes to the cost function does depend on parameters $\boldsymbol{\lambda}$. The values of these parameters must be known to completely specify the likelihood. In general and in our application in particular the values for the parameters are a-priori not known. Parameter estimation is, thus, an integral part of any approach for the solution of our segmentation problem.

In the Bayesian formulation the parameters are treated as random variables Λ with realizations $\boldsymbol{\lambda}$, that is in the same way as our observations Y and the labeling X are treated. We can therefore sum up the labeling and the additional parameters as a single random field $\Theta = \{X, \Lambda\}$ taking on the values $\theta = \{x, \boldsymbol{\lambda}\}$.

As shown in Section 3.6 our model depends on additional parameters through the likelihood function: $l(\boldsymbol{y} \mid x, \boldsymbol{\lambda})$. The a-posteriori probability must be extended to incorporate the additional parameters.

$$P(x \mid \boldsymbol{y}, \boldsymbol{\lambda}) \propto l(\boldsymbol{y} \mid x, \boldsymbol{\lambda})\, P(x) \tag{3.22}$$

The estimation of the unknown parameters $\boldsymbol{\lambda}$ is an important part of our segmentation problem. In the case, where the values of all random variables except the parameters are known the optimal parameter values $\boldsymbol{\lambda}^*$ can be estimated by using the maximum likelihood estimate [82, 116].

$$\boldsymbol{\lambda}^* = \arg\max_{\boldsymbol{\lambda}} \log l(\boldsymbol{y} \mid x, \boldsymbol{\lambda}) \tag{3.23}$$

For the application of image segmentation, however, the values for X_i are not known. The x_i, furthermore, are hidden, that means, they in principle cannot be observed. We are, thus, missing data — the unobservable x_i — to be able to estimate the parameters using Equation 3.23. Such a problem is called an incomplete data problem.

Some methods exist for solving the incomplete data problem. In the case where the gradient of the cost function can be computed the gradient descent method [164] can be used to determine both the labeling x and the parameters λ which locally minimize the cost function. The Expectation-Maximization (EM) method [43] describes another way of estimating the parameters which yield a maximum a-posteriori probability. Both approaches for the estimation of the labeling and the parameters will be described in conjunction with the segmentation approaches where they are employed (Secs. 3.8 and 3.9).

3.8 Active Contour Segmentation

Active contour segmentation approaches are a popular class of approaches for the segmentation of biomedical images [87]. Early work on active contours was done by Kass et al. [72], Malladi et al. [86] and Caselles et al. [29] to name just a few seminal papers.

The idea of active contour approaches is to solve the segmentation problem by computing the evolution of interfaces between the segmented regions. As an example, Kass et al. [72] described an active contour algorithm that evolves a contour represented by a 2D spline under the influence of internal and external forces. The authors present external forces that drive the contour towards regions with a large image gradient, that is, towards edges of objects in the image. In this way their active contour model is able to compute an edge based segmentation.

Another active contour approach also able to segment an input image by evolving a contour towards an image gradient was presented by Caselles et al. [29]. An important difference to the work of Kass et al. is how Caselles et al. represent the contour. Instead of using a parametric representation, like a spline, they used a non-parametric, implicit level set representation.

Level set functions are Lipschitz continuous functions from $\mathbb{R}^3 \rightarrow \mathbb{R}$ for 3D images, respectively, from $\mathbb{R}^2 \rightarrow \mathbb{R}$ for 2D images. The contours or interfaces are implicitly represented as the zero level set of these functions. Level set functions for the evolution of propagating fronts were first described by Osher and Sethian [100]. The level set representation has a number of advantages over an explicit, parametric representation. Its main advantage is that it naturally handles topological changes of the evolving surface [100]. Handling topological changes of parametric contours on the other hand is a tedious task requiring heuristic determination when the topology changes and explicit adaptation of the contour's topology. Furthermore, when using the level set representation the objective function for the segmentation can be parametrized on a fixed regular grid. This is an advantage over approaches where, for example, a surface is represented by triangular meshes, and where re-parametrization might be necessary during surface evolution when the size of the triangles becomes too large [30].

These advantages were also recognized by Chan and Vese [30] who proposed a level set based image segmentation approach that is able to deal with weak and spurious edges which pose a problem for edge based segmentation approaches. The authors later expanded their work by describing a multi-phase segmentation approach for the segmentation of the image domain into multiple regions at once [143]. In the latter paper it is shown that the proposed segmentation approach performs well for noisy images and images with weak edges. For these reasons one of the approaches for the segmentation of volume conductor models for EEG and MEG source analysis proposed in this thesis is based on the work from Vese and Chan [143].

For binary segmentation the partition $\mathcal{P}(\Omega)$ of the image domain Ω into two disjunct regions Ω_0 and Ω_1, so that $\Omega = \Omega_0 \cup \Omega_1$ and $\Omega_0 \cap \Omega_1 = \varnothing$, can be described using the sign of the level set function ψ_0 [30].

$$\Omega_0 = \{\boldsymbol{r} \mid \psi_0(\boldsymbol{r}) > 0\} \tag{3.24}$$

$$\Omega_1 = \{\boldsymbol{r} \mid \psi_0(\boldsymbol{r}) \leq 0\} \tag{3.25}$$

Equations 3.24 and 3.25 can be reformulated using characteristic functions $\chi_i, i \in \{0, 1\}$.

$$\Omega_i = \{\boldsymbol{r} \mid \chi_i(\boldsymbol{r}) = 1\} \tag{3.26}$$

The characteristic functions for the binary segmentation are defined using the Heaviside function $H(y)$ as

$$\chi_0(\boldsymbol{r}) = H(\psi_0(\boldsymbol{r})) \tag{3.27}$$
$$\chi_1(\boldsymbol{r}) = 1 - H(\psi_0(\boldsymbol{r})) \tag{3.28}$$

The key idea of Vese and Chan's [143] multi-phase level set approach is to represent the segmentation of the image domain in up to 2^n regions using n level set functions $\psi_i, i = 1, \ldots, n$. One way to do so is to identify a region by a characteristic combination of the signs of multiple level set functions.

First, a vector of level set functions and the vector Heaviside function are introduced.

$$\boldsymbol{\psi} = \begin{pmatrix} \psi_0 \\ \psi_1 \\ \cdots \\ \psi_n \end{pmatrix}, \qquad \boldsymbol{H}(\boldsymbol{\psi}) = \begin{pmatrix} H(\psi_0) \\ H(\psi_1) \\ \cdots \\ H(\psi_n) \end{pmatrix} \tag{3.29}$$

Given these we can define characteristic functions by introducing an unique constant vector $\boldsymbol{H}_i$ for each region i where $H_i^j \in \{0, 1\}$.

$$\chi_i = \begin{cases} 1 & \text{if } \boldsymbol{H}(\boldsymbol{\psi}) = \boldsymbol{H}_i \\ 0 & \text{else} \end{cases} \tag{3.30}$$

For n level set functions 2^n different constant vectors $\boldsymbol{H}_i$ can be constructed. Using n level set functions it is, thus, possible to encode $m = 2^n$ regions Ω_i.

The segmentation approach presented by Vese and Chan [143] is based on the optimization of an energy functional first proposed in the context of image segmentation by Mumford and Shah [92]. The active contours approach presented in this thesis will not use the Mumford and Shah energy functional. Instead, the segmentation problem will be described in the framework of Bayesian image analysis. However, the final objective functions are nearly identical using either the Mumford and Shah model or the Bayesian formulation.

In the context of Bayesian image analysis the segmentation problem can be defined as finding the maximum of the a-posteriori probability defined in Equation 3.14. For practical reasons we will try to minimize the negative logarithm of the a-posteriori probability which is equivalent to maximizing the a-posteriori probability. Equation 3.31 is the basic objective function for the active contours segmentation.

$$E(\mathcal{P}(\Omega)) = -\log p(\mathcal{P}(\Omega) \mid \boldsymbol{y}) = -\log l(\boldsymbol{y} \mid \mathcal{P}(\Omega), \lambda) - \log p(\mathcal{P}(\Omega)) \tag{3.31}$$

$p(\mathcal{P})$ here is an a-priori probability for the segmentation, and $l(\boldsymbol{y} \mid \mathcal{P}, \lambda)$ is the likelihood.

Different priors $p(\mathcal{P})$ can be specified to incorporate a-priori knowledge into the segmentation problem. The geometric interpretation of the level set function with the zero level set implicitly representing the interface between regions allows to define a-priori probabilities taking some geometric properties of the segmentation into account. Choosing the following prior, for example, favors segmentations with smaller interfaces between two adjacent regions:

$$p(\mathcal{P}(\Omega)) \propto \exp\left(-\epsilon |C(\mathcal{P}(\Omega))|\right) \tag{3.32}$$

$|C(\mathcal{P}(\Omega))|$ here denotes the area of the interface between the regions in the partition $\mathcal{P}$, and ϵ is a parameter determining the strength of the influence of this prior.

For the case of a binary segmentation the surface area of the interface between the two regions can be easily computed using the level set representation and the characteristic functions defined in Equations 3.27 and 3.28.

$$|C| = \int_\Omega \left(|\nabla \chi_0(\boldsymbol{r})| + |\nabla \chi_1(\boldsymbol{r})|\right) \, \mathrm{d}\boldsymbol{r} \tag{3.33}$$

For the segmentation into multiple regions we can use Equation 3.30 and write [143]

$$\begin{aligned} |C| &= \sum_{0 \le i < m} \int_\Omega |\nabla \chi_i(\boldsymbol{r})| \, \mathrm{d}\boldsymbol{r} && (3.34)\\ &\approx \sum_{0 \le i < n} \int_\Omega |\nabla H(\psi_i)| \, \mathrm{d}\boldsymbol{r} && (3.35)\\ &= \sum_{0 \le i < n} \int_\Omega \delta(\psi_i) \, |\nabla \psi_i| \, \mathrm{d}\boldsymbol{r} && (3.36) \end{aligned}$$

In the second step we approximated the sum of surface areas across each region by the sum of surface areas of the zero level set across each level set function. Due to this approximation the area of some interfaces might be taken into account more than once. However, this is still assumed to be a valid approximation in general [143].

Another prior on $\mathcal{P}$ applicable to binary image segmentation is given in Equation 3.37.

$$p(\mathcal{P}(\Omega)) \propto \exp\left(\nu V(\Omega_1)\right) \tag{3.37}$$

$V(\Omega_1)$ is the volume of the foreground region, and ν is a weighting factor. Using this prior favors segmentations with an as large as possible foreground region Ω_1.

Again using the level set representation of the binary segmentation the volume $V(\Omega_1)$ can be easily expressed as

$$V(\Omega_1) = \int_\Omega \left\{1 - H(\psi_0(\boldsymbol{r}))\right\} \, \mathrm{d}\boldsymbol{r} \tag{3.38}$$

In the previously presented Bayesian framework we can also incorporate a-priori tissue probability maps $P_c(\boldsymbol{r})$, $c = 1, \ldots, m$, where $P_c(\boldsymbol{r})$ denotes the probability that tissue region c occurs at location $\boldsymbol{r}$. Such a-priori probability maps could, for example, be derived from an empirically determined probabilistic tissue atlas. The logarithm of the a-priori probability for the complete

segmentation $\mathcal{P}(\Omega)$ is then computed from the logarithm of the probability maps by integration across the segmented regions.

$$\log p(\mathcal{P}(\Omega)) = \sum_{c=1}^{m} \int_{\Omega} \log P_c(\boldsymbol{r})\, \chi_c(\boldsymbol{r})\, \mathrm{d}\boldsymbol{r} \tag{3.39}$$

In addition to the prior $p(\mathcal{P})$ the likelihood function contributes to the objective function of our segmentation approach. A possible choice for the likelihood function is the multivariate Gaussian distribution (Eq. 3.20). The logarithm of the likelihood from Equation 3.21 can easily be rewritten using the level set functions, respectively, the derived characteristic functions.

$$\begin{aligned} \log l\,(\boldsymbol{y} \,|\, \mathcal{P}(\Omega), \lambda) = \sum_{c=1}^{m} \int_{\Omega} \Bigg\{ & -\frac{1}{2} \log\left((2\pi)^k \,|\boldsymbol{\Sigma}_c|\right) \\ & -\frac{1}{2}\, (\boldsymbol{y}(\boldsymbol{r}) - \boldsymbol{\mu}_c)^{\intercal}\, \boldsymbol{\Sigma}_c^{-1}\, (\boldsymbol{y}(\boldsymbol{r}) - \boldsymbol{\mu}_c) \Bigg\} \chi_c(\boldsymbol{r})\, \mathrm{d}\boldsymbol{r} \end{aligned} \tag{3.40}$$

Here, the sum across the discrete sites $i \in \{i \,|\, i \in \mathcal{S}, x_i = c\}$ from Equation 3.21 has been replaced by the integral across the region Ω_c which in turn was substituted by the integral $\int_{\Omega}\{\dots\}\chi_c(\boldsymbol{r})\,\mathrm{d}\boldsymbol{r}$.

Incorporating only the smoothness prior and employing the level set representation the objective function for the multiphase active contour segmentation (Eq. 3.31) can be written as follows.

$$\begin{aligned} E(\boldsymbol{y}, \mathcal{P}(\Omega), \lambda) \propto \epsilon & \sum_{i=1}^{n} \int_{\Omega} |\nabla H\,(\psi_i(\boldsymbol{r}))|\, \mathrm{d}\boldsymbol{r} \\ & + \sum_{c=1}^{C} \int_{\Omega} \Bigg\{ \frac{1}{2} \log\left((2\pi)^k \,|\boldsymbol{\Sigma}_c|\right) \\ & + \frac{1}{2}\, (\boldsymbol{y}(\boldsymbol{r}) - \boldsymbol{\mu}_c)^{\intercal}\, \boldsymbol{\Sigma}_c^{-1}\, (\boldsymbol{y}(\boldsymbol{r}) - \boldsymbol{\mu}_c) \Bigg\} \chi_c(\boldsymbol{r})\, \mathrm{d}\boldsymbol{r} \end{aligned} \tag{3.41}$$

The likelihood function in the a-posteriori probability depends in our case of a multivariate Gaussian likelihood function on unknown parameters $\boldsymbol{\lambda}$ so that the parameter values must be estimated simultaneously with the partition of the image domain. In summary, the segmentation problem of the active contours segmentation at this point is to optimize the objective function in Equation 3.41 with respect to the segmentation $\mathcal{P}(\Omega)$ represented by the level set functions ψ_i, and the parameters $\boldsymbol{\lambda} = \{\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1, \dots, \boldsymbol{\mu}_m, \boldsymbol{\Sigma}_m\}$.

The problem of optimizing the functional with respect to the shapes Ω_i is called a shape optimization problem [27]. A solution to this problem based on variational calculus was presented by Zhao et al. [164]. For an arbitrary perturbation ϕ of the level set functions ψ the Frechet derivative of the functional is computed. At the minimum of the functional this derivative has to be zero for any perturbation. This leads to a set of differential equations — sometimes called the Euler-Lagrange equations [30] — which are the necessary condition for a minimum of the objective function. The differential equations also define the gradient of the functional.

The optimization is performed in a dynamical scheme where the shapes are assumed to evolve

over an artificial time t so that $\Omega_i = \Omega_i(t)$, $\psi_i = \psi_i(t)$, and $E(\boldsymbol{y}, \mathcal{P}(\Omega(t)), \lambda) = E(t)$.

The gradient descent method is then used to actually minimize the objective function with respect to the level set functions ψ_i and the parameters $\boldsymbol{\mu}_c$ and $\boldsymbol{\Sigma}_c$. The basic principle of the method is to project the current solution $\psi_i^{(t)}$ along the direction of the gradient to obtain a new improved solution $\psi_i^{(t+1)}$.

Using a fixed time step Δt the level set functions ψ_i at the artificial time $(t+1)$ are in this way obtained by updating the current solution $\psi_i^{(t)}$ using the following update equation.

$$\psi_i^{(t+1)} = \psi_i^{(t)} - \Delta t\, \delta_\epsilon \left(\boldsymbol{U}_i^{\text{likelihood}} + \boldsymbol{U}_i^{\text{probability map}} + \boldsymbol{U}_i^{\text{inflation}} + \boldsymbol{U}^{\text{smoothness}} \right) \tag{3.42}$$

The $\boldsymbol{U}_i^*$ terms here stand for the terms originating from the likelihood, the a-priori probability maps from Equation 3.39, the inflation prior from Equation 3.37, and the smoothness prior from Equation 3.32, respectively. The terms can be expressed as

$$\boldsymbol{U}_i^{\text{likelihood}}(\boldsymbol{r}) = \sum_{c=1}^{n} \left(\frac{1}{2} \log\left((2\pi)^k \left| \boldsymbol{\Sigma}_c \right| \right) + \frac{1}{2} \left(\boldsymbol{y}(\boldsymbol{r}) - \boldsymbol{\mu}_c \right)^\intercal \boldsymbol{\Sigma}_c^{-1} \left(\boldsymbol{y}(\boldsymbol{r}) - \boldsymbol{\mu}_c \right) \right) S_c^i \tag{3.43}$$

$$\boldsymbol{U}_i^{\text{probability map}}(\boldsymbol{r}) = \sum_{c=1}^{n} \left(-\log P_c \right) S_c^i \tag{3.44}$$

$$\boldsymbol{U}_i^{\text{inflation}}(\boldsymbol{r}) = -\nu S_0^i \tag{3.45}$$

$$\boldsymbol{U}_i^{\text{smoothness}}(\boldsymbol{r}) = \epsilon \nabla \left(\frac{\nabla \psi_i^{(t)}}{|\nabla \psi_i^{(t)}|} \right) \tag{3.46}$$

The term S_c^i originates from the derivative of the characteristic functions χ_c. It is defined as

$$S_c^i = \left(2H_c^i - 1 \right) \prod_{\substack{j=1,\dots,n \\ j \neq i}} \left(2H_c^j H_\epsilon(\psi_j^{(t)}) - H_\epsilon(\psi_j^{(t)}) - H_c^j + 1 \right) \tag{3.47}$$

Regularized Heaviside and Delta functions H_ϵ and δ_ϵ had to be introduced here as described by Chan and Vese [30] so that the objective function is from C^1 which is a necessary condition to compute the derivative.

Updated values for $\boldsymbol{\mu}_c$ and $\boldsymbol{\Sigma}_c$ are determined by directly optimizing the objective function with respect to the parameters of the likelihood function. Equating the derivative of Equation 3.41 with respect to the mean intensities $\boldsymbol{\mu}_c$ to zero for each region Ω_c and solving for $\boldsymbol{\mu}_c^{(t+1)}$ yields

$$\boldsymbol{\mu}_c^{(t+1)} = \frac{\int_\Omega \boldsymbol{y}(\boldsymbol{r}) \chi_c^{(t+1)}(\boldsymbol{r})\, \mathrm{d}\boldsymbol{r}}{\int_\Omega \chi_c^{(t+1)}(\boldsymbol{r})\, \mathrm{d}\boldsymbol{r}} \tag{3.48}$$

Equation 3.48 amounts to computing the average intensity across region Ω_c.

In a similar way equations for the covariance matrices Σ_c can be derived.

$$\boldsymbol{\Sigma}_c^{(t+1)} = \frac{\int_\Omega \left(\boldsymbol{y}(\boldsymbol{r}) - \boldsymbol{\mu}_c^{(t+1)} \right) \left(\boldsymbol{y}(\boldsymbol{r}) - \boldsymbol{\mu}_c^{(t+1)} \right)^\intercal \chi_c^{(t+1)}(\boldsymbol{r})\, \mathrm{d}\boldsymbol{r}}{\int_\Omega \chi_c^{(t+1)}(\boldsymbol{r})\, \mathrm{d}\boldsymbol{r}} \tag{3.49}$$

In the gradient descent scheme the optimization of the level set functions ψ_i and of the parameters $\boldsymbol{\mu}_c$ and $\boldsymbol{\Sigma}_c$ are iterated until the algorithm converges.

In general it cannot be proven that a gradient based level set optimization as described in this section converges to a minimum of the cost function [27]. Evolution towards a stationary point, however, is certain as the method decreases the objective function (for a sufficiently small time step Δt) at every iteration.

The multi-phase level set segmentation procedure yields as results the final level set functions $\boldsymbol{\psi} = (\psi_0, \ldots, \psi_m)$, and the estimated parameters $\boldsymbol{\mu}_c$ and $\boldsymbol{\Sigma}_c$. A labeling of the image can be generated by evaluating the characteristic functions $\chi_i(\boldsymbol{\psi})$ for each voxel of the image grid. The resolution of the labeling is on the order of Δx which is the step size of the image grid.

A more accurate representation of the segmentation can be obtained by interpolating the interfaces between the segmented regions from the level set functions $\boldsymbol{\psi}$ using, for example, the Marching Cubes algorithm [83]. The active contours segmentation procedure has, thus, the potential to generate a segmentation with sub-voxel accuracy. The interfaces reconstructed from $\boldsymbol{\psi}$ will, in addition, be smooth and they will not exhibit the staircase-like artifacts which can be observed for interfaces reconstructed from binary images. This makes the reconstructed interfaces suitable for visualization purposes without any additional smoothing.

3.9 Markov Random Field Segmentation

The a-priori probability in Bayes formula (Eq. 3.14) allows to incorporate a-priori knowledge on the sought-after labeling. One way to define an informative a-priori probability distribution is to use an MRF model.

Definition 5 (Markov Random Field) *A random field X that fulfills the following two conditions is a Markov random field.*

$$P\left(x_i \,|\, x_{\mathcal{S}\setminus\{i\}}\right) = P\left(x_i \,|\, x_{\mathcal{N}_i}\right) \qquad \textit{Markovianity} \tag{3.50}$$

$$P(x) > 0, \ \forall x \in \mathcal{X} \qquad \textit{Positivity} \tag{3.51}$$

The first condition means that in an MRF the probability for the label at site i only depends on the labels at the sites i' in a certain neighborhood $\mathcal{N}_i$ of i. A further requirement for a family of random variables to be a MRF is the positivity, that is, that the probability of all configurations x is positive.

MRFs are always defined with respect to some neighborhood system.

Definition 6 (Neighborhood system) *The neighborhoods for all sites $i \in \mathcal{S}$ constitute a neighborhood system $\mathcal{N}$.*

$$\mathcal{N} = \{\mathcal{N}_i \,|\, i \in \mathcal{S}\} \tag{3.52}$$

The neighborhood relationship is defined by two conditions.

Definition 7 (Neighborhood relationship) *The relationship between sites i and i' is a neighborhood relationship if a site is not a neighbor of itself, and if the relationship is mutual.*

$$i \notin \mathcal{N}_i \tag{3.53}$$

$$i \in \mathcal{N}_{i'} \Leftrightarrow i' \in \mathcal{N}_i \tag{3.54}$$

Such a neighborhood relationship can be defined by using a suitable distance function. For two sites with lattice indices (i, j, k) and (i', j', k') a suitable distance function, for example, could be $d((i, j, k), (i', j', k')) = |i - i'| + |j - j'| + |k - k'|$.

$$\mathcal{N}_{(i,j,k)} = \{(i', j', k') \,|\, d((i, j, k), (i', j', k')) \leq 1\} \tag{3.55}$$

This will define the commonly used 6-neighborhood. Another commonly used neighborhood definition contains 26 neighbors for each site.

$$\mathcal{N}_{(i,j,k)} = \left\{(i', j', k') \;\middle|\; |i - i'| \leq 1 \wedge |j - j'| \leq 1 \wedge |k - k'| \leq 1\right\} \tag{3.56}$$

A useful theorem when working with MRFs is the equivalence between MRFs and Gibbs random fields (GRF).

Theorem 2 (Hammersley-Clifford Theorem [57, 82]) *If and only if a random field is an MRF then it also is a GRF.*

This means, that for each MRF an equivalent GRF exists and vice versa. A GRF is defined as follows.

Definition 8 (Gibbs Random Field) *If and only if the probability distribution of a random field X can be given through a function of the following form then this random field is a Gibbs random field.*

$$P(x) = \frac{1}{Z} \exp\left(-\frac{1}{T} U(x)\right) \tag{3.57}$$

Z is a normalization constant, and T is a parameter called temperature. The probability $P(x)$ only depends on the configuration x through the Gibbs energy function

$$U(x) = \sum_{c \in C} V_c(x) \tag{3.58}$$

with $V_c(x)$ being the clique potential.

The sum in Equation 3.58 runs across all possible cliques $c \in C$, and the clique potential $V_c(x)$ defines the potential of a certain clique c in the current configuration x.

Definition 9 (Clique) *A clique c is an ordered subset of one or more sites from $\mathcal{S}$. If the clique contains more than one site then any two different sites in the clique must be neighbors to one another.*

C_1 contains all single site cliques, that is, subsets containing only a single site i.

$$C_1 = \{i \,|\, i \in \mathcal{S}\} \tag{3.59}$$

Cliques with two or more sites might be possible for the chosen neighborhood system.

$$C_2 = \{(i, i') \mid i, i' \in \mathcal{S} \wedge i' \in \mathcal{N}_i\} \qquad \text{Pairwise cliques} \tag{3.60}$$

The set of all possible cliques is then made up of all single-site cliques, all pairwise cliques, and all cliques with more than two sites.

$$C = C_1 \cup C_2 \cup \ldots \tag{3.61}$$

For practical purposes it is of advantage to split up the Gibbs energy function into the contributions of single-site, pairwise and higher order cliques.

$$U(x) = \sum_{i \in \mathcal{S}} V_1(i, x_i) + \sum_{i \in \mathcal{S}} \sum_{i' \in \mathcal{N}_i} V_2\,(i, i', x_i, x_{i'}) + \ldots \tag{3.62}$$

$V_1(i, x_i)$ and $V_2(i, i', x_i, x_{i'})$ are the clique potentials for single-site and pairwise cliques.

Instead of specifying the MRF model directly through the conditional probabilities $P(x_i \mid x_{\mathcal{N}_i})$ it is of advantage to specify the random field through the clique potentials $V_c(x)$ via a GRF [82, 18]. The relation between the clique potentials and the conditional probability at site i is given by the following equation [82].

$$P(x_i|x_{\mathcal{N}_i}) = \frac{\exp\left(-\sum_{c \in C} V_c(x)\right)}{\sum_{x'_i \in \mathcal{L}} \exp\left(-\sum_{c \in C} V_c(x')\right)} \tag{3.63}$$

x' here is the configuration which is equal to x at all sites except at site i, where it is x'_i. $\mathcal{L} = \{0, \ldots, m-1\}$ is the set of labels.

The pairwise clique potentials in Equation 3.62 can depend on the relative positions of the sites i and i' to each other. In this case, the MRF is called anisotropic. Otherwise, when $V_2(i, i', x_i, x_{i'}) = V_2(i, x_i, x_{i'})$ the MRF is called isotropic. An isotropic MRF can still depend on the absolute position of the site i. This is a nonhomogeneous MRF as opposed to a homogeneous MRF where $V_2(i, x_i, x_{i'}) = V_2(x_i, x_{i'})$.

For the solution of our segmentation problem a labeling $x \in \mathcal{X}$ must be determined that maximizes the a-posteriori probability. For a finite lattice and a finite domain of labels $\mathcal{L}$ the space of possible configurations $\mathcal{X}$ also contains a finite number of different configurations. This is a combinatorial optimization problem. For typical problems, however, the number of different configurations is so large that an exhaustive search for the optimal configuration is computationally intractable. For the segmentation of a $256 \times 256 \times 256$ image into 9 different regions (labels) $9^{256} \approx 10^{244}$ configurations exist. Due to this reason, some other optimization procedure must be employed.

Geman and Geman [49] describe how the a-posteriori probability can be maximized using a simulated annealing (SA) algorithm. The SA algorithm successively visits each site $i \in \mathcal{S}$ where a new label is chosen by sampling from the local conditional probability. This means, that a new label is randomly drawn in such a way that the probability distribution for the random draw is equal to $P(x_i|x_{\mathcal{N}_i}, \boldsymbol{y}_i)$. This does not guarantee that the a-posteriori probability is increased at every iteration. It might happen, although it is less likely, that a label is chosen which decreases

the a-posteriori probability. This is a desired characteristic of the SA as it prevents the algorithm to only reach the next local minimum. During the annealing the temperature parameter T in the Gibbs distribution is decreased following an annealing schedule. For large values of T the Gibbs distribution is flat and maxima are shallow. This allows the SA algorithm to escape from local maxima and reach the global maximum. When decreasing the temperature the maxima become more pronounced and it is more likely that the algorithm is captured in a maximum. It can be proven that the SA algorithm converges to the global maximum with a probability of 1 if T is reduced according to the following annealing schedule [49].

$$T(k) \geq \frac{c}{\log(1+k)} \tag{3.64}$$

Here, $T(k)$ is the temperature for the k-th visited site, and c is a constant. The SA algorithm is run until T is lower than a specified threshold value. Following this schedule, however, would lead to a number of iterations which is too high for any application. In practice, less conservative schedules are used to keep the computation time within a reasonable limit.

A method for finding at least the next local minimum is the iterated conditional modes (ICM) method by Besag [19]. Again, every site i is visited following an arbitrary update order. To each visited site the label is assigned which maximizes the conditional probability $P(x_i \mid x_{\mathcal{N}_i}, \boldsymbol{y}_i)$ taking into account the current labels at the neighboring sites. This label is the mode of the conditional probability distribution which is why the algorithm is called iterated conditional modes method. The a-posteriori probability can be expressed as a product of the conditional probability for the label at site i and the joined conditional probability of the labels at all other sites.

$$P(x \mid \boldsymbol{y}) = P\left(x_i \mid x_{S\setminus i}, \boldsymbol{y}\right) \cdot P\left(x_{S\setminus i} \mid \boldsymbol{y}\right) \tag{3.65}$$

Each iteration of the ICM algorithm increases the first conditional probability and, thus, also the a-posteriori probability [19]. This means, that the algorithm converges to a maximum of the a-posteriori probability. An increase of the a-posteriori probability at each iteration of the algorithm also means that the convergence can only be local. Experiments show that the ICM algorithm rapidly converges to the next local maximum reaching the maximum after each site was visited only few times [19].

Estimating a configuration x which maximizes the a-posteriori probability is, however, only one part of the MRF segmentation. As for the active contours segmentation values for the parameters of the likelihood function also have to be determined. For the MRF segmentation we cannot compute the derivative of the cost function due to the discrete nature of the labeling x, and we, thus, cannot use the gradient descent method to simultaneously estimate the labeling and the parameters $\boldsymbol{\mu}_c$ and $\boldsymbol{\Sigma}_c$.

Instead, we will use the EM method first formally proposed by Dempster et al. [43]. An intuitive solution to an incomplete data problem would be to first choose an initial value for the unknown parameters $\lambda^{(0)}$. Using this initial guess it is then possible to obtain a MAP estimate for the labeling x, our missing data. Now, having the complete data $(\boldsymbol{y}, x)$ updated parameter values $\lambda^{(t+1)}$ can be computed using the maximum likelihood estimation from Equation 3.23.

This intuitive approach is formalized in the EM algorithm. The EM algorithm consists of two steps. The first step involves the computation of the conditional expectation of the log likelihood

for the complete data.

$$\text{E-step:} \qquad Q(\lambda|\lambda^{(t)}) = E\left[\log p(\boldsymbol{y}, x \,|\, \lambda)|\boldsymbol{y}, \lambda^{(t)}\right] \tag{3.66}$$

The EM algorithm is then completed by the M-step in which values for the parameters are determined which maximize the conditional expectation of the log likelihood.

$$\text{M-step:} \qquad \lambda^{(t+1)} = \arg\max_{\lambda} Q(\lambda \,|\, \lambda^{(t)}) \tag{3.67}$$

The E- and the M-step are performed iteratively until convergence.

In the case where the complete data likelihood is from the exponential family of distributions (Eq. 3.68) the E-step and the M-step take on a simpler form.

$$p(\boldsymbol{y}, x \,|\, \lambda) = b(\boldsymbol{y}, x) \exp\left[c(\lambda)^{\top} t(\boldsymbol{y}, x)\right] / a(\lambda) \tag{3.68}$$

$t(\boldsymbol{y}, x)$ here is a sufficient statistic. A sufficient statistic is a function of the complete data that provides all the information that is necessary to estimate the parameters λ [43, 91]. The conditional expectation of the logarithm of the complete data likelihood from the exponential family can be written as

$$Q(\lambda \,|\, \lambda^{(t)}) = E\left[\log(b(\boldsymbol{y}, x) \,|\, \boldsymbol{y}, \lambda^{(t)}\right] + c(\lambda)^{(t)} E\left[t(\boldsymbol{y}, x) \,|\, \boldsymbol{y}, \lambda^{(t)}\right] - \log(a(\lambda)) \tag{3.69}$$

In the M-step the logarithm of the conditional expectation is maximized with respect to λ. This requires that the conditional expectation is first computed in the E-step. In Equation 3.69 we can ignore the first term as it does not depend on λ and is, thus, irrelevant for the optimization in the M-step. To fully specify the complete data likelihood only the sufficient statistic $t(\boldsymbol{y}, x)$ must be determined. Therefore, the E-step in the case of a complete data likelihood from the exponential family of distributions consists of estimating $t(\boldsymbol{y}, x)$.

$$t(\boldsymbol{y}, x)^{(t+1)} = E\left[t(\boldsymbol{y}, x) \,|\, \boldsymbol{y}, \lambda^{(t)}\right] \tag{3.70}$$

For our application, the prior can be described by a Gibbs distribution, and our likelihood function is given by a Gaussian distribution so that the complete data likelihood is from the exponential family of distributions. We are, thus, dealing with the special case discussed above.

The complete data $(\boldsymbol{y}, x)$ itself is a sufficient statistic for the chosen complete data likelihood. Our task in the E-step is, thus, to estimate the complete data which is equivalent to estimating the labeling x since the image data $\boldsymbol{y}$ is already given. Due to the difficulties in actually computing the expectation of x we are instead estimating the x which maximizes the a-posterior probability using Besag's ICM algorithm [19].

Dempster et al. [43] prove in their paper for the EM algorithm that the likelihood $L(\lambda) = \log l(\boldsymbol{y} \,|\, \lambda)$ converges to some value L^* if the likelihood function is bounded. If $Q(\lambda^{(t+1)}, \lambda^{(t)})$ is continuous in both parameters then L^* is a stationary value of L [159]. Under some additional constraints which are fulfilled for our choice of likelihood function [159] this stationary value is a local maximum. In summary, the EM algorithm finds a local maximum of the logarithm of the likelihood. It, thus, solves the parameter estimation for the incomplete data problem in the same

sense as the maximum likelihood estimation (Eq. 3.23) for the case where the complete data is present.

4 Influence of Skull Segmentation Inaccuracies on EEG Source Analysis

4.1 Introduction

In Chapter 2 the advantages of using individual, realistic volume conductor models for EEG source analysis have been described. In summary, it was stated, that the accuracy of source analysis can be considerably improved by employing an individual, realistic model of the head for solving the EEG forward problem [61, 50, 111, 53, 153, 154, 1, 88]. Commonly, individual, realistic head models are created by first segmenting the different tissues based on imaging data of the subject's head, and then assigning previously published conductivity values to the segmented tissue regions.

In practice, the segmentation of the head tissues is not a trivial task. One of the reasons for this is that most of the time MRIs are the only available imaging modality providing information on the individual anatomy of the subject's head. Furthermore, in many cases the available MRIs are not optimally suited for volume conductor segmentation. Practically used MRIs, for example, often suffer from a low SNR, especially in the inferior head region, or they are affected by artifacts, such as the water-fat-shift [149] or the shading artifact [133, 132, 131]. Another problem is the low contrast between the skull and the surrounding tissues if only a T1-weighted MRI is available. This renders especially the segmentation of the skull problematic.

Due to these difficulties we have to expect inaccuracies in the geometry of the segmented skull and it may even be necessary to apply simplifications to the skull geometry. At the same time the skull plays a special role in modeling the volume conductor for EEG source analysis [55, 120, 39]. The skull is situated between the sensors and the source space and its conductivity is an order of magnitude lower than that of the surrounding tissues. Therefore, skull geometry inaccuracies might have a potentially large influence on the solution of the forward problem and, thus, on EEG source analysis in general.

Holes and skull thickness misspecifications are one type of skull geometry inaccuracies that might be expected when trying to segment the skull from MRI data. The intricate internal structure of the skull bone (Sec. 3.1) poses a problem for any segmentation procedure. In combination with noise, a low image resolution on the order of magnitude of the compact bone thickness, and water-fat-shift artifacts the segmentation algorithm might not be able to discern the thin compact bone layer, and it might fail to differentiate between cancellous bone and the surrounding muscle or brain tissue. This would cause a segmented skull with underestimated thickness or, in extreme cases, even holes. An overestimation of the skull thickness might occur in T1-weighted MRIs where there is very low contrast between CSF and the close-by compact bone.

Given the problems of segmenting the skull from MRI data one might also a-priori conclude that it is not possible to accurately segment some features of the skull geometry. In this case, these features might be approximated by applying some simplifications. As MRIs show nearly

no contrast between compact bone and air it will not be possible to determine the shape and position of air-filled cavities, for example, the sinuses, from an MRI. One possible solution is to simplify the skull geometry by modeling the air-filled cavities as compact bone. Simplifications could also make sense in the area of the skull base. Due to the especially low SNR in the inferior head and neck region it might not be possible to accurately segment the complicated geometry of the skull base. This problem could be addressed by approximating the skull at its base as a layer of constant thickness. In the same sense it might not be feasible to try to segment the different head tissues in the inferior head and neck region. Here, simplifying this region as an isotropic region with a single conductivity value could be a solution. As an alternative, the model could be cut immediately below or at a short distance below the skull to avoid the problems segmenting the inferior head and neck region.

The special role of the skull for volume conductor modeling in EEG source analysis is widely recognized. For this reason, several previous publications exist investigating single aspects of how skull geometry inaccuracies influence source analysis. Bénar et al. [11] studied in BEM models how the presence of burr holes in the skull affects source reconstruction. They find localization errors of up to 20 mm depending on the location and orientation of the simulated source. The influence of skull holes on the EEG was also investigated by Vanrumste et al. [142]. In their work the authors similarly find large localization errors when skull holes are not incorporated into the volume conductor model. A few other publications on the influence of skull holes on the EEG exist which all come to similar results [141, 81, 98].

In addition, previous studies [36, 117, 160] have dealt with the effects of locally under- or overestimating the skull thickness. The downward extension of the volume conductor model was studied before by Bruno et al. [25]. It was found that cutting the model along a plane intersecting the inferior part of the skull results in large errors for the solutions of the forward and inverse problem.

Each of the previous studies only investigated a single or only few types of skull defects. Thus, the results from these works cannot directly be compared to each other as the study setups differed considerably. One aspect in which the studies differed is the detail of the employed volume conductor models. In most studies simple spherical or realistically shaped three layer head models are used, and only few studies employed more realistic and more detailed models. Further aspects, in which previous investigations differ are, for example, the used conductivity values and the number and placement of the electrodes.

In the presented computer simulation study the influence of a wide variety of skull geometry inaccuracies on EEG source analysis will be investigated. These inaccuracies include skull holes, local errors in skull thickness, modeling air-filled sinuses as compact bone, the downward extension of the head model, and the simplification of the inferior skull, respectively, the inferior skull and scalp as layers of constant thickness. Reference EEG data is generated using a detailed and anatomically plausible head model in combination with high-resolution 1 mm geometry-adapted hexahedral FE meshes and the accurate FEM [118, 158]. Forward simulation and source reconstruction errors are assessed in test models representing the investigated skull geometry inaccuracies, and for probe sources densely distributed throughout the whole brain volume. By investigating the wide variety of skull defects in a common study setup it is possible to compare their influences, and to draw guidelines on how the skull geometry shall be modeled for the generation of accurate volume conductor models.

Using the results from this study it is also possible to link the accuracy of the segmentation approaches proposed in Chapters 7 and 8 to the performance of volume conductor models based on these approaches in EEG source analysis. If, for example, an error in skull thickness is observed in one of the segmentation results then from the results of this study a prediction can be made with regard to the influence of the observed segmentation inaccuracy on EEG source analysis.

The study presented here was published by the author of this thesis as first author in [77]. The presentation in this chapter mainly follows the presentation in the published article.

4.2 Materials and Methods

4.2.1 Study Setup

A detailed and anatomically plausible volume conductor model was constructed for this study. This model serves as a reference model for our simulations, and the data simulated in this model is our reference data. The reference data is assumed to be representative for real EEG data as it would be measured during an actual measurement. From the reference model test models were derived representing the investigated skull geometry inaccuracies.

The investigations done in the presented study can roughly be divided into two parts. In the first part the influence of the skull defects on the solution of the forward problem were studied. To do this, first reference data was generated by simulating the EEG potentials in the reference model for a large set of probe sources. For the same set of probe sources the potentials in each of the test models were computed. The obtained simulated potentials were then compared for each source individually. Error measures which were used for the comparison are described below.

In the second part of the presented study the influence of the skull defects on source reconstruction was considered. Again, for a set of probe sources reference data was computed by solving the forward problem in the reference volume conductor model. This reference data was then reconstructed using some kind of inverse method, and while doing so, the forward problem associated with the source reconstruction was solved in the test model. Source reconstruction was done individually for each probe source and test model. The differences between the reference and the reconstructed source positions were evaluated.

Forward and inverse errors were computed for probe sources distributed throughout the whole brain. Thus, error maps can be plotted to illustrate the dependency of the error on the location of the probe source inside of the volume conductor model and relative to the skull geometry inaccuracy. Errors for probe sources with different orientations but with the same positions are averaged to enable a clearer visualization of the error distribution. Probe sources were distributed on a regular cubic grid. In the following, errors are presented for slices of the regular cubic grid. To aid the interpretation of the error maps the slices of the maps are overlaid onto slices of the corresponding test model.

In addition, the errors were analyzed using descriptive statistics to obtain information about the maximal and mean errors. The proportion of affected sources is also assessed.

4.2.2 Construction of Reference and Test Models

The reference model was constructed based on several MRIs of the same subject. In detail, a T1-weighted MRI with selective water excitation (3D TFE sequence, TR = 9.32 ms, TE = 4.45 ms), a T1-weighted MRI with minimal water-fat-shift (3D TFE sequence, TR = 4.05 ms, TE = 1.81 ms, bandwidth = 777 Hz), and a T2-weighted image (2D SE sequence, TR = 5460 ms, TE = 60 ms) were used. MRIs were recorded on a Siemens 3T MRI Scanner at an original resolution of $1.17 \times 1.17 \times 1.17\,\mathrm{mm}^3$. During preprocessing the images were resampled to an isotropic $1.0\,\mathrm{mm}^3$ resolution. An affine registration approach based on the mutual information registration measure [151] was utilized to align the T1- and T2-weighted images.

Segmentation of the multimodal image data was done in a semi-automatic fashion. Only those tissue types were differentiated during segmentation, that differed substantially. In detail the segmentation of the reference model was done as follows. A so-called brain mask, marking the intracranial space, was extracted from the T2-weighted image by thresholding using a manually chosen threshold value. The surface of the brain mask was subsequently smoothed by means of applying suitable morphological operations and fitting a deformable surface model to it. The T1-weighted image recorded with water-selective excitation was then masked with this brain mask. Unsupervised classification using the software tool *FSL Fast* [163] was performed on the masked image to identify the tissues CSF, GM and WM.

Unsupervised classification was also performed on the parts of the T1-weighted image that were not masked by the brain mask. All voxels in a class with low gray value were then selected to construct a first approximate mask of the cranial skull. This first iteration of the cranial skull mask was subsequently cleaned up by applying morphological operations, in particular a closing operation and filling of holes. Additional manual corrections, for example, manually adding the openings for the optical nerves, completed the cranial skull mask. The cancellous bone layer was segmented from the T1-weighted image with minimal water-fat shift masked by the cranial skull mask using a simple thresholding procedure. Accurately segmenting the cancellous bone layer is possible only from MRIs which are not or merely minimally affected by the water-fat-shift artifact. In MRIs showing a large water-fat shift the signal from the cancellous bone layer, which is mainly consisting of fatty tissue, is displaced and the position of the diploe layer cannot be estimated accurately. The cancellous bone mask was then masked with a by 2 mm eroded cranial skull mask to guarantee a minimum thickness for the compact bone layers.

The bony structures and air-filled cavities in the inferior head and neck region (i.e., a region in the inferior part of the head and at a certain minimum distance to the segmented brain mask) were segmented in a similar way by first selecting voxels with a low gray value from the initial classification. Morphological operations were then performed to clean up the structures to be segmented. In addition, extensive manual corrections were applied. Some inferior structures, for example, the vertebrae and the sinus cavities, were even labeled completely by hand.

Up to this point the tissues brain, CSF, and compact and cancellous bone were identified. The remaining not yet identified voxels were assigned to the tissues scalp, muscle, fat and soft tissue. First, fatty tissue was segmented by selecting all bright voxels above a certain threshold in the T1-weighted image with minimal water-fat shift. Muscle and scalp tissue coincided and were, thus, identified with separate classes from the initial unsupervised classification. A minimum thickness of 2 mm was enforced for the scalp and the superior muscle layer by erosion of the

outer head surface, respectively, by dilation of the outer skull mask. The remaining voxels in the inferior part of the head were labeled as soft tissue.

The individual tissue masks as described above were finally combined to obtain the labeled volume of our reference model at 1.0 mm^3 resolution. Sagittal slices taken at regular intervals of the labeled volume can be seen in Figure 4.1d. In addition, a rendering of some selected surfaces of the reference models is shown in Figure 4.1a.

The investigated skull geometry inaccuracies and simplifications were represented by a series of test models (TMs) which were derived from the reference model. An overview of the used TMs is given in Table 4.1. In addition, renderings of selected TMs are shown in Figure 4.2.

The first test model series (Figure 4.2a) was derived from the reference model by introducing holes with a varying diameter in the skull. Holes were located in the left-temporal bone. The tissue inside of the holes was modeled with the same conductivity as muscle tissue.

A second test model series was used to study the effect of over- or underestimating the skull thickness. In the models of this series the skull thickness was increased, respectively, decreased by a varying amount in a circular region of 20 mm diameter. Figure 4.2b depicts TM 2f where the skull thickness is decreased.

The influence of ignoring the air-filled sinus cavities is investigated using the TM 3a and 3b. In these all sinuses except the frontal sinuses (TM 3a), respectively, all sinuses without exception (TM 3b) are modeled as filled with compact bone. A rendering of TM 3a can be seen in Figure 4.2c.

The models in the fourth series represent different ways of extending the volume conductor model downwards (see Figure 4.2d for a rendering of TM 4a). In TM 4a to 4c the model was cut along an axial plane at varying distances to the occipital hole of the skull. TM 4d was derived from TM 4c by applying further simplifications. In TM 4d all sinuses were modeled as bone and the inferior head and neck region was modeled as an isotropic, homogeneous region with a single conductivity value. The inferior head and neck region was defined for this purpose as consisting of all positions at a minimum distance of 30 mm from any brain tissue, and below a plane passing approximately through the nasion and the inion.

In TM 5a to 5c the inferior skull, which approximately corresponds to the skull base, was modeled as a layer of constant thickness. Thickness of the inferior skull was varied between 4 (TM 5a) and 8 mm (TM 5c). A rendering of TM 5b is shown in Figure 4.2e.

The sixth series of test models was derived from TM 5b. In TM 6a to 6c not only the inferior skull but also the inferior scalp was approximated by a layer of constant thickness. For the constant scalp thickness values from 4 to 8 mm were chosen. The inferior skull thickness was not varied but kept constant at 6 mm. Figure 4.2f shows a depiction of TM 6b.

4.2.3 Construction of Geometry-Adapted Hexahedral FEM Volume Conductor Models

In our simulation study the FEM was employed to solve the forward problem in the detailed, realistic volume conductor model. The FEM requires a FE mesh discretizing the volume conductor model. For the present study geometry-adapted hexahedral meshes were chosen which were constructed as follows. In a first step a regular hexahedral FE mesh was constructed by simply identifying each non-background voxel with a hexahedral element. Then, in a second

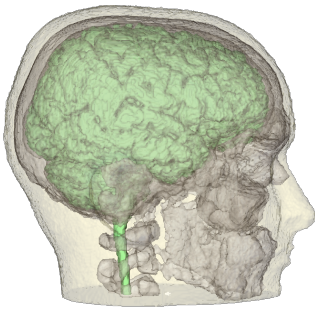

(a) Rendering of selected compartments of reference model. Only scalp, skull and brain surfaces are rendered.

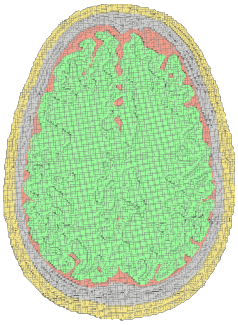

(b) Axial cut through FE mesh of reference model. Only the tissues scalp, skull, CSF and brain are differentiated here to improve clarity.

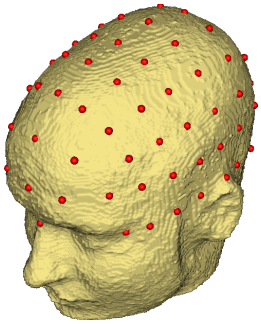

(c) The EEG configuration with 79 electrodes on top of the scalp surface of the reference model.

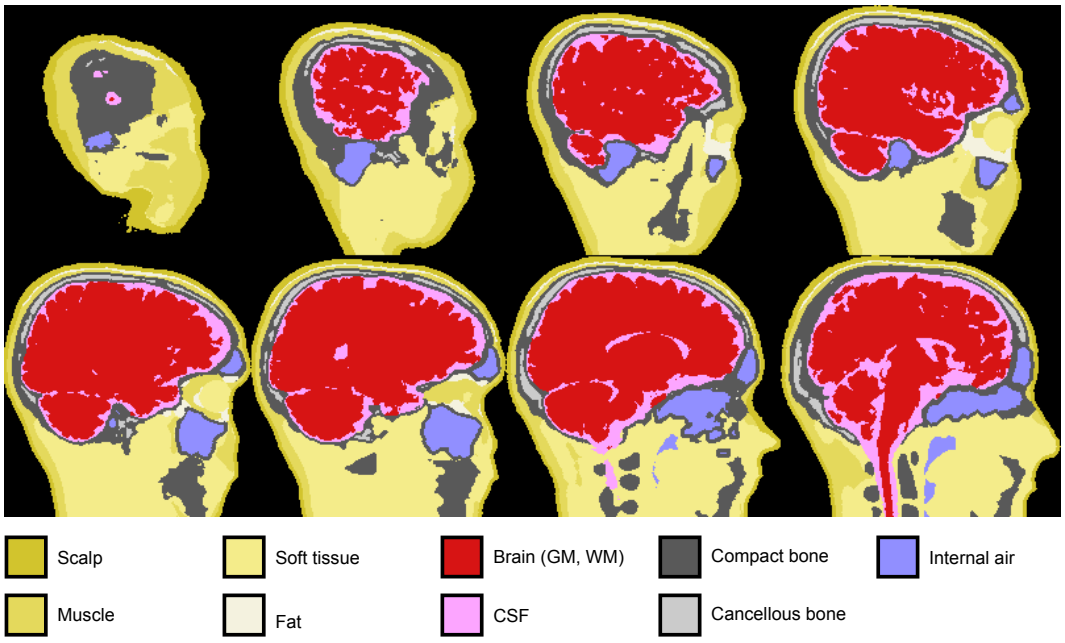

(d) Equidistant sagittal slices of the labeled MRI which is the basis for the construction of the reference model. GM and WM are not differentiated here because they were modeled with the same conductivity value.

Figure 4.1: Visualizations of the reference model and the used electrode configuration.

Table 4.1: Details on the test models used in the present study.

Test models 1- Skull hole	**Test models 2 - Skull thickness**	**Test models 3 - Sinuses**	**Test models 4 - Downward extension**	**Test models 5 - Simplified inferior skull**	**Test models 6 - Simplified inferior skull and scalp**
Hole in left-temporal skull bone; hole diameter is varied	Skull too thick or too thin in circular skull region of 20 mm diameter	Ignoring air-filled sinuses and modeling them as compact bone instead	Model cut along axial plane at varying distances below skull	Inferior skull approximated as layer of constant thickness	Inferior skull and scalp approximated as layers of constant thickness; inferior skull thickness is here kept constant at 6 mm
1a: 2 mm hole **1b:** 6 mm hole **1c:** 10 mm hole	**2a:** 2 mm thicker **2b:** 4 mm thicker **2c:** 6 mm thicker **2d:** 2 mm thinner **2e:** 4 mm thinner **2f:** 6 mm thinner	**3a:** All sinuses except frontal sinuses modeled as compact bone **3b:** All sinuses without exception modeled as compact bone	**4a:** Cut 0 mm below the skull **4b:** Cut 20 mm below the skull **4c:** Cut 40 mm below the skull **4d:** As TM 4c with additional simplifications	**5a:** 4 mm skull thickness **5b:** 6 mm skull thickness **5c:** 8 mm skull thickness	**6a:** 4 mm scalp thickness **6b:** 6 mm scalp thickness **6c:** 8 mm scalp thickness

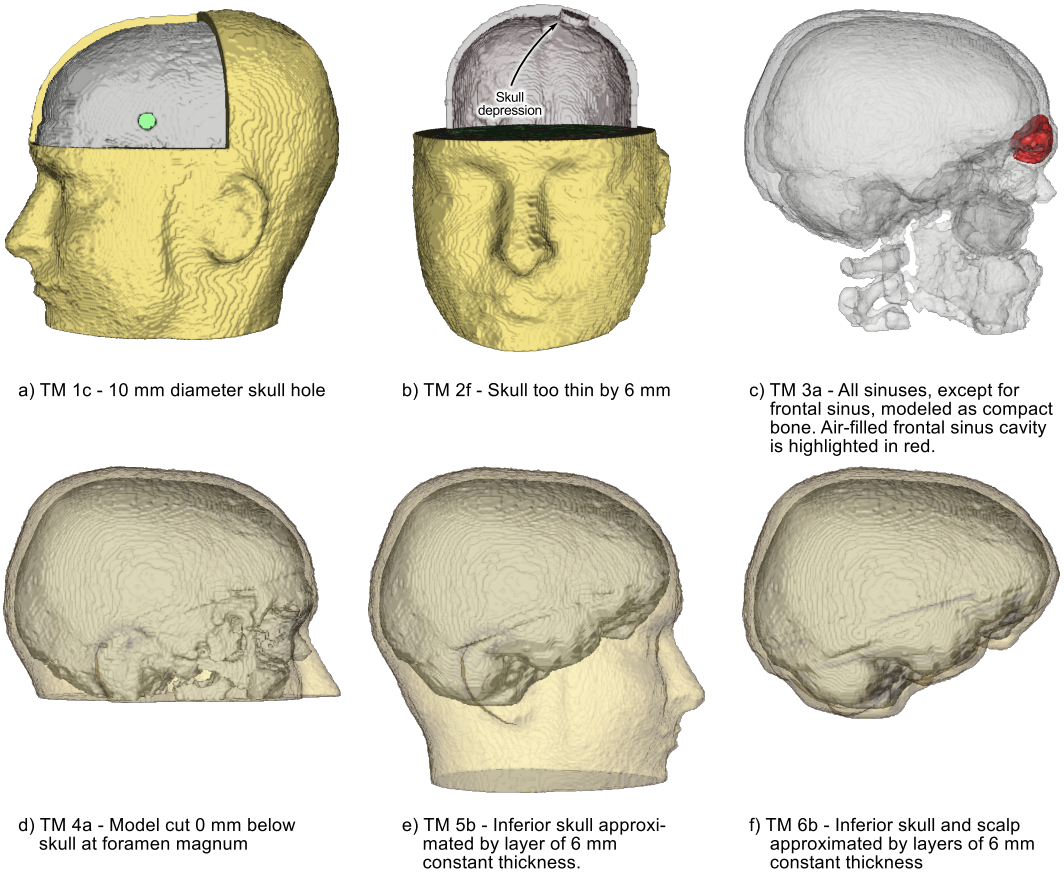

Figure 4.2: Renderings of selected test models used in the present study.

step, FE nodes at the interfaces between two tissues were shifted to smooth the stair-step like interfaces of the regular hexahedra mesh. Shifting was done using the so-called node-shift approach [158]. The interfaces in the resulting FE mesh were smooth, and they in general better approximated the head geometry. In this way an FE mesh with around 3.56 million nodes and around 3.47 million elements was obtained.

Tetrahedral meshes are besides hexahedral meshes also commonly used for FE modeling. Tetrahedral mesh construction is often done using surface-based approaches where a description of all tissue interfaces must be provided as input. A disadvantage of these surface-based approaches is that they require that some conditions for the input surfaces must be fulfilled. The surfaces must, for example, not be self-intersecting and they must not intersect each other. These conditions are non-trivial, and constructing such surfaces is difficult in practice.

In comparison, creating a high-resolution geometry-adapted hexahedral mesh is easier while still providing a very good accuracy especially in combination with the node-shift approach as it was shown by Wolters et al. [158]. The ease of mesh construction and the high accuracy were the main reasons for choosing geometry-adapted hexahedral meshes for the present study.

The FE representation of the volume conductor model is completed by assigning isotropic conductivity values to each element. All the assigned tissue conductivity values and the sources

from which the values were taken are listed in Table 4.2.

Tissue	Conductivity (S/m)		Tissue	Conductivity (S/m)	
Scalp	0.43	[61]	Cancellous bone	0.021	[4]
Muscle	0.11	[61]	Internal air	0.00001	[61]
Fat	0.04	[61]	CSF	1.79	[8]
Soft tissue	0.17	[61]	GM	0.33	[61]
Compact bone	0.005	[4]	WM	0.33	[61]

Table 4.2: Conductivity values used for the reference and test models in the present study.

4.2.4 Error Measures

The differences between the forward simulations in the reference and in the test models were quantified using the following error measures. The relative difference measure (RDM) [89] was used to express the differences in the topography of the EEG potentials. It is defined as follows:

$$\mathrm{RDM} = \left\| \frac{\phi^{\mathrm{ref}}}{||\phi^{\mathrm{ref}}||} - \frac{\phi^{\mathrm{test}}}{||\phi^{\mathrm{test}}||} \right\| \tag{4.1}$$

ϕ^{ref} and ϕ^{test} here denote the reference, respectively, the test potentials, and $\|\cdot\|$ means the L_2 norm. The RDM quantifies differences in the topography of the potentials only. It is not affected by a scaling of the magnitudes. The RDM ranges between 0 and 2, with 0 being its optimal value. The magnitude error (MAG) [89] is defined to measure the difference in the total magnitudes of the reference and test potentials.

$$\mathrm{MAG} = \frac{||\phi^{\mathrm{test}}||}{||\phi^{\mathrm{ref}}||} \tag{4.2}$$

The optimum value for the MAG is 1.

The results of a single source reconstruction study can in principle be evaluated by looking at the localization, the magnitude and the orientation error. In the presented study, only the localization error is regarded. It is computed as the Euclidean distance between the reference source position and the reconstructed source position. In addition, the mislocalization is considered. It is defined as the vector pointing from the reference to the reconstructed source position.

To get a single value quantifying the size of an influence on EEG source analysis the proportion of affected sources was defined for the forward solution and for source reconstruction. Threshold values were chosen from experience above which errors can be regarded as non-negligible. For the forward solutions a RDM value of 0.1, and for the localization error a value of 5 mm were set. Using these threshold values the proportion of affected sources can be computed as the number of sources with errors above the threshold value divided by the total number of probe sources. A larger proportion of affected sources will be interpreted as a larger influence of the currently investigated skull geometry defect.

4.2.5 Probe Sources, Electrode Configuration and Inverse Method

Reference data for the forward simulations was simulated for a set of probe sources distributed on a regular, cubic 4 mm grid. Sources oriented in x-, y- and z-direction were placed at each node of the grid. The magnitude of each source was set to 1 nAm.

For the source reconstruction study probe sources were distributed on a 10 mm, regular, cubic grid. A coarser grid resolution as compared to the forward simulations was chosen here to keep the computational demand for the inverse computations in reasonable bounds. The source reconstruction errors caused by the skull geometry defects might strongly depend on the orientation of the probe source. Therefore, sources with ten different orientations regularly sampling one hemisphere were placed at each node of the grid. The magnitude of the sources was again set to 1 nAm.

In the following, the forward and inverse errors are averaged for all sources at the same location but with different orientations. This makes it easier to visualize the errors and partly removes the dependency from the dipole orientation.

All probe sources for the forward and inverse studies were placed inside of the brain compartment of the reference volume conductor model. In addition, it was required that the probe sources have a minimum distance to the surrounding CSF and bone tissue, so that the FE node closest to the probe source only belongs to FE labeled as brain. This condition is called the *Venant* condition. It must be fulfilled to avoid unrealistic source modeling and numerical problems for the Venant dipole modeling approach [80].

All simulations presented here were carried out for an electrode configuration with 79 electrodes distributed across the scalp surface following the international 10-10-system [95]. The electrodes together with the scalp surface are depicted in Figure 4.1c.

For the second part of the presented work source reconstruction was performed using the robust Goal Function Scan method. To achieve a high reconstruction accuracy the scans were performed on a 1 mm scanning grid covering the entire source space. The node position for which the highest goal function value has been computed was interpreted as the reconstructed source position. To avoid an inverse crime [69] the cubic grid was shifted by 0.5 mm in each direction. In this way, the location of the probe sources do not coincide with the scanning grid. This also means that a perfect reconstruction is not possible. Finally, it was taken care that the Venant condition is fulfilled also for all nodes of the scanning grid.

4.2.6 Computation Platform

All simulations were carried out on a regular PC equipped with an *Intel Core 2 Quad Q6600* processor and 8 GB of RAM using the *SimBio* toolbox [130]. Mesh generation was performed using the *VGrid* software [60]. Visualization, for example, of error maps and rendered test models were done in *SCIRun* [125].

4.3 Results

4.3.1 Forward Simulations

To assess the influences of the skull geometry defects on the forward simulations the potentials for a set of probe sources were simulated in the reference and each of the test models. The simulated reference and test potentials were then compared directly. RDM and MAG error maps are shown in Figures 4.3 and 4.4, respectively.

When introducing *skull holes* of varying diameter into the skull considerable RDM errors can be found only for probe sources in the immediate vicinity of the hole (Fig. 4.3a). MAG errors for sources close to the skull hole are well above 1.0 (Fig. 4.4a) indicating that a hole in the skull leads to an overestimation of the EEG potentials for these sources.

Simulations in TM 2a to 2f have shown that an *over- or underestimation of the skull thickness* in a defined region caused non-negligible RDMs only in a small area close to the skull defect (Fig. 4.3b). The distribution of the MAG errors for these test models (Fig. 4.4b) is very similar to the distribution of the RDMs. For models TM 2d to 2f, in which the skull is modeled too thin, we can make out a small tendency to overestimate the simulated potentials.

In the third series of test models either all *sinuses* except the frontal sinuses or all sinuses without exception were modeled as filled with compact bone. When ignoring all sinuses except the frontal sinuses considerable RDM errors could only be observed for few inferior sources close to the skull base where the large sinuses are located (Fig. 4.3c). When also ignoring the frontal sinuses non-negligible errors can be observed in addition for frontal and temporal sources close to where the frontal sinuses in the reference model were located (Fig. 4.3d). In TM 3b a minor overestimation tendency can be made out for these sources (Fig. 4.4d).

The test models of the fourth series were constructed to study the *downward model extension*. Considerable RDM and MAG errors were found for TM 4a, 4b, and 4d (Fig. 4.3e,f). When cutting the model directly below the skull large errors could be observed for sources in the cerebellum and in a parts of the temporal lobes. Cutting the model at a distance of 20 mm from the skull already decreases the errors, so that non-negligible errors are found mainly for cerebellar sources. The MAG error maps for TM 4a and 4b (Fig. 4.4e) hint at an overestimation of the simulated potentials in these models. In TM 4d considerable errors occur for only a limited number of sources in the temporal and frontal regions (Fig. 4.3f).

In TM 5a to 5c the *inferior skull* was approximated as a layer of constant thickness. This simplification resulted in large RDM and MAG errors along the skull base (Fig. 4.3g). Errors were most prominent in the area of the occipital hole, in the temporal lobes close to the skull openings for the optical nerves, and in the vicinity of the frontal sinuses, where the skull is especially thick. The MAG errors in Figure 4.4g indicated that the potentials simulated in TM 5a to 5c are smaller as compared to the reference data.

In the sixth test model series the *inferior skull and scalp* were approximated as layers of constant thickness. The observed distribution of RDM errors for these test models resembled the error distributions of TM 5a to 5c but errors were generally larger (Fig. 4.3). The MAG error maps showed a different picture. In TM 6a, the model with the thinnest constant scalp thickness, MAG errors were above 1 for nearly all sources. Increasing the inferior scalp thickness the MAG error distribution became more heterogeneous. At 8 mm inferior scalp thickness the simulated

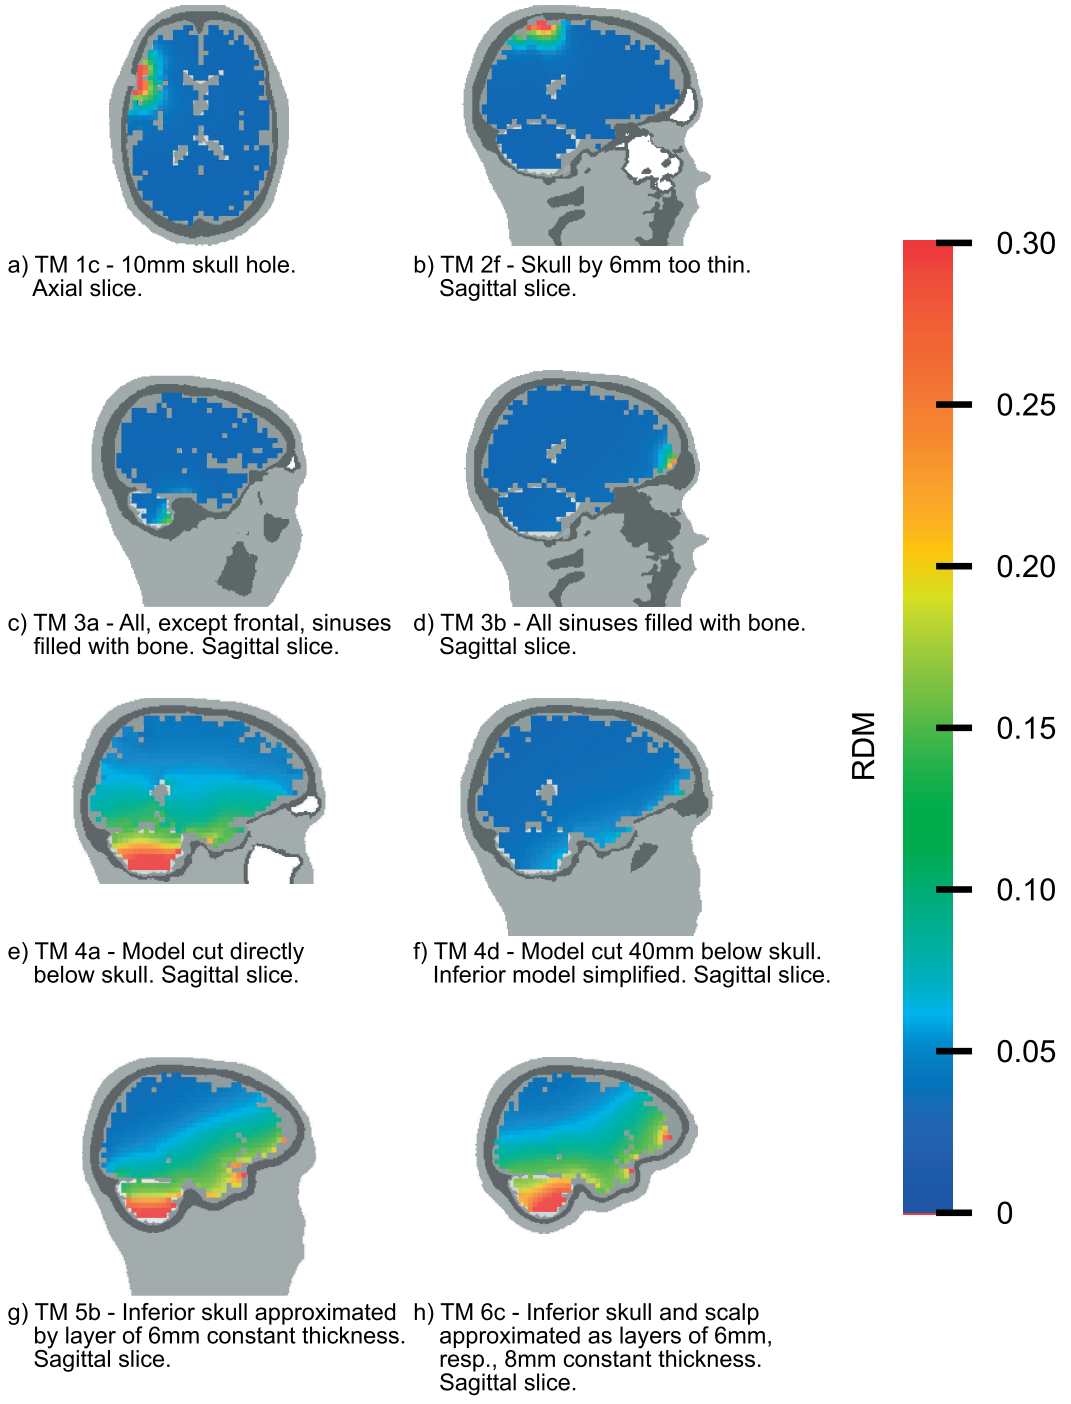

Figure 4.3: RDM error maps for selected test models. Errors were averaged across sources located at the same position but with different orientations.

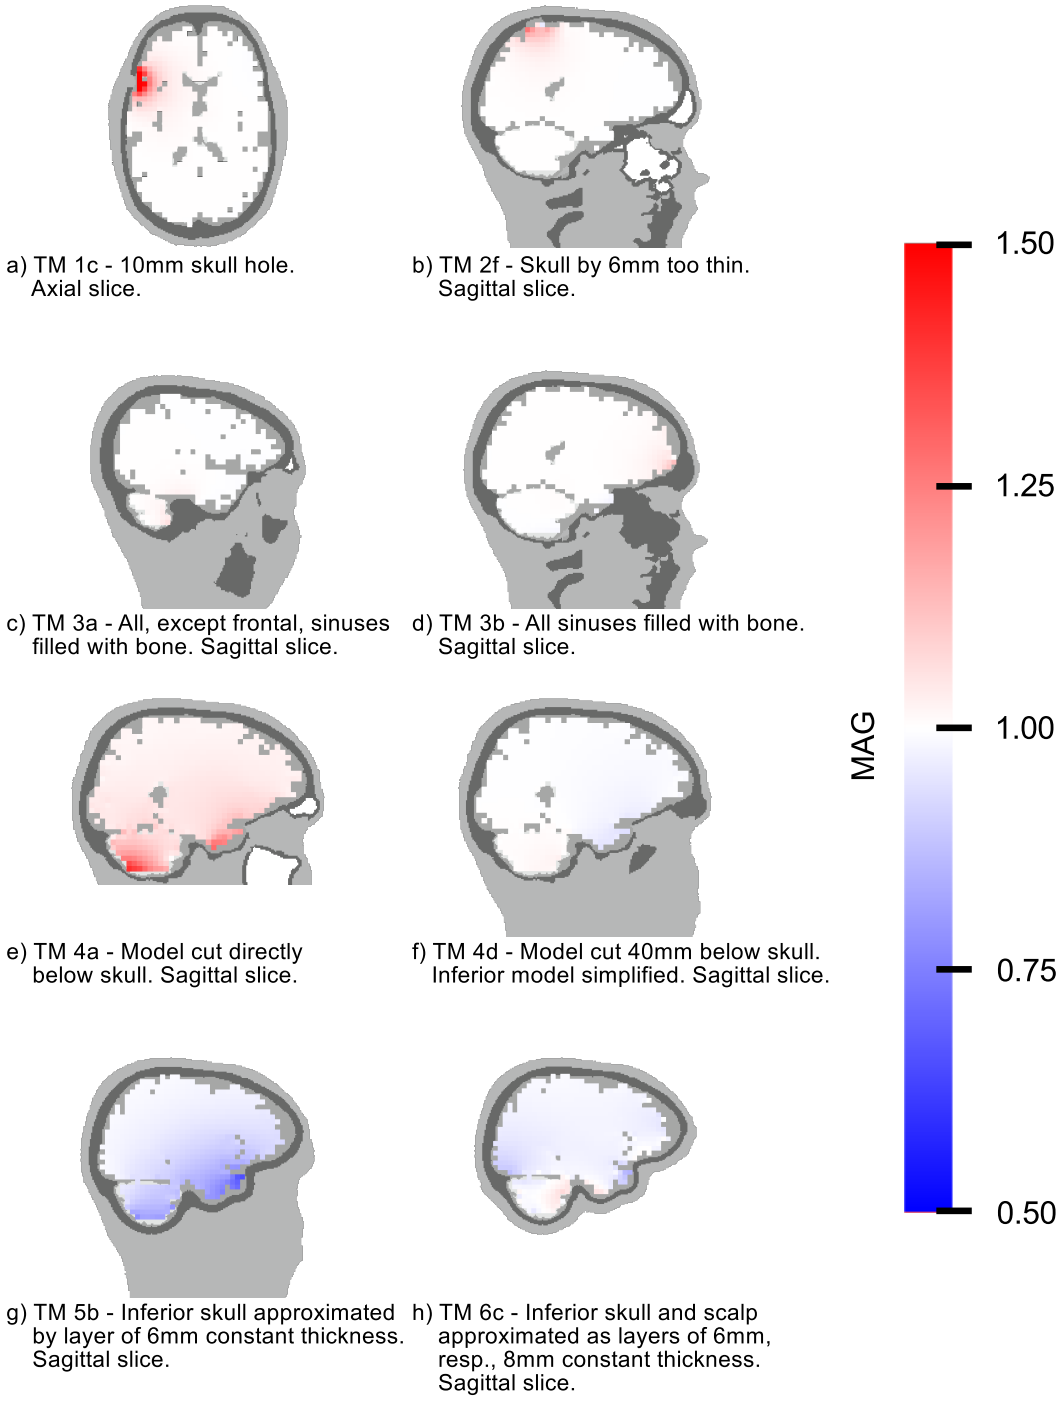

Figure 4.4: MAG error maps for selected test models. Errors were averaged across sources located at the same position but with different orientations.

potentials were overestimated for probe sources in some regions, for example, in parts of the temporal lobe and in the cerebellum, while the potentials were underestimated in other regions, like the remaining part of the temporal lobes and in occipital regions. For TM 6c it can be observed that the effects on the MAG of simplifying the inferior skull and of simplifying the inferior scalp partly canceled each other out. As a result the MAG became negligble in some areas along the skull base, for example, in temporal regions.

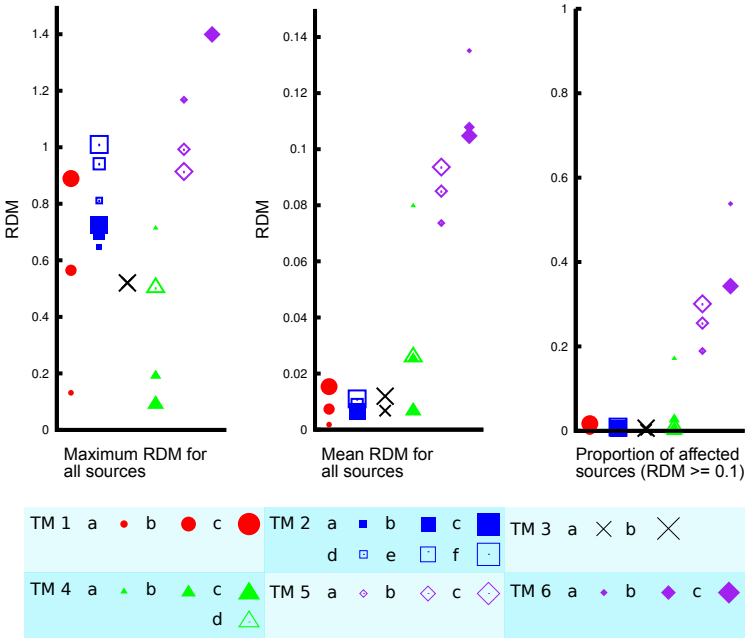

Figure 4.5: RDM error characteristics for all probe sources.

To sum up the errors for the different skull geometry inaccuracies the maximum and mean RDM errors across all probe sources, and the proportion of affected sources were computed. These values are plotted in Figure 4.5. The error characteristics were computed without averaging across the dipole orientations first. As all skull geometry inaccuracies have been studied in the same simulation setup Figure 4.5 allows to directly compare the influence sizes of the investigated inaccuracies.

First the test models inside of each series are compared. Increased errors can be observed for increased skull hole diameter (TM 1a-c), increased local error in skull thickness (TM 2a-f), and increased inferior skull thickness (TM 5a-c). Decreased errors were found for increased distance of the cutting plane to the skull (TM 4a-c), and for increased constant thickness of the inferior scalp layer (TM 6a-c).

Regarding the comparison of the different skull geometry inaccuracies it was found that a skull hole with only 2 mm diameter (TM 1a) caused the overall smallest forward errors. The low mean error for this test model can be explained by the very local nature of the skull defect. Considerable RDM errors were only found for very few probe sources in the immediate vicinity of the hole (Fig. 4.3a). Similarly, low mean RDM errors can also be found for further test models, where the skull geometry defect was of a local nature, like TM 3a and b where sinuses were

modeled as compact bone, and TM 2a-f with locally over- or underestimated skull thickness. The average RDM error for the model cut 40 mm below the skull can also be considered as negligible. Mean RDMs for the model cut 20 mm below the skull (TM 4b), and the model with additional inferior simplifications (TM 4d) were only slightly larger. In contrast, large averaged forward errors were found for the model cut directly below the skull (TM 4a), and the models with simplified inferior skull (TM 5a-c), respectively, simplified inferior skull and scalp (TM 6a-c). For these models also more than 17% of the sources were affected. This means, that the RDM for at least 17% of the probe sources was above 0.1.

4.3.2 Source Reconstruction

In the second part of the present study the influence of the skull segmentation inaccuracies on source reconstruction was investigated in single source scenarios. The observed localization errors are presented in this section. Figures 4.6 and 4.7 show the localization error maps, respectively, the maps of mislocalization tendencies for selected test models.

In the first test model series a non-negligible localization error could be observed only for the models with *skull holes* of at least 6 mm in diameter (Fig. 4.6a). In these models the sources tend to be reconstructed at positions farer away from the skull hole as can be seen in Figure 4.7a.

Only small localization errors were found for all models with an locally *over- or underestimated skull thickness* (Fig. 4.6b) The few probe sources with a minor localization error were located in the region close to the skull defect. Due to the small size of the localization errors a mislocalization tendency could only be made out for TM 2f where the sources tend to be reconstructed at position deeper inside the brain, that is, away from the skull thickness misspecification.

Modeling the *sinuses* as compact bone resulted in localization errors for probe sources in the vicinity of the sinus cavities. For TM 3a where all sinuses except the frontal sinuses were ignored noticeable errors only occured for sources in the region of the cerebellum (Fig. 4.6c). When ignoring all sinuses without exception, as it was done in TM 3b, additional localization errors are found for a few frontal sources in the immediate neighbourhood of the ignored frontal sinuses (Fig. 4.6d). These localization errors reached values of up to 9 mm. The affected sources in TM 3b were reconstructed at positions deeper inside the brain (Fig. 4.7c). For the affected sources in the cerebellum no consistent mislocalization tendency can be made out in Figure 4.7d.

In the fourth test model series (*downward model extension*) large localization errors were found when cutting the model directly below the skull. These errors were mainly located in the cerebellum and the frontal lobes (Fig. 4.6e). Cutting the model at 20 mm distance resulted in lower localization errors with a majority of the affected probe sources located in the cerebellum. In the model cut at a distance of 40 mm from the skull negligible errors below 4 mm were observed for all sources in the cerebrum. Considerable errors were only found for few sources in the area of the cerebellum. Applying additional simplifications to the model as done in TM 4d resulted in considerable localization errors for a minority of probe sources in occipital and frontal regions (Fig. 4.6f). In TM 4a and 4b affected sources in the cerebrum were consistently reconstructed too far superior (Fig. 4.7e). For cerebellar sources the mislocalization tendency in these models was not clear.

Large localization errors were found for the models with an *simplified inferior skull* (TM 5a

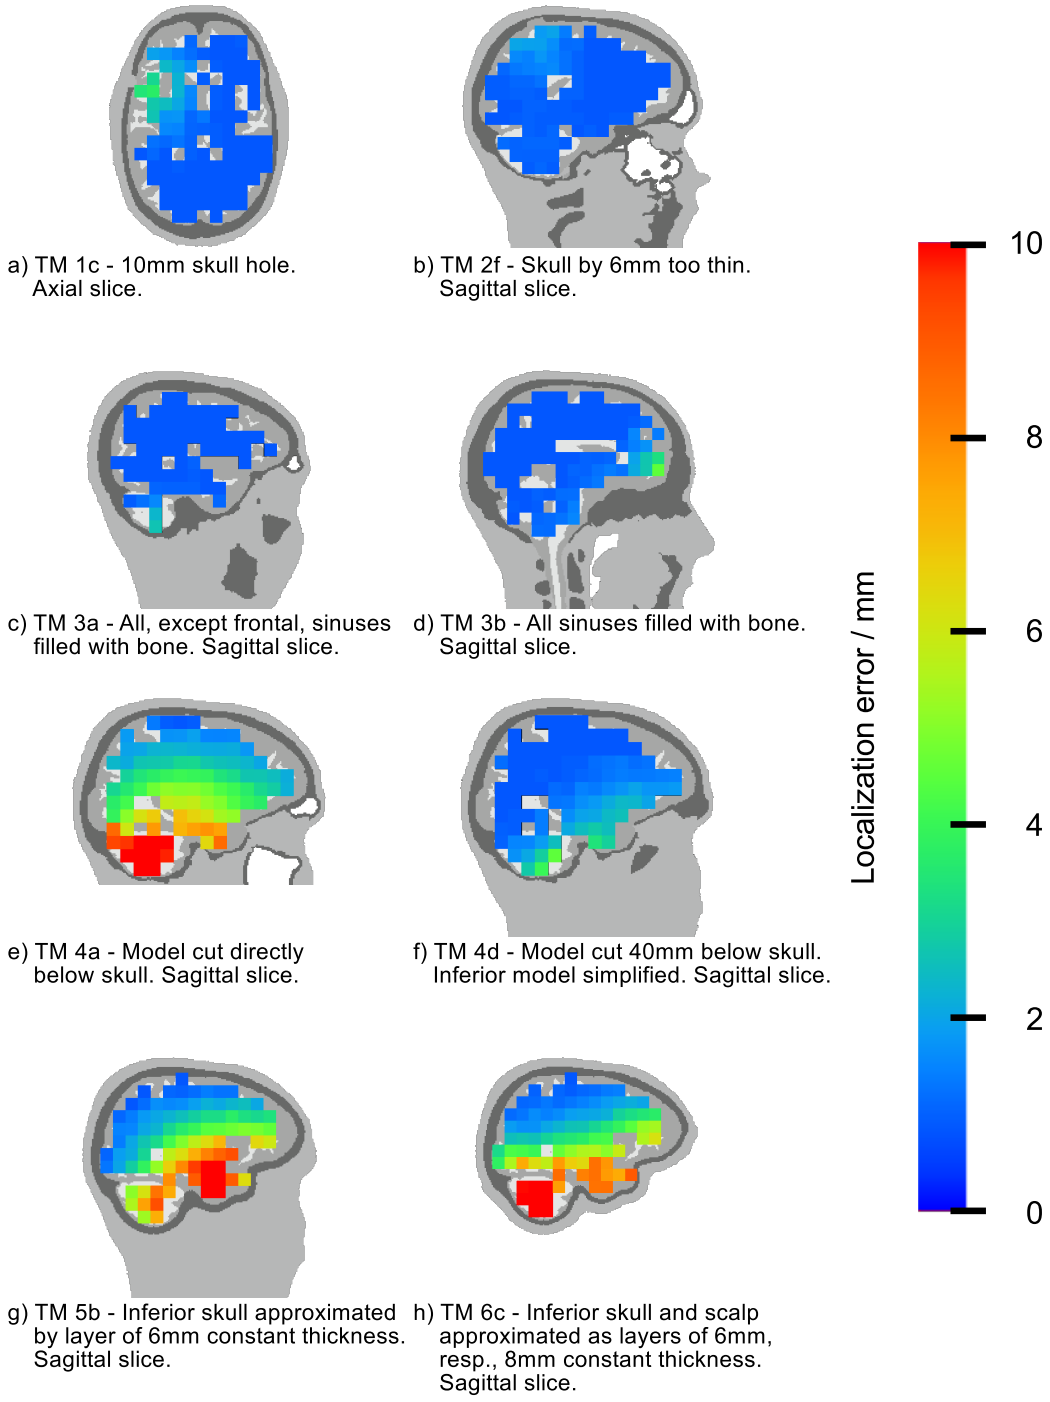

Figure 4.6: Localization error maps for selected test models. Errors were averaged across sources located at the same position but with different orientations.

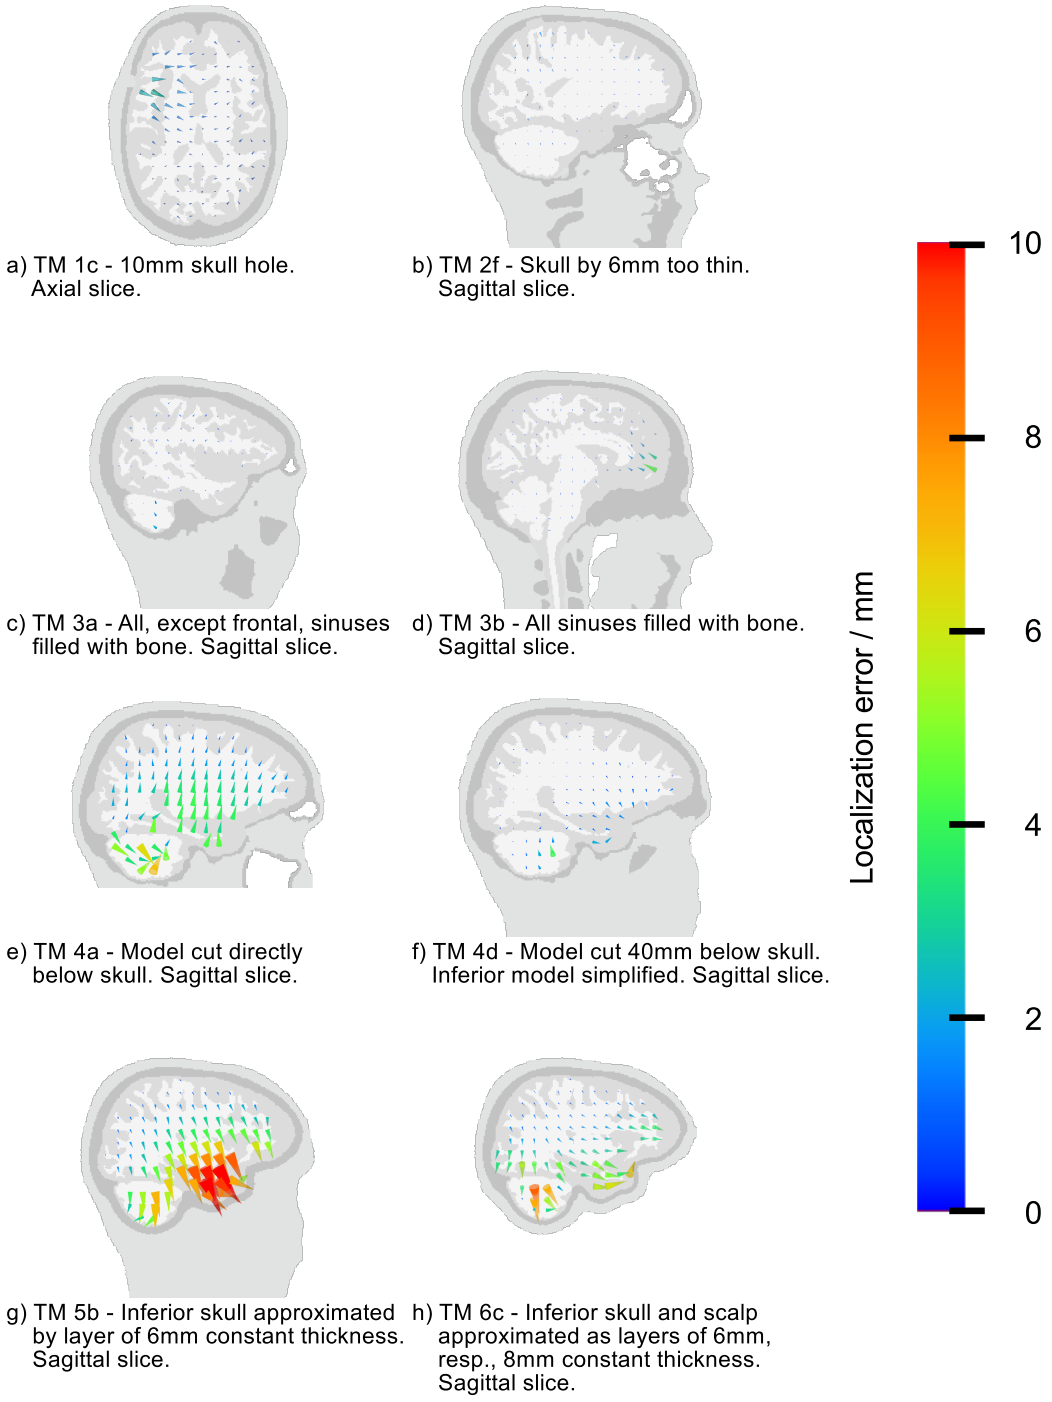

Figure 4.7: Mislocalization tendency maps for selected test models. Mislocalization vectors were averaged across sources located at the same position but with different orientations.

to 5c). Probe sources close to the skull base were reconstructed with the largest errors of up to 24 mm but probe sources at some distance to the skull base also showed considerable localization errors (Fig. 4.6g). All affected sources were consistently mislocalized towards the skull base (Fig. 4.7g).

In the sixth set of test models it was investigated how *simplifying the inferior skull and scalp* as layers of constant thickness influences source analysis. Regarding the influence on source reconstruction large localization errors were observed especially for sources close to the skull base (Fig. 4.6h). It was not possible to make out a consistent mislocalization tendency across TM 6a to 6c. In TM 6a sources were still consistently mislocalized towards more superior positions. When increasing the constant thickness of the inferior scalp, however, this tendency changed. Figure 4.7h shows that in TM 6c frontal and temporal sources were reconstructed at more anterior locations while sources in occipital regions were reconstructed towards inferior positions. In TM 6c for some sources in the cerebellum only small mislocalization tendencies could be observed (Fig. 4.7h) while the average localization errors for these sources were very large (Fig. 4.6h). This indicates that the mislocalizations for cerebellar sources with different orientations are not consistent so that the mislocalization vectors for these sources partly cancel out when averaging across dipole orientations for visualization.

The source reconstruction errors for all test models are summed up in Figure 4.8 where the mean and maximum localization errors, as well as the proportion of affected sources are plotted. The mean and maximum localization errors were computed without averaging the errors across dipole directions beforehand.

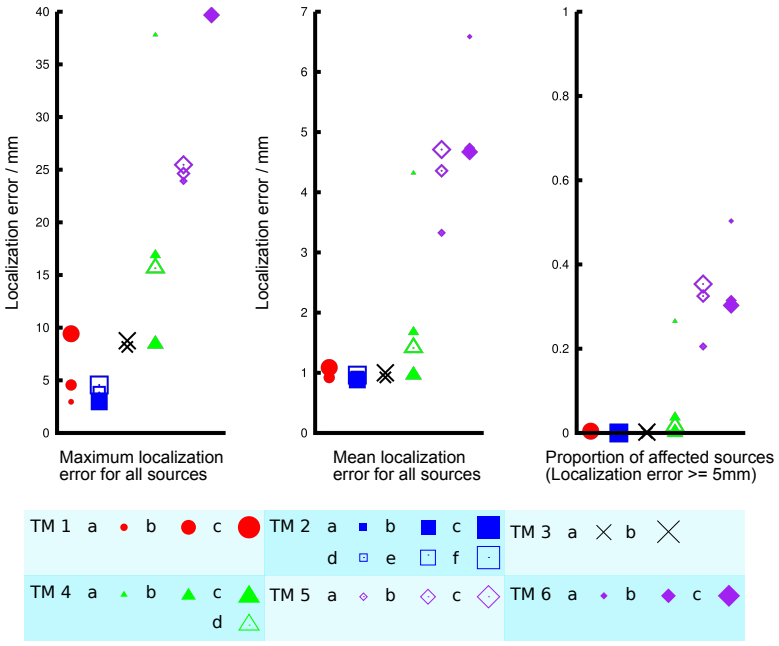

Figure 4.8: Localization error characteristics for all probe sources.

When comparing the source reconstruction errors between models in the same series of test models nearly the same tendencies were found as reported before for the forward errors.

Regarding the localization errors averaged across all probe sources it could be observed that the mean errors are close to the minimum achievable localization error for all models in the first three sets of test models (TM 1a-c, 2a-f, 3a,b). The mean errors for models TM 4b to 4d are also small staying well below 2 mm. Large errors were however observed for the model cut directly below the skull (TM 4a), and for the models where either the inferior skull, or the inferior skull and scalp were approximated as layers of constant thickness. In these models average localization errors reached values ranging from 3.2 mm to more than 6.5 mm. In addition, more than 20% of the probe sources in these models were mislocalized by at least 5 mm, that is, they can be regarded as affected by the respective skull geometry inaccuracy.

4.4 Discussion and Conclusion

In the present study influences were investigated of a wide range of skull geometry inaccuracies on the EEG forward solution and on source reconstruction. The spatial distribution of the forward and inverse errors were presented in error maps and the mislocalization tendencies were described. Finally, all studied influences were compared.

Summing up the presented results it was found that a local skull defect, that is, one which only effects a small part of the skull, caused a local error distribution where only sources close to the defect are affected. The lowest errors were observed for the models with small and local skull geometry inaccuracies. These were the model with a skull hole of only 2 mm diameter (TM 1a) and the model where the skull thickness is over- or underestimated by 2 mm in a circular skull region (TM 2a,d). Introducing changes to the geometry across larger parts of the model, for example, when simplifying the inferior skull or when cutting the model directly below the skull, lead to larger errors and a more extended error distribution. In test models 4b,c and 4d the model was cut at distances from 20 mm to 40 mm below the skull. Additional simplifcations were applied to the inferior head and neck region in TM 4d. The rather large changes in the model geometry, especially for TM 4d, only lead to minor errors for these models. This shows that even large modifications of the model can have only small influences if the modifications affect model regions far away from the source space and the sensors.

A carefully constructed reference model provides the basis for the presented simulation study. From the experience of the author the used reference model can be regarded as detailed and anatomically plausible. Yet, some limitations of the model have to be discussed.

MRI data was the only source of anatomical information for the construction of the reference model. Due to the properties of MRI it must be expected that some aspects of the reference model are not entirely accurate. As common MRIs show virtually no contrast between compact bone and air-filled cavities, it must be expected that there is some uncertainty regarding the shape and position of the sinuses in the reference model. An overestimation of the size of the sinus cavities would lead to larger errors in TM 3a and 3b. In addition, some simplifications regarding the geometry of the skull base were necessary for the construction of the reference model. The skull base has dozens of skull openings through which nerves and blood vessels pass. The reference model used in this study, however, only incoporates the foramen magnum and the openings, through which the eye nerves pass. It can be assumed, that incorporating additional openings would lead to larger errors for the test models with a simplified inferior

skull layer (TM 5a-c, 6a-c) where the openings in the skull base are effectively ignored.

Another aspect that has to be discussed is the uncertainty of the used conductivity values. For most tissues conductivity values that have been reported in the literature vary strongly between studies and even between samples of the same study. This is especially true for the conductivity of the skull [4, 97]. A higher value for the skull conductivity, which is closer to the conductivities of the surrounding tissues, would have lead to smaller errors for nearly all investigated skull defects. A smaller skull conductivity, on the other hand, would amplify nearly all studied influences. The test models where the air-filled sinus cavities are modeled as compact bone are the only exception to this rule. For these models, errors would decrease for a lower, and increase for a higher skull conductivity. These assumptions are supported by additional simulations that were carried out using a by 28% higher compact bone conductivity (0.0064 S/m instead of 0.005 S/m). Using the higher compact bone conductivity slightly smaller forward and inverse errors were found. The proportion of affected sources in TM 6c, for example, only changed from 30.3% to 27.0%. The authors, therefore, believe that the results of the present study are representative even if the exact skull conductivity values are not known.

Previous publications have described comparable studies on the influence of skull defects on EEG source analysis. Bénard et al. [11] investigated how the EEG forward solution is affected when ignoring a skull hole introduced into the skull, for example, by surgery. The authors report that forward errors were large in the proximity of the hole and decreasing for sources at a larger distance to the hole. This is in agreement with the RDM maps presented in Figure 4.3a. The source reconstruction results reported by Bénard et al. cannot directly be compared to the inverse errors shown, for example, in Figures 4.6 or 4.7. This is due to the fact that Bénard et al.'s study differs from the study presented in this work in one important aspect. Bénard et al.'s motivation is to investigate how skull holes introduced by, for example, surgery affect the source analysis. In their study design the model incorporating the hole is the reference model, while the model with the intact skull is the test model. In the design of the study presented here the opposite is assumed. The skull in the reference model is intact, while it has an additional hole in the test model. It is, nevertheless, possible to compare the general mislocalization tendency. Bénard et al. report that sources are reconstructed at positions more towards the skull hole when ignoring the skull hole. This agrees with the present study where it was observed that sources are mislocalized towards deeper positions away from the hole when the skull was modeled with an additional hole (4.7).

Another study investigating the influence of skull holes on source reconstruction is that of Vanrumste et al. [142]. The error maps presented in their work agree very well with the maps presented in Figure 4.3a. The maximum localization error reported by Vanrumste et al., however, is approximately 5 mm for a skull hole of 20 mm diameter, which is clearly smaller than the maximum error of 9 mm found in the study presented here for a skull hole of only half the diameter. A reason for this could be the differences between the two studies regarding the conductivity values for the skull and the skull hole. Vanrumste et al. chose a value of 16 for the ratio of skull hole conductivity to skull conductivity which is considerably smaller than the ratio of 22 which was chosen for the study at hand.

Cuffin discussed in his work [36] the influence of local errors in the skull thickness on source reconstruction reporting localization errors well below 1 cm. Similar results with a maximal mislocalization of less than 5 mm (Fig. 4.8) were obtained in the study presented here.

Some work on the downward extension of the volume conductor model was done by Bruno et al. [24]. In this work the authors report large forward errors when cutting the model along a plane that intersects the inferior skull, and non-linearly decreasing errors when cutting the model at larger distances from the skull. These observations are consistent with the averaged RDMs for TM 4a-c as plotted in Figure 4.5.

Finally, the model errors revealed by the present study shall be put into the context of other model errors that were previously studied in the literature. Dannhauer et al. [39], for example, investigated how modeling errors with regard to the internal structure of the skull bone affect source analysis. In their work Dannhauer et al. report large localization errors of 10 mm and more for sources throughout large areas of the source space when approximating the three-layered skull bone with a single isotropic layer. This error is on par with the errors that were observed in the study at hand when cutting the model directly below the skull (TM 4a), or when either simplifying the inferior skull (TM 5a-c) or the inferior skull and scalp (TM 6a-c).

Yvert et al. [162] studied how large the source reconstruction errors are that have to be expected when using simplified spherical models instead of realistic head models for EEG source analysis. Errors of $4 - 6$ mm for superior sources and $15 - 25$ mm for inferior sources were observed. These errors are approximately on the same order of magnitude as the errors observed for the models in the fifth and sixth set of test models.

In the present study a wide range of skull geometry inaccuracies were modeled in always the same study setup. In this way, it is possible to compare the influences of the different skull defects and, thus, the following guidelines can be formulated for the construction of accurate volume conductor models for EEG source analysis. Volume conductor models should not be cut immediately below the skull but they should be extended downwards at least 20 mm or even better 40 mm below the skull. Approximating the inferior head and neck region as a homogeneous region using a single isotropic conductivity is acceptable. It should be avoided to approximate the inferior skull or scalp as layers of constant thickness. The complicated geometry of the skull base should instead be modeled as accurately as possible. Skull holes larger than 2 mm and errors in skull thickness should be avoided if possible. The local errors caused by the latter skull defects might, however, be negligible in many applications.

The above guidelines can also serve as criteria to evaluate the performance of an automatic segmentation procedure. When observing, for example, skull holes in the automatic segmentation results then it can be expected that considerable localization errors occur in the vicinity of the hole when performing source analysis with the volume conductor model based on the segmentation. The performance of the two segmentation approaches which are presented later in this work will be also discussed in the context of these criteria.

5 Influence of Interior CSF Compartments on EEG Source Analysis

5.1 Introduction

In the previous chapter skull geometry inaccuracies in volume conductor models were investigated as an aspect influencing the accuracy of EEG source analysis. Another aspect of interest in this context is the modeling of the CSF filled spaces in the head. Realistic volume conductor models of the human head in practice at least differentiate the tissues scalp, skull, and brain. Previous studies have shown that also incorporating the CSF into these models can increase precision for EEG source analysis [110]. The CSF in the human head is mainly found in the ventricles, but also in deep sulci and in the space between the brain surface and the inner bone.

For source analysis it is generally assumed that sources must be placed inside of the brain compartment, and must not be placed in the CSF. Incorporating the interior CSF spaces, thus, has an influence on the space of allowed source positions. In particular, modeling the ventricles will lead to a source space which has large holes and an overall more complex geometry as compared to a simpler model not differentiating between brain and CSF.

A complex source space with holes and concavities can cause problems for certain inverse methods. The classical dipole fit approach [122], for example, might get ”stuck” due to the concavities of the source space. As a consequence the globally best fitting source position might not be found.

A possible solution to these problems is to use volume conductor models which are simplified with respect to the CSF. In a simplified model the CSF could be ignored either completely, or only interior CSF filled spaces, like the ventricles and deep sulci, could be ignored. It has to be expected that these model simplifications cause errors in the solution of the EEG forward problem and, thus, also for source reconstruction. The aim of the study presented in this chapter is to quantify these errors and to investigate if they are acceptable for our application.

The design of this study closely resembles that of the study presented in Chapter 4. Reference data was simulated in a reference model. Test models with partially or completely ignored CSF are constructed and the potentials simulated therein were compared to the reference potentials. In addition, source reconstruction was performed on the reference data while solving the associated forward problem in the simplified test models to assess the influence of the simplifications on the inverse problem.

The present study was published by the author of this thesis as first author in [75]. The presentation in this chapter closely follows this publication.

5.2 Materials and Methods

5.2.1 Study Design

Reference data for the simulation study was generated using a reference volume conductor model. The reference model is detailed and anatomically plausible, so that the reference data can be regarded as representative for data as it would be measured during an actual experiment.

From the reference model two test models were derived that represent simplifications with regard to the CSF. The first test model (TM A) does not differentiate between brain and CSF, so that the complete intracranial cavity is modelled as brain tissue. In the second test model (TM B) only interior CSF filled spaces are ignored by modelling them with the same conductivity as brain tissue. These interior CSF filled spaces included the ventricles and deep sulci. The CSF layer between the brain surface and the inner bone is still accounted for in TM B.

EEG potentials were computed for a set of probe sources in the test models and in the reference model. By directly comparing the simulated potentials it is possible to assess the influence of partially or completely ignoring the CSF on the solution of the EEG forward problem.

To investigate the influence on the inverse problem reference data for a set of probe sources was computed in the reference model. This reference data was then reconstructed using a suitable source reconstruction method. The forward problem associated with source reconstruction was thereby solved in the simplified test models. As a suitable source reconstruction method here the robust goal function scan method was employed. A fine 1 mm scanning grid was used to achieve a high spatial resolution for the solution of the inverse problem.

The same sets of forward and inverse probe sources, as well as the same electrode configuration were used for this study as were used for the study on the influence of skull geometry inaccuracies (Chap. 4). Forward and inverse errors were quantified using the error measures described in Section 4.2.4. Error maps of forward (RDM and MAG) and inverse errors were plotted to illustrate their spatial distribution. For this purpose errors were averaged across probe sources at the same position but with different orientations.

5.2.2 Construction of Reference and Test Models

As a reference model here the detailed and anatomically plausible model was used that was also used for the study described in the previous chapter. An in-depth description of the reference model construction can be found in Section 4.2.2.

Two test models were derived from the reference model. TM A was constructed by re-labeling all CSF voxels in the reference segmentation as brain. In TM B only those CSF voxels were relabeled that were situated in interior CSF filled spaces, like the ventricles or deep sulci. This was done as follows. A mask was created containing all voxels in the reference segmentation that were marked as brain. A dilation operation with a large dilation radius (15 mm) was then applied to the mask. Holes inside of the dilated mask which correspond to interior CSF spaces were filled. Subsequently, the mask was eroded by the same radius as used for the dilation. The final labeled image for TM B was constructed by re-labeling all CSF voxels as brain that were located inside of the previously generated mask.

5.3 Results

Comparing the forward solutions in the reference and test models resulted in the RDM and MAG error maps presented in Figures 5.1a and 5.1b, respectively. Figure 5.2 shows histograms of the RDM, MAG and localization errors to sum up the distributions of the errors across all sources.

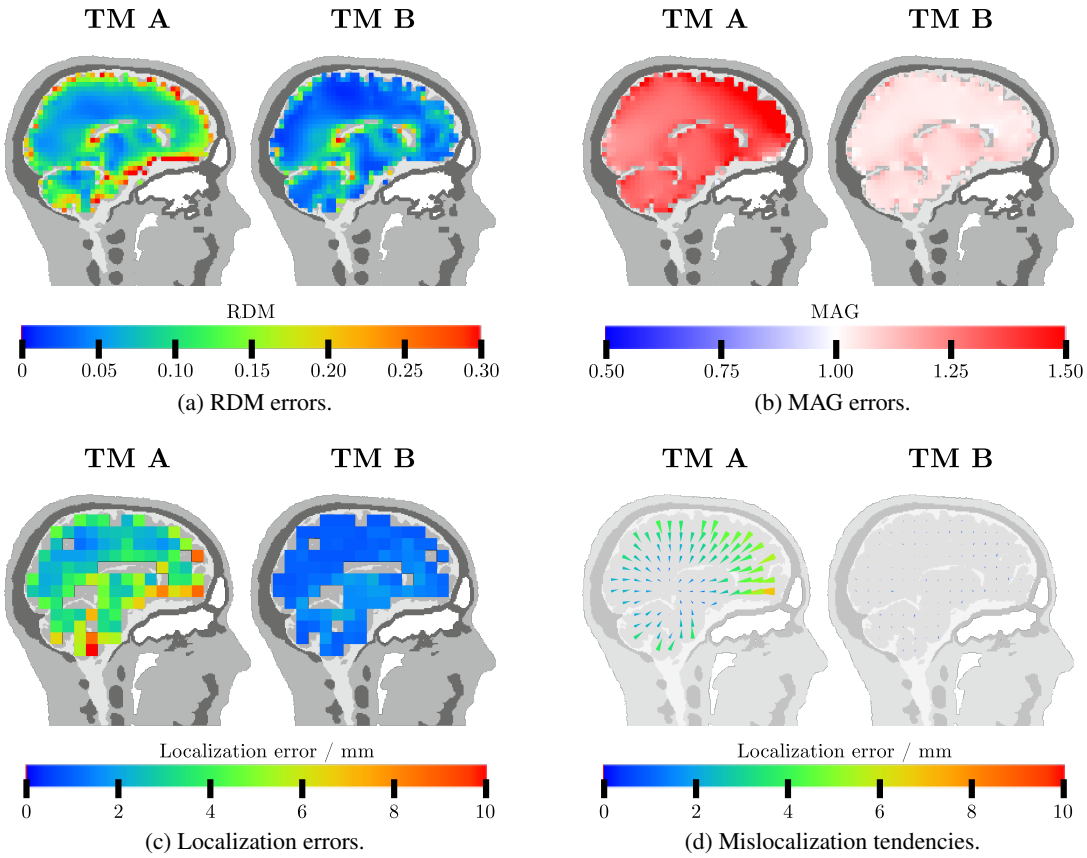

(a) RDM errors. (b) MAG errors. (c) Localization errors. (d) Mislocalization tendencies.

Figure 5.1: Error maps showing forward and inverse errors for simplified CSF models.

Large RDM errors (> 0.1) could be observed in TM A for superficial sources and sources close to the interior CSF compartments (Fig. 5.1a). The substantial RDM errors for sources close to, for example, the ventricles were also seen in TM B. Superficial sources, however, were in general less affected in TM B than in TM A (Fig. 5.2).

The MAG error maps for TM A (Fig. 5.1b) draws a clear picture. Sources throughout the complete source space consistently showed MAG errors that were clearly larger than 1. This means that surface potentials for most sources were substantially overestimated in TM A. For TM B it could also be observed that the MAG error was above 1 for most sources. The MAG in TM B, however, was everywhere notably closer to the optimal value of 1 than in TM A.

Figures 5.1c and 5.1d show the localization errors, respectively, the mislocalization tendencies caused by completely or partially ignoring the CSF. For TM A notable localization errors

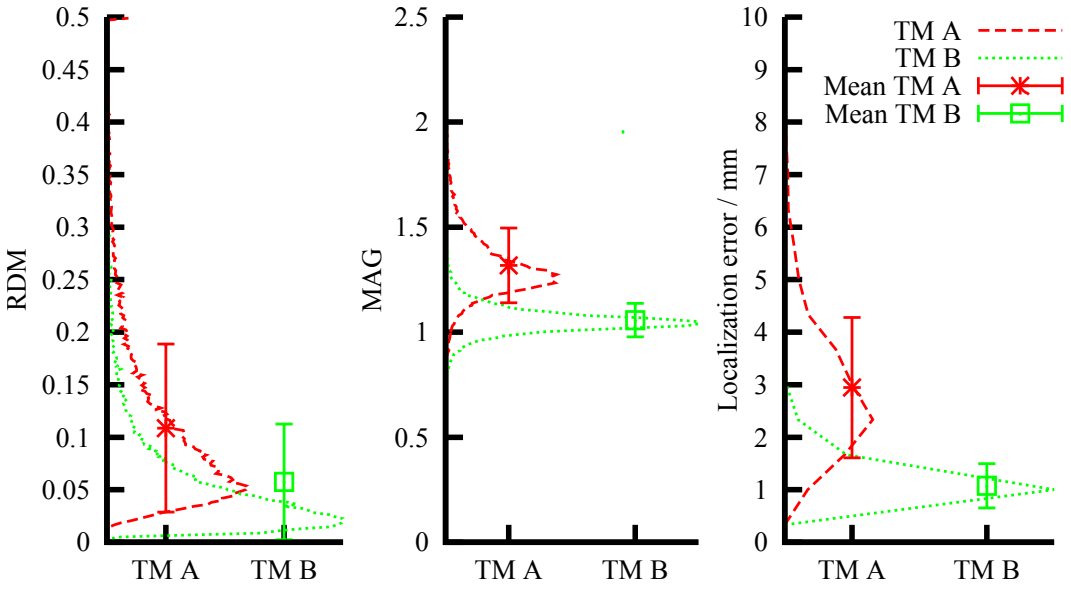

Figure 5.2: Histograms of RDM, MAG and localization errors for TM A and B.

up to 4 mm and more could be seen for frontal sources and sources along the skull base. The mislocalization tendencies indicated that sources are reconstructed too far into the depth of the head when completely ignoring the CSF. Observed source reconstruction errors when only partially ignoring the CSF as in TM B were small (< 2 mm) for most sources (Fig. 5.2). Due to the small size of the localization errors in TM B no mislocalization tendency could be made out in Figure 5.1d.

5.4 Discussion and Conclusion

In summary, the presented results showed large forward and inverse errors when completely ignoring the CSF. Only ignoring interior CSF filled compartments in comparison lead to considerably smaller errors. Source reconstruction errors were particularly small so that the influences of interior CSF spaces might be negligible for most applications.

In the present study only single source source scenarios were investigated. Inverse errors in source scenarios with two or more sources might be considerably larger due to the instability of source reconstruction for multiple sources.

In the context of the presented results the work of Ramon et al. [110] is of importance. Ramon et al. also investigated how EEG source analysis is affected when completely ignoring CSF. They report average localization errors of $2 - 3$ mm depending on the dipole orientation for sources in the area of the motor cortex. Similar localization errors were found in the present study for sources in the some regions of the cortex, for example, in the area of the motor cortex and the occipital cortices. However, also larger errors of 4 mm and more were observed, for example, for frontal sources. A reason for this difference might be the presence of relatively large CSF filled spaces in frontal areas. In general, errors when completely ignoring the CSF are expected

to be larger in the vicinity of large CSF filled spaces, and smaller when there is no CSF nearby.

The large magnitude errors which could be observed when ignoring the CSF completely are in good agreement with a recent study of Rice et al. [114]. In their work the authors investigate how a different positioning of the head can lead to a difference in the local thickness of the CSF layer. They observe that the CSF layer in the occipital region is 30% thinner when the subject is lying on his or her back instead of sitting upright. This also leads to a relative increase of the EEG magnitude by on average 80%. This corresponds to the large MAG values we found in TM A (Fig. 5.1b).

The errors for TM B reported in this study, in particular the marginal single source reconstruction errors, finally lead to the conclusion that internal CSF filled spaces, like ventricles or deep sulci, can safely be ignored while still allowing a relatively high accuracy for EEG source analysis. Simplifed head models that resemble TM B can, thus, be used to avoid the problems of complex source space geometries for some inverse methods.

6 A Pipeline for MR Image Processing in the Context of EEG Source Analysis

6.1 Introduction

The focus of this thesis is on our newly developed methods for the segmentation of volume conductor models from MRI data. Before the proposed methods can be applied to the data, however, some preprocessing is necessary. The pipeline for preprocessing the MRI data, which was implemented into *BESA MRI* [17] software as part of the dissertation work, is described in this chapter. In addition, the pipeline also delivers further anatomical information from the MRI data which can be used for the purpose of EEG source analysis.

The pipeline consists of reading the T1- and optional T2-weighted MR data, and transforming them into AC-PC space which is defined by landmarks set by the user.

In the next step, the optional T2-weighted MRI is aligned with the T1-weighted image. The mutual information registration measure [152, 85] and the Nelder-Mead downhill simplex algorithm [94] were used to estimate the optimal affine transformation which registers the T2- to the T1-weighted MRI.

A possible shading artifact is removed by applying our modified version of the AFCM algorithm [105] to the data. The AFCM algorithm and our modifications improving the accuracy and robustness are described below (Sec. 6.2).

A representation of the individual scalp surface can be used in EEG source analysis for the visualization of the scalp potentials and for aligning the electrode coordinates to the individual head surface. For this reason, a method for the reconstruction of the scalp surface from the individual MRI was implemented (Sec. 6.3). This method also delivers a scalp reference surface which is used during the volume conductor segmentation (Chaps. 8 and 9).

A surface representation of the individual subject's cortex can serve as a source space in EEG source analysis, that is, it can define the space where sources of the EEG signal are expected. The reconstruction of a suitable cortex surface was implemented in the *BESA MRI* software (Sec. 6.4). As a side-product the method for the reconstruction of the cortex surface also yields a WM reference surface for the volume conductor segmentation.

As a final step in the pipeline, after the inhomogeneity corrected MRI, and the scalp and WM reference surfaces have been prepared the volume conductor segmentation can be run (Chaps. 7 and 8).

6.2 Intensity Inhomogeneity Correction

The intensity inhomogeneity or shading artifact (Sec. 3.2.2) is very common in MRIs. Due to this artifact spatially smooth variations can be observed in the gray values of homogeneous tis-

sue regions. These spatially varying gray values affect the performance of tissue segmentation procedures that rely on the assumption of homogeneous intensity regions, for example, the cortex segmentation described in Section 6.4. For this reason, our pipeline includes a correction step in which the intensity inhomogeneity is removed. The implemented procedure for this step is based on the AFCM clustering method proposed by Pham and Prince [105]. For a previous application of the AFCM algorithm on MRI data of the human head see the work of Wolters ([155], Chap. 2.5). Here, we extended the AFCM approach to incorporate a-priori information to improve robustness and reliability of the estimation of the inhomogeneity field. A summary of the implemented method is given in Algorithm 1.

Algorithm 1 The extended AFCM algorithm for intensity inhomogeneity correction.

Require: $\boldsymbol{y}$: image data (only one modality).
Require: c: number of classes.
Require: n: maximum number of iterations.

Initialize evenly distributed class centroids μ_k
Apply pre-correction method to initialize multiplier field $\boldsymbol{m}$
$count \leftarrow 0$ ▷ Reset iteration counter

repeat
    Update membership values u_i^k according to Eq. 6.10
    Update class centroids μ_k according to Eq. 6.9
    Estimate multiplier field m_i by iteratively solving Eq. 6.11
    $count \leftarrow count + 1$ ▷ Increase iteration counter
until $(count \geq n)$
$y_i \leftarrow y_i \, (m_i)^{-1}$ ▷ Correct image intensities

The general idea of the AFCM is to cluster the data points y_i, that is, the gray values of the MRI, into c clusters. Clustering is done in a fuzzy way, meaning, that a voxel is not strictly assigned to a single cluster. Instead membership values u_i^k $\left(\sum_{k=1}^{c} u_i^k = 1\right)$ are determined for each voxel describing, in some sense, how probable it is that the voxel belongs to the k-th cluster. The centroids μ_k of the clusters are a-priori not known. A spatially smoothly varying multiplier field m_i is introduced to account for the intensity inhomogeneity. This correction field is multiplied to the centroids so that there will effectively be a different centroid μ_k at each site i of the image grid. The introduction of the multiplier field is based on the assumption that the shading artifact acts on the true gray values y_i^* in the form of a multiplicative field.

$$y_i = y_i^* \cdot 1/m_i + \text{noise} \tag{6.1}$$

The fuzzy clustering done by the AFCM approach can be expressed as the optimization of a suitable cost function. For the classical AFCM as proposed by Pham and Prince [105] two terms contribute to the cost function. The first term J_{cluster} is related to the distance of the data points to the class centroids, and the second term J_{mult} accounts for the smoothness of the multiplier

field .

$$J_{\text{AFCM}} = J_{\text{cluster}} + J_{\text{mult}} \tag{6.2}$$

The overall distance of the data points to the cluster centroids is measured in the classical AFCM algorithm as the weighted sum of the squared distances of the data points to the corrected class centroids. The class centroids are corrected by the spatially varying multiplier field and the squared membership functions are used as weighting factors.

$$J_{\text{cluster}} = \sum_i \sum_{k=1}^{c} \left(u_i^k\right)^2 \|y_i - m_i \mu_k\|^2 \tag{6.3}$$

The second term J_{mult} guarantees a smoothly and slowly varying multiplier field. It can be thought of as a regularization of the cost function to avoid high-frequency spatial oscillations of the multiplier field due to, for example, image noise.

$$J_{\text{mult}} = \lambda_1 \sum_i \|\nabla m_i\|^2 + \lambda_2 \sum_i \|\nabla (\nabla m_i)\|_F^2 \tag{6.4}$$

Here ∇ stands for a gradient operator, λ_1 and λ_2 are regularization parameters, and $\|\cdot\|_F$ is the Frobenius matrix norm. Also note that in the last term the Frobenius norm is computed for the gradient of the gradient of the multiplier field. Using the discrete forward difference operator with respect to the r-th coordinate D_r the multiplier field term of the cost function can also be expressed as follows [105]:

$$J_{\text{mult}} = \lambda_1 \sum_i \sum_{r=1}^{3} (D_r * m)_i^2 + \lambda_2 \sum_i \sum_{r=1}^{3} \sum_{s=1}^{3} (D_r * D_s * m)_i^2 \tag{6.5}$$

In this equation $(D_r * m)_i$ denotes the convolution of the forward difference operator D_r with the discretized multiplier field m evaluated at node i.

Our contribution is to extend the cost function of the original AFCM algorithm [105] by introducing a weighting factor p_i^k in the first term.

$$J_{\text{cluster}} = \sum_i \sum_{k=1}^{c} \left(u_i^k\right)^2 \|y_i - m_i \mu_k\|^2 \left(p_i^k\right)^{-1} \tag{6.6}$$

The inverse weighting factor p_i^k can be defined using a-priori probability maps. In our implementation we derive the p_i^k from an empirically determined WM probability map.

$$p_i^k = \begin{cases} p_i^{\text{WM}} & \text{if } k = c \\ \frac{1-p_i^{\text{WM}}}{c-1} & \text{else} \end{cases} \tag{6.7}$$

In regions where $p_i^c = p_i^{\text{WM}} \approx 1$ a large weight is put on the terms

$$\left(u_i^k\right)^2 \|y_i - m_i \mu_k\|^2, \qquad k = 1, \ldots, c-1$$

To minimize these the membership values u_i^k for the clusters $k = 1, \ldots, c-1$ must become small. In this way, the voxels in the regions with $p_i^c \approx 1$ are much more likely clustered into the c-th cluster. This is intended by us.

Our WM probability atlas was empirically determined by generating WM segmentations for 66 T1-weighted MRIs using the *FreeSurfer* software [48]. The WM segmentations were transformed to the Talairach space. For each voxel in Talairach space it was then counted in how many of the segmentations WM occurs. A probability map was obtained by normalizing the counted occurrences.

In a similar way, an additional weighting factor w_i only depending on the location i can be introduced to focus the clustering on a certain region of interest (ROI). This can be used to compute a clustering of the image data that best fits, for example, in the brain region.

The extended cost function (Eq. 6.8) is optimized in an iterative scheme successively with respect to the membership functions u_i^k, the class centroids μ_k, and the multiplier field m_i.

$$J_{\mathrm{AFCM}} = \sum_i \sum_{k=1}^{c} \left(u_i^k\right)^2 \|y_i - m_i\,\mu_k\|^2 \left(p_i^k\right)^{-1} + \lambda_1 \sum_i \|\nabla m_i\|^2 + \lambda_2 \sum_i \|\nabla\left(\nabla m_i\right)\|_F^2 \tag{6.8}$$

Closed formulas for the class centroid and membership values that optimize the cost function can be derived by setting the partial derivative of the cost function with regard to these variable equal to zero and solving for the sought after variable. Incorporating the a-priori probability p_i^k the following formulas are obtained:

$$\mu_k = \frac{\sum_i \left(p_i^k\right)^{-1} \left(u_i^k\right)^2 m_i y_i}{\sum_i \left(p_i^k\right)^{-1} \left(u_i^k\right)^2 m_i^2} \tag{6.9}$$

$$u_i^k = \frac{\|y_i - m_i\mu_k\|^{-2}\, p_i^k}{\sum_{l=1}^{c} \|y_i - m_i\mu_k\|^{-2}\, p_i^k} \tag{6.10}$$

Computing the partial derivative of the cost function with respect to the multiplier field value m_i results in a matrix equation (Eq. 6.11).

$$\boldsymbol{W}\,\boldsymbol{m} + \lambda_1 H_1 * \boldsymbol{m} + \lambda_2 H_2 * \boldsymbol{m} = \boldsymbol{f} \tag{6.11}$$

$$\text{with:} \quad \boldsymbol{W} = \mathrm{diag}(\boldsymbol{w}) \tag{6.12}$$

$$w_i = \sum_{k=1}^{c} \left(u_i^k\right)^2 \mu_k^2 \left(p_i^k\right)^{-1} \tag{6.13}$$

$$f_i = \sum_{k=1}^{c} \left(u_i^k\right)^2 \mu_k\, y_i \left(p_i^k\right)^{-1} \tag{6.14}$$

$$H_1 = \sum_{r=1}^{3} \left(D_r + \check{D}_r\right) \tag{6.15}$$

$$H_2 = \sum_{r=1}^{3} \sum_{s=1}^{3} \left(\left(D_r * D_s\right) + \left(\check{D}_r * \check{D}_s\right)\right) \tag{6.16}$$

For the definition of the differential operators H_1 and H_2 the mirrored forward difference opera-

tors $\check{D}_r$ are used.

The dimensions of $\boldsymbol{m}$ are typically on the order of several hundred thousands. A direct inversion of Equation 6.11 is, thus, not feasible. An estimate for the multiplier field $\boldsymbol{m}$ can still be computed by using an iterative solver. In our implementation a Jacobi-scaled CG solver [20] was employed. For this kind of solver in combination with the AFCM algorithm appropriate convergence rates have been demonstrated [155].

We choose a non-parametric representation for the multiplier field. The field is defined by its values on a regular cubic grid. The multiplier field is, however, not represented on the same grid on which the image data is defined. The multiplier field is computed on a coarse 4 mm grid while the image data is given on a fine 1 mm grid. The 4 mm resolution for the multiplier field is sufficient because the multiplier field is assumed to only slowly vary across space, and the lower resolution substantially speeds up the computations. The multiplier field values on the fine grid which are needed in Equations 6.9 and 6.10 are obtained by linear interpolation from the coarse grid. In Equations 6.13 and 6.14 we are summing up across all fine grid nodes i' which are in the domain Ω_i of the coarse voxel i.

$$w_i = \sum_{i' \in \Omega_i} \sum_{k=1}^{c} \left(u_{i'}^k\right)^2 \mu_k^2 \left(p_{i'}^k\right)^{-1} \tag{6.17}$$

$$f_i = \sum_{i' \in \Omega_i} \sum_{k=1}^{c} \left(u_{i'}^k\right)^2 \mu_k y_{i'} \left(p_{i'}^k\right)^{-1} \tag{6.18}$$

The AFCM is initialized by choosing initial values for the class centroids and the multiplier field. In our implementation the initial class centroids were evenly distributed across the gray value range. The simplest choice for the multiplier field is $m_i = 1, \forall i$. A better initial guess of the multiplier field might increase the speed of convergence and improve the robustness of the AFCM algorithm.

In some cases using the AFCM a non-optimal intensity inhomogeneity correction was observed in the inferior part of the head when there was a strong decay of signal intensity from superior to inferior. We, therefore, implemented a simple pre-correction procedure in which we try to estimate the component of the multiplier field which is responsible for this strong signal decay in inferior direction. In this procedure the user places a few marker points at different z-coordinates into the brain stem in the MRI. The WM tissue in a certain diameter around these marker points is then segmented by simply thresholding the intensity values. We assume a multiplier field which is varying in z-direction only. A spline function $f(z)$ is fitted to the intensities of the segmented WM voxels and their z-coordinates. The multiplier field at voxel i with coordinates (x, y, z) is then initialized as $m_i = f(z_{\max})/f(z)$, where $z_{\max}$ is the maximum z-coordinate of all data points. A faster convergence and a more robust classification could be observed when initializing the multiplier field for the AFCM algorithm using the above procedure.

The main result obtained by this method is the intensity inhomogeneity corrected MRI. It is computed by dividing the gray values y_i of the original image by the multiplier field values m_i.

The performance of the implemented intensity inhomogeneity correction method was evaluated by inspecting the intensity histograms and by direct inspection of the corrected images. The shading artifact is typically most obvious in the WM region (Fig. 6.1a). It causes a blurring of the

Figure 6.1: Typical results of the implemented inhomogeneity correction procedure. Results were obtained for $\lambda_1 = 1$, $\lambda_2 = 150$, and 20 iterations. a) Original T1-weighted MRI with noticeable shading artifact; b) Intensity histogram of original MRI; c) Resulting multiplier field $\boldsymbol{m}$; d) Resulting intensity inhomogeneity corrected image; e) Intensity histogram of corrected image; f) Hard classification obtained from clustering algorithm.

intensity distribution so that the expected histogram peaks of the distinct tissues cannot be seen in the image histogram (Fig. 6.1b). After applying our implemented correction method nearly no inhomogeneity can be made out in the WM region, and clearly distinct intensity peaks are observable in the intensity histogram of the corrected MRI (Fig. 6.1d and e). The resulting multiplier field is very smooth, as it was desired, and the resulting classification is plausible (Fig. 6.1c and f).

In some cases the incorporation of the WM a-priori probability atlas substantially improved the inhomogeneity correction (Fig. 6.2). Using the classical AFCM without incorporating the a-priori probabilities yielded an unsatisfactory result (Fig. 6.2b) for an input T1-weighted MRI which was too dark in the center of the brain (Fig. 6.2a). By incorporating the a-priori probabilities a better result with a more homogeneous WM could be obtained (Fig. 6.2e). The improvement becomes even more obvious when looking at the intensity histograms and at intensity profiles along a line in the WM region (Fig. 6.2c and f). In particular, the WM peak is much more pronounced in the histogram of the MRI that was corrected taking the a-priori probability into account. The intensity profiles of the original image and the image corrected using the classical AFCM show a strong intensity decay by approximately 30 %, respectively, 14 % in the center of the brain. This decay can not be seen in the intensity profile of the image processed with our correction procedure.

In summary, we implemented a new method for the correction of intensity inhomogeneity artifacts based on the AFCM algorithm [105, 155]. The classical AFCM algorithm was extended to incorporate a-priori probability maps to aid the clustering. The method was tested on a large number of data sets. For all data sets a clear reduction of the intensity inhomogeneity was observed, and for most data sets the inhomogeneity in the resulting corrected image was hardly noticeable. The incorporation of the WM a-priori probabilities was shown to improve the robustness and reliability of the method. We conclude that the presented method is very robust and efficient, and, thus, suitable for correcting the intensity inhomogeneity artifacts in a preprocessing pipeline for MRI data.

6.3 Scalp Surface Reconstruction

In our pipeline a scalp surface is reconstructed from a given T1-weighted MRI and an optional T2-weighted MRI. The reconstructed scalp surface shall be accurate and smooth. To be visually more pleasing the bottom of the reconstructed surface should be cut along a straight plane. Finally, the surface must consist of a single connected component to allow mapping of the scalp potentials.

Pseudo-code for the reconstruction of the scalp surface is presented in Algorithm 2. The reconstruction is done using a multi-resolution active contour segmentation approach as described in Section 3.8. In a first step, an initial solution is created by thresholding the T1-weighted MRI and by subsequently applying morphological operations to the binarized image (Algorithm 3). An opening operation in combination with selection of the biggest connected component is performed to remove small islands of background (BG) voxels from the mask which were initially included due to noise. Holes inside of the mask are then filled by applying a closing operation in combination with explicit filling of holes. The cleaning of the mask is finalized by an opening

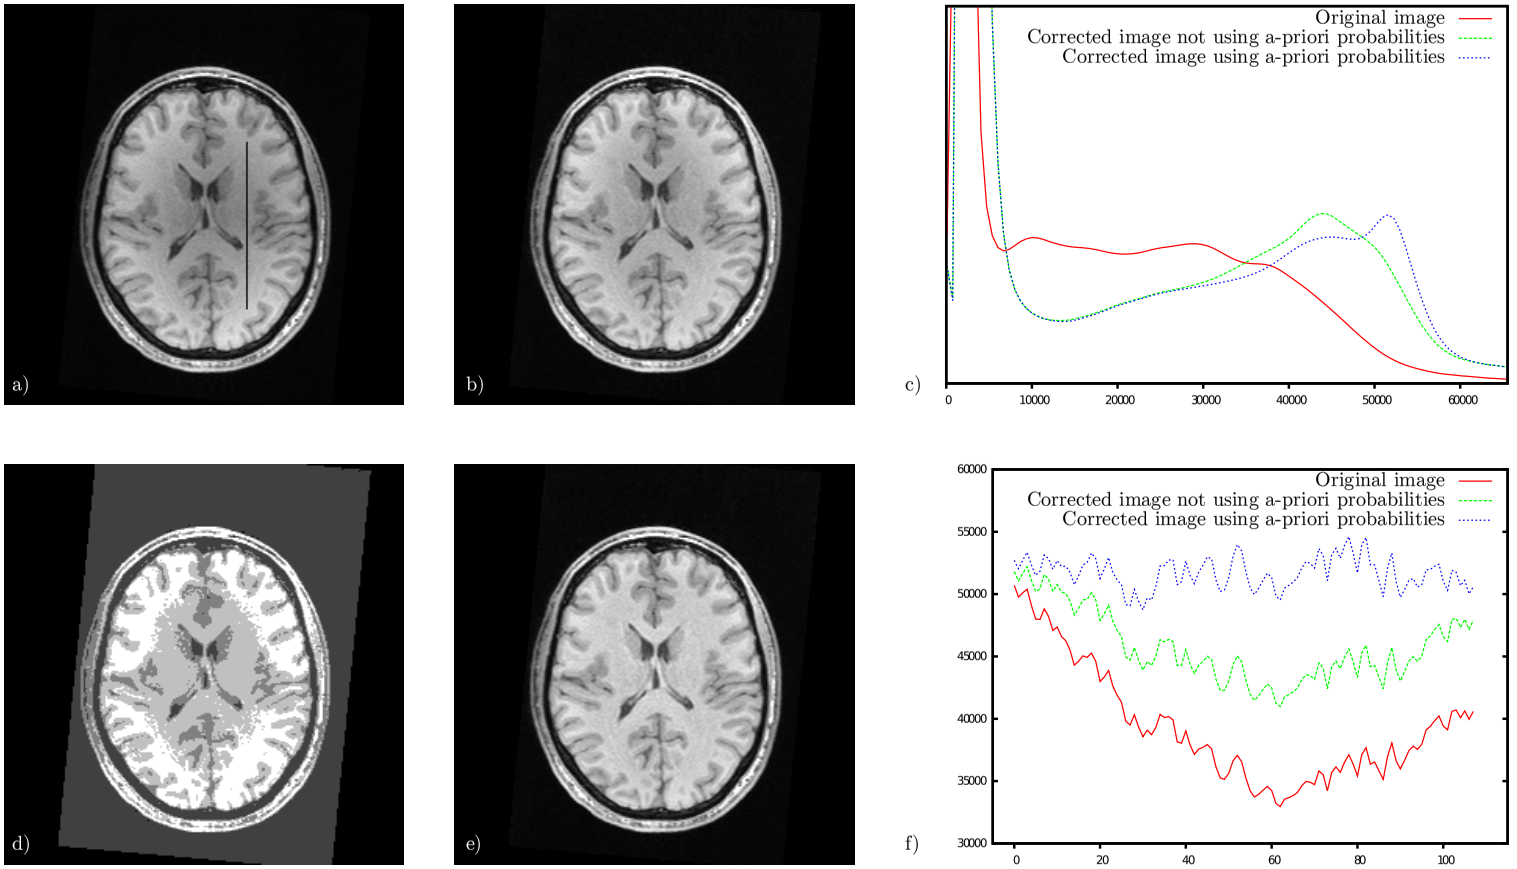

Figure 6.2: Improvement of the inhomogeneity correction by incorporating a-priori tissue probability maps. Results were obtained for $\lambda_1 = 1$, $\lambda_2 = 150$, and 20 iterations. For both approaches, the original AFCM and our modified approach, the same parameters (e.g., solver convergence criteria) were used. a) Original T1-weighted MRI with noticeable shading artifact. The black line indicates the line along which the intensity profiles in f) are plotted; b) Resulting intensity inhomogeneity corrected image without using the a-priori probabilities; c) Intensity histograms of original image, and corrected images with and without using the a-priori probabilities; d) Imperfect hard classification resulting from the intensity inhomogeneity correction without incorporating the a-priori probabilities; e) Resulting intensity inhomogeneity corrected image using the a-priori probabilities f) Intensity profiles plotted along a posterior to anterior line in the WM (see a)).

operation and selection of the biggest connected component. The mask is then down-sampled to a resolution of $4 \times 4 \times 4\,\mathrm{mm}^3$. This is the resolution on which the first iterations of the active contours segmentation will be performed. All voxels below a user-defined plane are finally removed from the mask to obtain a scalp surface which is cut along a straight plane. An example for the coarse initial solution $M^{\text{start, coarse}}$ is presented in Figure 6.3d.

The actual segmentation of the scalp surface is performed using a multi-resolution active contours segmentation approach (Sec. 3.8). A Gaussian distribution is chosen for the likelihood function. As a prior only the smoothness prior from Equation 3.32 is incorporated to achieve a smooth surface reconstruction. At each iteration of the active contours segmentation first the level set functions are updated according to Equation 6.19. The time step Δt for the gradient descent is chosen sufficiently small to guarantee numerical stability.

$$\psi^{(t+1)} = \psi^{(t)} - \Delta t\, \delta_\epsilon \left(\boldsymbol{U}^{\text{likelihood}} + \boldsymbol{U}^{\text{smoothness}} \right) \tag{6.19}$$

Next, the mean μ_c^m and standard deviation σ_c^m of the gray values inside of the scalp and the BG are updated. For the sake of simplicity, it is assumed for multiple image modalities that the intensities in the different modalities are mutually independent. As parameters for the likelihood we, thus, have a mean gray value and standard deviation for each region c and each modality m.

$$\mu_c^{m(t+1)} = \frac{\int_\Omega y^m(\boldsymbol{r})\, \chi_c(\boldsymbol{r})\, \mathrm{d}\boldsymbol{r}}{\int_\Omega \chi_c(\boldsymbol{r})\, \mathrm{d}\boldsymbol{r}} \tag{6.20}$$

$$\sigma_c^{m(t+1)} = \frac{\int_\Omega \left(y^m(\boldsymbol{r}) - \mu_c^{m(t+1)} \right)^2 \chi_c(\boldsymbol{r})\, \mathrm{d}\boldsymbol{r}}{\int_\Omega \chi_c(\boldsymbol{r})\, \mathrm{d}\boldsymbol{r}} \tag{6.21}$$

The active contours segmentation is performed first on a coarser 4 mm and then on a finer 2 mm level. The previously created initial solution is used to initialize the coarse segmentation. In the next step, the segmentation result obtained on the coarser level M^{coarse} is sampled up and used to initialize the segmentation at the finer $2 \times 2 \times 2\,\mathrm{mm}^3$ resolution. The segmentation at the finer level yields a scalp mask M^{fine} and an implicit representation of the scalp surface through the level set function ψ_0, both at 2 mm resolution. Segmentation at the coarse and fine resolution levels is run for a fixed number of iterations n^{coarse}, respectively, n^{fine}. Exemplary segmentation results at 4 mm and 2 mm resolution can be seen in Figures 6.3b and e.

One of the requirements for the reconstructed scalp surface is that it is cut at the bottom along a user-defined plane. To do this the level set function is directly modified and certain constraints are applied to the propagation of the zero level set representing the scalp surface. First, the sign of the level set function is adapted so that all voxels below the plane are assigned to the BG region ($\psi_0 > 0$). The level set values for voxels directly adjacent to the cut plane are set to the negative or positive distance of these voxels to the cut plane. The negative distance is chosen if the voxels are marked in the initial scalp mask. Otherwise the positive distance is chosen. In this way, the scalp surface implicitly represented by the level set function follows the cut plane as closely as possible. The level set values for the voxels directly adjacent and below the plane are then frozen, that is, they are not updated during the iterative segmentation procedure.

Some post-processing steps are finally applied to the segmentation result. The level set repre-

sentation of the segmented scalp has the advantage that it is more accurate than a representation using a binary mask. By means of, for example, the Marching Cubes algorithm [83] it is possible to interpolate the position of the segmented surface with an accuracy which is higher than the resolution of the regular grid on which the level set function is defined. During the post-processing we are, therefore, manipulating the level set function directly to preserve this accuracy while we use morphological operations on a binary mask to decide where we have to manipulate the level set function. During post-processing first holes in the binary scalp mask are filled and the biggest connected component is selected. The sign of the level set function is then modified to reflect the changes in the binary mask.

A triangular parameterization of the scalp surface is extracted from the level set representation utilizing the widely-used Marching Cubes algorithm. For an exemplary scalp surface reconstructed in this way see Figure 6.3c.

Although we are ensuring a certain smoothness of the segmented surface through the incorporated priors in some cases staircase-like artifacts are still observed for the reconstructed scalp surfaces. To reduce these artifacts an additional mesh based smoothing procedure [138] is performed on the surface. The effect of the additional smoothing gets obvious when comparing Figures 6.3c and f.

The volume conductor segmentation described in detail below also needs the reconstructed scalp surface as a reference surface. For this purpose, the level set function ψ_0 at 2 mm resolution which is implicitly representing the scalp surface is again up-sampled and then binarized to obtain a 1 mm resolution scalp mask. This scalp mask is used as input for the volume conductor segmentation.

The scalp surface reconstruction procedure described in this section was applied to a large number of T1-weighted MRIs with varying image characteristics. The results were inspected visually. For all inspected data sets an accurate and visually pleasing scalp surface could be reconstructed. The presented procedure depends on two parameters which can be adjusted by the user. These are the gray value threshold for the generation of the initial solution and the number of iterations for the segmentation on the coarse resolution level. For the inspected data sets the performance of the procedure was robust with respect to these parameters. Default parameters could be chosen so that accurate results were obtained for nearly all data sets. For few MRIs only parameter values had to be adapted. The run-time of the complete scalp surface reconstruction, including post-processing and frequent reconstruction of triangular meshes to visualize intermediate results, is approximately 4 min 20 s for $n^{\text{coarse}} = 300$ and $n^{\text{fine}} = 150$. A large fraction of this time is spent for the active contours segmentation, in particular, on the fine resolution level. We consider this run-time to be absolutely noncritical especially as the scalp surface reconstruction can usually be run unsupervised because user intervention is necessary only in the rarest cases.

6.4 Cortex Surface Reconstruction

The reconstruction of an approximate cortex surface is another task of the segmentation pipeline presented here. The cortex surface is intended to define the allowed space of sources for source reconstruction using certain inverse methods (e.g., minimum norm estimation [66]). Physiolog-

Algorithm 2 Reconstruction of the scalp surface from MRI data.

Require: $\boldsymbol{y}$: image data. T1-weighted MRI (T2-weighted MRI optional).
Require: y': threshold for T1-weighted image.
Require: n^{coarse}, n^{fine}: number of iterations at coarse and fine resolution levels.

$M^{\text{start, coarse}} \leftarrow$ CreateInitialSolution$(\boldsymbol{y}, y')$
InitializeSegmentation$(M^{\text{start, coarse}})$ ▷ Initialize from start mask
$M^{\text{coarse}} \leftarrow$ RunSegmentation$(\boldsymbol{y}, n^{\text{coarse}})$
$M^{\text{start, fine}} \leftarrow$ SampleUp(M^{coarse}) ▷ Sample up results to finer resolution
InitializeSegmentation$(M^{\text{start, fine}})$ ▷ Initialize from segmentation result at coarser resolution
$M^{\text{fine}} \leftarrow$ RunSegmentation$(\boldsymbol{y}, n^{\text{fine}})$
$M^{\text{fine}} \leftarrow$ PostProcessing(M^{fine})
Reconstruct scalp surface from ψ_0 using Marching Cubes algorithm [83]
Smooth scalp surface [138]

function InitializeSegmentation(M^{init})
    Initialize ψ_0 from M^{init}
    Set level set values at voxels adjacent to cut plane equal to distance to plane (see text)
    Freeze level set values for voxels adjacent to and below of cut plane
end function

function RunSegmentation$(\boldsymbol{y}, n)$
    $i \leftarrow 0$ ▷ Reset iteration counter
    while $i < n$ **do**
        Update ψ_0 according to Eq. 6.19
        Update μ_c^m according to Eq. 6.20
        Update σ_c^m according to Eq. 6.21
        $i \leftarrow i + 1$ ▷ Increase iteration counter
    end while
    Generate segmentation mask M^{seg}
    return M^{seg}
end function

function PostProcessing(M^{seg}, ψ_0)
    $M^{\text{seg,mod}} \leftarrow \text{fill}(M^{\text{seg}})$ ▷ Fill holes in binary mask
    $M^{\text{seg,mod}} \leftarrow \text{big}(M^{\text{seg,mod}})$ ▷ Select biggest connected component
    for $i \in \{i \mid M_i^{\text{seg}} \neq M_i^{\text{seg,mod}}\}$ **do**
        $\psi_0(\boldsymbol{r}_i) \leftarrow -\psi_0(\boldsymbol{r}_i)$
    end for
    return $M^{\text{seg,mod}}, \psi_0$
end function

Algorithm 3 Generating an initial solution for the scalp surface reconstruction.

1: **function** CREATEINITIALSOLUTION($\boldsymbol{y}$)
2: M^{start} = threshold($\boldsymbol{y}$) ▷ Binarize image
3: M^{start} = dil(big(ero(M^{start}))) ▷ Opening with additional biggest connected component labeling
4: M^{start} = ero(fill(dil(M^{start}))) ▷ Closing with additional filling of holes
5: M^{start} = open(M^{start}) ▷ Final opening ...
6: M^{start} = big(M^{start}) ▷ ... and biggest connected component labeling
7: Downsample M^{start} to 4 mm resolution
8: Cut M^{start} below user-defined plane
9: **return** M^{start}
10: **end function**

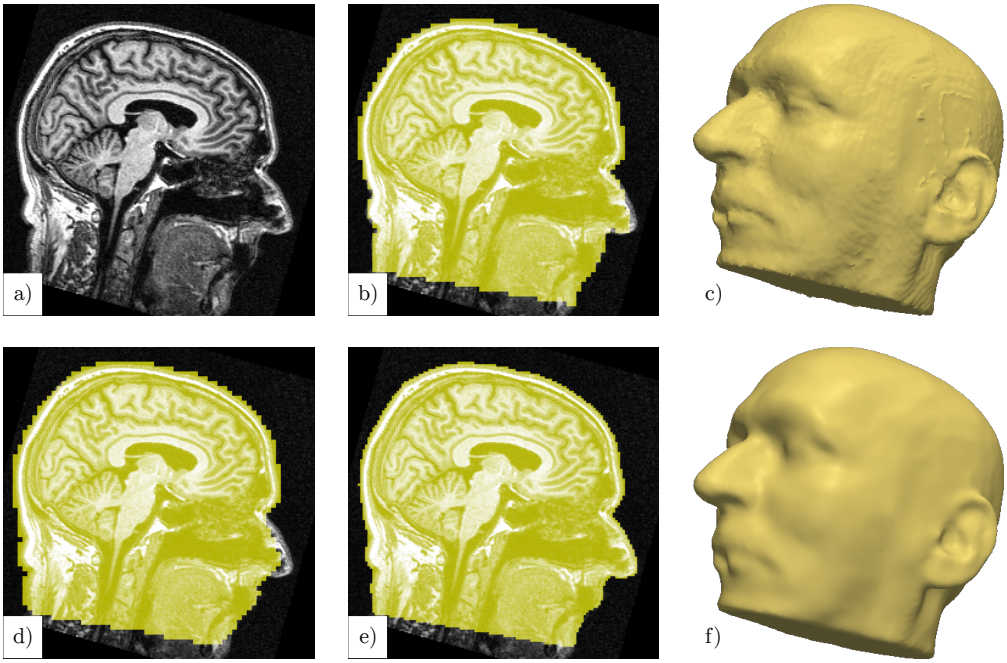

Figure 6.3: Intermediate and final results of the scalp surface reconstruction described in Alg. 2. (a) T1-weighted image; (b) coarse segmentation result M^{coarse} (Line 3); (c) rendering of reconstructed scalp surface (Line 8); (d) coarse initial solution $M^{\text{start, coarse}}$ (Line 1); (e) fine segmentation result M^{fine} (Line 6); (f) rendering of reconstructed scalp surface with additional surface based smoothing (Line 9).

ically sources are expected in the cortex layer of the brain. The cortex surface shall, thus, be reconstructed somewhere between the WM and the GM surface. Furthermore, the sources of interest in many applications are located in the neocortex. The reconstructed surface therefore must not include the cerebellum and other sub-cortical structures.

Algorithm 4 Overview of the cortex surface reconstruction procedure.

Require: $\boldsymbol{y}$: T1-weighted image
Require: T^{clean}: template of subcortical structures, the cerebellum and the ventricles
Require: y': Threshold for brain mask creation

$M^{\text{brain}} \leftarrow \textsc{CreateBrainMask}(\boldsymbol{y}, y')$
$(P_{\text{WM}}, M^{\text{WM}}) \leftarrow \textsc{DoClassification}(\boldsymbol{y}, M^{\text{brain}})$ ▷ Perform MRF-MAP segmentation of brain tissues
$M^{\text{WM}} \leftarrow \textsc{SmoothAndInflateWMSurface}(M^{\text{WM}}, P_{\text{WM}})$ ▷ Perform smoothing and inflation of WM surface

The implemented cortex surface reconstruction is sketched in Algorithm 4. Our procedure essentially consists of the following steps. First, morphological operations in combination with a pre-defined brain template are used to isolate the brain in the MRI from surrounding bone and soft tissues. In the next step, a voxel-based MRF classification as described in Section 3.9 and in some more detail below is performed to segment the image with the isolated brain into the four tissue regions WM, GM, CSF and BG. The classification approach employed here is similar to that presented by Zhang et al. [163]. As output of the classification tissue probabilities $P_c(\boldsymbol{r})$ at each voxel are obtained. Using pre-defined templates the probability maps are modified so that the cerebellum and subcortical structures are excluded from the WM. From these probabilities the WM surface can be extracted, which is then optimized in a level set framework with respect to the tissue probabilities (Eq. 3.39) and a smoothness prior (Eq. 3.32). In this way, a smoothed WM surface is extracted. This surface is in a next step inflated by computing the propagation of the zero level set under the influence of only the smoothness and the inflation prior (Eq. 3.37). This, finally, yields a smooth cortex surface approximately situated between the GM and WM surfaces.

The procedure of removing non-brain tissue before the actual segmentation starts is detailed in Algorithm 5. The brain mask is created by first thresholding the T1-weighted image using the threshold value y'. During the scalp segmentation the averaged gray values inside and outside of the scalp surface were computed. The threshold value y' is derived from these averaged gray values to account for possible variations in the gray value characteristics of the MRI. A standard brain template is then transformed to the individual AC-PC coordinate system and the thresholded MRI is masked with the transformed brain template. The template was constructed by aligning T1-weighted MRIs of 66 different subjects in Talairach space. For each of the MRIs the brain and the subcortical structures were segmented using the FreeSurfer software [48]. The brain template was finally constructed as the union of the segmented brains of all 66 subjects. In this way, we can safely cut away most of the non-brain tissue in the individual MRI by masking it with the brain template. Finally an opening operation combined with selection of the biggest connected component and a closing operation in combination with explicit filling of holes are

applied to the brain mask. An example of a brain masked T1-weighted image generated in this way is shown in Figure 6.4a.

Algorithm 5 Creating a brain mask M^{brain} to remove non-brain tissue.

Require: $\boldsymbol{y}$: input T1-weighted image.
Require: y': threshold value for initial thresholding.
Require: T^{brain}: pre-computed brain template.

1: **function** CREATEBRAINMASK($\boldsymbol{y}, y'$)
2: $M^{\text{brain}} \leftarrow \text{threshold}(\boldsymbol{y}, y')$ ▷ Binarize image
3: $M^{\text{brain}} \leftarrow M^{\text{brain}} \cap T^{\text{brain}}$ ▷ Apply standard brain template
4: $M^{\text{brain}} \leftarrow \text{dil}(\text{big}(\text{ero}(M^{\text{brain}})))$ ▷ Opening with additional selection of biggest connected component
5: $M^{\text{brain}} \leftarrow \text{ero}(\text{fill}(\text{dil}(M^{\text{brain}})))$ ▷ Closing with additional filling of holes
6: **end function**

An MRF classification as described in Section 3.9 is then performed on the brain masked MRI. For the likelihood function a Gaussian distribution inside of each tissue region with mean μ_c and standard deviation σ_c is assumed. An isotropic MRF model is specified to account for the layered structure of the tissues WM, GM and CSF. Pairwise clique potentials are chosen so that only transitions from WM to GM, from GM to CSF, and from CSF to BG are allowed but direct transitions, for example, from WM to CSF are forbidden. The MRF classification finally yields a WM tissue probability map from which a binary WM mask (Figure 6.4b) can be derived.

In the next step, the binary WM mask is cleaned up to remove those regions of the WM surface that are not part of the neocortical surface. Subcortical structures and the cerebellum are removed and the ventricles are filled up. The associated WM probability map is changed accordingly by setting the probability to 0, respectively, 1 for all voxels that were removed from or added to the WM mask. The removal of the subcortical structures and the cerebellum, and the filling of the ventricles is done using a standard template. The template is transformed to the individual AC-PC coordinate system. If a voxel is marked in the transformed template as belonging to a subcortical structure or the cerebellum then it is removed from the WM mask. If a voxel is marked as belonging to the ventricles then it is added to the WM mask. Finally, any remaining holes in the WM mask are filled. Figure 6.4c presents the WM mask from Figure 6.4b after cleaning up.

The standard template was generated in a very similar way as the brain template used during the brain extraction. The subcortical structures, the cerebellum, and the ventricles were segmented from MRIs of 66 different subjects aligned in Talairach space. Segmentation was performed using the FreeSurfer software [48]. Histograms were then computed by counting in how many subjects a voxel in Talairach space belongs to subcortical structures, the cerebellum, or the brain stem. By thresholding the histogram maps binary masks were generated for each of the structures. The binary masks were finally combined into one template image.

The cleaned WM probability map and the corresponding binary WM mask are used to initialize the level set based active contour approach. The active contour approach is described in

Algorithm 6 Smoothing and inflation of the initial WM surface using level set based surface evolution.

Require: : M^{WM}: binary WM mask as obtained by previous classification
Require: : $P_{\mathrm{WM}}(\boldsymbol{x})$: WM probability map as obtained by previous classification
Require: : n^{smooth}: number of iterations for smoothing
Require: : n^{inflate}: number of iterations for inflation

function SmoothAndInflateWMSurface($M^{\mathrm{WM}}, P_{\mathrm{WM}}(\boldsymbol{x})$)
  Initialize $\psi_0^{(0)}$ from M^{WM}
  $i \leftarrow 0$ ▷ Reset iteration counter
  while $i < n^{\mathrm{smooth}}$ **do** ▷ Smooth WM surface
    Update ψ_0 according to Eq. 6.22
  end while
  while $i < n^{\mathrm{inflate}}$ **do** ▷ Inflate WM surface
    Update ψ_0 according to Eq. 6.23
  end while
  Reconstruct cortex surface from ψ_0 using Marching Cubes algorithm [83]
end function

detail in Section 3.8. Algorithm 6 gives an overview of the procedure. For smoothing and inflating the WM surface we ignore the likelihood function and evolve the zero level set taking into account only one or multiple priors. In a first phase, the WM probability map is considered as a prior. In addition, the smoothness prior from Equation 3.32 is incorporated. With these priors the zero level set evolves to a smoothed version of the WM surface. In a second phase, only the smoothness prior and the inflation prior (Eq. 3.37) are taken into account. Update equations for the smoothing and the inflation can be derived from Equation 3.42 by only taking the relevant terms into account.

$$\psi_i^{(t+1)} = \psi_i^{(t)} - \Delta t\, \delta_\epsilon \left(\boldsymbol{U}_i^{\text{probability map}} + \boldsymbol{U}^{\text{smoothness}} \right) \qquad \text{Smoothing} \tag{6.22}$$

$$\psi_i^{(t+1)} = \psi_i^{(t)} - \Delta t\, \delta_\epsilon \left(\boldsymbol{U}_i^{\text{inflation}} + \boldsymbol{U}^{\text{smoothness}} \right) \qquad \text{Inflation} \tag{6.23}$$

Level set evolution is computed for a fixed number of iterations. During these iterations the WM surface moves outwards, that is, it is inflated, while remaining smooth. The number of iterations was empirically determined so that we end up with a smooth cortex surface approximately lying between the WM and GM surface. The level set computations are carried out on a $1 \times 1 \times 1$ mm grid to ensure a good spatial resolution of the resulting cortical surface. Renderings of an original, a carefully smoothed, and a final, smoothed and inflated cortex surface can be seen in Figures 6.4d, e and f, respectively.

The presented cortex segmentation procedure was tested on T1-weighted MRIs from a large number of different subjects and the resulting cortex surfaces were inspected visually. For all subjects a visually pleasing smooth cortical surface could be reconstructed. Removal of subcortical structures and filling of the ventricles worked well for nearly all images. The accuracy of the segmentation was found to be good with only few exceptions where some thin gyri or

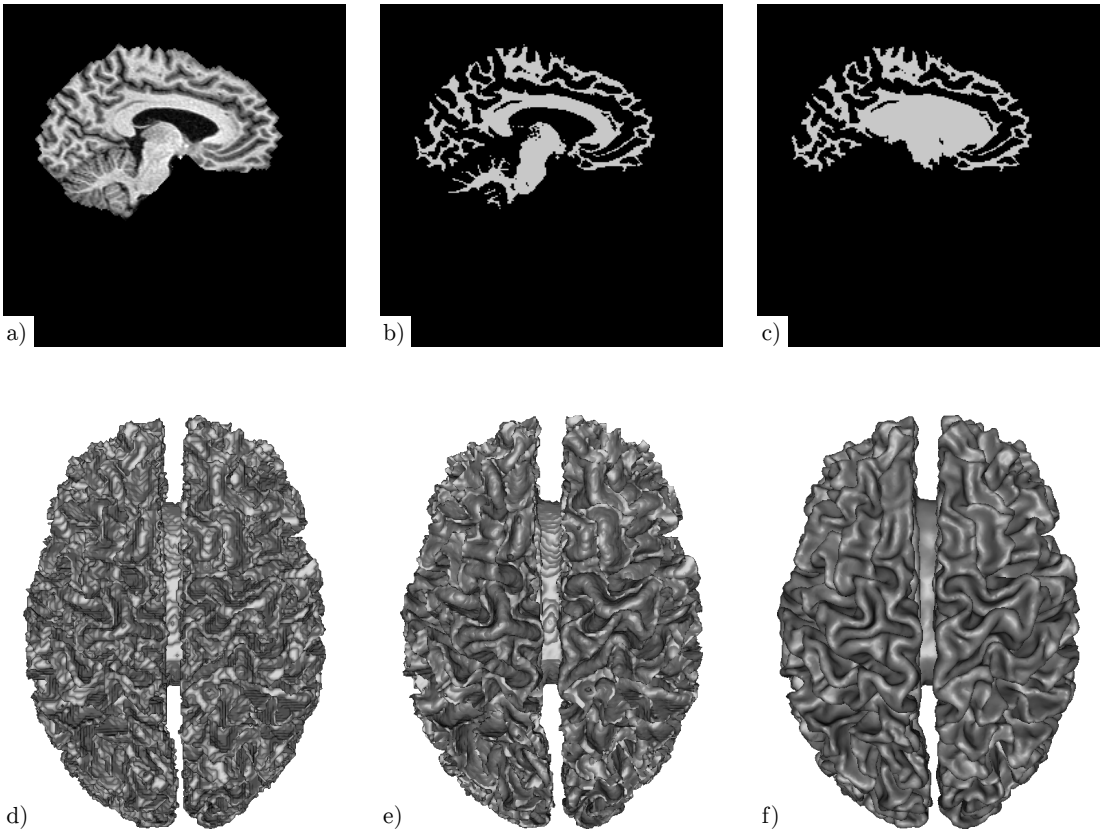

Figure 6.4: Intermediate results from the cortex surface reconstruction. (a) T1-weighted image masked with brain mask; (b) first WM classification M^{WM}; (c) WM mask with cerebellum and subcortical structures removed and filled ventricles; (d) initial reconstructed WM surface; (e) WM surface after first smoothing; (f) final inflated and smoothed cortex surface.

narrow sulci were not included in the final cortex surface. Some parameters might influence the results of the cortex reconstruction. For the creation of the brain mask, for example, these are the threshold value or the parameters for the closing and opening operations. As a further example, the amount of inflation is specified by the number of iterations n^{inflate} during which the WM surface is propagated under the influence of the inflation prior. Our tests on a large number of different MRIs, however, showed that default parameters can be chosen so that the segmentation results are acceptable for all images with only very few exceptions. User intervention to adjust the parameters is, thus, not necessary and the procedure can run unsupervised.

In our approach we combined a voxel-based MRF segmentation approach similar to the one presented by Zhang et al. [163] with level set based surface evolution. The level set based surface evolution was added to the procedure to be able to smooth and inflate the WM surface to obtain a suitable cortex surface. Priors on the geometry of the reconstructed surface (e.g., smoothness or inflation priors) can naturally be incorporated into level set based approaches as it was shown in Section 3.8. Representing these priors in voxel-based approaches is hardly possible.

For some applications it is desirable that the cortex surface has a spherical topology. At the moment, the above procedure cannot guarantee that the final surface has the correct topology. It should, however, be easily possible to extend the procedure to correct and preserve the topology. Bazin and Pham's approach [9] can be used to correct the topology of the WM mask obtained from the classification. During the inflation the level set evolution can be constrained to preserve the now correct topology of the WM mask [58].

The MRF classification and the level set computations on a fine 1 mm grid are computationally demanding. A typical cortex surface reconstruction, thus, lasts approximately 13 min. Nevertheless, we regard these run-times as sufficiently fast because the algorithm is intended to be run without user supervision and intervention.

7 A Multiphase Active Contours Segmentation Approach for the Generation of Volume Conductor Models

7.1 Introduction

The active contours based segmentation introduced in Section 3.8 will be the basis for the volume conductor segmentation approach presented in this chapter. There are several reasons why we think it is suitable for segmenting the tissues relevant for EEG source analysis from MRI data.

One reason is that it is able to simultaneously segment multiple regions. It is, thus, in principle able to at once segment the MRI into the tissues scalp, skull, CSF and brain that should also be distinguished in volume conductor models of the human head for EEG source analysis (Sec. 2.4). This avoids problems like overlap or vacuum which might occur when composing the final segmentation from multiple seperately segmented masks.

A further advantage of the active contours segmentation approach is its ability to incorporate the geometric smoothness prior. Smoothness is a desirable characteristic for the solutions of our segmentation problem. The inner and outer skull boundaries in the region of the skull cap, as an example, are rather smooth and have a low curvature.

Some recent work [5] using the multiphase active contours segmentation proposed by Chan and Vese [143] for the delineation of brain tissues from MRI data reassures us that their segmentation approach is a good basis for our method for volume conductor segmentation.

7.2 The Segmentation Algorithm

In our approach we will use Chan and Vese's multiphase level set framework [143] to segment the head into a fixed number of regions. Experiments will be carried out at the end of this chapter to investigate the ideal number of regions to be differentiated during the segmentation.

The presented approach will incorporate the smoothness prior from Equation 3.32. The smoothness prior makes the segmentation robust against noise in the MRI data. In addition, an a-priori probability distribution, for example, derived from a probabilistic atlas can be included using the prior from Equation 3.39.

Taking the above mentioned priors into account the a-posteriori probability is optimized with respect to the level set functions according to the following update equation:

$$\psi_i^{(t+1)} = \psi_i^{(t)} - \Delta t\, \delta_\epsilon \left(\boldsymbol{U}_i^{\text{likelihood}} + \boldsymbol{U}_i^{\text{probability map}} + \boldsymbol{U}_i^{\text{smoothness}} \right),\ i = 1, \ldots, n \qquad (7.1)$$

The update matrix is composed from the terms associated with the likelihood $\boldsymbol{U}_i^{\text{likelihood}}$ (Eq. 3.43), the a-priori probability maps $\boldsymbol{U}_i^{\text{probability map}}$ (Eq. 3.44), and the smoothness prior $\boldsymbol{U}_i^{\text{smoothness}}$

(Eq. 3.46). Δt is the time step of the gradient descent method. It is determined in dependence of the update matrices in such a way that the values of ψ_i at the grid nodes do not change by more than the constant parameter $\eta \leq \min(\Delta x, \Delta y, \Delta z)$ where $\min(\Delta x, \Delta y, \Delta z)$ is the minimum voxel size in x-, y-, and z-direction. By determining the time step in this way the Courant-Friedrichs-Lewy (CFL) condition (see, e.g. [99]) is fulfilled which is a necessary condition for the convergence of finite difference methods.

We assume that the MRI intensities for the distinguished tissues are homogeneous corrupted by Gaussian noise. The gray value distribution in region c could, thus, be described using a Gaussian distribution with means $\boldsymbol{\mu}_c$ and a covariance matrix $\boldsymbol{\Sigma}_c$. In the case of multiple image modalities for the sake of simplicity and a more efficient estimation of the gray value parameters we assume the intensities in the different image modalities to be independent of each other. In this way the multivariate Gaussian distribution can be factorized into multiple univariate distributions (Eq. 3.19) with means μ_c^m and standard deviations σ_c^m where m is the current image modality. The gray value parameters $\boldsymbol{\mu}_c^{(t+1)}$ and $\boldsymbol{\sigma}_c^{(t+1)}$ at iteration $(t+1)$ are determined using the current estimate for the level set functions according to the following equations.

$$\mu_c^{m(t+1)} = \frac{\int_\Omega y^m(\boldsymbol{r}) \chi_c(\boldsymbol{r}) \, \mathrm{d}\boldsymbol{r}}{\int_\Omega \chi_c(\boldsymbol{r}) \, \mathrm{d}\boldsymbol{r}} \tag{7.2}$$

$$\sigma_c^{m(t+1)} = \frac{\int_\Omega \left(y^m(\boldsymbol{r}) - \mu_c^{m(t+1)}\right)^2 \chi_c(\boldsymbol{r}) \, \mathrm{d}\boldsymbol{r}}{\int_\Omega \chi_c(\boldsymbol{r}) \, \mathrm{d}\boldsymbol{r}} \tag{7.3}$$

To improve the convergence of the segmentation and to make it more robust against non-optimal initial solutions segmentation is performed on multiple resolution levels. All input data are sub-sampled to a coarser resolution. On the coarsest level then a fixed number of iterations of the active contours segmentation approach is performed. This yields a segmentation represented by the level set functions ψ_i at the resolution of the coarsest level. The level set functions are then up-sampled to the next higher resolution level using nearest-neighbour or trilinear interpolation. Further iterations of the active region segmentation are performed on the current resolution level starting from the up-sampled segmentation obtained at the coarser level. After a fixed number of iterations the resulting level set functions are again up-sampled. This procedure is repeated until we have performed a sufficient number of iterations on the finest level which then yields the final segmentation result.

For all our level set computations we are employing the narrow band approach [3]. This means, we are initializing the level set functions only for voxels up to a certain distance to the zero level set. All computations are then also only carried out inside of this narrow band around the tissue interfaces. The narrow band is reinitialized around the current region boundaries whenever the zero level set comes too close to the edge of the narrow band.

The complete proposed approach runs as follows (Alg. 7): An initial solution is generated from a scalp and a WM reference surface by applying morphological operations and combining the resulting masks (Sec. 8.5.1). Next, the initial solution, the input MRIs and a-priori probability maps (if available) are sub-sampled once or multiple times to generate input data at the coarser resolution levels.

Using n level set functions a maximum of $m = 2^n$ regions can be encoded. Here, the

number of regions m is given and the number of required level set functions is determined as $n = \text{ceil}(\log_2(m))$. The level set functions ψ_i, $i = 1, \ldots, n$ are then initialized based on the initial solution to signed distance functions using the fast marching method (FMM) [128]. More details on the encoding of the regions using the n level set functions are given below (Sec. 7.3).

Next, the initial gray value parameters $\boldsymbol{\mu}_c^{(0)}$ and $\boldsymbol{\sigma}_c^{(0)}$ are estimated from the sub-sampled initial solution and the image data according to Equations 7.2 and 7.3.

At each resolution level l starting at the coarsest one now a fixed number of iterations $t_l^{\max}$ of the active contours segmentation is performed.

In each of these iterations, first, the update matrix for each level set function is computed. The time step for the gradient descent is then chosen to fulfill the CFL criterion to guarantee convergence of the gradient descent method. Updated level set functions $\psi_i^{(t+1)}$ are obtained by adding the update matrices multiplied with δ_ϵ and Δt to the previous estimate $\psi_i^{(t)}$.

Next, new estimates $\boldsymbol{\mu}_c^{(t+1)}$ and $\boldsymbol{\sigma}_c^{(t+1)}$ for the gray value parameters are determined taking into account the updated level set functions $\psi_i^{(t+1)}$ in the form of the characteristic functions χ_c.

After having finished $t_l^{\max}$ iterations and in the case that we do not yet have reached the finest resolution level the level set functions ψ_i are up-sampled using nearest-neighbour or trilinear interpolation to serve as an initial solution at the next finer resolution level.

From the final level set functions obtained at the finest resolution level the segmented image can be reconstructed by assigning to each voxel i the region $c \in 1, \ldots, m$ for which $\chi_c(\boldsymbol{r}_i) \neq 0$. Surface representations of the region boundaries can be interpolated from the level set functions using, for example, the Marching Cubes algorithm [83].

7.3 Implementation Details

An important aspect of the multiphase active contours segmentation approach is how exactly we define the characteristic functions χ_c. Each region c is encoded by a certain combination of the signs of the level set functions expressed by the unique constant vector $\boldsymbol{H}_c$ (Eq.3.30). In early experiments we observed that the boundary between two regions does not converge well towards the true region boundary if the characteristic vectors for the two regions differ in more than one element. This means, that for the region boundary to propagate across a voxel the sign of two or more level set functions has to change at the voxel in the same time step. This hinders the propagation of the region boundary and, thus, also a good convergence.

To enable a good convergence we, thus, try to encode the regions that shall be segmented in such a way that the characteristic vectors $\boldsymbol{H}_c$ only change in one element for neighbouring regions. A possible way to encode the regions in the four-, and five-tissue segmentation are shown in Figure 7.1. For the segmentation into five head tissue regions it was not possible to encode the regions in such a way that the characteristic vectors of neighbouring tissues only differed in one element. For a voxel to pass from brain to CSF the sign of two level set functions has to change.

Algorithm 7 Active contours based segmentation approach.

Require: $\boldsymbol{y}$: (Multi-modal) image data
Require: $M^{\text{ref. scalp}}$, M^{WM}: Scalp and WM reference masks
Require: r: Number of resolution levels
Require: m: Number of regions

Generate initial labeled image $x^{(0)}$ (Sec. 8.5.1)
Sub-sample input data $(r-1)$ times
$n \leftarrow \text{ceil}(\log_2(m))$ ▷ Determine number of level set functions
Initialize $\psi_i,\ i = 1, \ldots, n$ at coarsest level using FMM [128]
Estimate $\boldsymbol{\mu}_c^{(0)}$ and $\boldsymbol{\sigma}_c^{(0)}$ (Eqs. 7.2 and 7.3)
for $l \leftarrow 0$ to $(r-1)$ **do** ▷ Loop across resolution levels
    $t \leftarrow 0$ ▷ Reset iteration counter
    repeat
        for $i \leftarrow 0$ to $(n-1)$ **do** ▷ Loop across level set functions
            Compute $\boldsymbol{U}_i = \boldsymbol{U}_i^{\text{likelihood}} + \boldsymbol{U}_i^{\text{smoothness}} + \boldsymbol{U}_i^{\text{probability map}}$
        end for
        $\Delta t_i = \eta^{-1} \max |\delta_\epsilon U_i(\boldsymbol{r})|, i = 0, \ldots, n-1$ ▷ Determine time step for gradient descent
        $\Delta t = \min(\Delta t_i),\ i = 0, \ldots, n-1$
        for $i \leftarrow 0$ to $(n-1)$ **do**
            $\psi_i^{(t+1)} \leftarrow \psi_i^{(t)} + \Delta t\ \delta_\epsilon\, \boldsymbol{U}_i$ ▷ Apply update matrices
        end for
        Estimate $\boldsymbol{\mu}^{(t+1)}$ and $\boldsymbol{\sigma}_c^{(t+1)}$ (Eqs. 7.2 and 7.3)
        $t \leftarrow t+1$ ▷ Increase iteration counter
    until $t \geq t_l^{\max}$
    if $l < (r-1)$ **then** ▷ If finest resolution level is not yet reached
        Up-sample level set functions ψ_i
    end if
end for
$x_i \leftarrow c \mid \chi_c(\boldsymbol{r}_i) \neq 0,\ i \in \mathcal{S}$ ▷ Reconstruct labeled image x

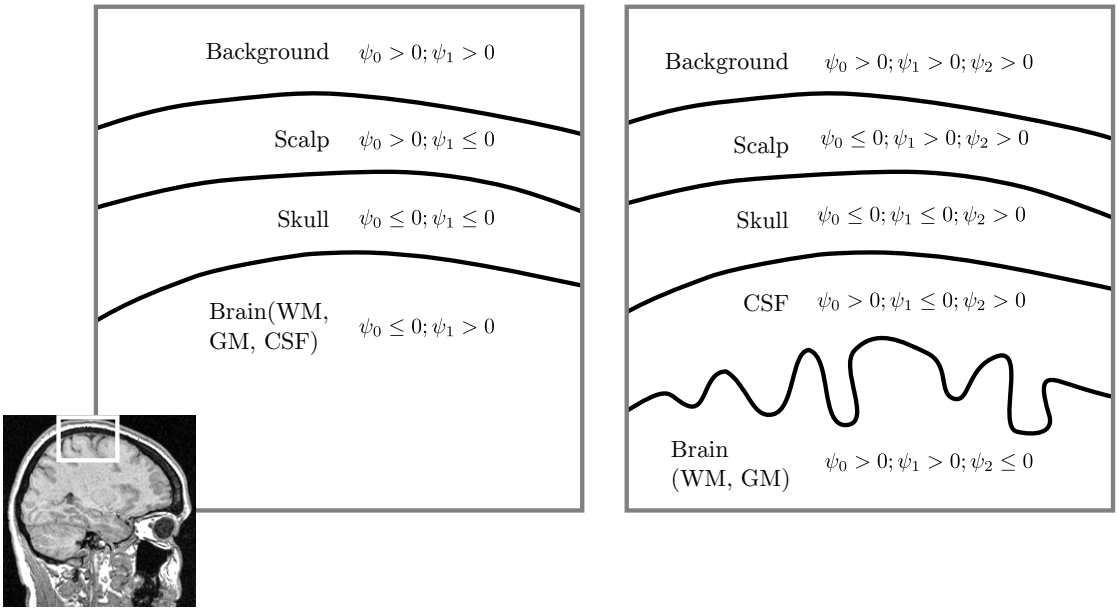

Figure 7.1: Illustration of how the head tissue regions are encoded using the signs of two or more level set functions. The cases of four and five regions segmentation are presented.

7.4 Experiments

Experiments on real data were performed to determine some parameters for the active contours segmentation (the number of regions), and to demonstrate the effectiveness of the method for the segmentation of the skull and other tissues relevant for EEG source analysis.

All experiments in this section were performed with the following parameters unless stated otherwise. The smoothness parameter was chosen as $\epsilon = 1$, the multi-resolution segmentation was performed at 2 mm and 1 mm resolution levels. $t_0^{\max} = 1000$ and $t_1^{\max} = 500$ iterations were performed at the coarse and at the fine level, respectively. For all segmentation experiments except for the first one five different regions, that is, BG, scalp, skull, CSF and brain, are differentiated. The pipeline described in Chapter 6 was used to preprocess the input MR images.

To decide on the optimal number of regions we applied the active contours approach to segment a data set containing a T1- and a T2-weighted MRI into four, five, and eight regions (Figs. 7.3b, c, and d). On first sight, there are problems with the delineation of the inner skull for the four region segmentation. The segmented regions correspond well to the actual tissue distributions in the case of the five region segmentation. The eight region segmentation is able to differentiate GM, WM, and CSF. However, it is apparently not able to distinguish accurately between compact and cancellous bone, and between muscle or scalp tissue and SCT.

A closer look at the skull outlines for the segmentation into different numbers of regions confirms the issues of the four region segmentation with the delineation of the inner skull (Fig. 7.2). Furthermore, it can be noted that the skull segmentation accuracy of the eight region segmentation lacks in some areas of the head. In the area of the interhemispheric fissure, for exam-

ple, the outer skull boundary is reconstructed at the interface between the skull and the dura mater (Fig. 7.2b).

From these observations we conclude that the optimal number of regions for the volume conductor segmentation of the human head using the active contours approach is five. The most important reason for this decision were the observed skull segmentation errors for the four and eight region segmentations. For the five regions segmentation, in contrast, no major inaccuracies in delineating the skull were apparent.

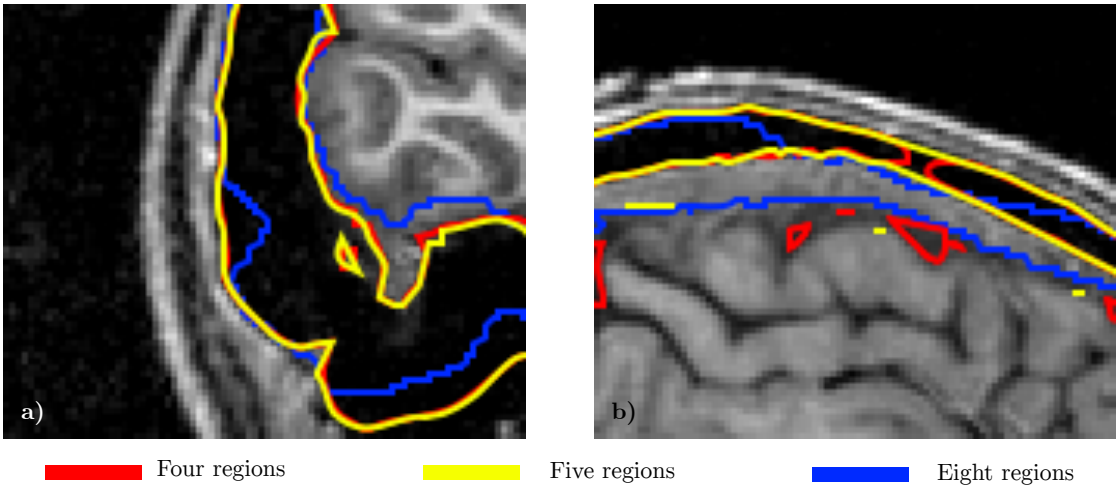

Figure 7.2: Details of the skull outlines for the active contours segmentation into four, five, and eight regions. The location of the enlarged regions is outlined by the white squares in Figure 7.3a.

The skull is of particular importance for volume conductor models to be used in EEG source analysis (Chap. 4). We, therefore, inspected the skull segmentation results obtained by the active contours segmentation approach for a data set containing only a T1-weighted MRI and for a data set containing a T1- as well as a T2-weighted MRI. Segmentation on both data sets yielded a skull segmentation that was found to be very accurate in the region of the skull cap (Figs. 7.4 and 7.5). Only for the most frontal part of the skull a small inaccuracy could be observed where the segmented skull does not propagate into the small, highly convex part of the frontal bone. Otherwise no substantial inaccuracies could be noted. Holes were not present in the skull cap of the segmented skull, and from visual inspection inaccuracies in the determination of the skull thickness are on the scale of one voxel, that is, 1 mm.

In the area of the skull base the segmentation is less accurate. Especially for the T1-only data set the outer skull boundary does not too closely follow the complicated geometry of the skull base (Fig. 7.4). For the segmentation based on a T1- and a T2-weighted image the segmentation of the skull base is more accurate. Still the outer skull boundary does not seem to propagate far enough in regions where the skull base is very thick (Fig. 7.5), for example, in the areas where sinuses are embedded into the skull bones.

Our proposed approach is in principle able to simultaneously segment the head into any number of tissue regions. Here, we are differentiating the five tissue regions which are most relevant

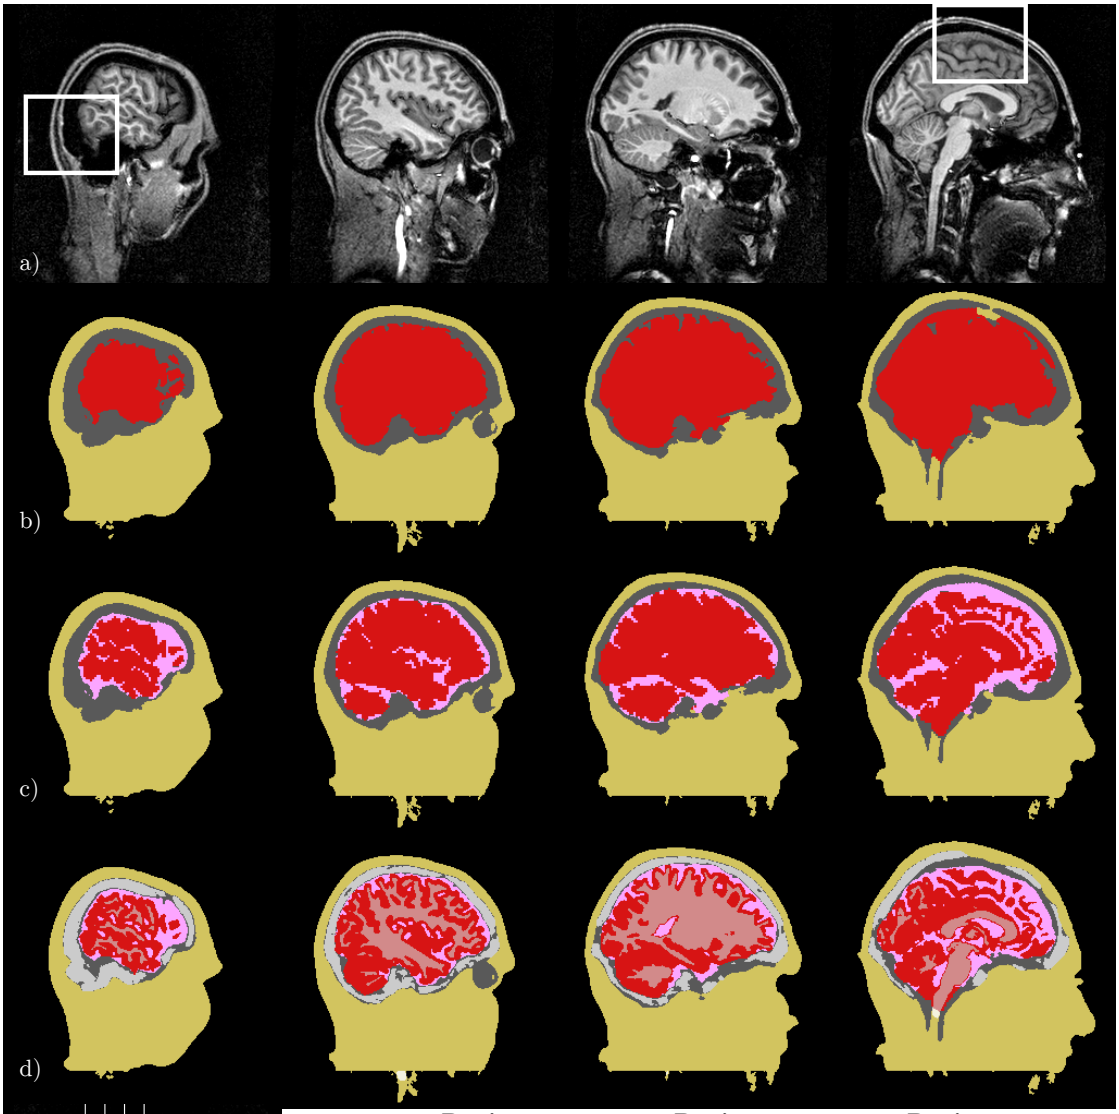

	4 Regions	5 Regions	8 Regions
	BG	BG	BG
	Scalp	Scalp	Muscle/scalp
			Subcutaneous tissue
	Skull	Skull	Compact bone
			Cancellous bone
		CSF	CSF
	Brain	Brain	Gray matter
			White matter

Figure 7.3: Results of the active contours segmentation approach for segmentation into b) four, c) five, and d) eight regions. Segmentation was performed based on a T1- and T2-weighted image. The T1-weighted image is depicted in a). The coronal slice in the bottom rows indicates the position of the sagittal slices.

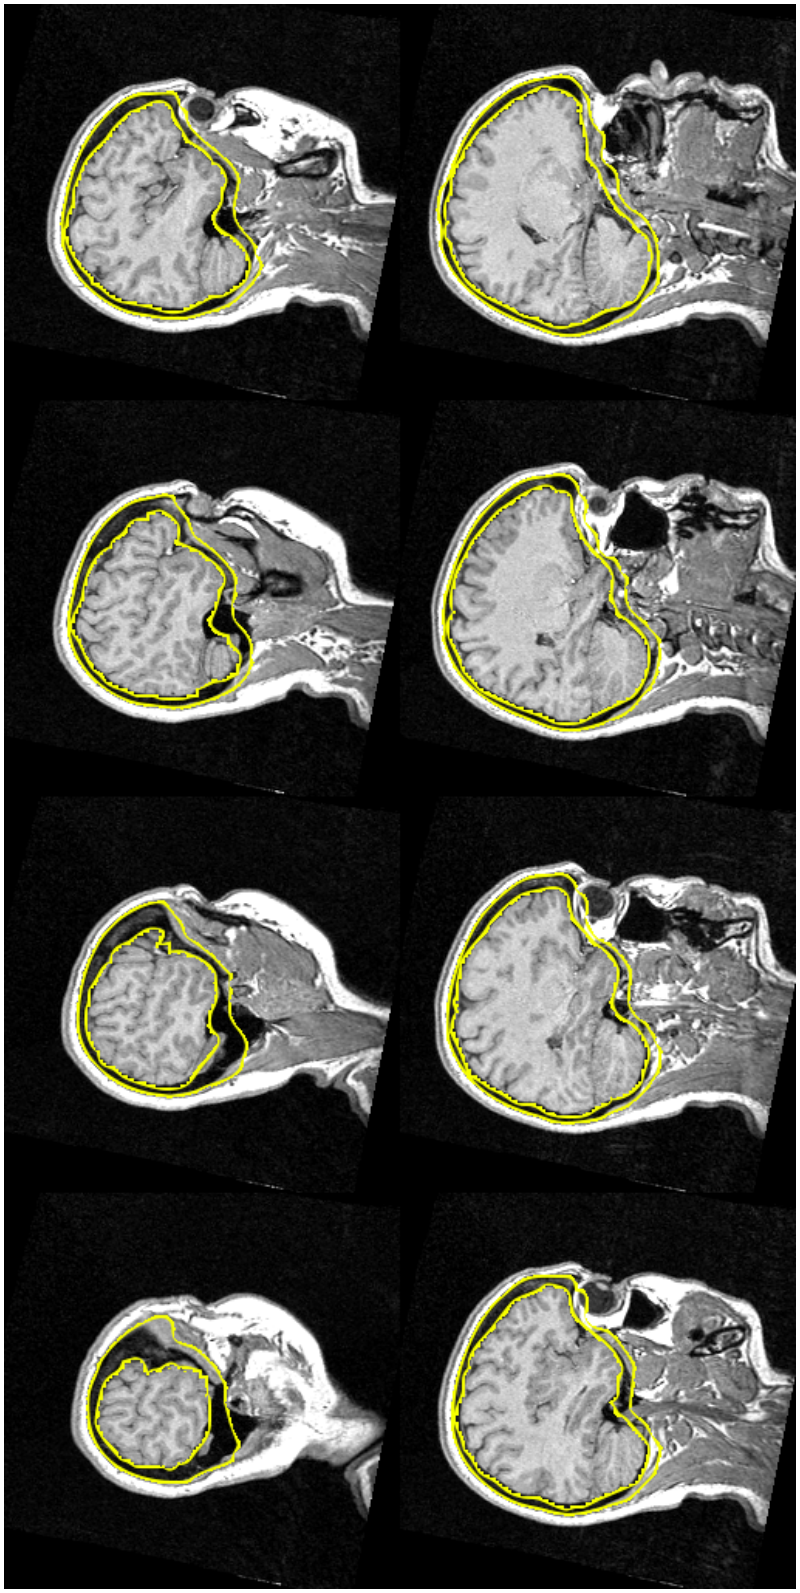

Figure 7.4: Exemplary segmentation results using the active contours approach from T1-weighted MRI data only. Outlines of the segmented skull region in yellow are overlaid onto sagittal slices of the T1-weighted MRI.

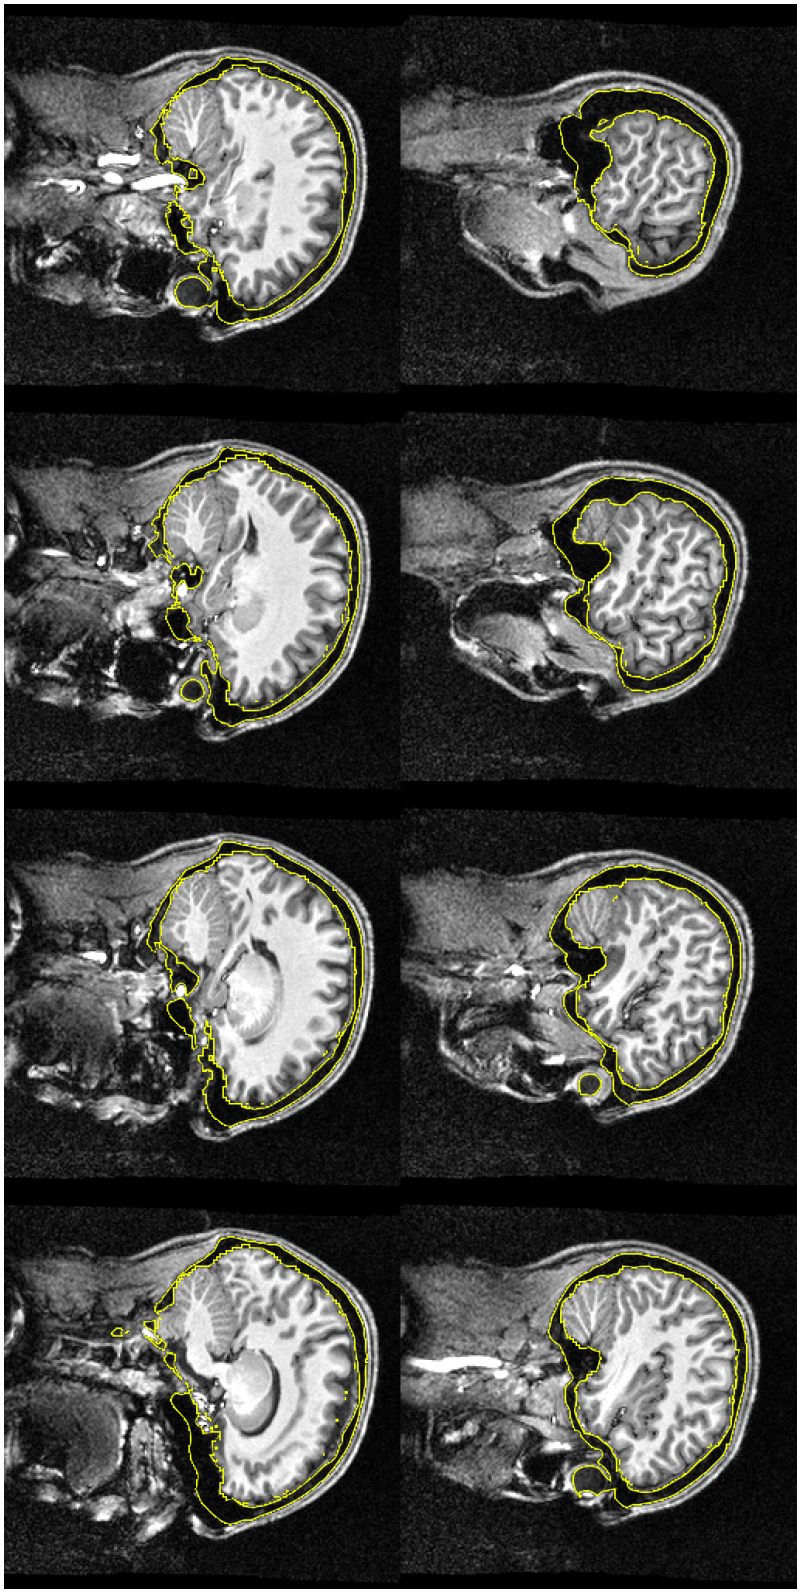

Figure 7.5: Exemplary segmentation results using the active contours approach from T1- and T2-weighted MRI data. Outlines of the segmented skull region in yellow are overlaid onto sagittal slices of the T1-weighted MRI.

for EEG source analysis. The accuracy of the scalp segmentation was found to be good even when only a T1-weighted image was available (Fig. 7.6). The effect of the smoothing prior was, nevertheless, notable. Reconstructed scalp surfaces were smooth and unaffected by noise. On the other hand, small, irregular features of the scalp were smoothed out, for example, the characteristic depression of the scalp surface above the nose.

With respect to the remaining regions it can be observed that the interface between the CSF and brain is very rough and that it does not accurately delineate the brain surface (Figs. 7.6 and 7.7). This is especially evident in the sagittal view of the first data set where two prominent superior CSF filled sulci are segmented as brain. Overall the brain and CSF segmentation accuracy are, thus, substantially worse than the accuracy of the skull and scalp segmentations.

The course of the mean gray values against the number of iterations indicates that the algorithm converges as expected towards a segmentation of the image domain according to our formulation of the segmentation problem (Fig. 7.8). Only the curves for the T1 intensities inside of the brain and the CSF do not show much change. This is also due to the problems with our choice for the encoding of the regions using the signs of the multiple level set functions.

For the segmentation we assumed that the image to be segmented is homogeneous with added Gaussian noise inside of each of the segmented regions. To test if this assumption is valid we examine the distribution of the gray values of the T1-weighted image inside of each region (Fig. 7.9). We observe that the distributions clearly deviate from the Gaussian distribution. This is most obvious for the gray values in the scalp and brain regions. The latter distribution is even multimodal with maxima around 17000 and 27000.

7.5 Discussion and Conclusion

The optimal number of regions was determined to be five. Segmentation into four and eight regions yielded less accurate skull segmentation results. The eight regions segmentation failed to differentiate the tissues compact and cancellous bone, and muscle or scalp and SCT. Thus, the advantage of a potentially more detailed segmentation in the eight regions case cannot compensate for the less accurately reconstructed skull.

The skull segmentation accuracy was by visual inspection found to be very accurate in the five regions segmentation. It might still be advantageous to enhance the presented segmentation approach using additional a-priori information on the anatomy of the human head, for example, in the form of a probabilistic atlas.

Comparing the segmentations for the data set with a T1-weighted image only and the data set with a T1- and a T2-weighted image it was found that incorporating an additional T2-weighted image increases the segmentation accuracy.

With respect to the segmentation of CSF and brain it was observed that the interface between brain and CSF was less accurately delineated. The reason for this is most likely the chosen encoding of the regions by the level set functions as discussed in Section 7.3.

This is also reflected in the plot of the mean gray values across the iterations of the algorithm. For the regions BG, scalp and skull this plot shows the convergence of the gray values which indicates that the algorithm works as expected. For the CSF and brain regions, however, only very small changes in the mean gray values can be seen.

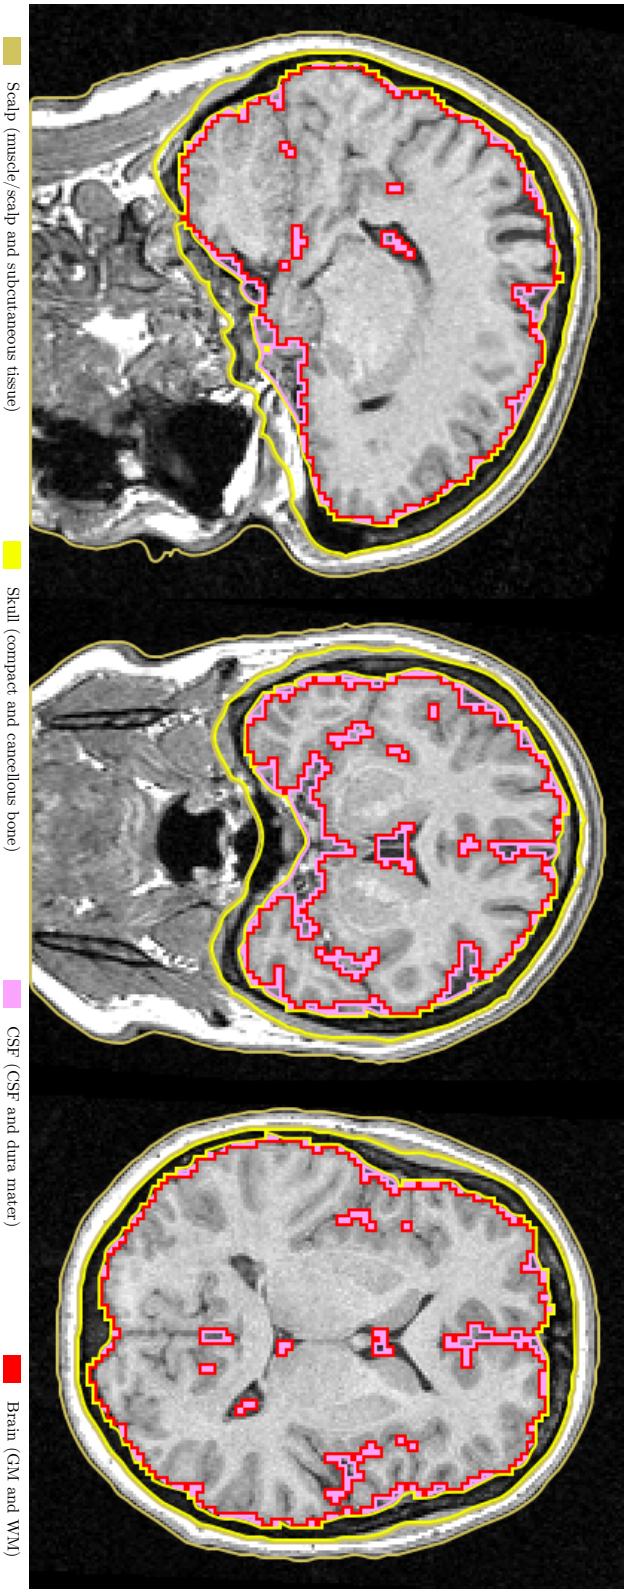

Figure 7.6: Exemplary segmentation results using the active contours approach from T1-weighted MRI data only. Outlines of the segmented tissue regions most relevant for the application in EEG source analysis (scalp, skull, CSF, and brain) are shown. Outlines are overlaid onto orthogonal slices of the T1-weighted MRI.

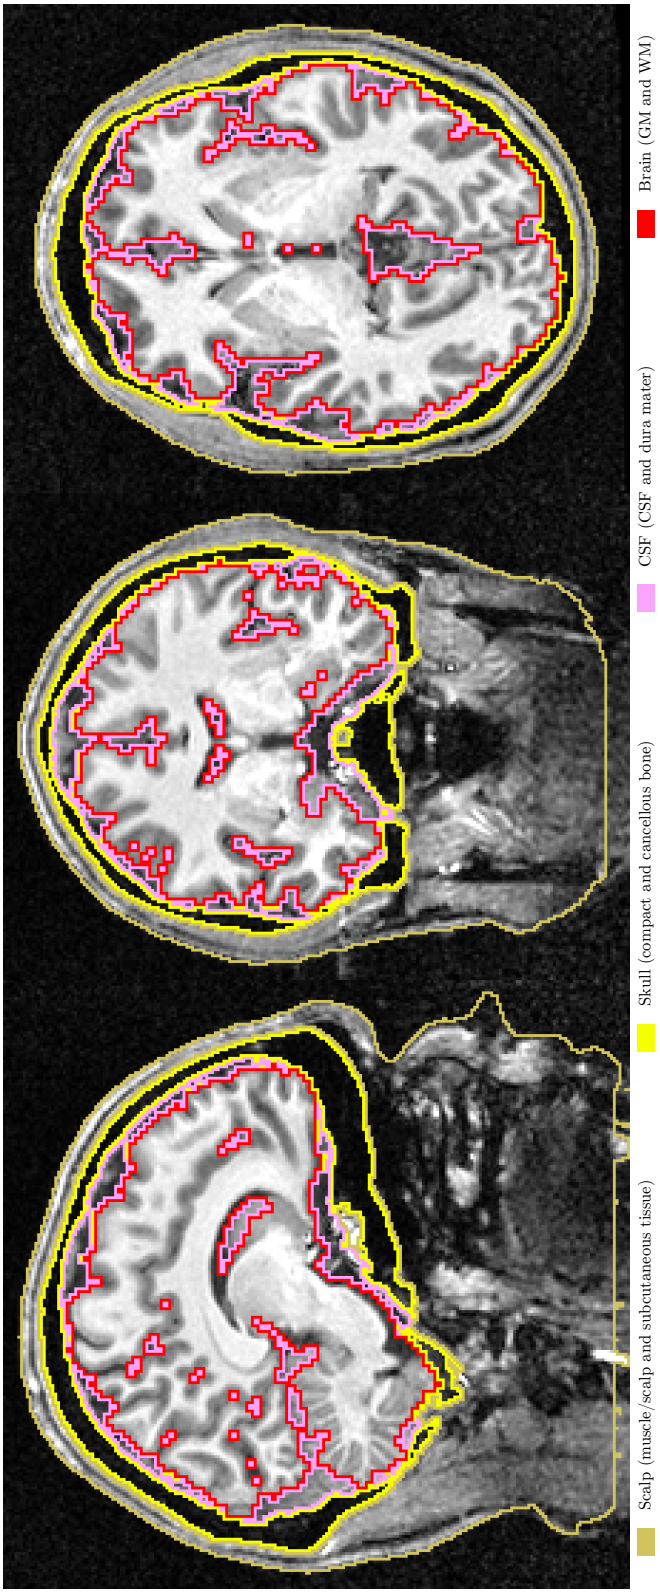

Figure 7.7: Exemplary segmentation results using the active contours approach from T1- and T2-weighted MRI data. Outlines of the segmented tissue regions most relevant for the application in EEG source analysis (scalp, skull, CSF, and brain) are shown. Outlines are overlaid onto orthogonal slices of the T1-weighted MRI.

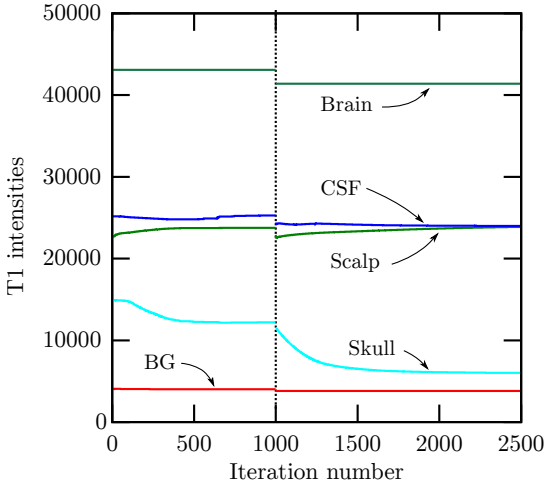

Figure 7.8: Plot of the mean gray values $\mu_{c,1}$ for the T1-weighted image versus the performed iterations of the active contours segmentation. Segmentation is performed on the coarse resolution level for the first 1000 iterations. After that the segmentation is continued on the fine level. The discontinuity in the plot is caused by this transition.

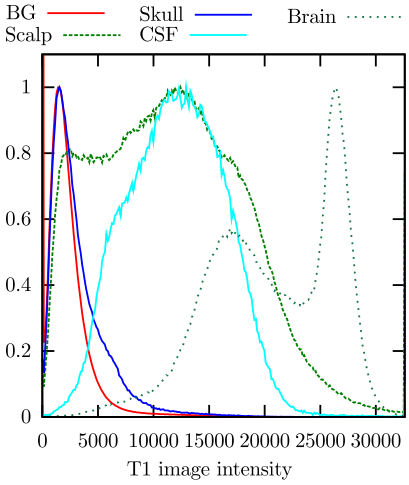

Figure 7.9: Histogram of the T1-image intensities for the segmented regions obtained by the five regions active contours segmentation. Histogram values are normalized.

Although the five regions active contours segmentation produced good results with respect to the skull segmentation there is a drawback. It is apparent from the histograms of the intensities inside of the final segmented regions that the intensities inside of each region are not distributed following a Gaussian distribution. This means that our image model where we assumed a Gaussian distribution might be inappropriate. The reason for this is that there are more than five tissues with differing gray value characteristic in the human head. When differentiating only five regions, for example, compact bone which is always dark is not distinguished from cancellous bone which might appear very bright.

In summary, it can be stated that the active contours segmentation approach was able to accurately segment the skull from MRI data. The CSF and brain segmentation was less good. It might, thus, be necessary for the construction of EEG head models to combine the segmented scalp and skull masks with CSF and brain masks obtained using a different, more suitable procedure.

8 A Markov Random Field Based Segmentation Approach for the Generation of Volume Conductor Models

8.1 Introduction

The problem of segmenting the human head from commonly available MRIs is a rather hard one. A large number of different tissue types with overlapping gray value distributions have to be differentiated (Fig. 3.7a), and the MRIs are often obstructed by noise and other imaging artifacts. For this reasons we cannot assume, that an accurate segmentation is possible taking only the observed MRI gray values into account, and additional information about the segmented image must be incorporated.

In this chapter we, therefore, propose an approach for the segmentation of volume conductor models for EEG source analysis from MR data that incorporates both information on the individual head anatomy contained in MR images, and a-priori knowledge on the anatomy of the human head.

Image information and a-priori knowledge are combined in a Bayesian framework. An MRF model is used to encode our knowledge on the structure of the head tissues. A good segmentation is then found by searching for the labeling x that produces the maximum a-posteriori probability (MAP). This MAP estimation is a widely used approach in Bayesian image analysis (e.g., [49, 166]).

In the following sections we will describe the components of our segmentation approach, and demonstrate its effectiveness. In Section 8.2 it is detailed how the MRF model is specified so that the a-priori probability $P(x)$ best represents the anatomy of the human head. Our choice for the likelihood function, which defines the link between the image data and the individual anatomy, is presented in Section 8.3. The complete algorithm is described in Section 8.4, and implementation details are given in Section 8.5. At the end of the chapter experiments on real MRI data are presented (Sec. 8.6), and the results of the experiments are discussed (Sec. 8.7).

8.2 A Markov Random Field Model to Incorporate A-priori Knowledge on the Anatomy of the Human Head

In the proposed segmentation approach, contextual information is exploited to incorporate a-priori knowledge about the anatomy of the head. The anatomy of the superior part of the head can be described as consisting of multiple tissue layers. Starting at the WM surface we find consecutive layers of GM, CSF, dura mater, compact bone, cancellous bone, muscle and SCT. In the inferior part of the head the geometry does not adhere so strictly to a layered structure but

starting at the basal WM we can observe consecutive layers of GM, possibly CSF, dura mater, compact bone and other soft tissues.

In such a layered structure regarding a point belonging to a certain region it can be expected that not all other regions occur equally likely in the point's neighborhood. Regarding, for example, a point in the compact bone compartment we expect to find compact bone, cancellous bone, muscle or dura mater in its neighborhood but we do not expect to find GM or WM, or points belonging to the BG. In the same sense we can give an a-priori probability for finding a certain region at a given point by looking at the regions occurring in its neighborhood. For example, for a point in which neighborhood only GM occurs there is a high a-priori probability that this point belongs to GM, WM, or CSF, but only a minor probability, that it belongs to tissues, like muscle or cancellous bone.

In a layered structure, additional information can be gathered by also considering where the neighboring regions occur relative to some given reference surfaces. We will discuss as an example a layered structure with consecutive regions $\Omega_1, \ldots, \Omega_n$ which is confined on one side by the reference surface A, so that A is adjacent to Ω_1 but not to Ω_2 (Fig. 8.1). Point $\boldsymbol{r}$ is our point of interest and we assume that region Ω_2 occurs frequently in the part of $\boldsymbol{r}$'s neighborhood, which is closer to A than $\boldsymbol{r}$ is. Not taking into account the position of the neighbors relative to $\boldsymbol{r}$ and A we would expect, that $\boldsymbol{r}$ equally likely belongs to Ω_1, Ω_2 or Ω_3. When now also incorporating the information that the neighbors belonging to region Ω_2 are closer to A, we would expect that $\boldsymbol{r}$ more likely belongs to Ω_2 or Ω_3.

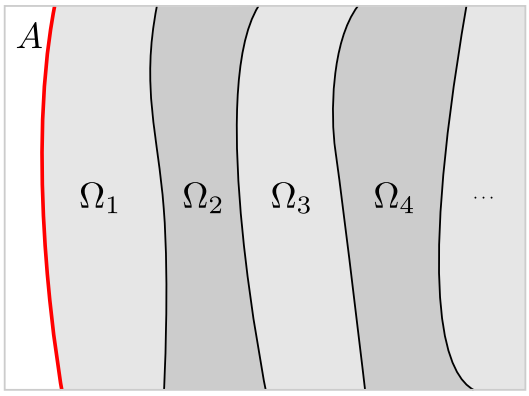

Figure 8.1: Illustration of layered structure of regions Ω_1 to Ω_n with reference surface A.

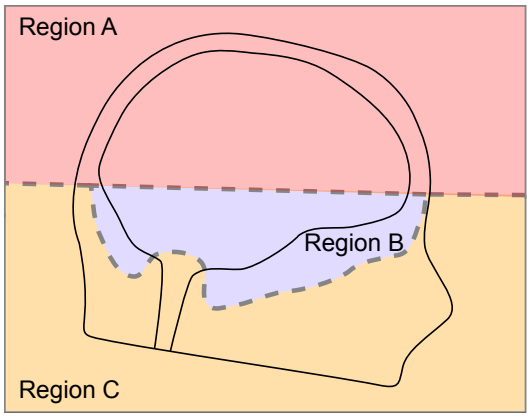

Figure 8.2: Head regions with different structures of head tissues. Region A: superior head and skull cap; region B: skull base (without foramen magnum); region C: inferior head and neck.

In our application we use the scalp and WM surfaces as reference surfaces because there exist several proven methods to accurately extract these from MR images [46, 22, 1, 38]. The image processing pipeline for EEG source analysis which is described in Chapter 6 is also able to deliver accurate segmentations of the scalp and WM surfaces.

The information about the layered anatomy of the head that was described informally above is incorporated into our algorithm by means of a nonhomogeneous, anisotropic MRF model.

A-priori probabilities $P(x_i|x_{\mathcal{N}_i})$ for finding the label x_i at site i given the labels at neighboring voxels $x_{\mathcal{N}_i}$ are specified to reflect the anatomy of the human head. The definition of the conditional probabilities $P(x_i|x_{\mathcal{N}_i})$ is done by employing the equivalence of MRFs and GRFs, and by specifying heuristic pseudo probabilities for the occurrence of different head tissues at neighboring sites (Sec. 3.9). Equation 3.63 gives the relationship between the conditional probabilities $P(x_i \mid x_{\mathcal{N}_i})$ and the clique potentials V_c of the Gibbs random field.

For the specification of the clique potentials we introduce pseudo transition probabilities $P_{x_i,x_{i'}}(i,i')$ between the different head regions. They are related to the pairwise clique potentials as given in the following equation.

$$V_2(i,i',x_i,x_{i'}) = -\ln\left(P_{x_i,x_{i'}}(i,i')\right) \tag{8.1}$$

Furthermore, for the single-site cliques we define the clique potentials using a-priori probabilities only depending on the location of the site i.

$$V_1(i,x_i) = -\ln\left(P^{\text{atlas}}_{x_i}(i)\right) \tag{8.2}$$

Using the pseudo transition probabilities and the single-site a-priori probabilities the conditional probabilities $P(x_i|x_{\mathcal{N}_i})$ can be defined as in Equation 8.3.

$$P(x_i|x_{\mathcal{N}_i}) = \frac{P^{\text{atlas}}_{x_i}(i)\ \prod_{i'\in\mathcal{N}_i} P_{x_i,x_{i'}}(i,i')}{\sum_{x'_i} P^{\text{atlas}}_{x'_i}(i)\ \prod_{i'\in\mathcal{N}_i} P_{x'_i,x_{i'}}(i,i')} \tag{8.3}$$

The pseudo transition probabilities are now chosen in such a way that the configuration x which best resembles the anatomy of the head has highest probability $P(x)$. From Equation 3.57 we can see that a high probability $P(x)$ can only be reached by achieving a low energy $U(x)$. This in turn means that the clique potentials $V_2(i,i',x_i,x_{i'})$ have to be small and that the transition probabilities for a reasonable clique configuration must be large. As an example, cancellous and compact bone tissue occur in the head neighbouring each other. For $x_i = \{\textit{Cancellous bone}\}$ and $x_{i'} = \{\textit{Compact bone}\}$, thus, the transition probability $P_{x_i,x_{i'}}(i,i')$ must be large, so that the clique potential $V_2(i,i',x_i,x_{i'})$ is small.

Regarding its anatomy (Sec. 3.1) the head can be divided into three regions (Fig. 8.2). In these regions the MRF model will be specified by different transition probabilities to reflect the different structure of the head tissues. In the superior part of the head above a virtual plane passing approximately through the nasion and the inion — region A — the head tissues are arranged in a layered structure between the scalp and WM reference surfaces. The region directly adjacent to the skull base shall be region B. In region B we find a layered structure of tissues relative to the WM reference surface. The third region — region C — makes up the remaining, inferior part of the head. Here, no layered structure relative to any reference surfaces exists but still a certain structure can be observed.

We will start by describing how we specified the MRF model in region C. In region C we do not find the layered structure with respect to some reference surface as we find it in regions A and B. Thus, our MRF model in the inferior head and neck region is isotropic and the clique

potentials do not depend on the location of the neighbouring site i' relative to i:

$$V_2(i, i', x_i, x_{i'}) = V_2(i, x_i, x_{i'})$$

Nevertheless, the anatomy in region C follows a certain structure in which, for example, cancellous bone is always enclosed in compact bone tissue, and WM is always covered by GM. A directed adjacency graph, like the graph shown in Figure 8.3a, can be used to illustrate the structure of the head tissues for the inferior head and neck region. The vertices in the graph represent the different head tissues and the edges represent allowed transitions from one into another tissue.

Transition probabilities between x_i and $x_{i'}$ for the isotropic case can now be defined based on whether a transition from x_i to $x_{i'}$ is allowed:

$$P^C_{x_i,x_{i'}} = P^{\text{iso}}_{x_i,x_{i'}} = \begin{cases} c_{\text{self}} & \text{if } x_i = x_{i'} \\ \epsilon & \text{if transition from } x_i \text{ to } x_{i'} \text{ is forbidden} \\ \frac{1-c_{\text{self}}-\epsilon'}{n'} & \text{if transition from } x_i \text{ to } x_{i'} \text{ is allowed} \end{cases} \tag{8.4}$$

c_{self} is a constant specifying to which extent configurations are favoured where two neighbouring sites have the same label; ϵ is a very small constant; n' is the number of sites $i' \in \mathcal{N}_i$ that have labels $x_{i'}$ for which transitions into x_i are allowed; and ϵ' is the sum of transition probabilities for all neighbouring sites having labels for which transitions are forbidden.

In region B we assume that there is a certain layered structure of head tissues relative to the WM reference surface. This is modeled by choosing clique potentials that depend on the position of the neighbouring site i' relative to site i and to the reference surface. Different clique potentials are used whether the neighbour i' is closer to the reference surface than i is, whether it is farer away, or whether i' and i are at approximately the same distance. As it was done for region C pseudo transition probabilities $P^B_{x_i,x_{i'}}$ are used to define the clique potentials according to Equation 8.1.

$$P^B_{x_i,x_{i'}}(i, i') = \begin{cases} P^{\text{from, WM}}_{x_i,x_{i'}} & \text{if } \quad \Delta d^{\text{WM}}(i, i') > +d' \\ P^{\text{to, WM}}_{x_i,x_{i'}} & \text{if } -d' > \Delta d^{\text{WM}}(i, i') \\ P^{\text{iso}}_{x_i,x_{i'}} & \text{if } -d' \leq \Delta d^{\text{WM}}(i, i') \leq +d' \end{cases} \tag{8.5}$$

$$\Delta d^{\text{WM}}(i, i') = d^{\text{WM}}(i) - d^{\text{WM}}(i') \tag{8.6}$$

$d^{\text{WM}}(i)$ is the distance of site i from the WM reference surface, d' is a constant determining when i and i' are considered to be equally far away from the reference surface, and $P^{\text{from, WM}}_{x_i,x_{i'}}$ and $P^{\text{to, WM}}_{x_i,x_{i'}}$ are pseudo transition probabilities for the cases, where i is farer away from the WM reference surface than i', respectively, where i is closer to the WM reference surface than i'.

To explain how the pseudo transition probabilities $P^{\text{from, WM}}_{x_i,x_{i'}}$ and $P^{\text{to, WM}}_{x_i,x_{i'}}$ were defined we again consider a layered structure (Fig. 8.1). We assume that i' lies in Ω_3 and i is farer away from A than i' so that $d^A(i) - d^A(i') > +d'$. This means, when moving from i' to i we are moving away from our reference surface A. In a layered structure we expect that we in this case consecutively pass through the regions Ω_1, Ω_2, ... Thus, when propagating from i' to i

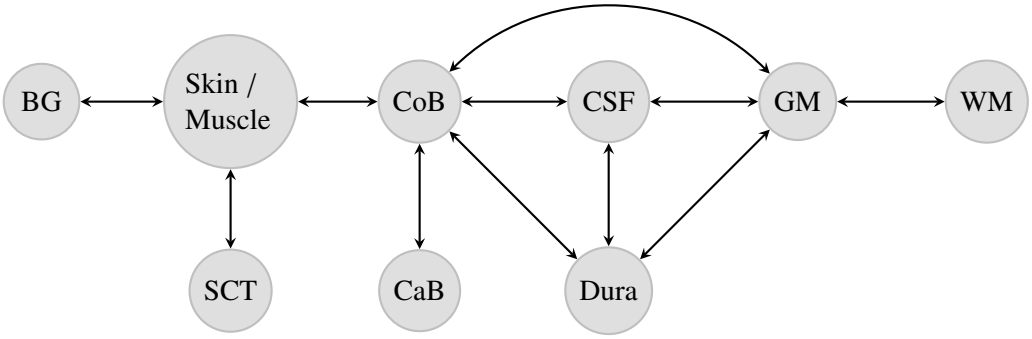

(a) Adjacency graph for region C - inferior head and neck

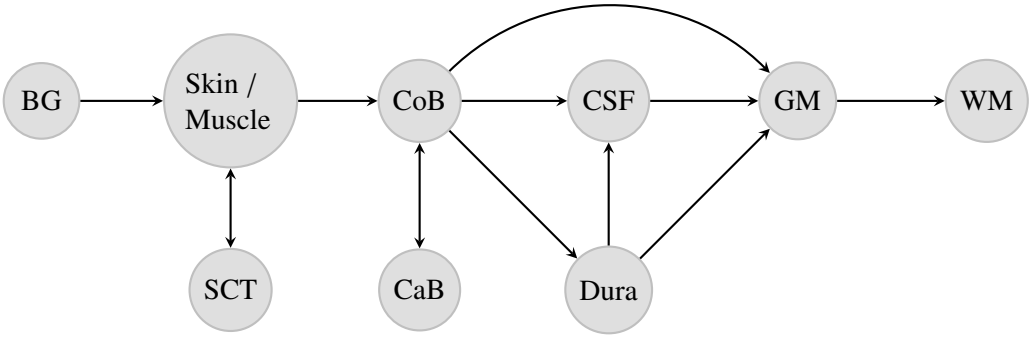

(b) Adjacency graph when propagating from scalp to WM reference surface.

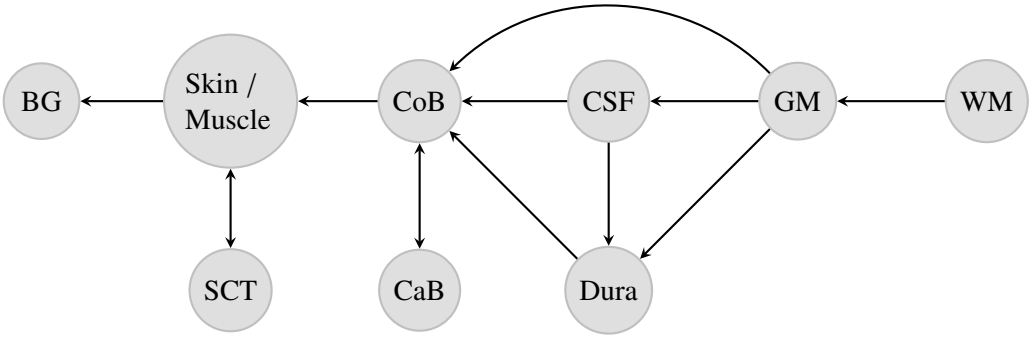

(c) Adjacency graph when propagating from WM to scalp reference surface.

Figure 8.3: Adjacency graphs illustrating definition of pseudo transition probabilities. BG: background; SCT: subcutaneous tissue; CoB: compact bone; CaB: cancellous bone; CSF: cerbrospinal fluid; Dura: Dura mater; GM: gray matter; WM: white matter.

we expect to find regions Ω_3 or Ω_4 at i, but not Ω_2 or Ω_1. As in the isotropic case we then say that the transition from Ω_3 to Ω_4 is allowed when moving away from the reference surface, while the transition from Ω_3 to Ω_2 or Ω_1 is forbidden. Forbidden and allowed transitions for the layered structure of the head tissues when propagating from the WM to the scalp reference surface can be illustrated using the directed graph in Figure 8.3c. The definition of the pseudo transition probabilities $P^{\text{from, WM}}_{x_i,x_{i'}}$ is then done analogously to Equation 8.4. $P^{\text{to, WM}}$ for the case that $d^{\text{WM}}(i) - d^{\text{WM}}(i') < -d'$, that is, when moving towards the WM reference surface is defined accordingly. A graph for this case illustrating the allowed and forbidden transitions in the layered structure of the head tissues is depicted in Figure 8.3b.

In region A — the superior head region — we find a layered structure of head tissues inbetween the WM surface and the scalp surface. As in region B the MRF model is anisotropic and the clique potentials depend on the position of i' relative to i and to the reference surfaces. The pseudo transition probabilities specifying the clique potentials can be defined similarly as in region B. The only difference is that in region A two reference surfaces are present. Thus, first transition probabilities $P^{A,\text{ WM}}_{x_i,x_{i'}}$ and $P^{A,\text{ Scalp}}_{x_i,x_{i'}}$ with respect to both of the reference surfaces individually are defined. $P^{A,\text{ WM}}_{x_i,x_{i'}}$ is defined exactly as in Equation 8.5, while $P^{A,\text{Scalp}}_{x_i,x_{i'}}$ is defined as follows:

$$
\begin{aligned}
P^{A,\text{ Scalp}}_{x_i,x_{i'}}(i) &= \begin{cases} P^{\text{from, Scalp}}_{x_i,x_{i'}} & \text{if} \quad \Delta d^{\text{scalp}}(i,i') > +d' \\ P^{\text{to, Scalp}}_{x_i,x_{i'}} & \text{if} \quad -d' > \Delta d^{\text{scalp}}(i,i') \\ P^{\text{iso}}_{x_i,x_{i'}} & \text{if} \quad -d' \leq \Delta d^{\text{scalp}}(i,i') \leq +d' \end{cases} \\
&= \begin{cases} P^{\text{to, WM}}_{x_i,x_{i'}} & \text{if} \quad \Delta d^{\text{scalp}}(i,i') > +d' \\ P^{\text{from, WM}}_{x_i,x_{i'}} & \text{if} \quad -d' > \Delta d^{\text{scalp}}(i,i') \\ P^{\text{iso}}_{x_i,x_{i'}} & \text{if} \quad -d' \leq \Delta d^{\text{scalp}}(i,i') \leq +d' \end{cases}
\end{aligned}
\tag{8.7}
$$

$\Delta d^{\text{scalp}}(i,i')$ is defined analogously to Equation 8.6. In the last step we exploited that for voxels between the scalp and WM surface moving towards the scalp surface is essentially the same as moving away from the WM reference surface and vice versa. Thus, $P^{\text{from, Scalp}}_{x_i,x_{i'}} = P^{\text{to, WM}}_{x_i,x_{i'}}$ and $P^{\text{to, Scalp}}_{x_i,x_{i'}} = P^{\text{from, WM}}_{x_i,x_{i'}}$.

The two pseudo transition probabilities depending on different reference surfaces are then combined to a weighted average with the weighting factor depending on the distances to both reference surfaces.

$$
P^{A}_{x_i,x_{i'}} = w_{\text{WM}} P^{A,\text{ WM}}_{x_i,x_{i'}} + w_{\text{scalp}} P^{A,\text{ Scalp}}_{x_i,x_{i'}} \tag{8.8}
$$

$$
w_{\text{WM}} = 1 - \frac{d^{\text{scalp}}(i)}{d^{\text{scalp}}(i) + d^{\text{WM}}(i)} \tag{8.9}
$$

$$
w_{\text{scalp}} = 1 - \frac{d^{\text{WM}}(i)}{d^{\text{scalp}}(i) + d^{\text{WM}}(i)} \tag{8.10}
$$

8.3 The Likelihood

In our segmentation approach volumetric medical images, for example, T1- or T2-weighted MRIs, provide information on the individual anatomy of the subject's head. In the context of Bayesian image analysis, these possibly multimodal images are our observations. For each site i we observe a (multimodal) intensity $\boldsymbol{y}(i)$.

The likelihood $l(\boldsymbol{y} \mid x)$ gives the probability with which the given labeling x generates the observed image.

For the segmentation of T1- and T2-weighted MRIs in our proposed segmentation approach, we assume that the intensity in each tissue region is constant with added Gaussian noise. In the work of Gudbjartsson and Patz [51] it was shown that this is a reasonable assumption for images with a not too low SNR. For a single modality the likelihood to observe y_i given the label $x_i = c$ is then equal to a Gaussian probability distribution with mean μ_c and standard deviation σ_c.

$$l(y_i \mid x_i, \mu_c, \sigma_c) = \frac{1}{\sqrt{2\pi}\sigma_c} \exp\left\{-\frac{1}{2}\left(\frac{y_i - \mu_c}{\sigma_c}\right)^2\right\} \tag{8.11}$$

When segmenting multimodal image data the likelihood can be expressed by a multivariate Gaussian distribution (Eq. 3.17) with $\boldsymbol{\mu_c}$ containing the mean intensities in region c for each of the image modalities and Σ_c meaning the covariance matrix for region c.

It has to be noted that the parameters for the Gaussian distributions, that is the mean gray values and covariance matrices, are not known a priori. Different MR recording sequences or only slightly different sequence parameters will lead to differences in the observed intensities for the head tissues. Thus, the parameters of the gray value distributions must be estimated simultaneously while determining the optimal labeling.

8.4 The Segmentation Algorithm

As already mentioned above, we are trying to find the segmentation x, and the parameters $\boldsymbol{\mu}_c$ and Σ_c which maximize the Bayesian a-posteriori probability from Equation 3.14. In our approach optimization is performed by means of an algorithm based on the Expectation-Maximization [43] and the ICM algorithm by Besag [19] (Sec. 3.9).

For the sake of simplicity, we first assume that the parameters of our Gaussian gray value distributions are known. Our task is then to maximize the a-posteriori probability with respect to x. In this case, the ICM algorithm can be used to maximize the a-posteriori probability. The algorithm runs as follows. Starting at an initial solution $x^{(0)}$ we construct a new possible segmentation $x^{(1)}$ by assigning the mode of the conditional probability distribution $P(x_i \mid \boldsymbol{y}_i, x_{\mathcal{N}_i})$ (Eq. 8.3) to site i.

$$x_i = \arg\max_c P(X_i = c \mid \boldsymbol{y}_i, x_{\mathcal{N}_i}) \tag{8.12}$$

This is done for all sites $i \in \mathcal{S}$ following a random update order.

Using the update rule in Equation 8.12 the local conditional probability $P(x_i \mid \boldsymbol{y}_i, x_{\mathcal{N}_i})$ increases monotonically whenever a label is replaced by the conditional mode. When the observations of the image data $\boldsymbol{y}$ are conditionally independent of each other then this means that also the

posterior probability $P(x \mid \boldsymbol{y})$ increases at each iteration of the ICM algorithm. Convergence to a local maximum of the probability distribution is therefore guaranteed [19].

The parameters of the gray value distributions are, contrary to our above assumption, not known beforehand. Our algorithm must, therefore, simultaneously estimate these parameters. For such types of problems in many applications algorithms based on the EM algorithm [43] are used. In an EM algorithm parameters and the sought-after labeling x are alternately optimized.

During each iteration first the E-step is performed. In the E-step the conditional expectation of the log-likelihood for the complete data is computed (Eq. 3.66). This essentially means, that we have to obtain an estimate for the missing data which is the labeling x. In our approach, we estimate the labeling x in the E-step using the ICM algorithm by Besag [19] as described above.

Having obtained an updated estimate for the labeling x now the M-step is carried out. In the M-step the conditional expectation of log likelihood for the complete data is optimized with respect to the parameters of the gray value distributions $\boldsymbol{\lambda} = \{\boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c \mid c \in \mathcal{L}\}$. For a complete data likelihood from the exponential family of distributions the conditional expectation is given in Equation 3.69. In this way two formulas are obtained that describe how the parameters $\boldsymbol{\mu}_c$ and $\boldsymbol{\Sigma}_c$ have to be updated to maximize $Q(\boldsymbol{\lambda} \mid \boldsymbol{\lambda}^{(t)})$.

$$\boldsymbol{\mu}_c^{(t+1)} = \frac{\sum_{i \in \mathcal{S}} P(X_i = c \mid \boldsymbol{y}_i, \boldsymbol{\lambda}^{(t)})\, \boldsymbol{y}_i}{\sum_{i \in \mathcal{S}} P(X_i = c \mid \boldsymbol{y}_i, \boldsymbol{\lambda}^{(t)})} \tag{8.13}$$

$$\boldsymbol{\Sigma}_c^{(t+1)} = \frac{\sum_{i \in \mathcal{S}} P(X_i = c \mid \boldsymbol{y}_i, \boldsymbol{\lambda}^{(t)}) \left(\boldsymbol{y}_i - \boldsymbol{\mu}_c^{(t+1)}\right)\left(\boldsymbol{y}_i - \boldsymbol{\mu}_c^{(t+1)}\right)^{\mathsf{T}}}{\sum_{i \in \mathcal{S}} P(X_i = c \mid \boldsymbol{y}_i, \boldsymbol{\lambda}^{(t)})} \tag{8.14}$$

Algorithm 8 Markov Random Field based segmentation approach.

Require: $\boldsymbol{y}$: Multimodal image data
Require: $t^{\max}$: Maximum number of iterations
Require: $M^{\text{ref. scalp}}$, M^{WM}: Scalp and WM reference surfaces
1: Preprocessing (Sec. 8.5)
2: Generate initial labeling $x^{(0)}$ (Sec. 8.5.1)
3: Choose initial gray value parameters $\boldsymbol{\Sigma}^{(0)}$ and $\boldsymbol{\mu}^{(0)}$ (Sec. 8.5.2)
4: $t \leftarrow 0$ ▷ Reset iteration counter
5: **repeat**
6: E-step: $x_i^{(t+1)} = \arg\max_c P(X_i = c \mid \boldsymbol{y}_i, \boldsymbol{\lambda}^{(t)}, x_{\mathcal{N}_i}),\ \forall i \in \mathcal{S}$
7: M-step: Compute $\boldsymbol{\mu}_c^{(t+1)}$ and $\boldsymbol{\Sigma}_c^{(t+1)}$ according to Eqs. 8.13 and 8.14
8: $t \leftarrow t + 1$ ▷ Increase iteration counter
9: **until** $(x_i^{(t)} = x_i^{(t-1)}, \forall i \in \mathcal{S})\ \vee\ (t \geq t^{\max})$
10: Post-processing (Sec. 8.5)

The complete segmentation approach now runs as follows (Alg. 8): First, an initial labeling $x^{(0)}$ is computed based on the scalp and WM reference surfaces (Sec. 8.5.1), and the initial gray value parameters are set to heuristically determined initial values (Sec. 8.5.2). Next, the E-step of the EM algorithm is performed, and a new estimate for the labeling $x^{(t+1)}$ is computed by running one iteration of the ICM algorithm, that is, for each voxel following a random update order the

label is replaced by the conditional mode of the local probability $P(X_i = c \,|\, \boldsymbol{y}_i, \boldsymbol{\lambda}^{(t)}, x_{\mathcal{N}_i})$. In the M-step then new estimates $\boldsymbol{\mu}_c^{(t+1)}$ and $\boldsymbol{\Sigma}_c^{(t+1)}$ for the gray value parameters which maximize the conditional expectation of the log-likelihood of the complete data are computed. The E- and the M-step are repeated until either the labeling did not change during the last iteration ($x_i^{(t)} = x_i^{(t-1)}, \forall i \in \mathcal{S}$), or until a fixed maximum number of iterations is reached. Finally, some optional post-processing operations are applied to the labeling x (Sec. 8.5.3).

8.5 Implementation Details

8.5.1 Generation of Initial Solution

The proposed algorithm as described in Algorithm 8 requires an initial labeling, here also called a start mask. This start mask is used during the first E-step while computing a new labeling using the update step from the ICM algorithm. As the ICM and EM algorithm only converge locally ideally a start mask shall be constructed which already resembles as good as possible the final solution. In the following, it will be explained how an initial solution is generated for the MRF segmentation approach.

The segmentation approach requires a scalp and a WM reference surface and additionally distance maps to both reference surfaces as input for the definition of the anisotropic MRF model (Sec. 8.2). These reference surfaces and the associated distance maps will be the basis for the generation of the start mask.

A detailed description of the start mask generation is given in Algorithm 9. Morphological operations, specifically dilation (dil(x)) and erosion (ero(x)) operations, are applied to the reference surfaces to consecutively create masks for all head tissues. Finally, all individual tissue masks are combined to create an initial solution to the segmentation problem.

Dilation and erosion of the binary reference scalp mask $M^{\text{ref. scalp}}$ and the reference WM mask M^{WM} are performed by employing thresholding of the maps of distances to the reference surfaces. Dilation of the scalp reference surface by a distance d^{dil} can, for example, be performed by thresholding the scalp distance map from $-d^{\text{dil}}$ to 0. Note, that the distance with respect to the scalp reference surface is computed in such a way that the distance is positive on the inside of the scalp reference surface and negative on the outside. Dilation and erosion of other masks where no distance map is available beforehand is performed based on a simpler Chamfer distance transform [68].

The generation of the initial solution begins by generating nested binary masks for the tissues WM, GM, CSF, dura mater and skull. Each tissue mask is created by dilation of the previous mask. The GM mask, for example, is generated by dilation of the WM mask and the CSF mask, in turn, is a dilated GM mask. A scalp region mask is computed by erosion of the scalp reference surface. The SCT and cancellous bone regions which are enclosed inside of the scalp, respectively, the skull region are constructed by erosion of the scalp, respectively, the skull mask. The foramen magnum is a major opening at the base of the skull. To incorporate this hole into the skull of the initial solution a pre-defined template mask is used which defines the approximate region of the foramen magnum. Finally, all masks are combined according to the scheme illustrated in Figure 8.4. Examples of the above tissue masks and of a final start mask are shown in Figures 8.5 and 8.6, respectively.

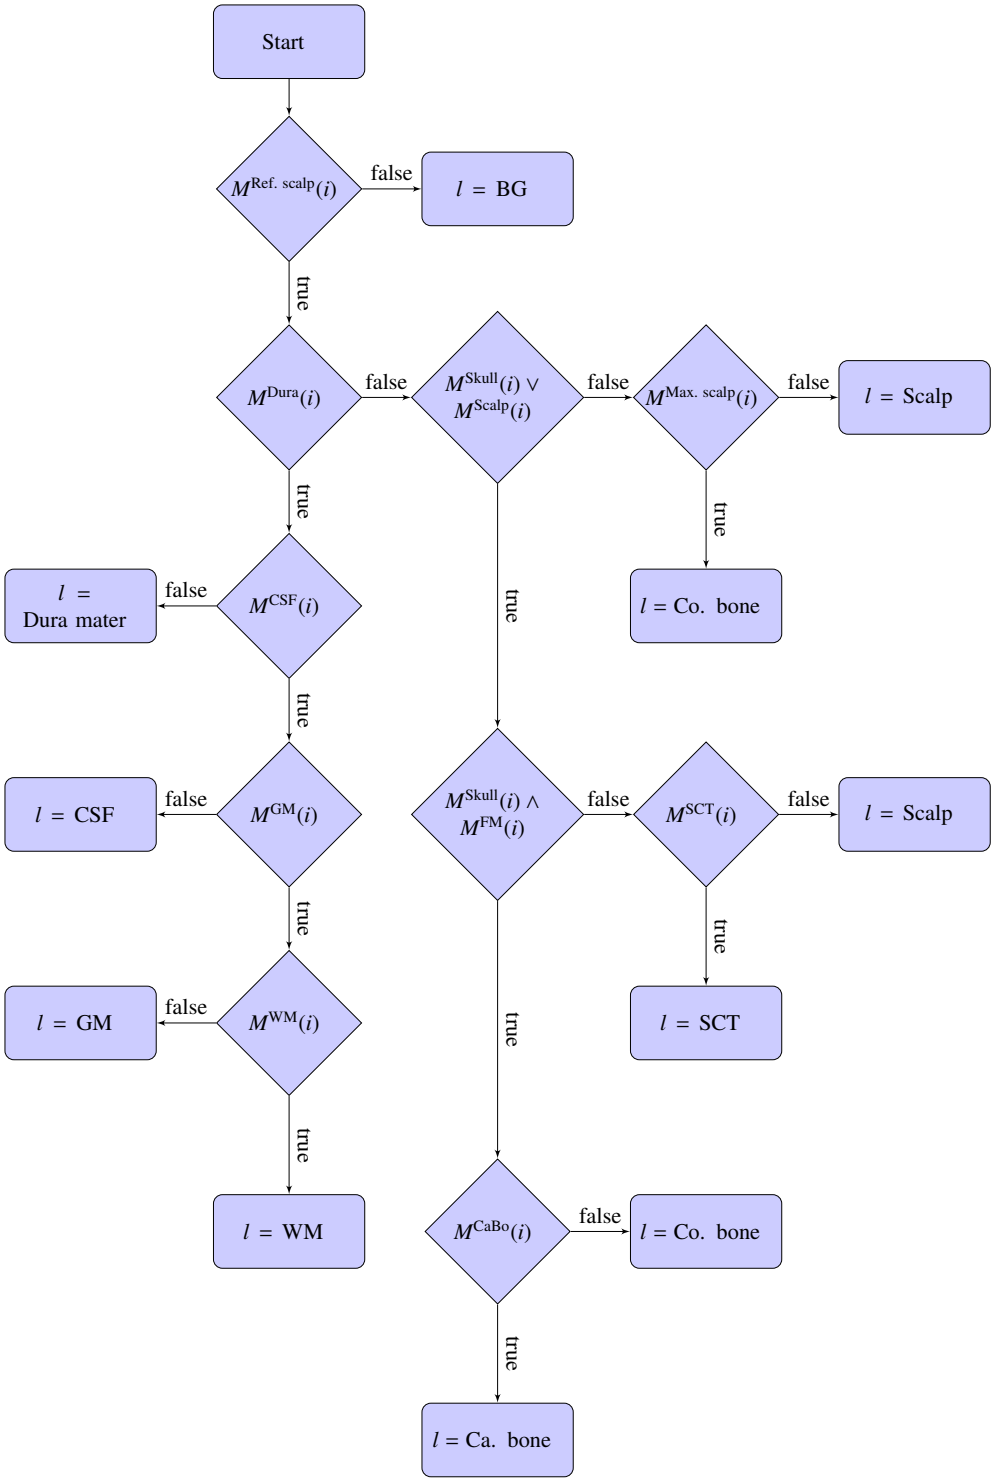

Figure 8.4: Composition of the initial solution from the individual tissue masks. Compact and cancellous bone are abbreviated as co. bone and ca. bone, respectively. FM means foramen magnum.

Figure 8.5: A sagittal slice of intermediate masks generated during construction of start mask.

Algorithm 9 Generation of an initial solution for the volume conductor segmentation.

Require: $M^{\text{ref. scalp}}$, M^{WM}: Scalp and WM reference surfaces
Require: Maps of distances to reference surfaces
Require: M^{FM}: Template based foramen magnum mask
Require: $M^{\text{inf. head}}$: Inferior head region (skull base) mask

1: $M^{\text{GM}} = \text{dil}(M^{\text{WM}}, 4\,\text{mm})$ ▷ GM mask
2: $M^{\text{CSF}} = \text{ero}(\text{dil}(M^{\text{GM}}, 20\,\text{mm}), 19\,\text{mm})$ ▷ CSF mask
3: $M^{\text{dura}} = \text{dil}(M^{\text{CSF}}, 1\,\text{mm})$ ▷ Dura mater mask
4: $M^{\text{outer skull}} = \text{dil}(M^{\text{dura}}, 12\,\text{mm})$ ▷ Outer skull boundary mask
5: $M^{\text{skull}} = \left(M^{\text{outer skull}} \cap M^{\text{ref. scalp}}\right) \cap \neg M^{\text{dura}}$ ▷ Skull mask
6: $M^{\text{scalp}} = \text{ero}(M^{\text{ref. scalp}}, 5\,\text{mm}) \cap \text{M}^{\text{ref. scalp}}$ ▷ Scalp mask
7: $M^{\text{max. scalp}} = \text{ero}(M^{\text{ref. scalp}}, 12\,\text{mm}) \cap \neg\text{M}^{\text{inf. head}}$ ▷ Maximum scalp mask for superior head region
8: $M^{\text{SCT}} = \text{ero}(M^{\text{scalp}}, 1\,\text{mm})$ ▷ SCT mask
9: $M^{\text{CaBo}} = \text{ero}(M^{\text{skull}}, 2\,\text{mm})$ ▷ Cancellous bone mask
10: Combine individual tissue mask according to Fig. 8.4.

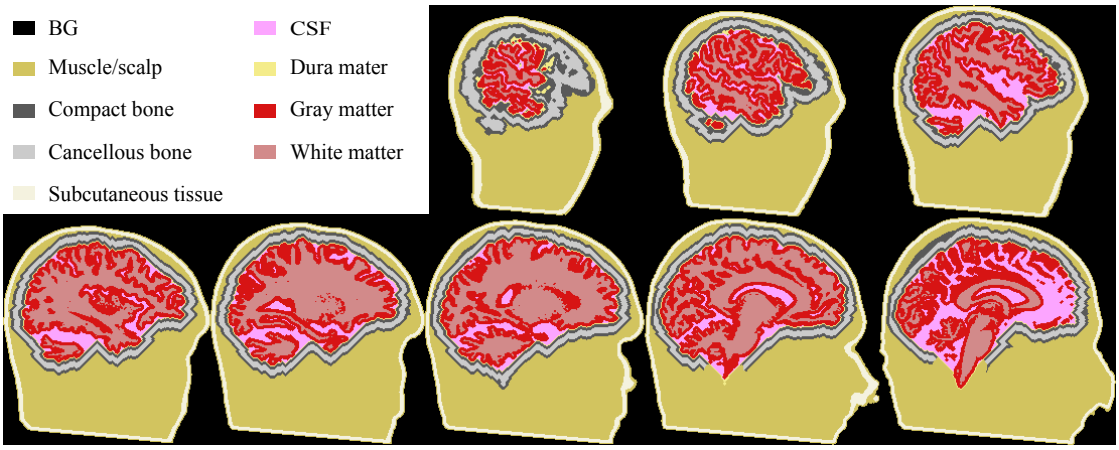

Figure 8.6: Montage of sagittal slices of the final start mask.

8.5.2 Initial Gray Values Parameters

For the EM algorithm initial gray value parameters — mean gray values $\boldsymbol{\mu}_c$ and covariance matrices $\boldsymbol{\Sigma}_c$ for each region c — are required for the definition of the likelihood during the first E-step. Due to the only local convergence of the EM algorithm it is advisable to choose initial gray value parameters which are as close as possible to the true parameters. This is especially true for the mean gray values $\boldsymbol{\mu}_c$. For this reason, we heuristically determined the mean gray values for each tissue and for different types of MRIs. How this was done is described in this section. In addition, it is explained how we choose the best set of initial gray value parameters for the current T1-weighted MRI without knowing a-priori to which type of MRI it belongs.

A total of four data sets containing a T1- and a T2-weighted MRI and seven data sets containing only a T1-weighted MRI were selected. All data sets were preprocessed in *BESA MRI* [17]. Preprocessing included transformation of the image to AC-PC space, resampling to $1.0 \times 1.0 \times 1.0\,\mathrm{mm}^3$ isotropic resolution, and intensity inhomogeneity correction. During the inhomogeneity correction the intensity of the T1-weighted MRIs were rescaled, so that the WM region was on average equally bright in all images.

In our segmentation approach we are in total differentiating eight different tissues. Thus, we need as initial parameters eight vectors of gray values $\boldsymbol{\mu}_c$. To obtain estimates for these mean gray value parameters we manually marked voxels from each tissue and in each of the data sets. Voxels were marked in every approximately fifth to tenth sagittal slice. Voxels at the interface between two tissues were not marked as their gray values might be affected by partial volume averaging.

Having marked a substantial number of voxels in each tissue region and for each data set we can now obtain an estimated intensity distribution for each individual tissue and MRI. Exemplary intensity distributions are shown in Figure 3.7a. Given these distributions it is now easy to compute an estimate for the mean gray values. Averaged gray values for each of the tissue regions and the different T1- and T2-weighted images can be seen in Figures 8.7a and 8.7b, respectively.

In our investigation we included T1-weighted MR images that were measured with and without fat suppression. This can be observed in the mean intensities shown in Figure 8.7a. The mean intensities of SCT and cancellous bone are much lower in the T1-weighted MRIs recorded with fat suppression as they are in the images recorded without fat suppression. Due to this large differences in the initial gray values we must determine two sets of mean gray values: one set for images with fat suppression and the other set for images without fat suppression. The final sets of average intensities for each tissue are obtained by averaging the mean intensities across all MRIs that are of the same kind.

For T2-weighted images we do not differentiate between different kinds of MRIs and, thus, average across all available T2-weighted images.

As we do not in all circumstances know if the T1-weighted MRI to be segmented was measured with or without fat suppression we must determine which set of initial gray value parameters to use. We do this, by computing the likelihood according to Equation 3.20 for each of the initial gray value parameters sets given the initial solution to our segmentation problem (the start mask, see Sec. 8.5.1). The set of initial parameters which yields the maximum likelihood is then selected. The likelihood computation is limited to a certain ROI in which we assume that

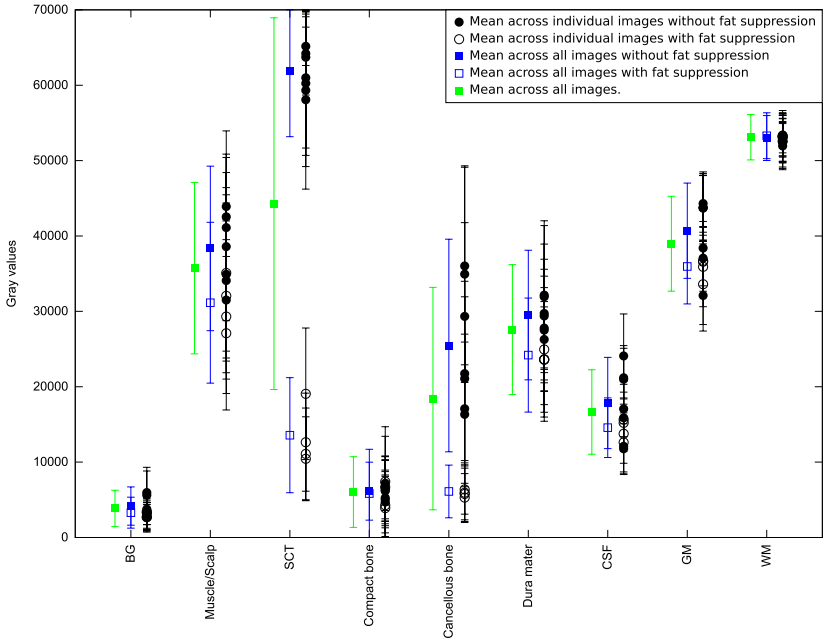

(a) Mean gray values in each of the tissue regions and for different T1-weighted MRIs.

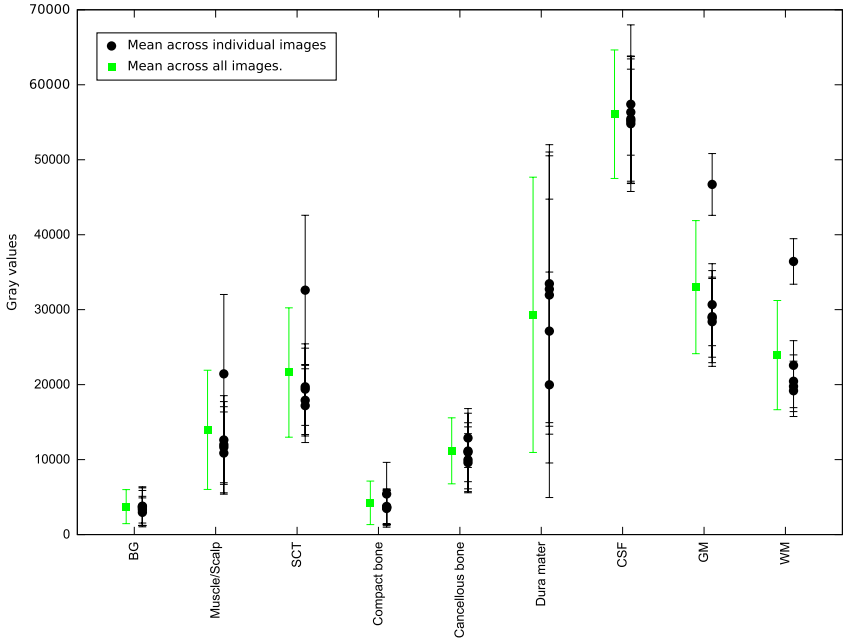

(b) Mean gray values in each of the tissue regions and for different T2-weighted MRIs.

Figure 8.7: Mean gray values for T1- and T2-weighted MRIs. Error bars indicate one standard deviation.

the start mask best fits the actual solution. We chose the superior head as ROI.

8.5.3 Additional Constraints on Allowed Segmentations

Besides the constraints defined by the MRF model we might want to incorporate additional constraints on the allowed segmentations to further improve the reliability of the algorithm. To do this, the proposed approach offers two mechanisms for incorporating additional a-priori knowledge.

The first approach utilizes a so-called safe mask. It is assumed that the labeling of all voxels which are marked in the safe mask is known with high certainty. So from the beginning fixed labels are assigned to all these voxels and their labels are not updated during the E-step of our algorithm.

The second mechanism through which constraints can be applied is the a-priori probability $P_{x_i}^{\text{atlas}}$. This mechanism can work similarly to employing a safe mask, for example, when setting

$$P_{x_i}^{\text{atlas}}(i) = \begin{cases} 1 - \epsilon' & \text{if } x_i = c \\ \epsilon & \text{if } x_i \neq c \end{cases}.$$

Here, ϵ is a very small positive number and $\epsilon' = \sum_{x_i \neq c} \epsilon$. In such a case, essentially only the label c would be allowed at site i. Moreover, using the a-priori probability it is also possible to incorporate a constraint where a certain label c is forbidden at site i by setting $P_c^{\text{atlas}}(i) = \epsilon$.

In the following we will shortly discuss some of the constraints which can optionally be applied in our segmentation approach. A minimum scalp thickness can be guaranteed by marking all voxels in the safe mask that are at a distance less than the minimum thickness from the reference scalp surface. These voxels can simply be found by thresholding the map of distances to the reference scalp surface which is available as input.

Another reasonable assumption one can make is that the input reference surfaces are true and can be trusted. This then means that every voxel outside of the scalp reference surface must belong to the BG and every voxel inside of the WM reference surface must belong to WM. This is easily enforced using the safe mask approach. Furthermore, assuming that the scalp reference surface is true we do not expect to find any voxels belonging to the BG inside of the scalp surface. Segmentations satisfying this assumption can be found when modifying the a-priori probability for BG in such a way that it is virtually zero everywhere inside of the reference scalp surface:

$$P_{\text{BG}}^{\text{atlas}}(i) = \epsilon, \forall i \in M^{\text{ref. scalp}}$$

8.5.4 Optional Post-Processing Steps

Some constraints cannot be applied a-priori because their formulation depends on the actual segmentation itself which is a-priori not known. Therefore, it is possible to enforce some additional, optional constraints in a post-processing step which is performed after the actual segmentation. A few of these constraints are described below.

Our study on the influence of skull geometry inaccuracies on EEG source analysis (Chap. 4) has shown that the inferior part of the head and neck can be modeled as a homogeneous region

with a conductivity equal to that of soft tissue while introducing only negligible errors into source analysis. We cannot completely trust our segmentation in the inferior part of the head due to the typically low SNR and the complicated anatomy in that region. Thus, it is useful to also simplify the segmentation result by labeling all voxels in the inferior part of the head and at a certain minimum distance from the brain as scalp. This is done by first creating a brain mask M^{brain} covering the regions that were segmented as GM or WM in the final segmentation result. The brain mask is then dilated with a radius equal to the minimum distance of the simplified region to the brain $d^{\text{min}} = 40\,\text{mm}$: $M^{\text{brain}} = \text{dil}(M^{\text{brain}}, d^{\text{min}})$. To only include voxels in the inferior part of the head the brain mask is intersected with the mask approximately defining the inferior part of the head: $M^{\text{brain}} = M^{\text{brain}} \cap M^{\text{inf. head}}$. Finally, all non-background voxels marked in this mask are labeled as scalp.

Another reasonable constraint is that the skull in the superior part of the head has a minimum thickness. This constraint is hard to apply a-priori through one of the above described mechanisms as the skull boundaries and the position of the skull are not known before the segmentation is performed. Therefore, it is implemented as a post-processing step. One possibility to ensure a minimum skull thickness is to mark all voxels up to a maximum distance from the segmented inner bone surface as skull. This approach has the disadvantage that it might produce inaccurate results if the segmented inner bone surface is erroneous. Similarly, enforcing a minimum skull thickness by erosion of the segmented outer bone surface is not an ideal solution either. In our approach a minimum skull thickness is, thus, enforced by dilating voxels on the medial plane of the segmented skull. This is the best compromise between dilating the segmented inner bone surface and eroding the segmented outer bone surface.

First, voxels on the medial plane are determined based on an algorithm described by Pudney [108]. According to this algorithm voxels are one after another removed from the skull mask if they fulfill two conditions. The first condition is that the topology of the skull mask must not change if the voxel is removed [15]. And the second condition is that the voxel must not be lying at the very edge of the medial plane [108]. Voxels are removed beginning with voxels at the surface of the skull object and then propagating towards the interior of the skull. We only propagate up to a distance of half the minimum skull thickness into the skull. In a final step, we create a mask with only those voxels that were not removed, because they did not fulfill the two conditions for removal. These are voxels on the medial plane of the skull.

The mask is then dilated by half the minimum skull thickness to obtain a minimum skull mask. The segmentation result is corrected by relabeling as compact bone all non-skull voxels that are lying in the superior part of the head and that are marked in the minimum skull mask. Figure 8.8 shows examples of a mask containing voxels on the medial plane and the resulting minimum skull mask.

8.6 Experiments

The proposed segmentation approach was applied to real MRI data. Data sets contained either a T1-weighted image only or a T1- and a T2-weighted MRI. In a first validation of the method the obtained segmentation results were inspected visually. Our findings from the inspection and some exemplary segmentation results for illustration are presented in this section.

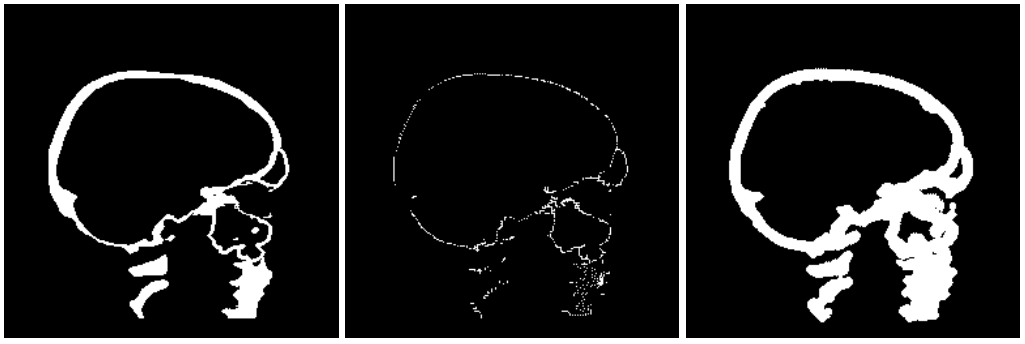

(a) Original skull mask.

(b) Mask with voxels on the medial plane.

(c) Post-processed skull mask with minimum thickness. A large value for the minimum thickness was chosen here for demonstration purposes.

Figure 8.8: Sagittal slices showing the original skull mask, the intermediate mask with selected voxels on the medial plane, and the final corrected skull mask. Please note, that a large value for the minimum skull thickness was chosen to illustrate the effect of the post-processing step. Contrary to the description in the text, here, post-processing was applied to the complete skull.

The segmentation results were generated by iterating for a fixed number of iterations $t^{\max} = 25$. The weighting for the MRF prior was set to $1/T = 0.02$. An initial labeling and initial parameters of the gray value distributions were determined according to the procedures from Section 8.5.2. Additional constraints ensuring a minimum scalp thickness and preserving the reference surfaces in the segmentation results were enforced (Section 8.5.3). The input reference surfaces were generated using the pipeline described in Chapter 6. The same pipeline was also used to preprocess the input MR images.

We already mentioned that it is not possible even for multimodal input to segment the large number of head tissues only taking the gray values into account. This is illustrated by the results obtained by the likelihood classification (Fig. 8.9, middle row). In the likelihood classification only the likelihood is evaluated to find the optimal segmentation. We observe that some tissues are quite accurately differentiated, for example, WM and GM. In other regions, however, the segmentation completely fails. The fatty SCT, for example, is labeled as WM, and most of the skull is labeled as BG. The proposed MRF classification (Fig. 8.9, bottom row), on the other hand, is able to distinguish between neighboring tissues, for example, GM and WM as good as the likelihood classification, but the classification result also accurately represents the anatomy. In particular, compact bone is not mislabeled as BG and the bright fatty tissue is distinguished from the equally bright WM.

During the visual inspection of the segmentation results we put particular focus on the skull segmentation accuracy because the skull plays a special role in volume conductor models for EEG source analysis. We find that using the proposed MRF segmentation approach the skull is in most cases very accurately segmented (Fig. 8.10 and 8.11). Some minor classification errors are, however, observed. In some cases, it might happen that parts of the muscle or scalp layer directly adjacent to the outer skull are misclassified as cancellous bone. The skull will,

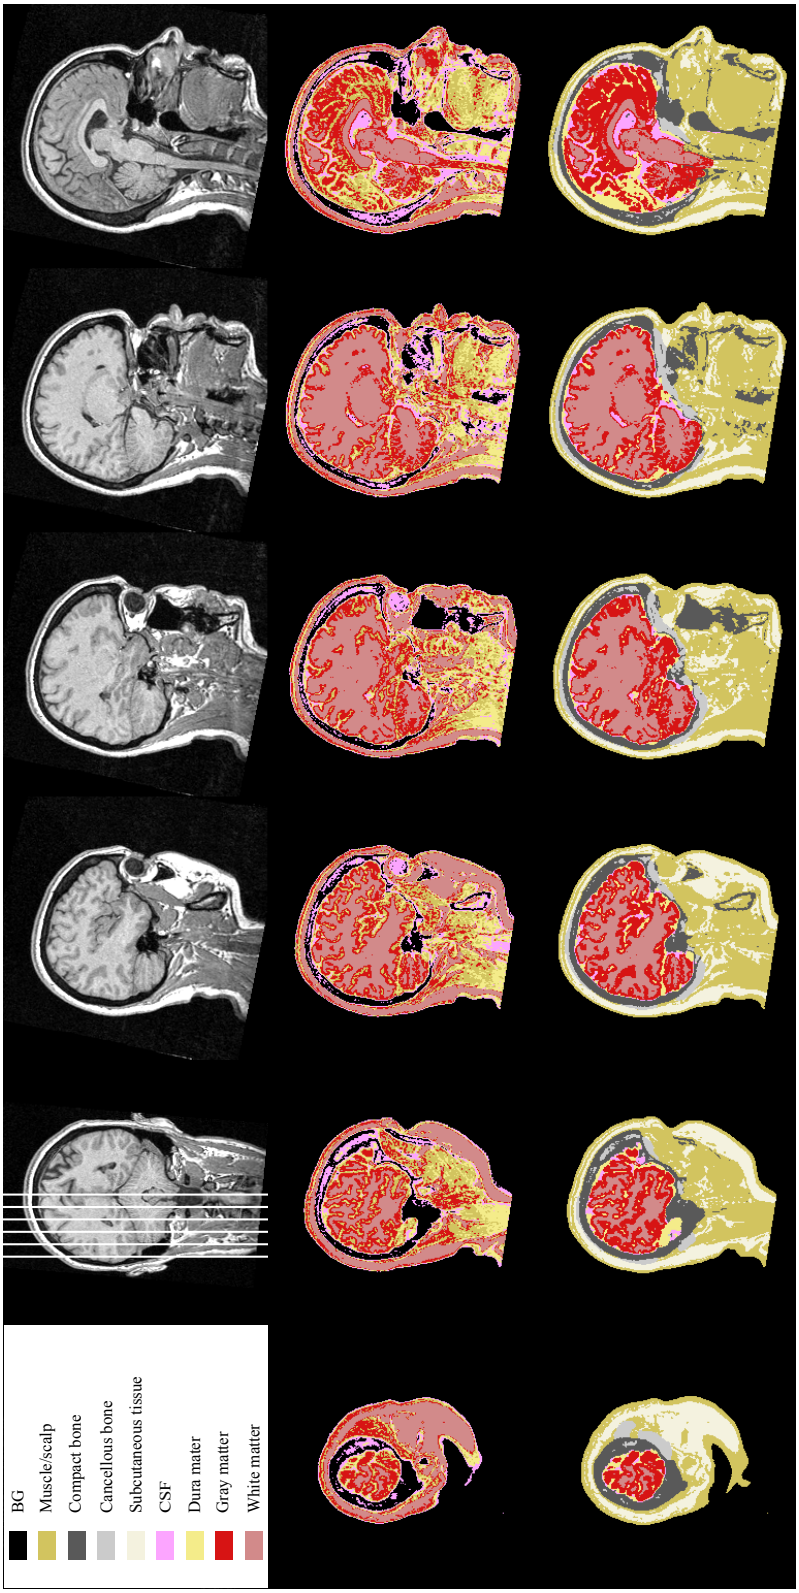

Figure 8.9: Sagittal slices of a T1-weighted MRI (top), the result of the likelihood classification (middle) and the proposed MRF classification (bottom). The coronal slice in the top row shows the position of the selected slices. The tissue legend is shown at the top left.

thus, extend to parts of the inner scalp. Additionally including a T2-weighted image into the classification improves the skull segmentation accuracy (Fig. 8.11).

For the application of EEG source analysis it is also of interest to see how accurate the tissue regions CSF, scalp and brain are reconstructed. By visual inspection the segmentation of these regions appears to be very accurate. In particular, the CSF is accurately segmented even when only a T1-weighted MRI is available (Fig. 8.12). When a T2-weighted MRI is available in addition the accuracy can even be improved (Fig. 8.13).

Our approach in principle distinguishes a total of eight tissues (and BG). With respect to these tissues it can be observed that WM and GM can be accurately segmented (Fig. 8.14).

Compact and cancellous bone are also treated as different tissues. The accuracy of the cancellous bone segmentation is, however, not optimal (Fig. 8.15). Anatomically the spongiosa layer in the area of the skull cap should be continuous interrupted only by the sutures of the skull bone. In our segmentation experiment on real T1- and T2-weighted MRI data only fragments of the true cancellous bone were correctly classified.

Convergence of the iterative MRF segmentation approach is fast. During the first iterations the segmentation result changes a lot while after approximately 25 iterations only few voxels change their labels (Fig. 8.16). The same is true for the sum of the log likelihood and the Gibbs energy $U(x)$ (Eq. 3.62) which rapidly converge towards their final values.

8.7 Discussion and Conclusion

In our experiments we found that the likelihood classification taking only the gray values into account was able to segment the input MRI data into a given number of gray value classes. Due to the overlap of the gray value distributions of the different head tissues this classification, however, does not say much about the actual spatial distribution of the head tissues. The a-priori knowledge on the anatomy of the human head which is incorporated into the proposed segmentation approach via the MRF model solves this problem. The MRF segmentation results clearly and accurately reflect the actual distribution of the head tissues.

Although the skull segmentation was overall very accurate some minor classification errors were observed. In some cases the skull extended to parts of the inner scalp. This classification error might occur in T1-weighted images with a noticeable water-fat-shift (Fig. 8.10). Here, the signal from the fatty SCT layer is displayed as shifted relative to the signal from the remaining scalp layer. In the MRI this results in a thin dark layer where the SCT was supposed to be followed by the bright layer of the shifted SCT. The local gray value profile of a voxel in the SCT resembles that of a cancellous bone voxel. In this situation the MRF prior is not able to resolve the ambiguity between SCT and cancellous bone and misclassification might happen. We will show in the next chapter that an a-priori information from a suitable tissue atlas is able to resolve this ambiguity.

Our results have shown that segmentation accuracy can be increased by incorporating T2-weighted MRI data as another image modality. The strong contrast between the CSF and the surrounding tissues in the T2-weighted MRI especially improves the delineation of the inner skull boundary but also the overall accuracy profits from the complementary information contained in the additional image modality.

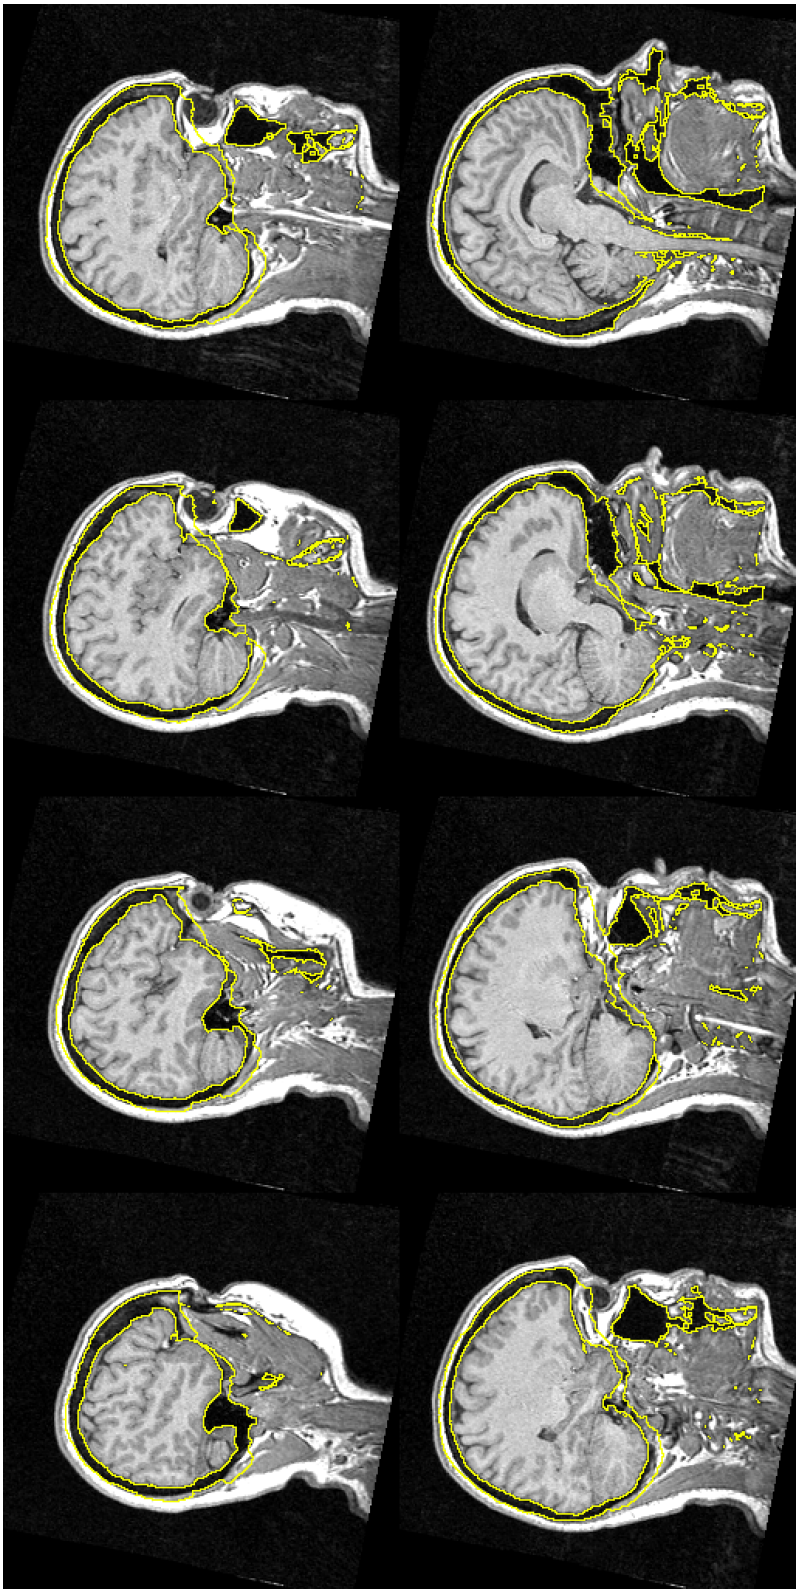

Figure 8.10: Exemplary segmentation results using the MRF approach from T1-weighted MRI data only. Outlines of the segmented skull region (compact and cancellous bone) in yellow are overlaid onto sagittal slices of the T1-weighted MRI.

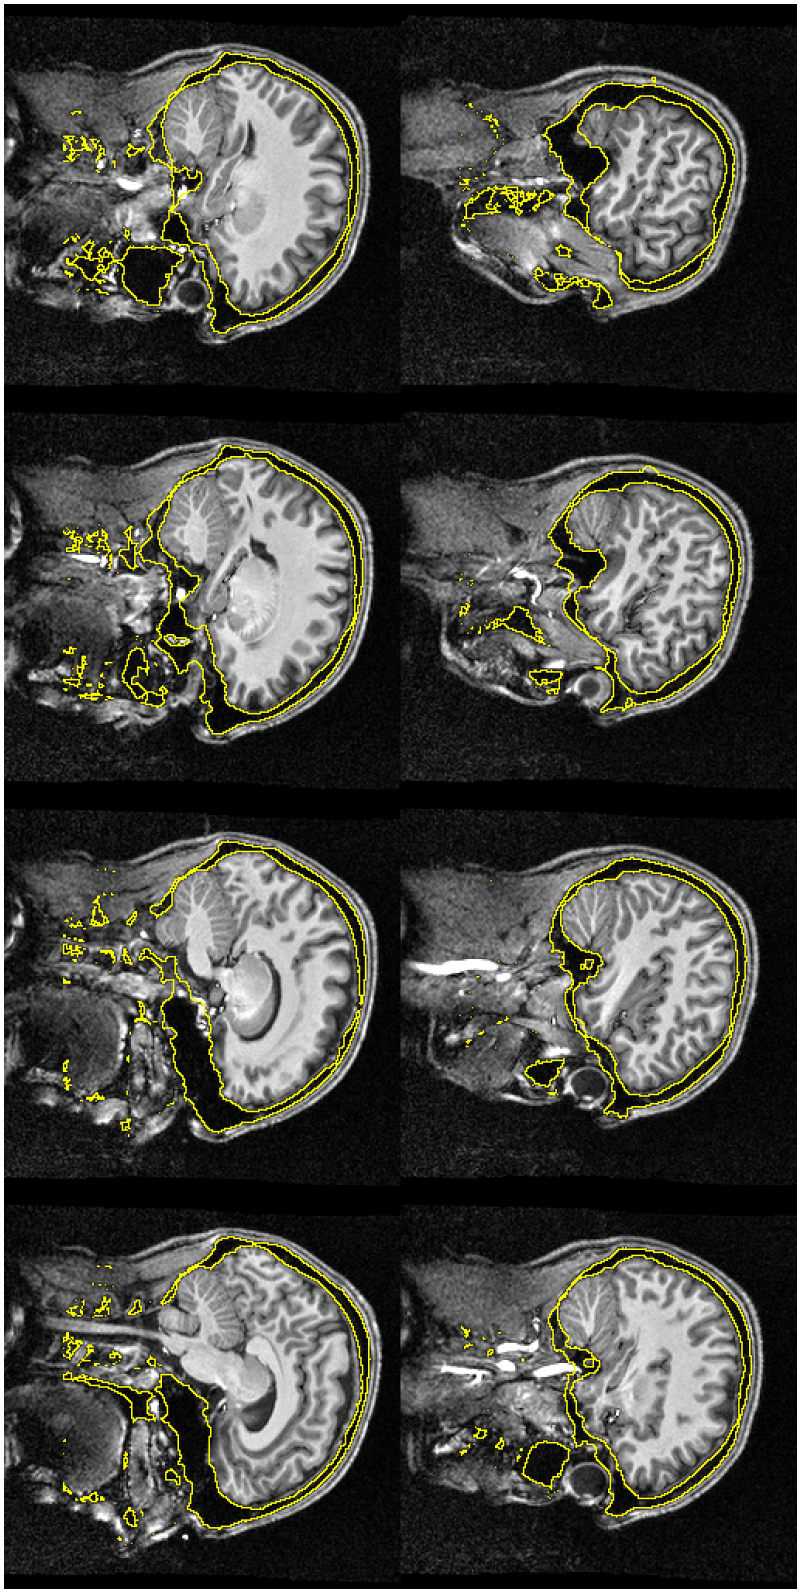

Figure 8.11: Exemplary segmentation results using the MRF approach from T1- and T2-weighted MRI data. Outlines of the segmented skull region (compact and cancellous bone) in yellow are overlaid onto sagittal slices of the T1-weighted MRI.

Figure 8.12: Exemplary segmentation results using the MRF approach from T1-weighted MRI data only. Outlines of the segmented tissue regions most relevant for the application in EEG source analysis (scalp, skull, CSF, and brain) are shown. Outlines are overlaid onto orthogonal slices of the T1-weighted MRI.

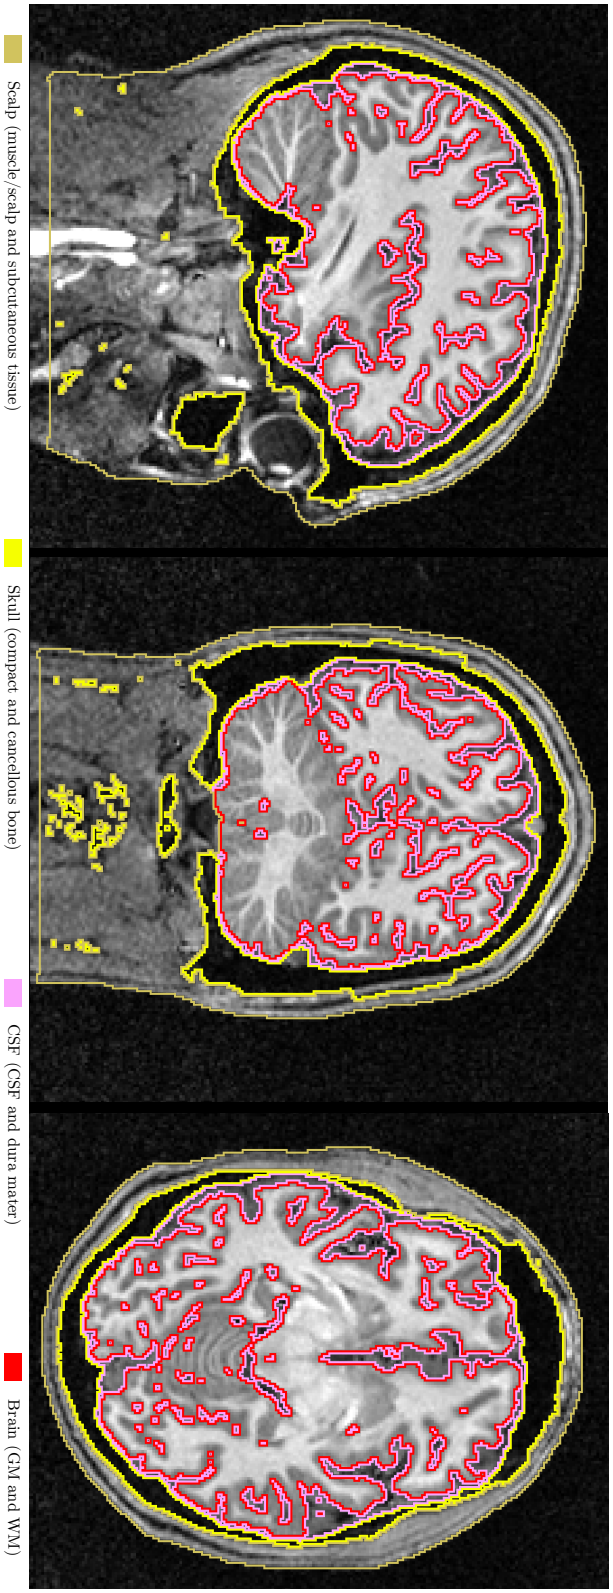

Figure 8.13: Exemplary segmentation results using the MRF approach from T1- and T2-weighted MRI data. Outlines of the segmented tissue regions most relevant for the application in EEG source analysis (scalp, skull, CSF, and brain) are shown. Outlines are overlaid onto orthogonal slices of the T1-weighted MRI.

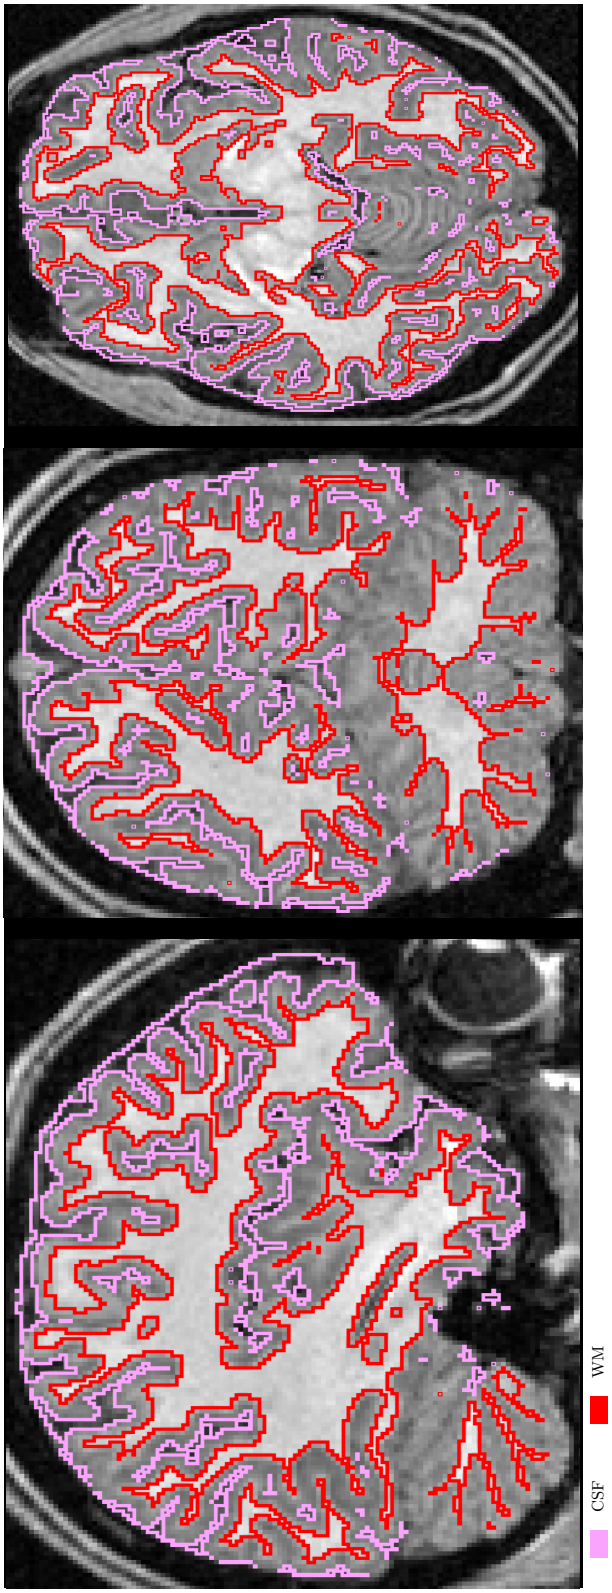

Figure 8.14: Exemplary segmentation results using the MRF approach from T1- and T2-weighted MRI data. Outlines of the segmented WM and CSF regions are shown. Outlines are overlaid onto orthogonal slices of the T1-weighted MRI.

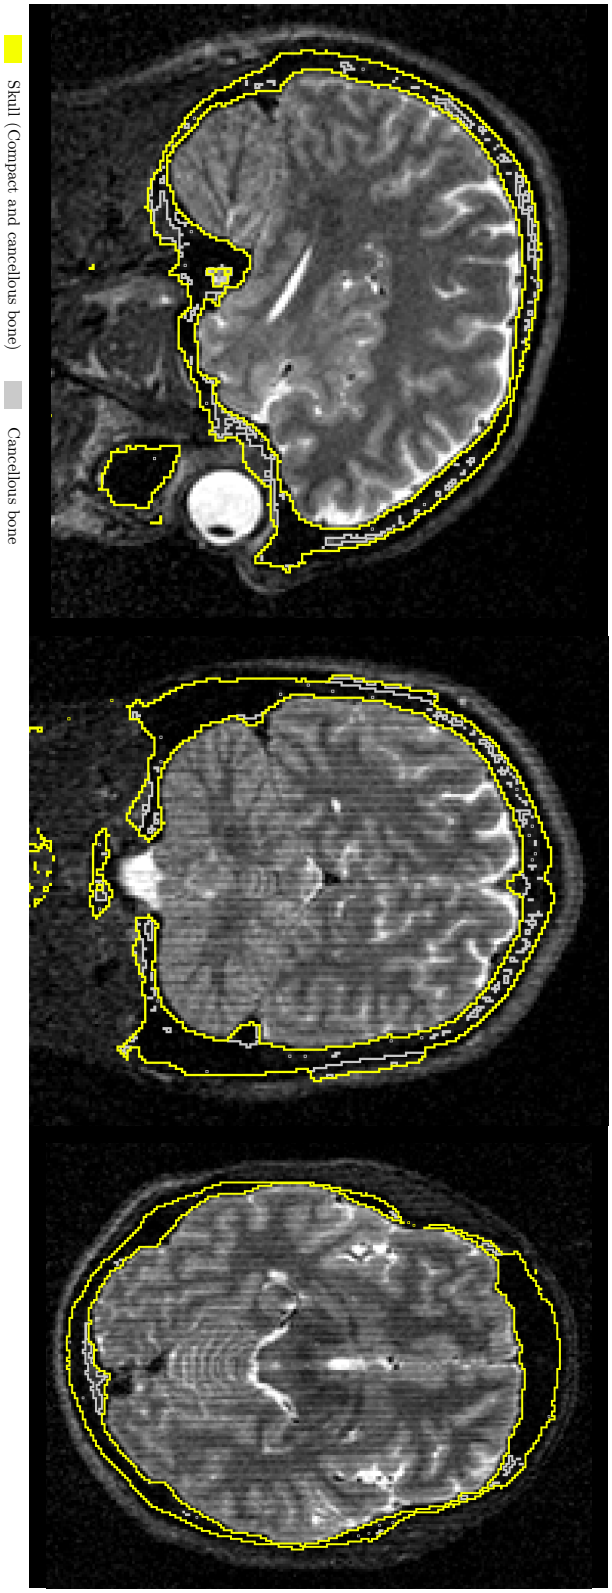

Figure 8.15: Exemplary segmentation results using the MRF approach from T1- and T2-weighted MRI data. Outlines of the segmented skull (compact and cancellous bone) and cancellous bone regions are shown. Outlines are overlaid onto orthogonal slices of the T2-weighted MRI.

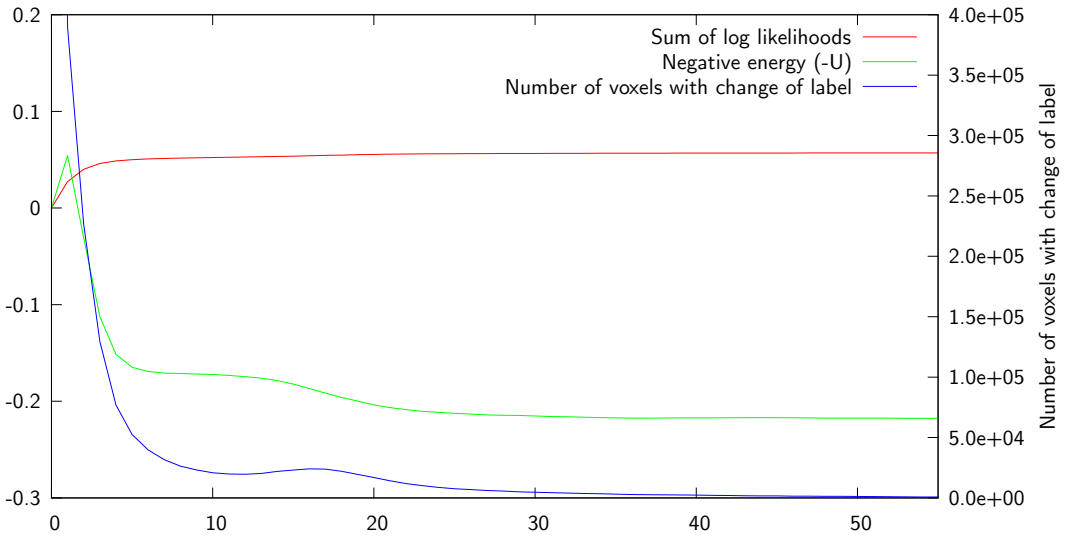

Figure 8.16: The sum of the log likelihood, the negative energy $-U(x)$ and the number of voxels which changed their labels during the iterations of the MRF segmentation. The sum of the log likelihood and the negative energy $-U(x)$ are given as normalized relative changes with respect to the respective values at the first iteration.

Chapter 4 investigated the influences of skull segmentation inaccuracies on EEG source analysis. We can relate the findings from that chapter to the observed skull segmentation results presented here. No erroneous holes were observed so that errors as in the first test model series are not expected. We must, however, expect the minor forward and inverse errors which were seen for the test models with skull thickness misspecifications, and for test models where the sinuses were ignored. The skull base in the segmentation provided by our proposed approach seems to be reasonably well approximated. Models constructed from the segmentation will, thus, most likely not exhibit the severe influences on source analysis of simplifying the inferior skull, or skull and scalp.

Overall the results showed that all tissues relevant for EEG source analysis could be distinguished with good accuracy. In addition, the segmentation approach also yielded an accurate WM and GM segmentation. Although it is not commonly done, this information might be exploited in volume conductor models.

According to our recent study [39] it is beneficial to distinguish between compact and cancellous bone in EEG head models. The segmentation with respect to these tissues, however, was not optimal in the presented results. This might be caused by the bad SNR and the sub-optimal recording sequences used to acquire the MRIs. The T1-weighted image (Fig. 8.11) was, for example, recorded with fat-suppression so that it did not contain any information to facilitate the differentiation between the compact bone and the fatty cancellous bone. The T2-weighted image was measured without fat suppression. It was, however, so noisy that it was still difficult to distinguish the layers of the skull. Better results for the segmentation of the skull spongiosa are expected for data sets with a larger SNR.

With respect to the convergence it was surprising at first glance that the negative energy $-U(x)$ of the labeled image x decreases and settles at an value around 20% smaller than for the initial state (Fig. 8.16). This can be explained with the initial labeling which was constructed according to the same anatomical assumptions on which the MRF model is based. The initial configuration was, thus, already fitting very well to the MRF model and any changes in the segmentation due to the likelihood function only decreased the fit to the model.

In conclusion, judging by the visual inspection the proposed segmentation approach produced accurate segmentation results which seem to be suitable as a basis for the construction of EEG head models. The good accuracy and stability of the method was achieved by incorporating the MRF model which encoded a-priori information on the anatomy of the head. Some remaining classification errors that were observed might be solved by exploiting additional a-priori information.

9 An Atlas-based Anatomical Prior for Head Tissue Segmentation

9.1 Introduction

Using the MRF model described above we are already incorporating important a-priori information on the anatomy of the head. Nevertheless, in some cases we still observe ambiguities and it is not clear how a voxel shall be labeled. In a T1-weighted image measured with fat suppression (Fig. 3.6a), for example, it is not possible to unambiguously label the voxels in the SCT region. Based on the MRF model and the likelihood voxels in this region might either be classified as SCT or as compact bone. In this chapter, an approach is described to generate atlas-based a-priori probability maps. It is shown that incorporating these maps helps resolving ambiguities during the classification.

In our approach first local tissue probability distributions depending on the distances to some reference surfaces are empirically constructed. These local probability distributions constitute our tissue atlas. Individual a-priori probability maps are then computed by projecting the probability distributions to individual reference surfaces. In the following, the creation of the atlas and the projection of the probabilities will be described in more detail. Examples of individual a-priori probability maps generated by our approach will be presented. In addition, we will describe how the a-priori probability maps can be incorporated into the two segmentation approaches presented in Chapters 7 and 8. Finally, segmentation results obtained using the two proposed segmentation approaches with and without incorporating the a-priori probability maps will be compared to show the efficacy of the generated atlas-based a-priori probability maps.

9.2 Generation of the Atlas

This section describes the generation of the tissue atlas. Details on the underlying data will be given. The construction of the reference surfaces and the computation of the distance maps will be discussed, and finally the computation of the local tissue probability mass functions will be presented.

Basis for our tissue atlas are reference segmentations from the BrainWeb data base [6]. In total 20 labeled MRIs of subjects aged between 24 and 37 years were used. The generation of the labeled MRIs is described in [6]. In summary, the labeling was done as follows. T1-, T2- and PD-weighted MRIs with very good SNR and a magnetic resonance angiography (MRA) were recorded for each of the subjects. The images were preprocessed, registered to each other and up-sampled to $0.5 \times 0.5 \times 0.5\,\mathrm{mm}^3$ resolution. A minimum distance fuzzy classification was then applied to the multimodal image data. The classification result was combined with scalp, inner skull and brain masks generated using the BET [134] and BET2 [67] software. The

BrainWeb labels	Atlas labels	BrainWeb labels	Atlas labels
BG	BG	Bone marrow	Skull
Skin / Muscles	Scalp	Dura mater	CSF / dura mater
Muscles	Scalp	CSF	CSF / dura mater
Fat	Scalp	GM	Brain
Tissue around fat	Scalp	WM	Brain
Skull	Skull	Blood vessels	Ignored

Table 9.1: Correspondences between tissues in the labeled BrainWeb images and the tissues differentiated in the proposed atlas.

blood vessels were extracted by thresholding the MRA using a manually chosen threshold value. Finally labeled volumes differentiating eleven tissues (and BG) were obtained. The following twelve regions were differentiated: BG, GM, WM, CSF, skull, bone marrow, dura mater, fat, tissue around the fat, muscles, skin/muscles, and blood vessels.

Our atlas distinguishes between five tissue regions. These regions are BG, scalp, skull, CSF/dura mater and brain. Distinguishing between a larger number of regions will have no significant benefit, due to the limited spatial resolution of our atlas. The labeled images from the BrainWeb data base are therefore first relabeled to the five tissues considered in our atlas. The correspondences between tissues in the BrainWeb images and the tissues in our atlas are listed in Table 9.1. Note, that the blood vessels cannot meaningfully be mapped to one of the tissue regions that are incorporated into the atlas. Thus, voxels labeled as blood vessels are ignored during the generation of the atlas, that is, they are treated as if they do not exist.

In the next step, reference surfaces are reconstructed from each of the original reference images (i.e., the reference images differentiating between twelve tissues and regions). As reference surfaces we choose the WM and the scalp surface. These surfaces are well suited as reference surfaces for our atlas because they confine the tissue regions of interest (namely scalp, skull, CSF, and GM) on two opposing sides. The distances to these reference surfaces are, thus, very informative with regard to the occurrence of the relevant tissues in our segmentation problem. A scalp reference mask is constructed by marking all non-background voxels and filling the holes in the resulting mask. By selecting all WM voxels in the labeled image an initial WM reference mask is constructed. Holes in the WM mask are also filled. The distances from each voxel to both of the reference surfaces are then computed using the FMM [128].

Our tissue atlas consists of local tissue probability mass functions depending on these distances. Using the absolute distances to the reference surfaces as parameters on which the tissue probabilities depend can, however, be problematic. This is due to the large differences in head size for different subjects, especially, when considering the heads of children as compared to the heads of adults. Furthermore, our reference data is based on MRIs of adult subjects only, so that the empirically determined probability distributions depending on the absolute distances will not be able to faithfully reflect the tissue probabilities for the much smaller head of a child. As a solution to this problem, we are also estimating tissue probability distributions depending on distances to the reference surfaces, which are normalized in a particular manner.

We chose to normalize the distances to the reference surfaces to the local distance between

the WM and scalp surfaces. This choice is based on the assumption that the relation between the total thicknesses of the layers scalp, skull, CSF, and GM for two different subjects is the same as the relation between the distances from WM and scalp in these two subjects.

To compute the normalized distances first the local distance between the WM and scalp reference surface has to be computed. This is done as follows: The distance from the WM surface is computed for all image voxels directly adjacent to the scalp reference surface using the FMM [128]. Adjacent here means that the voxel either is a background voxel with at least one neighbouring voxel inside the scalp reference surface or vice versa. These distance values are then projected to the vertices of a projection surface.

As a projection surface we are using the inner bone surface derived from a standardized, averaged model of the human head. The inner bone surface was chosen because it passes close by most voxels in the scalp and skull regions, which are of particular interest in our segmentation problem. The projection surface, as shown in Figure 9.1, consists of 642 vertices.

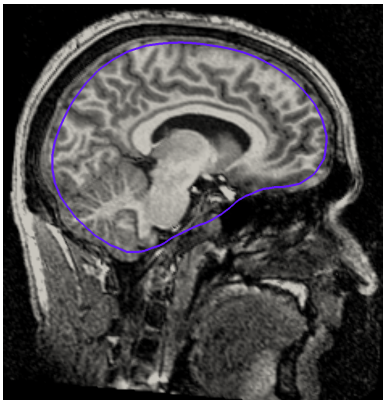

(a) Cut through the projection surface overlaid over a sagittal slice close to the mid-sagittal plane of a subject's T1-weighted MRI.

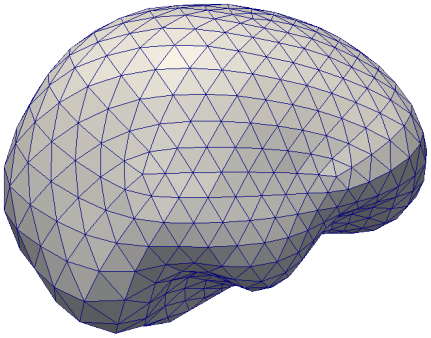

(b) Rendering of the projection surface.

Figure 9.1: The projection surface. The local tissue probability mass functions and the local thickness of the layers between scalp and WM are projected to this surface.

For a single vertex j then all voxels of the image grid are determined that fulfill two conditions. First, the voxel is directly adjacent to the scalp reference surface, and second, the voxel is closer to vertex j than to any other vertex of the projection surface. We now compute the average of the distance values for all voxels fulfilling these conditions to obtain the local distance from scalp to WM at vertex j. This local distance from the WM reference surface to the scalp reference surface can also be interpreted as the local thickness of the layers situated between scalp and WM. Normalization of the distances of a voxel to the two reference surfaces is done by dividing the absolute distance values by the local distance between the reference surfaces computed as described above.

The next step is the computation of the local tissue histograms. We will start by describing how this is done for a single reference data set (Fig. 9.2).

For each vertex of the projection surface it is counted how often each of the tissues occurs in

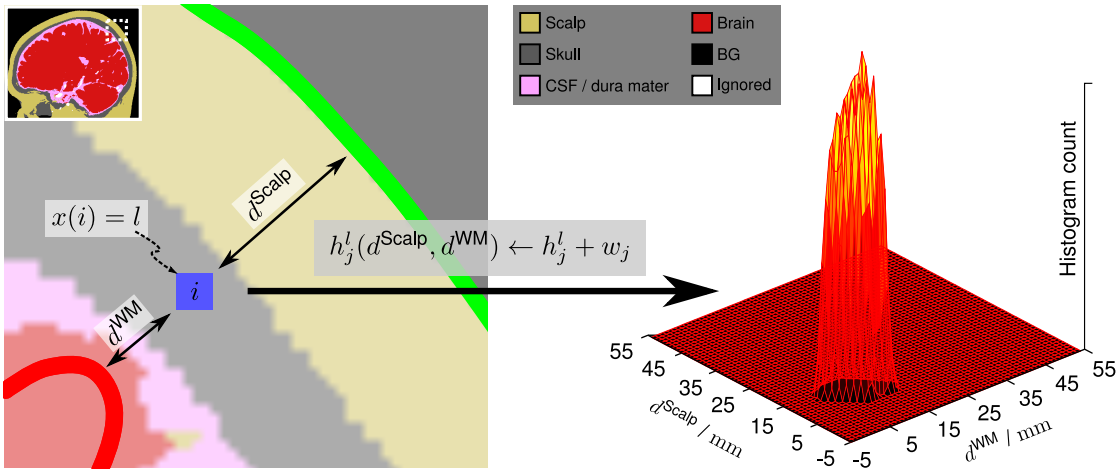

Figure 9.2: The generation of the tissue atlas for the head model segmentation. For each reference data set local histograms are computed for the occurrence of different tissues depending on the distances to the reference surfaces.

the vicinity of the vertex depending on the distances to the two reference surfaces. To do this, we successively consider each of the voxels inside of a certain distance range from both reference surfaces. We then determine the five vertices of the projection surface which are closest to the current voxel. Let us now assume that the currently considered voxel has the label l in the reference data set. We then increase the histogram counts $h_j^l(d^{\text{scalp}}, d^{\text{WM}})$ for each of the closest vertices of the projection surface. Histogram counts for vertex j are increased by a weighting factor w_j depending on the distance d_j of the current voxel to the vertex of the projection surface as defined in Equation 9.1.

$$w_j = \frac{1/d_j^2}{\sum_{k=1}^{5} 1/d_k^2} \tag{9.1}$$

The sum in the denominator runs across the five closest vertices $k = 1, \ldots, 5$. By increasing the histogram count not only for the single closest vertex but for the five closest vertices of the projection surface we obtain a smoother distribution of histogram values across the projection surface.

Histogram counts for all 20 reference data sets are finally summed up and histograms are normalized. We thusly obtain two-dimensional probability distributions $p_j^l(d^{\text{scalp}}, d^{\text{WM}})$ for each vertex j of the projection surface telling us the probability to find a voxel belonging to one of the differentiated tissues l at given distances d^{scalp} and d^{WM} to the scalp and WM reference surfaces, respectively.

The resulting empirically determined joint probability mass functions unfortunately are not defined for all pairs of $(d^{\text{scalp}}, d^{\text{WM}})$. This is due to the fact that for these pairs in the reference data sets there are simply no voxels located at the same time at a distance of d^{scalp} from the scalp reference surface and at a distance of d^{WM} from the WM reference surface. Imagine, for example, that the distance between the reference surfaces is approximately 20 mm in some area

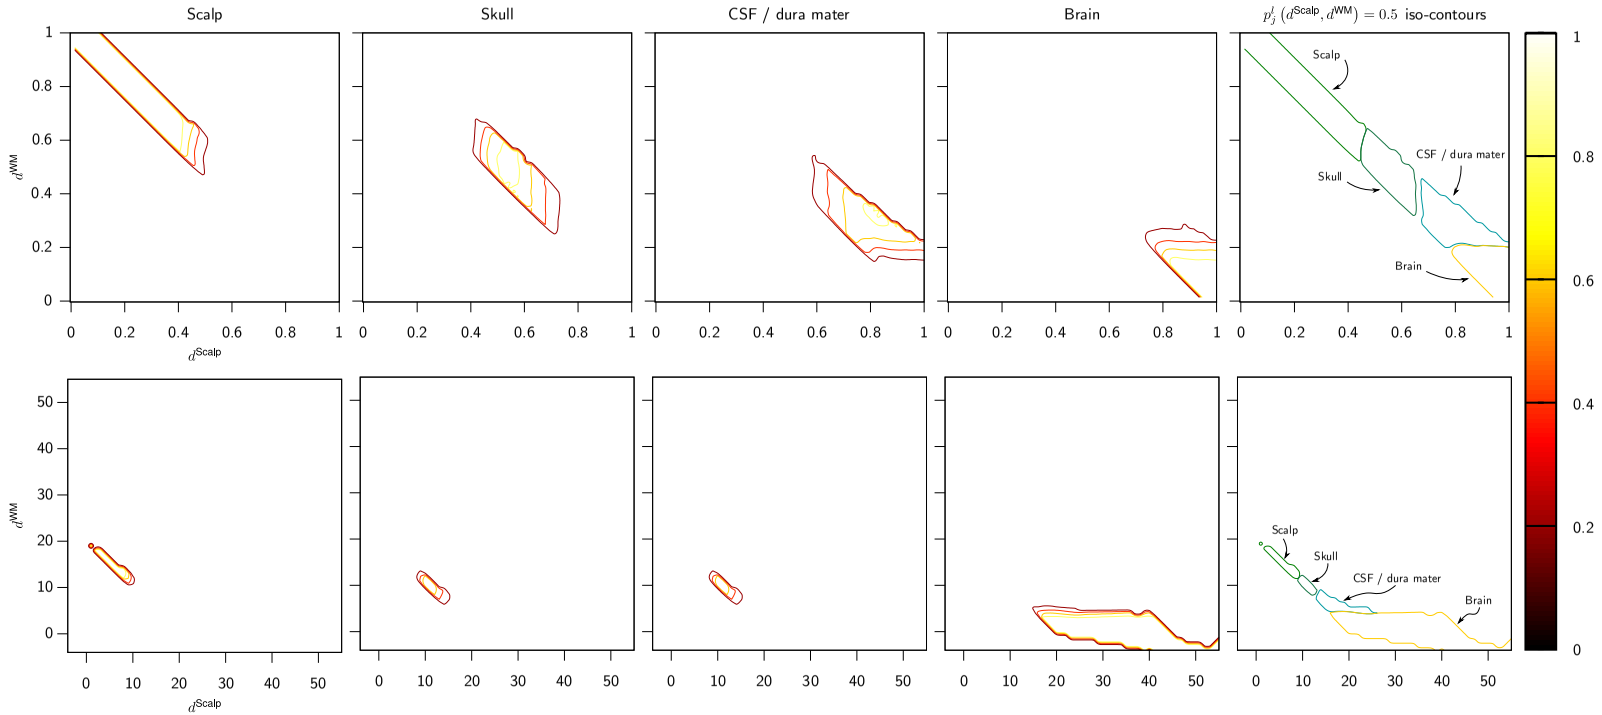

Figure 9.3: Contours of the probability mass functions depending on the normalized distances (top) and the absolute distances (bottom) to the scalp and WM reference surfaces. Contours are plotted at values of 0.2, 0.4, 0.6, and 0.8. The probability mass functions for a single vertex in the left parietal region is shown. In the last column the contour lines at $p_j^l = 0.5$ for the tissue probabilities of scalp, skull, CSF and brain are shown in one plot.

of the head. It is then not possible to find a voxel which is at a distance of 15 mm from both reference surfaces. For this point we, thus, cannot empirically determine the tissue probability mass function This problem of missing data is dealt with during the projection of the tissue probabilities to the individual subject's reference surfaces (Sec. 9.3).

Marginal probability mass functions only depending on the distance to one reference surface can be computed directly from the joint distributions.

The probability distributions for the tissues scalp, skull, CSF, and brain at one vertex in the region of the left parietal bone are shown in Figure 9.3. It can be seen that the empirical tissue probability distribution for this vertex is only defined in a narrow interval of values for the scalp and the WM distances. We also observe that the scalp and skull probabilities change continuously from close to one to zero when varying the scalp distance. They, however, change rather abrupt when varying the WM distance. The opposite is true for the brain probability. This shows that the WM distance does not contain much information on the tissue probabilities in the region close to the scalp surface. For a fixed small scalp distance the tissue probabilities are nearly constant with respect to the WM distance. The probabilities then drop to zero where they could not be determined empirically.

9.3 Generation of Individual A-Priori Probability Maps

Using the process described in Section 9.2 we generated an atlas describing how likely it is that certain head tissues occur depending on the distances to the scalp and WM reference surfaces. The atlas consists of local tissue probability mass functions at each vertex of the projection surface.

To aid the segmentation for an individual subject the atlas is now projected to the individual subject's reference surfaces to obtain tissue probability maps defined on the image grid. Figure 9.4 illustrates how the projection works.

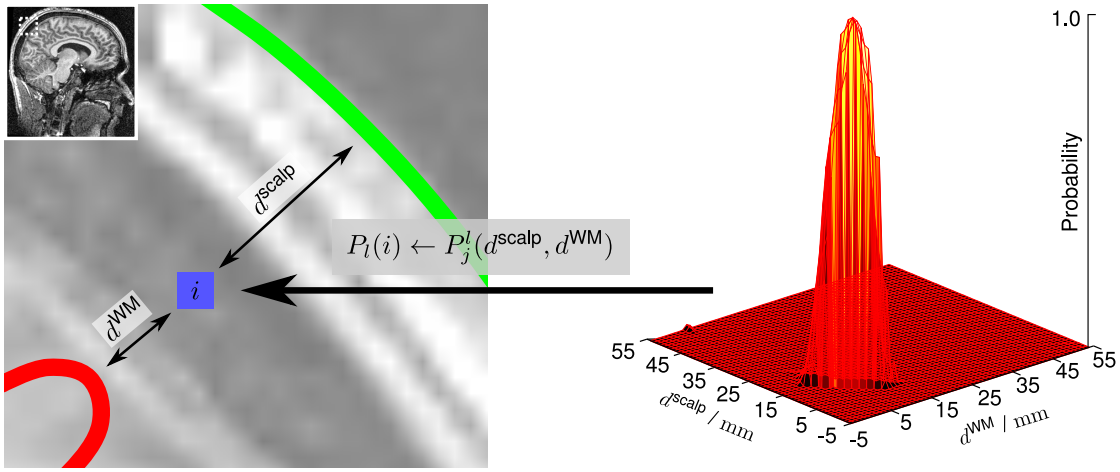

Figure 9.4: Projection of the atlas of the anatomy of the human head to the individual subject's reference surfaces.

Probabilities are projected to the image grid voxel by voxel. For each voxel, first, the distances to the individual reference surfaces d^{scalp} and d^{WM} are obtained. Next, the five vertices of the projection surface which are closest to the current voxel i are searched. Now, for each of the vertices j the tissue probabilities $P_j^l(d^{\text{scalp}}, d^{\text{WM}})$ for the distances to the reference surfaces at the current voxel are retrieved from the atlas. At this point it might happen that the empirically determined probability mass function for a vertex of the projection surface does not contain any information for the given distance pair $(d^{\text{scalp}}, d^{\text{WM}})$ as discussed at the end of Section 9.2. In this case, the projection vertex is simply ignored during the computation of the tissue probabilities at the current voxel. If none of the five closest vertices contains valid information then an uninformative tissue probability is set. If at least one vertex contains valid information then the final a-priori tissue probabilities are obtained by computing the weighted average of the probabilities for the closest projection vertices with valid data. For the weighting the same weighting factor as in Equation 9.1 is used.

In the area of the skull cap (region A in Fig. 8.2) the regions of interest (i.e., skull, scalp, CSF) form a layered structure which is encased by the two reference surfaces. The distances to the reference surfaces are thus very informative with regard to the occurrence of a certain head tissue. In the area of the skull base (region B in Fig. 8.2), however, the distance to the scalp reference surface is not very informative as it varies widely and is also influenced by, for example, how far the reconstructed scalp reference surface extends downwards. For this reason, the tissue a-priori probabilities in the area of the skull base are derived not from the joint probability mass functions depending on the distances to both reference surfaces, but from the marginal probability mass functions depending only on the distance to the WM reference surface. In the same sense, neither the distance to the WM reference surface, nor the distance to the scalp reference surface are informative in the inferior head and neck area (region C in Fig. 8.2). Thus, tissue probabilities for this area are not derived from the atlas.

9.4 Incorporation of A-Priori Probability Maps into Proposed Segmentation Approaches

Using the method outlined in the last section it is possible to project the probabilistic atlas based on the distances to two reference surfaces to an individual MRI. In this way, probability maps are obtained with probabilities $P^{\text{atlas}}(x_i = l)$ for the occurrence of the tissues and regions BG, scalp, skull, CSF and brain.

These probability maps can be incorporated into the active contours approach (Chap. 7, Eq. 8.2) and the MRF segmentation approach (Chap. 8, Eq. 7.1) for the segmentation of volume conductor models of the human head. Including the atlas-based a-priori probability into the approaches is straight-forward.

For the MRF approach the single-site clique potentials were set to the negative logarithm of an atlas-based probability (Eq. 8.2). The MRF approach distinguishes in total eight different tissues. It is, thus, not possible to directly identify the atlas-based probability described in this chapter with the probability in Equation 8.2. The labels distinguished in the atlas, however, are supersets of the more specific labels used for the MRF segmentation. As an example, the MRF uses the labels *compact bone* and *cancellous bone* which are subsets of the more generic *skull* label of

the atlas. Assuming that the probability distributions for all the labels in the same superset are uniform we can identify the probabilities for the labels in the subsets with the probability for the superset. In the MRF we will, for example, use the skull probability distribution to obtain a-priori probabilities for the occurrence of the compact and cancellous bone. A mapping from the atlas probabilities to the labels of the MRF approach is shown in Table 9.2.

MRF label	**Atlas label**	**MRF label**	**Atlas label**
BG	BG	Dura mater	CSF
Muscle / scalp	Scalp	CSF	CSF
SCT	Scalp	GM	Brain
Compact bone	Skull	WM	Brain
Cancellous bone	Skull		

Table 9.2: Mapping from the more specific labels used in the MRF segmentation approach to the supersets of labels used for the construction of the atlas.

The active contours segmentation approach also offers the possibility to incorporate a-priori probability maps. Equation 3.39 is the atlas-based prior formulated in the framework of the active contours segmentation. In this approach either four or five different regions are differentiated. If five regions are differentiated a one to one mapping from the projected probability maps to the regions is possible. As discussed above it is also possible to only use the four regions BG, scalp, skull and intracranial space, where the intracranial space consists of CSF, GM, and WM. In this case, the probabilities for the occurrence of CSF and brain are simply added up to obtain a-priori probabilities that a voxel belongs to the intracranial space.

9.5 Experiments

By projecting the atlas to the individual reference surfaces smooth probability maps are obtained. This can be seen from the rendering of the isosurface of the skull probability (Fig. 9.5). The holes in the rendered isosurface can be explained by the missing data for some pairs of scalp and WM distance. When comparing the projected probability maps to the underlying MRI of the subject we can also see that the $P_j^l = 0.5$ isolines agree well with the tissue region boundaries in the MRI. For the skull and scalp probability we also observe that the probability distribution is rather sharp in some areas, for example, in the occipital area at the back of the head. Please note, that the a-priori probability is not projected everywhere from the probabilistic atlas. The probabilistic atlas is only defined for absolute distances of -5 to 50 mm and normalized distances of 0 to 1 relative to the scalp and WM reference surfaces. This means, for example, that the probability inside of the WM reference mask ($d^{\mathrm{WM}} < -5$) is not computed from the atlas (Fig. 9.5). In these regions the a-priori tissue probabilities are either uninformative or other knowledge on the anatomy of the head is exploited (cf. Sec. 8.5.3).

For the active contours segmentation we find that incorporating the atlas-based a-priori probabilities increases the accuracy of the skull segmentation. Segmentation results without taking the atlas into account sometimes revealed minor mis-classifications where parts of the low-signal

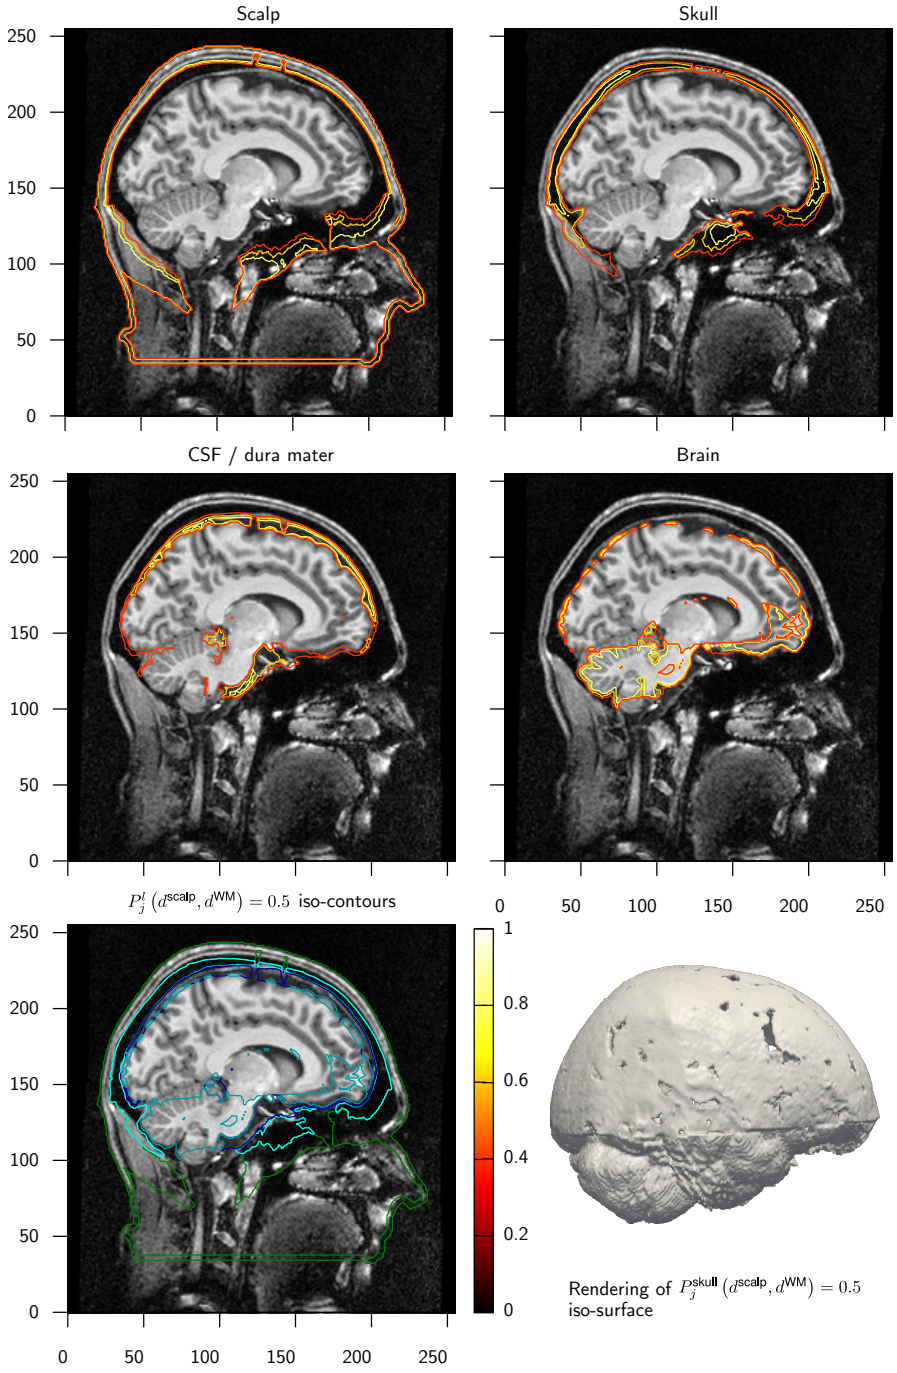

Figure 9.5: Probability maps generated by projecting the atlas to the individual reference surfaces of a subject. Probability distributions are indicated by isolines at 0.4, 0.6 and 0.8. In the lower left panel the isolines at 0.5 for the four tissues scalp (green), skull (light blue), CSF (blue), and brain (mint green) are presented. In addition, the lower right panel shows the rendering of the skull probability isosurface at 0.5.

SCT regions were falsely identified as skull (Fig. 9.6a). This does not happen when exploiting the anatomical a-priori information contained in the atlas as shown in Figure 9.6b.

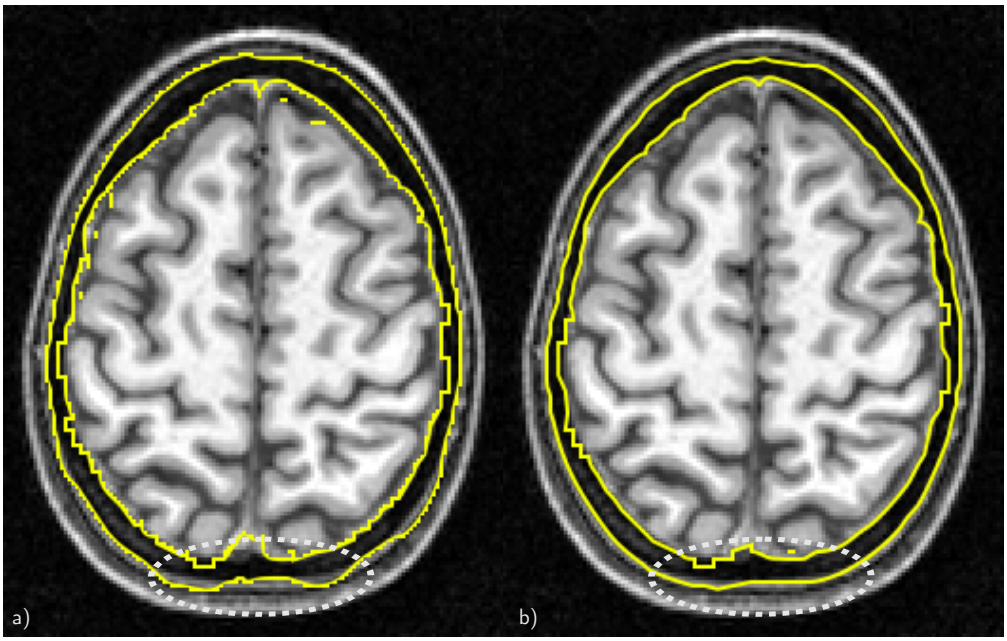

Figure 9.6: Comparison of a) the skull segmentation using the active contours segmentation approach described (cf. Chap. 7) and b) in addition employing the atlas-based a-priori probability maps. The skull outlines for both skull segmentations are shown overlaid over the same axial slice of the T1-weighted MRI of the subject.

For the MRF segmentation it was observed that in some cases parts of the subcutaneous scalp layer was misclassified as skull (Fig. 9.7a). The reason for this was the low signal of this layer in MRIs recorded with reduction of the fat signal. When incorporating the atlas-based a-priori information this misclassification does not occur. The SCT voxels are now correctly labeled as scalp (Fig. 9.7b).

9.6 Discussion and Conclusion

In this chapter, we described the construction and application of a tissue probability atlas for the segmentation of volume conductor models for EEG source analysis.

The resulting individual probability maps were found to fit well to the individual MRI (Fig. 9.5). This already indicates that the atlas will be able to facilitate the tissue segmentation. The projected maps were spatially smooth (see rendering of skull iso-probability surface in Fig. 9.5). This might be surprising at first. Discontinuities might be expected because the probabilities are projected from the sparse vertices of a coarse triangle mesh. The smoothness can be explained by the fact that we are not interpolating from a single vertex but from five different vertices. This avoids discontinuities.

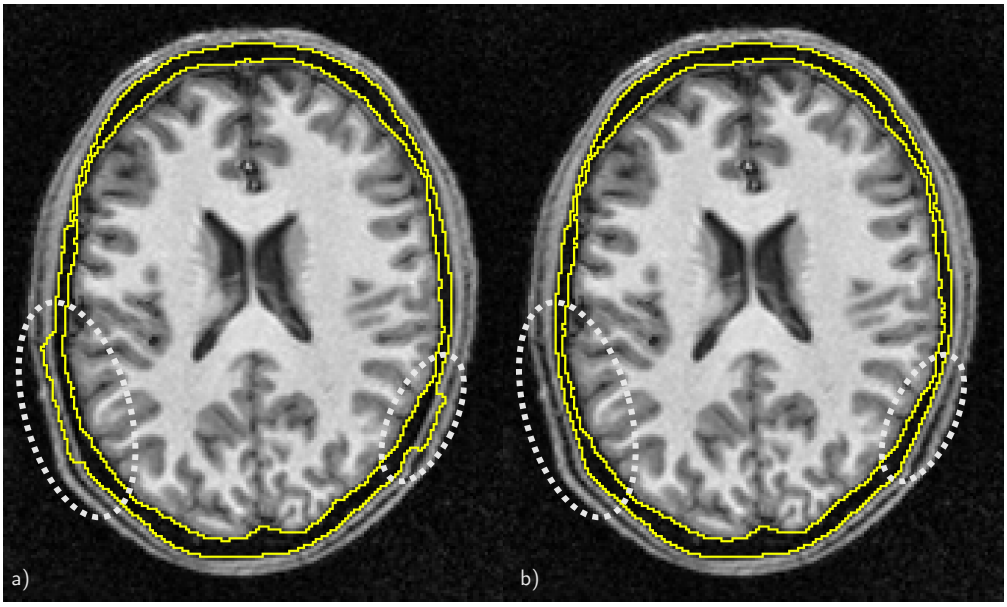

Figure 9.7: Comparison of a) the skull segmentation using the MRF segmentation approach described (cf. Chap. 8) and b) in addition employing the atlas-based a-priori probability maps. The skull outlines for both skull segmentations are shown overlaid over the same axial slice of the T1-weighted MRI of the subject.

A problem that occurred during the construction of the atlas was the presence of empty distance bins for which no data was available to heuristically determine the tissue probabilities. This led to some minor regions in the individual MRI where only an uninformative prior for the tissue probabilities could be used (Fig. 9.5). These minor regions did not have an observable influence on the accuracy of the segmentation results. Extending the set of labeled images from which the atlas was constructed might reduce the problem of empty bins.

When using the atlas-based a-priori probability maps a clear improvement of the skull segmentation accuracy could be observed for both proposed segmentation methods. Regarding the MRF segmentation approach without the atlas, that is, using only the MRF model some ambiguities could not be resolved. For example, in MRIs with fat reduction the ambiguity between muscle and cancellous bone, and SCT and compact bone (Fig. 9.7a). The atlas provides the information that the voxel close to the scalp surface more likely belongs to SCT than to cancellous bone. This resolves the ambiguity and the voxels are classified correctly (Fig. 9.7b). A similar improvement could also be observed for the active contours segmentation approach.

The atlas was constructed based on the labeled images of a collective of healthy subjects. It therefore represents the typical healthy anatomy of the human head. Due to illnesses or surgery, however, the anatomy of some patients might substantially deviate from this. When applied to MRIs of these patients it must be expected that the segmentation will not be as accurate as for subjects with a more typical anatomy. Lesions or other anatomical abnormalities will most likely not be segmented correctly.

Employing an empirical tissue probability distribution depending on the distances to certain

reference surfaces was also proposed by Kapur et al. [71] as a means to improve segmentation of MRI images. In contrast to our approach only a single, global probability distribution was used. In practice, Kapur et al.'s method would be less informative and , thus, not able to account for differences in, for example, skull thickness in different parts of the head.

Other approaches to incorporate atlas-based information into the segmentation exist (e.g., [31, 48, 103, 165, 45]). Many of these use non-linear registration algorithms to align the individual MRI with a template MRI which corresponds to a tissue atlas. Using the inverse transformation the information from the probabilistic atlas is then projected to the individual MRI. We decided to use a different approach for projecting the atlas to the individual MRI because of two reasons. First, it is known that non-linear coregistration works well on the brain [34] but it is not clear if it is also able to precisely align the scalp and skull regions. Gray value characteristics of the scalp and skull regions, for example, can differ substantially between different T1-weighted MRIs (fat reduction vs. no fat reduction) whereas the contrast between GM and WM will be similar in these images. This might cause problems for the projection of atlas-based a-priori knowledge on the anatomy using non-linear coregistration approaches in combination with only gray value based registration measures. And second, the effort of implementing a reliable method for non-linear registration is much larger than the implementation effort of our approach.

In the future the atlas might be further improved by incorporating more data. Additional labeled images for the construction of the atlas could be obtained by first using the proposed segmentation approach to generate already rather accurate segmentation results. These could then be manually corrected with not too much effort. The manually corrected images could finally be incorporated into the atlas.

Some open questions also remain with respect to the normalization of the distances during the generation and the projection of the atlas. Instead of normalizing the distances to the local WM to scalp distance other more suitable ways for the normalization might exist. Future work will be concerned with investigating these different ways of normalization.

In conclusion, we described the creation and projection of a probabilistic tissue atlas for the segmentation of EEG head models. In experiments on real data it could be shown that the a-priori information provided by the atlas is able to improve the skull segmentation accuracy. In this way, it enables us to build more accurate volume conductor models for EEG source analysis.

10 Validation of the Proposed Segmentation Approaches

10.1 Introduction

In the above chapters two approaches for the segmentation of volume conductor models from MRIs of the human head have been presented. Experiments on real data have already hinted at promising performance for both approaches. However, more thorough validations must be performed to be able to make a more meaningful statement regarding the accuracy of the proposed methods. In this chapter, validation studies are presented which aim at assessing the accuracy of the proposed segmentation approaches. Validation is done by comparing the automatic segmentation results to reference segmentations. We consider two types of reference segmentations: reference segmentations derived from manual segmentations and reference segmentations constructed from CT image data.

In the context of EEG source analysis the skull plays a special role as its conductivity is low, and as it is situated between the source space and the sensors. For EEG source analysis it is, therefore, especially important to accurately model the skull. For this reason, we will focus on investigating how accurately the skull is segmented from the input MRIs.

In the following sections, it will first be described how the proposed segmentation approaches were used to generate the segmentation results for the validation. In addition, we also describe the competing method that is included and how the competing method was used to generate skull segmentations. Next, the study design for the validation of the proposed segmentation approaches against manual reference segmentations is presented. The validation procedure using reference segmentations derived from CT images is described in the following section. Finally, results of the two validation studies are presented and discussed.

10.2 Generation of Automatic Segmentation Results

In our validation study we included results obtained using the active contours and the MRF segmentation approach. For both approaches segmentations were generated either using or not using the atlas based a-priori information. For comparison the skull segmentation provided by the widely-used BET2 tool [134, 67] was evaluated.

For all methods the T1- and T2-weighted images were intensity inhomogeneity corrected (Sec. 6.2), transformed to AC-PC-space, and resampled to $1 \times 1 \times 1\,\mathrm{mm}^3$ isotropic resolution.

To generate the active region contours results we ran the algorithm described in Chapter 7 for 700 iterations at a coarser $2 \times 2 \times 2\,\mathrm{mm}^3$ level and for 500 iterations at the resolution of the image data ($1\,\mathrm{mm}^3$). The smoothing parameter was chosen as $\epsilon = 1.0$ for all data sets. Based on our observations from previous experiments on real MRI data (Sec. 7.4) five regions (BG, scalp,

skull, CSF, brain) were differentiated. The segmentation was started from an initial solution constructed according to Section 8.5.1 and relabeled to five regions.

The MRF segmentation (Chap. 8) was employed to generate another set of head tissue segmentations. The MRF labeling was run for a fixed number of iterations ($t^{\max} = 25$). A value of $1/T = 0.2$ was chosen for the weighting of the MRF prior. The same start mask as for the active contour segmentation was used, and the initial parameters of the likelihood function were set to empirically determined values (Sec. 8.5.2). No post-processing was applied to the segmentation result but some a-priori constraints were enforced. These constraints were the minimum scalp thickness and the preservation of the input reference surfaces in the segmentation result (Sec. 8.5.3). The scalp and WM reference surfaces which are required as input for the MRF approach and for the atlas based a-priori probability maps were reconstructed in *BESA MRI* using automatic segmentation methods (Sec. 6.3 and 6.4).

For comparison skull segmentation masks were generated using the BET2 tool (Release 4.1, June 2008). From preliminary experiments with the tool it was obvious that the accuracy of the results strongly depends on two parameters which can be set by the user. Because of this we varied the chosen values for these parameters across the whole parameter range, and computed the skull segmentation mask for in total 99 different combinations of parameter values. In the results we will always list accuracies for the default parameter values and for those parameter values which yielded the highest accuracy.

10.3 Validation Against Manual Segmentations

For the segmentation task at hand human raters can be regarded as a benchmark for any algorithm. Still, it is known that a certain variability must be expected between the manual segmentations of different human raters [73, 107, 161]. An average segmentation generated from multiple manual segmentations is better suited to act as a reference than a manual segmentation from a single human rater. Furthermore, having multiple manual segmentations of the same data set allows us to estimate the inter-rater variability. The observed accuracy measures for the proposed segmentation approaches can then be evaluated in the context of the inter-rater variability.

In this validation study we, thus, collect multiple manual segmentations of the same data set from different human raters. From these then multiple reference segmentations are constructed. The results of the proposed segmentation approaches are compared to all of the reference segmentations and accuracy measures are statistically evaluated.

The validation was carried out on data sets of four different subjects. Two data sets included a T1- and T2-weighted MRI, while the other two data sets only included a T1-weighted MRI. Details on the MRIs are given in Table 10.1. Standard MRI sequences were used to record the MRIs. Only the T1-weighted images in the first two data sets (MS_A and MS_B in Tab. 10.1) had a notable special characteristic. These images were recorded with a water-excitation sequence (Sec. 3.2.1) which provides an efficient reduction of the signal from fatty tissue. In this way, problems arising from the water-fat shift artifact are avoided.

Manual segmentations from four different human raters were acquired. Raters had at least one year of experience in the field of realistic volume conductor modeling for EEG source analysis. Due to the large effort of manual segmentations it was not feasible to manually segment the

Parameter	MS_A		MS_B	
	T1	T2	T1	T2
Sequence	3D TFE	2D SE	3D TFE	2D SE
Resolution (mm^3)	1.17^3	1.17^3	1.17^3	1.17^3
Repetition time TR (ms)	9.32	3429	9.32	3429
Echo time TE (ms)	4.45	40	4.45	40
Flip angle (°)	9	90	9	90
Bandwidth (Hz)	224	777	224	777
Inversion recovery time (ms)	n.a.	n.a.	n.a.	n.a.
Fat reduction	water excitation	fat saturation	water excitation	fat saturation

	MS_C T1	MS_D T1
Sequence	MP-RAGE	3D FFE
Resolution (mm^3)	1.0^3	$0.46 \times 0.46 \times 1.0$
Repetition time TR (ms)	2200	30
Echo time TE (ms)	3.93	4.6
Flip angle (°)	15	25
Bandwidth (Hz)	130	648
Inversion recovery time (ms)	1200	n.a.
Fat reduction	none	none

Table 10.1: Details on MRIs for the four data sets used in this validation study. Columns with T1 as heading describe sequences for T1-weighted MRIs, those with T2 described sequences for T2-weighted MRIs.

complete head. Thus, the manual segmentation is limited to smaller ROIs randomly distributed throughout the head. In total, for each data set 15 ROIs of 40×40 pixels each are randomly placed in sagittal slices of the MRI. The centers of the square ROIs are restricted to lie within a distance of 20 mm from the scalp so that the ROIs cover a substantial part of the skull. In addition, ROIs are only placed in the superior part of the head as we are mainly concerned with the accuracy of segmenting the skull cap. As an example, Figure 10.1 shows a selection of labeled ROIs.

From the multiple acquired segmentations of the same data sets reference skull segmentations are constructed following the scheme illustrated in Figure 10.2. Note, that the same procedure can be used to generate reference segmentations for other tissues. First, from each of the manual segmentations a skull mask M_i is extracted. In the skull mask only those voxels are marked which are labeled as compact or cancellous bone in the manual segmentation. A reference skull mask R_i is then constructed from multiple manually segmented skull masks. To do this, each voxel that is marked as skull in a majority of the manually segmented masks is marked as skull in the reference mask. Hereby, the i-th reference mask R_i is based on all but the i-th manually segmented skull mask M_i. The i-th manual segmentation is left out for the construction of the i-th reference mask so that we can assess the accuracy of rater i by comparing rater i's labeling to the average of all other raters' labelings. In this way, we obtain as many reference segmentations as there are manual segmentations from different human raters.

The agreement between the reference skull masks R_i and the test skull masks T_j, that is, the skull masks generated by the automatic segmentation approaches, is measured using the Dice

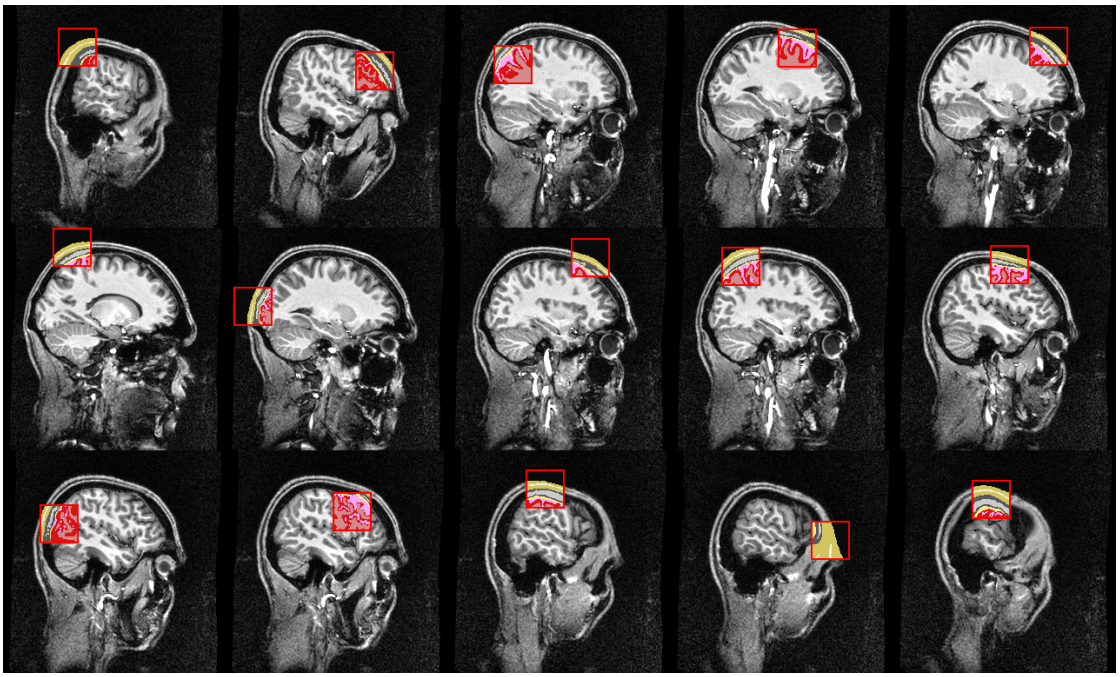

Figure 10.1: Exemplary labeled ROIs for data set MS_A blended over the T1-weighted MRI of that data set. ROIs are indicated by red squares. Dice coefficients were computed only for these ROIs. Voxels labeled outside of the ROIs were ignored.

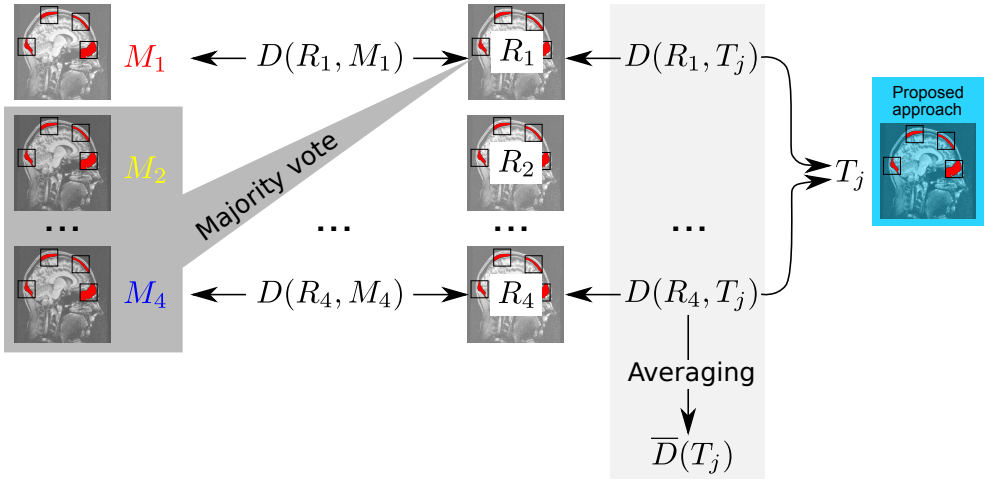

Figure 10.2: The design of the study validating the skull segmentation of the proposed approaches against multiple manual segmentations from different human raters. The construction of reference segmentations from multiple manual segmentations from different raters is shown. Furthermore, the computation of the difference metrics for the test masks and the manually segmented masks are illustrated.

coefficient [44].

$$D(R_i, T_j) = 2\,\frac{\left|R_i \cap T_j\right|}{\left|R_i\right| + \left|T_j\right|} \tag{10.1}$$

$|A|$ here means the volume of mask A computed by counting the marked voxels and multiplying with the volume of a voxel. The Dice coefficient can be interpreted as the overlap between the two masks relative to the average size of the masked regions. The optimal value of the Dice coefficient is 1, indicating a perfect agreement between the masks. For a finite size of the image grid a Dice coefficient of 0 means that there is no overlap between the two masks.

A test skull segmentation is compared to each of the reference skull segmentations. In this way we obtain for each test segmentation as many Dice coefficients as there are reference segmentations. Finally, these Dice coefficients $D(R_i, T_j)$ are averaged and the averaged Dice coefficient $\overline{D}(T_j)$ is interpreted as a measure for the accuracy of the skull segmentation obtained using the automatic segmentation approach. The validation of the segmentation approaches with respect to the accuracy of the other tissues is performed in an analogous way.

10.4 Validation Against CT Reference Segmentation

A lot of the difficulties in segmenting the human skull from MR images originate from properties of the MRIs, for example, the low contrast between skull and surrounding tissues, a low SNR, and other artifacts. CT images are much better suited for skull segmentation because the contrast between the skull and other tissues is very high and because CT images are mostly free from artifacts which might affect the segmentation (Sec. 3.2.4). Skull segmentations derived from CT images can, thus, be regarded as a reference to which we can compare our MRI based automatic segmentations.

In this validation study, we will investigate the accuracy of the proposed segmentation approaches by comparing the obtained results to a CT based reference segmentation.

We were able to obtain one data set including a CT image, a T1-weighted MRI, and a T2-weighted MRI of acceptable quality of a patient with a large tumor in the left hemisphere. CT data of a healthy subject was not available, and is in general very rare as CT scans are usually not performed on healthy subjects because of the exposure to ionizing radiation during the CT measurement. Slices of the available image data can be seen in Figure 10.3. Details on the images are given in the captions of the same figure.

The T1- and T2-weighted MRIs were preprocessed in *BESA MRI*. Preprocessing included intensity inhomogeneity correction (Sec. 6.2), transformation of the image into AC-PC space, coregistration of the T2-weighted to the T1-weighted MRI, and resampling to an isotropic resolution of $1.0\,\mathrm{mm}^3$.

Coregistration of the CT data to the T1-weighted MRI was a problem. Using a linear transformation with nine parameters (three translation, three rotation, and three scaling parameters), and using the mutual information [152, 85] registration measure no sufficiently accurate coregistration could be obtained. This is most likely due to the specific gray value characteristics of the T1-weighted image. In the T1-weighted MRI the intensity of the scalp is low, and the boundary between the outer skull and the scalp is not very well visible. This also effects the precise alignment of the CT image. This problem was partially solved by coregistering the CT

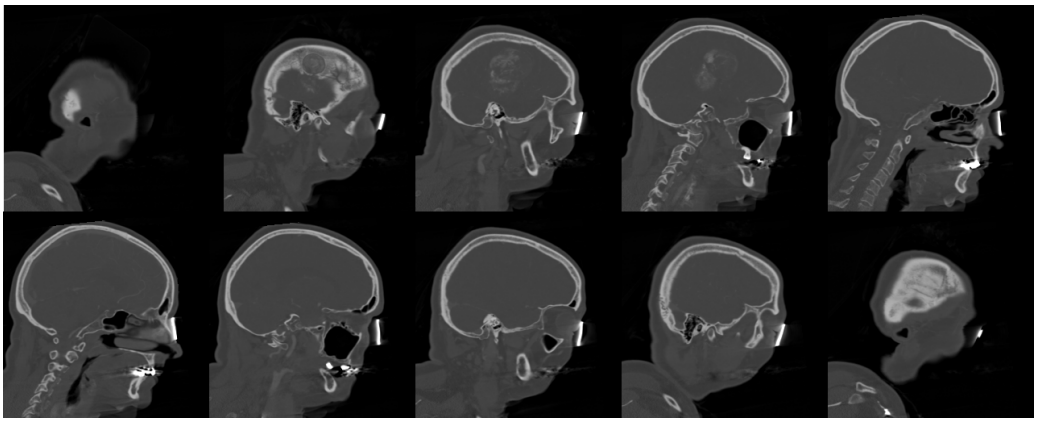

(a) CT image with original resolution of $0.5 \times 0.5 \times 0.7\,\mathrm{mm}^3$.

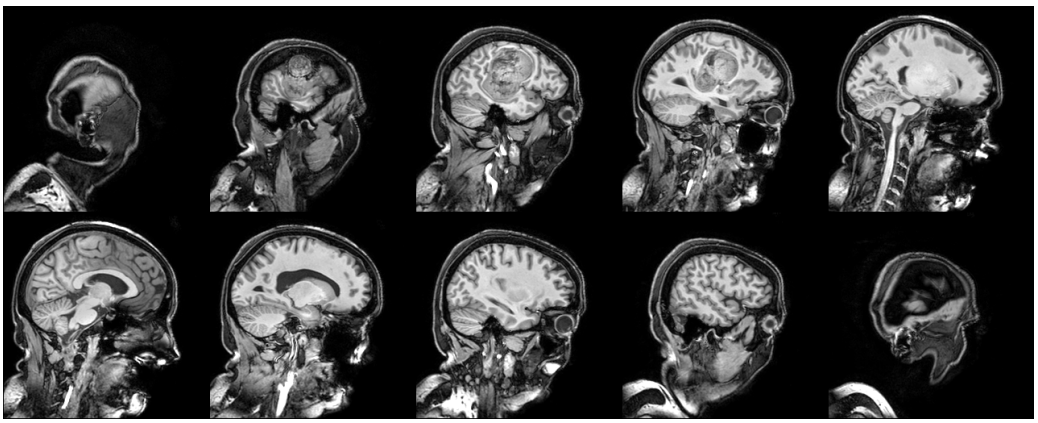

(b) T1-weighted MR image with original isotropic resolution of $(1.17\,\mathrm{mm})^3$. MRI was recorded using a 3D TFE sequence with selective water-excitation to avoid the water-fat-shift artifact.

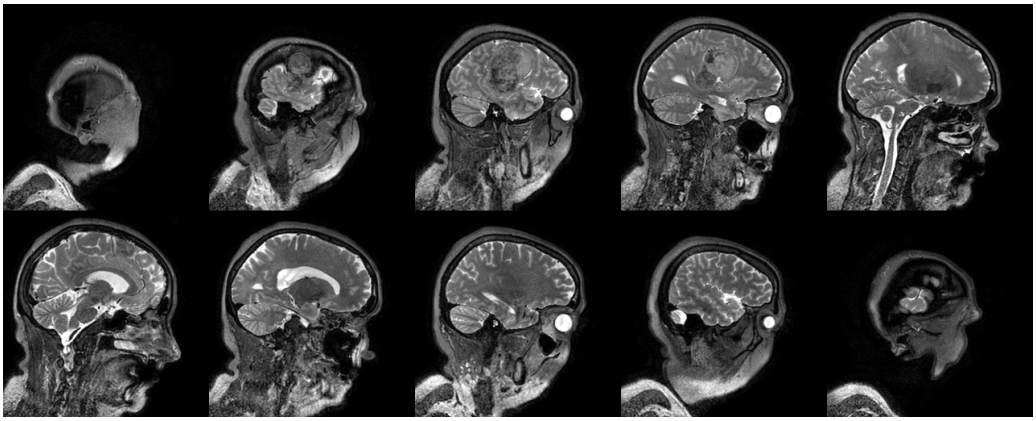

(c) T2-weighted MR image with original isotropic resolution of $(1.17\,\mathrm{mm})^3$.

Figure 10.3: Slices of the CT, T1-weighted, and T2-weighted images of the data set used for validation of the automatic segmentation against a reference segmentation based on the CT image.

image to a pre-segmented scalp mask derived from the T1-weighted image, and employing the sum of squared intensity differences as registration measure.

Another problem was that the posture of the patient during the CT measurement differed from that during the MRI recording. This lead to a sub-optimal coregistration when taking the complete target image into account for the computation of the registration criterion (cost function). Thus, during the coregistration we focused on a ROI ignoring the inferior head and neck regions. The registration criterion was computed only for the ROI. During the coregistration the images were resampled to the resolution of the target image. In this way, we obtained an improved coregistration result as verified by visual inspection.

The CT based reference skull segmentation was generated using simple thresholding on the coregistered CT image. The threshold value was chosen manually approximately half way between the intensities of soft tissue and compact bone. Obvious errors in the segmentation, for example, parts of the spongiosa that were darker than the threshold and that were thus mislabeled, were corrected manually. In addition, all air-filled cavities were added to the reference CT skull mask.

When comparing the test skull segmentation and the reference skull segmentation we are limiting ourselves to the superior part of the head, approximately the region of the skull cap. This is because the skull cap is more important for EEG source analysis because it is situated directly in-between the sources and the sensors.

Furthermore, we ignored for our comparison approximately two-thirds of the left half of the head where the large tumor in the patient's head is situated. Due to its large size the tumor distorts the normal anatomy of the head. Results obtained for the hemisphere with the pathological anatomy would, thus, not allow meaningful statements with respect to the segmentation accuracy in subjects with normal anatomy.

Comparisons between the test and reference skull segmentations are done by computing the Dice coefficients as defined in Equation 10.1.

10.5 Results

10.5.1 Validation Against Manual Segmentations

First, the results of the validation against the manually segmented images are presented. Table 10.2 lists the averaged Dice coefficients for the skull segmentation and the individual Dice coefficients for the four datasets (MS_A to MS_D). For the first two data sets accuracies are given for the segmentation based on the T1- and T2-weighted image, and for the segmentation based only on the T1-weighted image from the data set.

On average, the MRF approach incorporating the atlas-based prior performs best closely followed by the MRF approach without atlas information. For the active contours segmentation the mean Dice coefficients are marginally smaller than those for the MRF segmentations. It can be observed that for the active contours approach the overlap with the manual segmentations is increased by incorporating the atlas-based prior, too. The accuracy of the BET is on average less good, especially, when choosing the default parameter values. Only for MS_C BET performs best choosing either the default or the optimal parameters. The Dice coefficients for all

approaches stay below the inter-rater variability which can be interpreted as the average segmentation accuracy of a human rater.

Accuracies generally increased for all methods when incorporating both a T1- and a T2-weighted MRI. In the case of, for example, the MRF approach with atlas information the averaged skull segmentation accuracy for the data sets including two image modalities was 0.908 and, thus, larger than the average accuracy across all T1-only data sets (0.878). The improvement by using two modalities also gets obvious when looking at the voxels that were falsely classified as skull by the MRF approach (Fig. 10.4). When only incorporating a T1-weighted image some CSF voxels in the lower part of the ROI are mislabeled as skull. In the segmentation based on a T1- and a T2-weighted image these voxels are correctly classified.

In the masks of false positive skull voxels (Fig. 10.4) it can also be observed how using the atlas-based a-priori probability improves the skull segmentation. When not incorporating the atlas a lot of voxels actually belonging to the SCT are falsely classified as skull. This is due to the fat reduction sequence with which the MRI was recorded. The signal of the SCT is, thus, very low. When the atlas-based a-priori probabilities are used the number of voxels falsely classified as skull is much lower.

MRI characteristics also mattered for all segmentation methods. The computed dice coefficients varied between 0.826 (MS_D) and 0.943 (MS_A, T1 only) for the MRF with atlas information, between 0.802 (MS_D) and 0.939 (MS_A) for the active contours approach with atlas information, and between 0.684 (MS_B, T1 only) and 0.923 (MS_C) for the BET.

The MRF segmentation approach incorporating the atlas information proved to deliver on average the most accurate skull segmentation. We will, thus, have a closer look on the remaining falsely classified voxels for this method. To do so, the distance of the false positive (FP) and false negative (FN) voxels to the reference skull mask were computed. It was counted how often FP and FN voxels occurred at certain distances. The counts were normalized to the number of voxels in the reference mask.

The resulting histograms (Fig. 10.5) show more false negative than false positive voxels indicating that there is a trend of the MRF approach to underestimate the skull as compared to the segmentations produced by manual raters. For the data sets containing a T1- and a T2-weighted image (MS_A and MS_B in Fig. 10.5) only a very small number of voxels at a distance of more or equal than 2 mm from the reference skull mask were falsely classified ($< 2\%$ relative to the number of voxels in the reference skull mask). When only exploiting the information from a T1-weighted image (MS_C and MS_D in Fig. 10.5) more voxels at a larger distance to the reference skull are falsely classified. In relation to the number of voxels in the reference mask up to 12.3% of voxels at a distance of 2 mm or more are falsely classified, and this number drops to below 2% only for distances of 5 mm and more.

The MRF and the active contours segmentation approach are able to deliver a segmentation of the head into the four tissues scalp, skull, CSF and brain which are relevant for EEG source analysis. These tissues were also labeled by the manual raters. We are, therefore, able to assess the segmentation accuracy of the other tissue regions (Table 10.3). Note, that we only present the results for the MRF and active contours approaches incorporating the atlas-based a-priori information as those segmented the skull most accurately.

The MRF segmentation approach is able to segment the scalp and brain tissues at least as accurately as it is able to segment the skull. Brain segmentation accuracy is even on the same

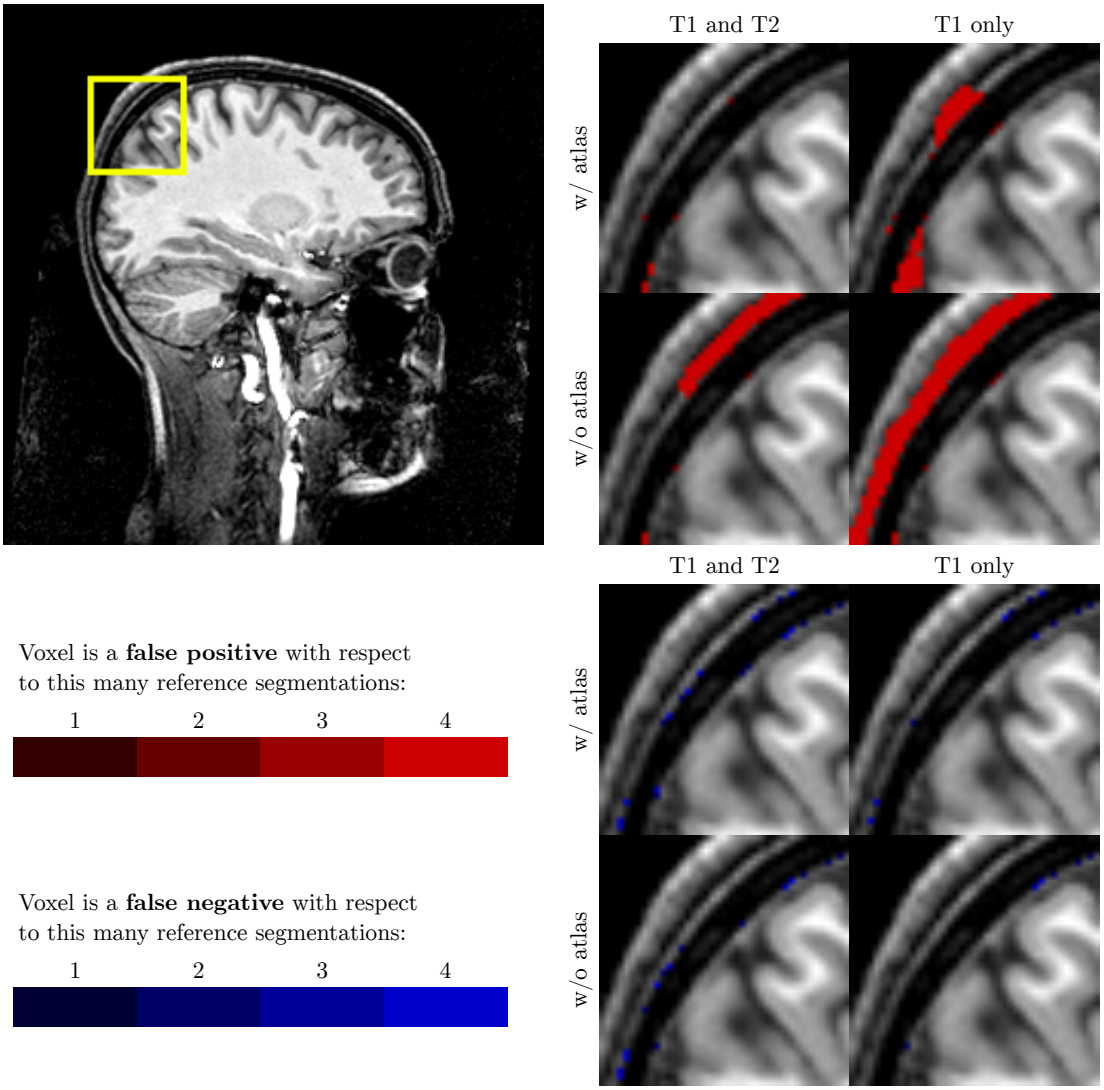

Figure 10.4: Difference masks showing false positives (top right) and false negatives (bottom right) for the MRF segmentation on data set MS_B. Results are shown for MRF segmentation based on T1- and T2-weighted images, and segmentation based on the T1-weighted image only. Furthermore, results using, respectively, not using the atlas-based a-priori probabilities are presented.

Table 10.2: Dice coefficients measuring the skull segmentation accuracy for our proposed segmentation approaches and another commonly used method (BET). Dice coefficients for the four data sets are given. In addition, the average Dice coefficients across all data sets are listed. Standard deviations are given in parentheses corresponding to the last digits of the quoted results. Here, MRF denotes the approach based on a MRF model of the human head, and CV (Chan-Vese) stands for the approach based on the multiphase active contours method.

Method	**MS_A**	**MS_A** (T1 only)	**MS_B**	**MS_B** (T1 only)
MRF w/ atlas	0.905(7)	0.943(7)	0.911(4)	0.895(3)
MRF w/o atlas	0.911(7)	0.908(1)	0.866(2)	0.745(7)
CV w/ atlas	0.939(7)	0.932(6)	0.828(5)	0.855(3)
CV w/o atlas	0.951(2)	0.894(4)	0.806(7)	0.747(8)
BET [67] (opt.)	0.910(4)	0.885(5)	0.765(7)	0.684(3)
BET [67] (def.)	0.781(4)	0.770(5)	0.698(9)	0.601(6)
Inter-rater	0.945(15)	*n.a.*	0.907(20)	*n.a.*

Method	**MS_C**	**MS_D**	**Average**
MRF w/ atlas	0.850(3)	0.826(3)	0.888(43)
MRF w/o atlas	0.901(2)	0.891(3)	0.871(64)
CV w/ atlas	0.858(2)	0.802(4)	0.869(55)
CV w/o atlas	0.910(3)	0.800(2)	0.851(78)
BET [67] (opt.)	0.923(3)	0.770(3)	0.823(97)
BET [67] (def.)	0.923(3)	0.652(3)	0.737(114)
Inter-rater	0.943(13)	0.964(7)	0.938(14)

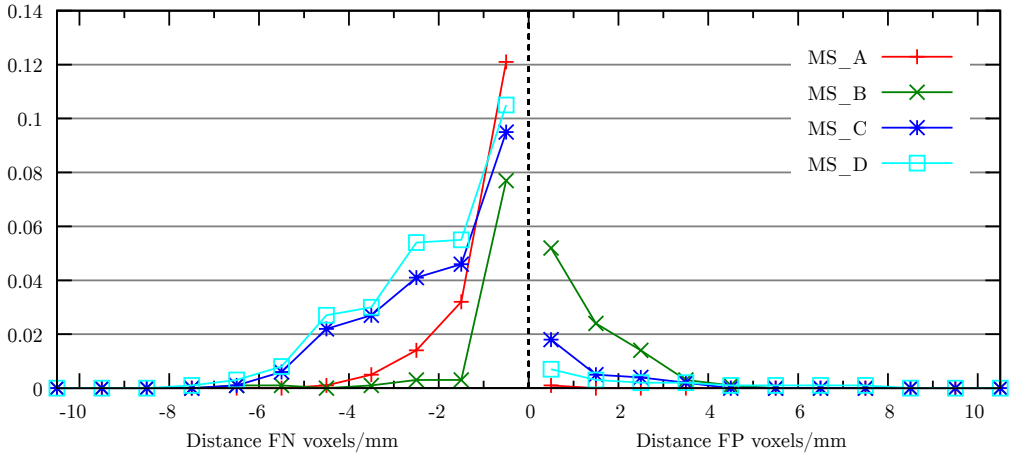

Figure 10.5: Normalized histograms of the distances from the false positive and false negative skull voxels to the reference skull mask for the MRF approach incorporating the probabilistic atlas. Results for MS_A and MS_B were based on the T1- and the T2-weighted MRI. Histogram counts are normalized to the number of voxels in the reference mask.

level as the inter-rater variability. The accuracy of the CSF segmentation is less good with an average Dice coefficient of 0.744. The smaller accuracy for the CSF segmentation can also be found for the manual raters. There the average relative overlap between the manual segmentation of one rater and the reference segmentation was 0.783 which is not much larger than the average overlap for the MRF segmentation.

For the active contours approach the accuracy of the segmented scalp masks is on the same level as that of the segmented skull. The brain segmentation accuracy is slightly lower. CSF is segmented only with a low Dice coefficient which is substantially below the inter-rater variability.

The results show that the MRF segmentation is more accurate than the active contours segmentation with respect to all tissue regions. Differences in accuracy are most pronounced for the CSF and brain regions.

In principle, the MRF segmentation approach yields a segmentation into even more than four tissue regions. During the segmentation a total of eight tissues are distinguished. The accuracy of this very detailed segmentation was investigated (Table 10.4) to see if it would make sense to incorporate more than four tissues into the head models for EEG source analysis.

The results show that compact bone, GM and WM could be segmented with an accuracy of more than 0.75 and the accuracy of the WM segmentation even surpassed the inter-rater variability. SCT and muscle tissue were labeled slightly less accurately. The average Dice coefficient for cancellous bone seems quite low at first but it is mainly due to the segmentation results for the T1-weighted images from data sets MS_A and MS_B. These images were recorded using a sequence suppressing the signal from fatty tissue. The different appearance of compact and cancellous bone in MR images nearly exclusively originates from this fat signal. The T1-weighted images, thus, do not contain the information to differentiate between cancellous and compact bone leading to low Dice coefficients. The average overlap for cancellous bone is 0.544

Table 10.3: Dice coefficients measuring the segmentation accuracy for the four tissue regions which are differentiated in our volume conductor models. Dice coefficients are given for the four investigated data sets separately and averaged across all data sets. In this table only the proposed segmentation approaches are included that incorporate the atlas-based prior. Standard deviations are given in parentheses corresponding to the last digits of the quoted results. Here, MRF denotes the approach based on a MRF model of the human head, and CV (Chan-Vese) stands for the approach based on the multiphase active contours method.

Data set	Scalp	Skull	CSF	Brain
MRF w/ atlas				
MS_A	0.911(5)	0.905(7)	0.775(18)	0.928(2)
MS_A (T1 only)	0.919(5)	0.943(7)	0.790(21)	0.912(2)
MS_B	0.898(3)	0.911(4)	0.817(8)	0.933(4)
MS_B (T1 only)	0.879(4)	0.895(3)	0.773(9)	0.913(2)
MS_C	0.905(4)	0.850(3)	0.578(9)	0.940(10)
MS_D	0.851(7)	0.826(3)	0.733(7)	0.904(14)
Average	0.894(25)	0.888(43)	0.744(86)	0.922(14)
CV w/ atlas				
MS_A	0.919(7)	0.939(7)	0.601(16)	0.862(6)
MS_A (T1 only)	0.911(5)	0.932(6)	0.608(18)	0.863(6)
MS_B	0.861(3)	0.828(5)	0.538(18)	0.875(9)
MS_B (T1 only)	0.864(3)	0.855(3)	0.593(20)	0.875(9)
MS_C	0.893(3)	0.858(2)	0.306(6)	0.878(12)
MS_D	0.778(7)	0.802(4)	0.475(5)	0.764(12)
Average	0.871(51)	0.869(55)	0.520(116)	0.853(44)
Inter-rater				
MS_A	0.957(10)	0.945(15)	0.839(20)	0.934(12)
MS_B	0.955(11)	0.907(20)	0.772(47)	0.919(24)
MS_C	0.949(14)	0.943(13)	0.648(87)	0.912(34)
MS_D	0.951(11)	0.964(7)	0.877(50)	0.907(58)
Average	0.951(5)	0.940(24)	0.783(100)	0.918(12)

when ignoring these two MRIs. The lowest overlap between the automatic and the reference segmentations can be observed for dura mater. It is nevertheless not substantially smaller than the inter-rater variability.

10.5.2 Validation Against CT Reference Segmentation

The Dice coefficients between the CT based reference skull segmentation and the results obtained from the investigated methods are listed in Table 10.5.

For the MRF approach we find Dice coefficients in the range between 0.843 and 0.863 depending on if the atlas information is incorporated, and depending on if only a T1-weighted image or if a T1- and a T2-weighted image were used. Incorporating the atlas information increased the accuracy when only using the T1-weighted image. In the case, where also a T2-weighted image is taken into account accuracy decreases when exploiting the information contained in the probabilistic tissue atlas.

The active contours segmentations agreed well with the CT reference skull mask as indicated by the Dice coefficients that ranged between 0.894 and 0.898. As for the MRF approach incorporating two image modalities increased the overlap. Incorporating the atlas-based a-priori probability maps did not lead to higher Dice coefficients. When incorporating a T1- and a T2-weighted image the Dice coefficient was even slightly reduced when utilizing the atlas.

Using the BET method Dice coefficients between 0.829 and 0.886 were achieved. Segmentation accuracy of the BET method was substantially better when both imaging modalities were used. Dice coefficients for the optimal choice of parameters were only slightly larger than the coefficients for the default parameters.

When only incorporating the T1-weighted image the MRF achieved the largest Dice coefficients (0.857) followed by the active contours approach (0.850). When additionally using a T2-weighted image the active contours approach resulted in the skull mask with the best agreement to the reference segmentation (0.898). In the multimodal case BET yielded the skull mask with the second highest Dice coefficients (0.886).

Some of the observed results are not easily explainable, for example, the result that the incorporation of the atlas decreases the Dice coefficients in some cases. For this reason, we will have a closer look at the CT reference segmentation and the difference masks for two selected approaches (Fig. 10.6).

When comparing the outline of the CT reference skull mask to the underlying T1-weighted image it can be observed that the CT and the MRI are not perfectly aligned (Fig. 10.6a). This is most obvious in the superior part of the head (Fig. 10.6b). Here, the CT reference skull mask overlaps a region which can be clearly identified as SCT in the MRI. In general, the outer boundary of the CT reference skull mask seems to be located too far towards the scalp surface.

When zooming in on the superior region we can see that the MRF segmentation approach using both image modalities and the atlas information yields an outer skull boundary that, judging by the underlying T1-weighted MRI, precisely delineates the skull from the thin scalp layer on the outside of the skull (Fig.10.6c). For the BET method we observe that the outer boundary is reconstructed closer to the interface between SCT and the skin layer which is too far superior (Fig. 10.6d). Nevertheless, due to the slight misalignment of the CT image the overlap

Table 10.4: Dice coefficients measuring the segmentation accuracy of the MRF segmentation approach incorporating the atlas-based prior for eight different tissues. Dice coefficients are given for the four investigated data sets separately and averaged across all data sets. Standard deviations are given in parentheses corresponding to the last digits of the quoted results.

Data set	Scalp/ muscle	SCT	Compact bone	Cancellous bone	Dura mater	CSF	GM	WM
MRF w/ atlas								
MS_A	0.653(9)	0.466(17)	0.716(17)	0.557(5)	0.276(51)	0.723(5)	0.838(5)	0.924(3)
MS_A (T1 only)	0.835(15)	0.640(22)	0.741(12)	0.047(2)	0.360(28)	0.574(9)	0.809(3)	0.924(3)
MS_B	0.720(3)	0.710(11)	0.847(11)	0.420(11)	0.621(29)	0.795(15)	0.824(8)	0.882(7)
MS_B (T1 only)	0.806(7)	0.776(7)	0.811(14)	0.078(12)	0.477(34)	0.612(8)	0.791(5)	0.882(7)
MS_C	0.626(75)	0.703(150)	0.755(13)	0.352(14)	0.294(26)	0.189(23)	0.729(18)	0.706(25)
MS_D	0.501(21)	0.770(24)	0.662(8)	0.850(3)	0.479(170)	0.550(11)	0.571(24)	0.591(24)
Average	0.690(124)	0.678(115)	0.755(66)	0.384(302)	0.418(132)	0.574(210)	0.760(100)	0.818(138)
Inter-rater								
MS_A	0.892(16)	0.824(68)	0.854(34)	0.846(25)	0.568(97)	0.854(14)	0.835(22)	0.903(15)
MS_B	0.851(24)	0.828(23)	0.791(29)	0.605(20)	0.542(74)	0.746(57)	0.810(47)	0.872(17)
MS_C	0.651(47)	0.552(65)	0.829(19)	0.728(38)	0.359(76)	0.544(61)	0.754(52)	0.704(54)
MS_D	0.805(8)	0.849(23)	0.854(36)	0.888(42)	0.499(51)	0.766(49)	0.791(91)	0.743(58)
Average	0.800(106)	0.763(141)	0.832(30)	0.767(127)	0.492(93)	0.727(131)	0.797(34)	0.805(97)

Table 10.5: Dice coefficients for the automatically segmented skull in comparison to a CT based reference skull segmentation. Dice coefficients are given for each of the two proposed segmentation methods, as well as for a competing method (BET). MRF denotes the approach based on a MRF model of the human head, and CV (Chan-Vese) stands for the approach based on the multiphase active contours method.

Method	Dice coefficient
MRF w/ atlas, T1 only	0.857
MRF w/o atlas, T1 only	0.843
MRF w/ atlas, T1 and T2	0.845
MRF w/o atlas, T1 and T2	0.863
CV w/ atlas, T1 only	0.850
CV w/o atlas, T1 only	0.849
CV w/ atlas, T1 and T2	0.884
CV w/o atlas, T1 and T2	0.898
BET [67] (T1 only, def.)	0.829
BET [67] (T1 only, opt.)	0.833
BET [67] (T1 and T2, def.)	0.883
BET [67] (T1 and T2, opt.)	0.886

between the CT reference skull mask and the BET skull mask is higher than the overlap between the reference and the MRF skull mask.

This observation is confirmed by the difference mask showing the voxels that were falsely classified as skull (false positives) or falsely classified as non-skull (false negatives). For the MRF approach we can see a thin layer of false negative voxels adjacent to the outer skull boundary (Fig. 10.6e), and only very few false positive voxels. For the BET the layer of false negative voxels is thinner and in some regions it even happens that voxels adjacent to the outer skull boundary are falsely identified as skull (Fig. 10.6f).

10.6 Discussion and Conclusion

Comparing the skull segmentation accuracy of all investigated methods we find that our two proposed approaches yield more accurate results for the validation against the manual segmentations than an in practice widely-used method [67]. Another advantage of the proposed methods is the relative independence of parameter values. It is true that our methods depend on parameters. The MRF approach, for example, depends on the parameter $1/T$ defining the strength of the influence of the MRF model. These parameters, however, were fixed for all data sets in our validation study, and the approaches still yielded good results. We may, thus, expect that they also yield good results for other data sets while using the same parameter values. The performance of the BET, however, strongly varied depending on the values chosen for the parameters. Knowing the reference segmentation we could optimize the parameters to obtain an optimal Dice coefficient of on average 0.823. Using the predefined default parameter values that are set in the BET2 software, however, the achieved skull segmentation accuracy was considerably smaller ($D = 0.737$). In any practical application, optimal parameters could only be chosen by visually

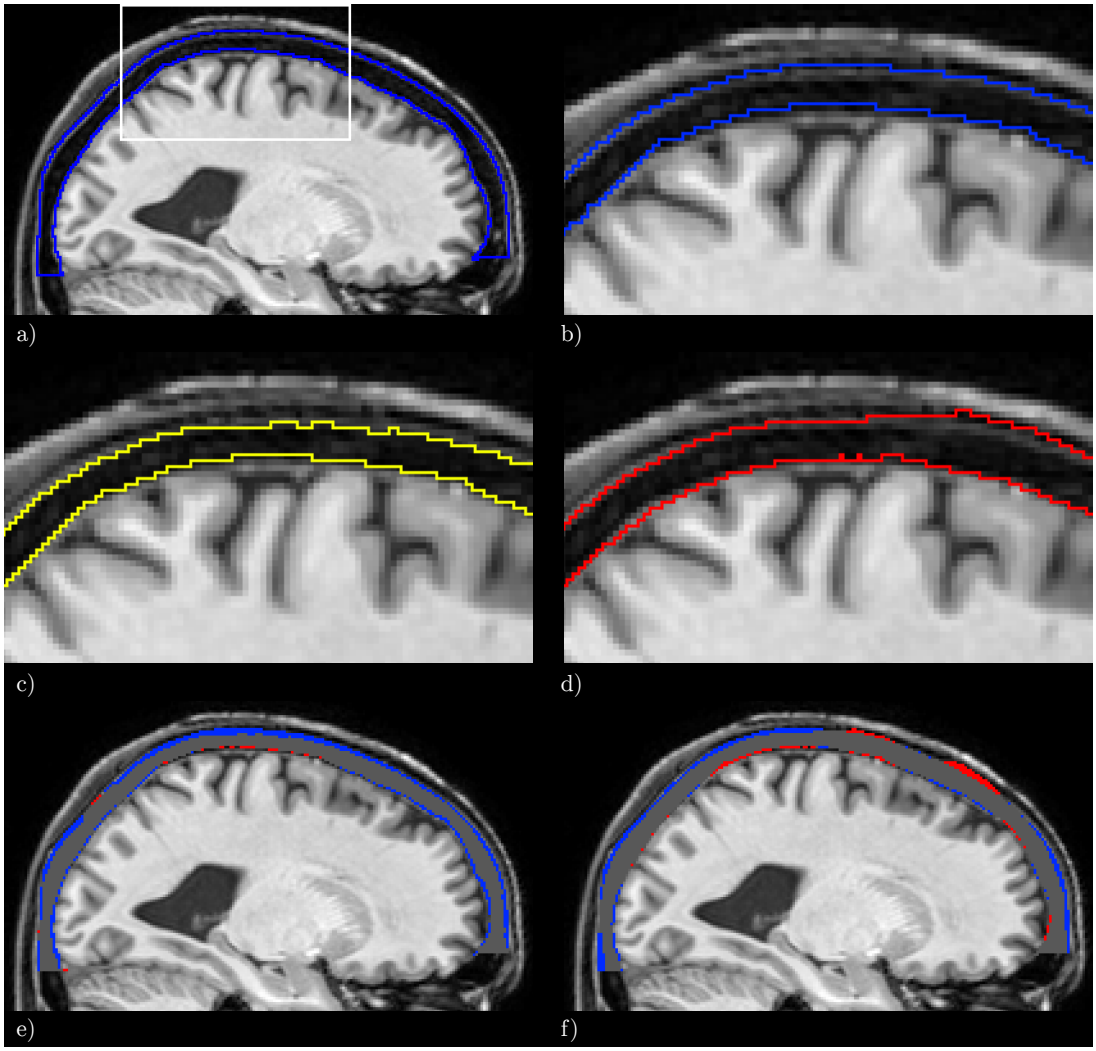

Figure 10.6: CT reference skull masks and difference masks for two selected segmentation methods. a) T1-weighted image with outline of CT reference skull mask. All voxels within the outline are marked in the skull mask; b) to d) Zoomed in ROI of the T1-weighted image with outlines of the CT reference skull mask (b), blue), the MRF skull mask with atlas based on T1- and T2-weighted images (c), yellow), and the BET skull mask based on T1- and T2-weighted images (d), red); e) and f) Difference masks for the MRF skull mask with atlas and the BET skull mask, respectively. Both skull masks were generated based on T1- and T2-weighted images. Red voxels are false positives, blue voxels are false negatives.

inspecting the segmentation results for a large number of parameter combinations which would be very labor intensive.

Regarding the skull segmentation accuracy all investigated methods stayed well below the accuracy of the human raters. The MRF approach was, however, able to segment the brain (consisting of GM and WM) with an on average higher accuracy than the human raters. In addition, the MRF approach yielded a CSF segmentation whose accuracy was nearly on the level of the accuracy achieved by the manual segmentations. In another aspect the performance of the MRF approach and the manual segmentation differed substantially. Using the MRF approach a segmentation of the complete head into eight different tissues can be obtained in far less than an hour of computation time. In contrast, a human rater needed approximately two to three hours to label only the 15 regions-of-interest in a single data set.

Some more insight could be gained from the validation of the segmentation results against manual segmentations. It was found that incorporating multiple image modalities in general increased the skull segmentation accuracy. Including a T2-weighted image in particular helped in correctly distinguishing CSF from skull but the additional anatomical information also improved the segmentation of the outer skull boundary.

Segmentation accuracy also seemed to be influenced by the MRI recording sequence and its parameters. The T1-weighted images of the first two data sets (MS_A and MS_B) were recorded with the same sequence, and the skull segmentation results for these images ($D = 0.919$) were clearly better than those for the other two MRIs ($D = 0.838$). The main reason for this performance difference is the fat reduction sequence used for recording the first two T1-MRIs. Using such a sequence cancellous bone appears nearly as dark as compact bone so that it can more easily be distinguished from the surrounding brighter tissues. In addition, by suppressing the signal from fatty tissue we are avoiding problems with the water-fat shift artifact. This artifact might otherwise lead to situations where the signal from the fatty SCT layer is shifted downwards and obstructs the outer skull boundary.

The validation presented in this chapter investigated how accurate the segmentation results of the proposed approaches are. It did, however, not show how large the errors would be if we built volume conductor models from the segmentations, and used these for EEG source analysis. We can get a rough estimate for the errors of the forward and inverse solution by relating the segmentation errors observed here to the skull geometry errors that were investigated in the sensitivity study in Chapter 4.

During visual inspection of the skull segmentation results we did not observe any erroneous holes. We, thus, do not expect the scalp potential and localization errors observed in the sensitivity study (Figs. 4.3a and 4.6a). Holes can in any case be avoided by using suitable post-processing steps (Sec. 8.5.4).

For the MRF segmentation approach incorporating the probabilistic atlas some false negative voxels were found at a distance of up to 4 mm from the reference skull mask (Fig. 10.5) when only using T1-weighted MRIs. This segmentation error approximately correlates with models TM 2e and 2f from the EEG sensitivity study where the skull thickness was decreased locally. For these models only minor errors were found. Localization errors, as an example, stayed below 5 mm for all sources. The mean localization errors were below 1 mm and, thus, negligible.

Another result of the sensitivity study was that skull and scalp should be modeled as accurately as possible even in the inferior part of the head. From visual inspection it can be concluded that

the segmentation results for the inferior head and neck region are not as good as for the region of the skull cap. Here, there is still some room for improvement of the proposed methods. Nevertheless, our segmentation improves substantially upon the test models TM 5a–c and 6a–c from the sensitivity study. The foramen magnum, for example, is correctly segmented as a skull hole by both of our segmentation approaches. This will already avoid most of the large forward and inverse errors found for sources at the skull base and in the cerebellum. In addition, in our segmentation methods inferior skull and scalp are not restricted to have constant thickness, and the model is not automatically cut close below the skull (cf. TM 4a and b).

Several previously published studies exist which are validating methods for other segmentation tasks, and which make use of the Dice coefficient to quantify the overlap between segmentations. The absolute values for the Dice coefficients reported in these studies cannot be compared to the values reported in this thesis as the Dice coefficient also depends on the geometry of the segmented object and on the ROI in which the segmentation differences are evaluated. When segmenting any object classification errors will most likely occur at the surface of this object. Voxels deep inside of the object at some distance to the surface will be reliably identified as belonging to the object. Due to this reason the segmentation accuracy measured by the Dice coefficient will be larger for compact objects where the ratio of the surface area to the volume is small. As an example, Shattuck et al. [129] studied the performance of skull-stripping algorithms, and they state that the best algorithm yields a Dice coefficient of 0.969 which is larger than the values reported in this thesis for any Dice coefficient. During skull-stripping the brain is segmented from non-brain tissue. The brain is a rather compact object so that high accuracies can be reached. In our segmentation problem, however, the object to be segmented is an only few mm thick shell, and, therefore, reported Dice coefficients are on a lower level.

In the second part of this chapter the skull segmentation accuracy is validated by comparing the segmentation results to a CT based reference skull mask. Such a validation study makes sense because the CT image is, due to its image characteristics (Sec. 3.2.4), ideally suited for an accurate segmentation of the skull. In our segmentation study, however, a problem occurred. The complicated gray value profile of the extra-cranial tissue in the investigated MRI made it impossible to accurately align the CT to the MR image using a gray value based coregistration (using, e.g., mutual information [151]). Even using a coregistration measure based on presegmented scalp masks only yielded an alignment that was suboptimal in the region of the scalp and skull (Fig. 10.6a). The CT based reference skull mask was, thus, overlapping with regions that were clearly identifiable as scalp in the MRI. This suboptimal alignment will have had an influence on the reported segmentation accuracies. We observed, for example, that the BET skull segmentation in some regions erroneously protrudes into the SCT region (Fig. 10.6d). Although, it is obvious from the MRI that the segmentation in these regions is wrong because of the misalignment of the CT the resulting Dice coefficient might still be larger than for a truly more accurate segmentation.

The uncertainty in the Dice coefficients caused by the imperfect alignment can be estimated by computing the Dice coefficients between two skull masks that are translated or scaled with respect to each other but otherwise identical. If we scale the original mask by 0.5% and compare the scaled to the original mask the resulting Dice coefficient decreases to 0.870. Similarly, the Dice coefficient decreases to 0.904 if we translate the skull mask by only 1 mm. These decreases are on the same scale as the differences between the Dice coefficients for the investigated meth-

ods. Due to the problem with the alignment of the CT image we cannot interpret the rather small differences for the different investigated methods.

A validation study against a CT-based reference segmentation comparable to the study presented here was performed by Wang et al. [147] to prove the accuracy of their segmentation approach. Wang et al. find overlaps of $D = 0.70$ and $D = 0.75$ for the method from [46], respectively, for their proposed method. Dice coefficients for our MRF segmentation approach are clearly larger reaching values up to 0.86 (Tab. 10.5). The difference of the Dice coefficients is so clear that it cannot be explained by the imperfect alignment of the CT image in our study alone. The comparison, thus, indicates that our segmentation approach yields more accurate results than that presented by Wang et al..

A limitation of our CT validation study is that only one data set was investigated. We were not able to obtain more than one suitable data set that included both MR images of reasonable quality and similarly good quality CT images. Such data sets are very rare as CT images are only recorded if there is a clear medical indication to expose the patient to the ionizing radiation during the CT measurement. And in the cases where a good CT is available the MRIs are often not suitable for segmenting the head because the field-of-view is limited, or the resolution in at least one direction is too low.

In conclusion, it can be stated that both proposed methods, the active contours and the MRF method, are able to segment the skull with high accuracy. By relating the segmentation accuracy to the geometric skull inaccuracies investigated in the sensitivity study from Chapter 4 we can conclude that the skull segmentation accuracy achieved by the proposed methods is sufficient for EEG source analysis.

Using the MRF approach, in addition, an accurate segmentation of the head into the four tissues most relevant for EEG head models can be obtained. As a further benefit the MRF approach is a fully automatic method, and it is free of parameters that need to be adapted by the user.

For these reason, we are convinced that the MRF approach using the atlas-based a-priori information is best suited for segmenting volume conductor models from MRI data.

We, furthermore, suggest that T1-weighted MRIs shall be used for segmentation that were measured using sequences with suppression of the fat signal as there was some indication that such MRIs allowed for the most accurate skull segmentations. In addition, if possible the complementary information from T2-weighted MRIs should be exploited by incorporating them into the segmentation.

11 Summary and Outlook

11.1 Summary

This thesis describes the problem of volume conductor segmentation in the context of EEG source analysis, it suggests criteria with respect to the required accuracy and level of detail of volume conductor models, it proposes two new methods for the automatic segmentation of those models, and it proves by validation against manual and CT-based reference segmentations that the proposed methods deliver accurate segmentations that are suitable for the construction of volume conductor models for EEG source analysis.

Chapter 2 introduces the reader to EEG source analysis which is the context of this thesis. It is described which role individual, realistic volume conductor models of the human head play. Furthermore, the literature on realistic EEG head models is reviewed, and requirements with regard to the level of detail of the model are stated.

Next (Chap. 3) we describe the problem of determining a segmentation of the human head from MRI data for the purpose of creating a realistic volume conductor model. Here, also the theoretical basis is presented for the two segmentation approaches which are proposed in this thesis.

After reviewing the literature on realistic volume conductor modeling for EEG source analysis some open questions remained with regard to the required level of detail and the required accuracy of the geometric representation. Two simulation studies were, thus, performed to answer some of the open questions. Both studies were published by the author of this thesis and his co-authors in peer-reviewed journals [75, 77].

The first study (Chap. 4) investigated how skull geometry inaccuracies influence EEG source analysis. Several different types of inaccuracies and simplifications were examined. From the results of the study criteria were derived stating how accurately the skull must be represented in volume conductor models for EEG source analysis.

In the second study (Chap. 5) we investigated how a certain simplification of the CSF in the EEG head model affects source analysis. Our results demonstrated that the simplification will be acceptable in most applications since the observed errors in the forward and inverse computations were mostly negligible. The investigated simplified CSF model solves the problem that occurs when using some inverse methods, for example, the dipole fit, in combination with a realistic source space which has a very complicated geometry. Our study, thus, opens up a possibility to also use these inverse methods in conjunction with realistic volume conductor models.

This thesis also provides the description of an alternative pipeline for preprocessing MRIs and for extracting anatomical information that can be used in EEG source analysis (Chap. 6). The pipeline includes steps for removing the shading artifact from MRIs, and for the reconstruction of scalp and cortex surface representations. In particular, we proposed a modified version of the

AFCM algorithm [105] taking into account a-priori information on the occurrence of WM tissue. We were able to demonstrate that the modified algorithm is more robust and more accurate than the original approach.

One of the major contributions of this thesis is the proposal of two approaches for the automatic segmentation of realistic head models from MRI data. The first approach (Chap. 7) is based on the multiphase active contours segmentation approach by Vese and Chan [143] in which the segmentation is represented by a combination of level set functions. Using a Bayesian framework an objective function is derived that takes into account the observed image data, a geometric smoothness prior, and optional a-priori information on the anatomy of the head, for example, from a probabilistic tissue atlas. An optimal segmentation which minimizes the objective function is found using the gradient descent method.

The second approach (Chap. 8) presented in this thesis uses an MRF model to encode prior knowledge about the layered structure of the head tissues. Again, the segmentation problem is formulated in a Bayesian framework that allows to combine the MRF prior and the agreement of the segmentation with the observed image data. An optimal segmentation is found by estimating the labeling and the likelihood parameters that maximize the a-posteriori probability. The simultaneous segmentation and parameter estimation is done employing an E-M type algorithm [43]. In the E-step the labeling is optimized using the ICM method [19].

Both new methods were applied to real MRI data. The resulting segmentations were inspected visually with particular focus on the skull, and it was found that both approaches are able to accurately delineate the skull. Only minor errors in the skull outlines were observed.

A possible means to avoid also these minor errors is to incorporate additional a-priori knowledge about the anatomy of the human head. Therefore, we devised a probabilistic tissue atlas specifically designed to enhance volume conductor segmentation (Chap. 9). The atlas consists of empirically determined local tissue probability distributions depending on the distances to the scalp and WM surfaces as reference surfaces. After computing the distances to the individual reference surfaces of a subject the probability distributions can be projected to the subject's MRI to obtain tissue probability maps. These maps can be used in conjunction with either of the two proposed segmentation approaches. Experiments on real MRI data showed that the atlas-based a-priori probabilities improve the skull segmentation accuracy in both cases.

In the final chapter (Chap. 10) we conducted two studies to validate the proposed methods for volume conductor segmentation. In addition, we assess their accuracy in comparison to the accuracy of a competing method and to the variability of the labeling from different manual raters. We found that the MRF segmentation approach utilizing the probabilistic atlas delivers the most accurate skull segmentation of all investigated methods. This approach was also shown to be able to accurately delineate all other tissues relevant for EEG source analysis, that is, scalp, CSF, GM, and WM.

In conclusion, the work presented in this thesis makes it substantially more feasible to use realistic volume conductor models for EEG source analysis or in other applications, for example, magnetoencephalography based source analysis, or the simulation of transcranial magnetic or direct current stimulation. It does so by resolving open questions with regard to the construction and design of volume conductor models, and more importantly by proposing an automatic segmentation method that allows to create accurate and detailed head models with nearly no effort at all. The utilization of realistic head models will in the end lead to improvements in EEG

source analysis with respect to reliability, and to the significance of the source analysis results.

11.2 Outlook and Further Work

The experiments with real MRI data and the validation studies have shown that our proposed segmentation approaches perform very well. Especially the MRF approach is able to deliver a detailed and accurate volume conductor segmentation. Still, some improvements might be possible to even further improve the accuracy of the approach.

Some recent publications [14, 62] have shown that accurate segmentation results in various applications can be obtained using label propagation methods that in some way combine the labels of multiple reference segmentations that have been non-linearly aligned with the individual MRI. Instead of directly creating a segmentation such methods could also be used to generate improved tissue probability maps which then enhance the MRF approach. The improved probability maps might in particular be more suited to represent the anatomy of the skull base than the current atlas is. In addition, the label propagation methods could provide a better initial solution. This might be beneficial since the optimization methods used in the MRF segmentation are only converging to local optima and the segmentation accuracy, thus, depends on the quality of the initial solution.

Another aspect where the MRF segmentation might be improved is the differentiation of compact and cancellous bone. The validation studies have shown that in the investigated data sets compact bone can not reliably be distinguished from cancellous bone. Differentiating these tissues in a volume conductor model will, however, be advantageous for EEG source analysis [39]. One possible improvement in this regard might be to also differentiate compact and cancellous bone in the probabilistic atlas. Improvements for the cancellous bone segmentation can also be expected when using optimized MRI sequences. In our experiments on data sets with a T1- and a T2-weighted image only the T2-weighted image contained information on the presence of cancellous bone tissue because the T1-weighted image was recorded with suppression of the fat signal. Unfortunately, the sequence that was used to measure the T2-weighted image was not optimal, thus, producing very noisy images. In future work different T2 sequences will be tested and their parameters will be optimized to obtain a T2-weighted image with low noise and a good contrast between compact and cancellous bone.

The inspection of the MRF segmentation results did not reveal any errors that can be unambiguously attributed to a sub-optimal optimization of the a-posteriori probability. Nevertheless, it might be that using a global optimization procedure, for example, the SA algorithm, would yield better results than the currently used ICM method which is known to converge to the next local maximum only. A future study might, thus, investigate the effects of different optimization procedures on the segmentation results.

An important aspect greatly influencing the segmentation accuracy is the quality of the input MRI data. In particular, a good segmentation will not be possible for MRIs that are affected by a strong water-fat shift artifact. In these images the inner and outer skull boundary will be obstructed by the fat signal from the cancellous bone and the SCT layer, respectively. Furthermore, the layered structure visible in an MRI affected by the artifact does not resemble the layered structure encoded in our MRF model. For example, in the case where the fat signal is

shifted from superior to inferior a thin, low intensity layer is visible at the boundary between the SCT and the skin (Fig. 3.5b). This is because the signal from the SCT which mostly consists of fat is shifted downwards. The low intensity layer does not correspond to any of the layers in the MRF model and might, thus, be misclassified, for example, as compact bone. A segmentation approach might be able to account for the water-fat shift artifact by using an image generation model that treats the signal from fatty and aqueous tissue separately. In this model the observed image is then the superposition of a water image and a shifted fat image. The density of fatty tissue and the parameters of the gray value distribution of the fat image would then have to be estimated by the segmentation algorithm.

Bibliography

[1] Z. A. Acar and S. Makeig. Neuroelectromagnetic forward head modeling toolbox. *Journal of Neuroscience Methods*, 190:258–270, 2010.

[2] Z. A. Acar and S. Makeig. Effects of forward model errors on EEG source localization. *Brain Topography*, 26(3):378–396, July 2013.

[3] D. Adalsteinsson and J. Sethian. A fast level set method for propagating interfaces. *Journal of Computational Physics*, 118(2):269–277, 1995.

[4] M. Akhtari, H. C. Bryant, A. N. Mamelak, E. R. Flynn, L. Heller, J. J. Shih, M. Mandelkem, A. Matlachov, D. M. Ranken, E. D. Best, et al. Conductivities of three-layer live human skull. *Brain Topography*, 14(3):151–167, 2002.

[5] E. D. Angelini, T. Song, B. D. Mensh, and A. Laine. Multi-phase three-dimensional level set segmentation of brain MRI. In *SPIE International Symposium, Medical Imaging*, pages 318–326, 2004.

[6] B. Aubert-Broche, M. Griffin, G. Pike, A. Evans, and D. Collins. Twenty new digital brain phantoms for creation of validation image data bases. *IEEE Transactions on Medical Imaging*, 25(11):1410 –1416, Nov. 2006.

[7] P. D. Ax, D. Weishaupt, and V. Kochli. *How does MRI work?. An Introduction to the Physics and Function of Magnetic Resonance Imaging*. Springer-Verlag Berlin and Heidelberg GmbH & Co. K, 3rd edition, Oct. 2003.

[8] S. Baumann, D. Wozny, S. Kelly, and F. Meno. The electrical conductivity of human cerebrospinal fluid at body temperature. *IEEE Transactions on Biomedical Engineering*, 44(3):220–223, Mar. 1997.

[9] P. Bazin and D. Pham. Topology preserving tissue classification with fast marching and topology templates. In *Information Processing in Medical Imaging: 19th International Conference, IPMI 2005, Glenwood Springs, CO, USA, July 10-15, 2005: Proceedings*. Springer, 2005.

[10] E. Bellon, E. Haacke, P. Coleman, D. Sacco, D. Steiger, and R. Gangarosa. MR artifacts: a review. *Am. J. Roentgenol.*, 147(6):1271–1281, Dec. 1986.

[11] C. G. Bénar and J. Gotman. Modeling of post-surgical brain and skull defects in the EEG inverse problem with the boundary element method. *Clinical Neurophysiology*, 113(1):48–56, 2002.

[12] P. Berg and M. Scherg. A fast method for forward computation of multiple-shell spherical head models. *Electroencephalography and Clinical Neurophysiology*, 90(1):58–64, Jan. 1994.

[13] H. Berger. Über das elektrenkephalogramm des menschen. *Archiv für Psychiatrie und Nervenkrankheiten*, 94(1):16–60, Dec. 1931.

[14] A. Bertelsen, A. Munoz-Barrutia, S. Tejada, C. Ortiz-de Solórzano, and D. Borro. Segmentation of the skull from MR T1 images using label fusion. In *Proceedings of the 32nd Annual International Conference of the IEEE EMBS*, pages 3121–3124, 2010.

[15] G. Bertrand. Simple points, topological numbers and geodesic neighborhoods in cubic grids. *Pattern Recognition Letters*, 15(10):1003–1011, Oct. 1994.

[16] O. Bertrand, M. Thevenet, and F. Perrin. 3D finite element method in brain electrical activity studies. In *Biomagnetic Localization and 3D Modelling*, pages 154–171. Report of the Dep. of Tech.Physics, Helsinki University of Technology, 1991.

[17] BESA GmbH. BESA MRI. http://www.besa.de, Nov. 2012.

[18] J. Besag. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, 36(2):192–236, Jan. 1974.

[19] J. Besag. On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society. Series B (Methodological)*, 48(3):259–302, 1986.

[20] D. Braess. *Finite Elemente*. Springer, 2007.

[21] R. Brette and A. Destexhe. *Handbook of Neural Activity Measurement*. Cambridge University Press, 1 edition, Aug. 2012.

[22] B. H. Brinkmann, T. J. O'Brien, M. A. Dresner, T. D. Lagerlund, F. W. Sharbrough, and R. A. Robb. Scalp-recorded EEG localization in MRI volume data. *Brain Topography*, 10(4):245–253, June 1998.

[23] V. Brodbeck, L. Spinelli, A. M. Lascano, M. Wissmeier, M.-I. Vargas, S. Vulliemoz, C. Pollo, K. Schaller, C. M. Michel, and M. Seeck. Electroencephalographic source imaging: a prospective study of 152 operated epileptic patients. *Brain*, 134(10):2887–2897, Oct. 2011.

[24] P. Bruno, F. Vatta, S. Mininel, and P. Inchingolo. Head model extension for the study of bioelectric phenomena. *Biomed Sci Instrum*, 39:59–64, 2003.

[25] P. Bruno, F. Vatta, S. Mininel, and P. Inchingolo. Referenced EEG and head volume conductor model: geometry and parametrical setting. In *Engineering in Medicine and Biology Society, 2004. IEMBS '04. 26th Annual International Conference of the IEEE*, volume 1, pages 833–836, 2004.

[26] H. Buchner, G. Knoll, M. Fuchs, A. Rienäcker, R. Beckmann, M. Wagner, J. Silny, and J. Pesch. Inverse localization of electric dipole current sources in finite element models of the human head. *Electroencephalography and Clinical Neurophysiology*, 102:267–278, 1997.

[27] W. Burger and M. J. Burge. *Digitale Bildverarbeitung*. 2005.

[28] T. Buzug. *Computed Tomography. From Photon Statistic to Modern Cone-Beam CT*. Springer, Berlin, 2008.

[29] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *International Journal of Computer Vision*, 22(1):61–79, 1997.

[30] T. F. Chan and L. A. Vese. Active contours without edges. *Image Processing, IEEE Transactions on*, 10(2):266–277, 2001.

[31] D. L. Collins, A. P. Zijdenbos, W. F. C. Baaré, and A. C. Evans. ANIMAL+INSECT: improved cortical structure segmentation. In A. Kuba, M. Šáamal, and A. Todd-Pokropek, editors, *Information Processing in Medical Imaging*, number 1613 in Lecture Notes in Computer Science, pages 210–223. Springer Berlin Heidelberg, Jan. 1999.

[32] M. Cook and Z. Koles. A high-resolution anisotropic finite-volume head model for EEG source analysis. In *28th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 2006. EMBS '06*, pages 4536 –4539, Sept. 2006.

[33] T. F. Cootes, C. J. Taylor, D. H. Cooper, J. Graham, et al. Active shape models-their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.

[34] W. Crum, L. Griffin, D. Hill, and D. Hawkes. Zen and the art of medical image registration: correspondence, homology, and quality. *NeuroImage*, 20(3):1425–1437, Nov. 2003.

[35] V. Csépe, C. Pantev, M. Hoke, S. Hampson, and B. Ross. Evoked magnetic responses of the human auditory cortex to minor pitch changes: localization of the mismatch field. *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section*, 84(6):538–548, Nov. 1992.

[36] B. N. Cuffin. Effects of local variations in skull and scalp thickness on EEG's and MEG's. *IEEE Transactions on Biomedical Engineering*, 40(1):42 – 48, 1993.

[37] B. N. Cuffin. EEG localization accuracy improvements using realistically shaped head models. *IEEE Transactions on Biomedical Engineering*, 43(3):299–303, Mar. 1996.

[38] A. M. Dale, B. Fischl, and M. I. Sereno. Cortical surface-based analysis: I. segmentation and surface reconstruction. *NeuroImage*, 9(2):179–194, Feb. 1999.

[39] M. Dannhauer, B. Lanfer, C. H. Wolters, and T. R. Knösche. Modelling of the human skull in EEG source analysis. *Human Brain Mapping*, 2011.

[40] J. de Munck. The potential distribution in a layered anisotropic spheroidal volume conductor. *Journal of Applied Physics*, 64:464, 1988.

[41] J. C. de Munck. A linear discretization of the volume conductor boundary integral equation using analytically integrated elements. *IEEE Transactions on Biomedical Engineering*, 39(9):986–990, 1992.

[42] E. M. Delfaut, J. Beltran, G. Johnson, J. Rousseau, X. Marchandise, and A. Cotten. Fat suppression in MR imaging: Techniques and pitfalls. *Radiographics*, 19(2):373–382, 1999.

[43] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38, Jan. 1977.

[44] L. Dice. Measures of the amount of ecologic association between species. *Ecology*, 26(3):297 – 302, 1945.

[45] J. Diedrichsen, J. H. Balsters, J. Flavell, E. Cussans, and N. Ramnani. A probabilistic MR atlas of the human cerebellum. *NeuroImage*, 46(1):39–46, May 2009.

[46] B. Dogdas, D. W. Shattuck, and R. M. Leahy. Segmentation of skull and scalp in 3-D human MRI using mathematical morphology. *Human Brain Mapping*, 26(4), 2005.

[47] F. Drechsler, C. Wolters, T. Dierkes, H. Si, and L. Grasedyck. A full subtraction approach for finite element method based source analysis using constrained delaunay tetrahedralisation. *NeuroImage*, 46(4):1055–1065, July 2009.

[48] 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(3):341–355, Jan. 2002.

[49] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6):721–741, Nov. 1984.

[50] N. G. Gencer and C. E. Acar. Sensitivity of EEG and MEG measurements to tissue conductivity. *Physics in Medicine and Biology*, 49:701–717, 2004.

[51] H. Gudbjartsson and S. Patz. The Rician distribution of noisy MRI data. *Magnetic Resonance in Medicine*, 34(6):910–914, 1995.

[52] D. Güllmar, J. Haueisen, M. Eiselt, F. Gießler, L. Flemming, A. Anwander, T. Knösche, C. Wolters, M. Dümpelmann, D. Tuch, et al. Influence of anisotropic conductivity on EEG source reconstruction: Investigations in a rabbit model. *IEEE Transactions on Biomedical Engineering*, 53(9):1841–1850, 2006.

[53] D. Güllmar, J. Haueisen, and J. R. Reichenbach. Influence of anisotropic electrical conductivity in white matter tissue on the EEG/MEG forward and inverse solution. a high-resolution whole head simulation study. *NeuroImage*, 51(1):145–163, May 2010.

[54] H. Hallez, B. Vanrumste, P. V. Hese, Y. D'Asseler, I. Lemahieu, and R. V. de Walle. A finite difference method with reciprocity used to incorporate anisotropy in electroencephalogram dipole source localization. *Physics in Medicine and Biology*, 50:3787–3806, 2005.

[55] M. Hämäläinen, R. Hari, R.Ilmoniemi, J. Knuutila, and O.Lounasmaa. Magnetoencephalography: theory, instrumentation, and applications to noninvasive studies of the working human brain. *Reviews of Modern Physics*, 65:413–497, 1993.

[56] M. Hämäläinen and J. Sarvas. Realistic conductivity geometry model of the human head for interpretation of neuromagnetic data. *IEEE Transactions on Biomedical Engineering*, 36(2):165–171, 1989.

[57] J. M. Hammersley and P. Clifford. Markov fields on finite graphs and lattices. Unpublished manuscript, 1971.

[58] X. Han, C. Xu, D. Tosun, and J. L. Prince. Cortical surface reconstruction using a topology preserving geometric deformable model. *Mathematical Methods in Biomedical Image Analysis, IEEE Workshop on*, 0:213, 2001.

[59] S. E. Harms, D. P. Flamig, K. L. Hesley, M. D. Meiches, R. A. Jensen, W. P. Evans, D. A. Savino, and R. V. Wells. MR imaging of the breast with rotating delivery of excitation off resonance: clinical experience with pathologic correlation. *Radiology*, 187(2):493, 1993.

[60] U. Hartmann and G. Berti. VGRID: A fast and robust octree-based 3D mesh generator for unstructured grids. [130], 2000-2003.

[61] J. Haueisen, C. Ramon, P. Czapski, and M. Eiselt. On the influence of volume currents and extended sources on neuromagnetic fields: A simulation study. *Annals of Biomedical Engineering*, 23(6):728–739, Nov. 1995.

[62] R. A. Heckemann, S. Keihaninejad, P. Aljabar, D. Rueckert, J. V. Hajnal, A. Hammers, and A. D. N. Initiative. Improving intersubject image registration using tissue-class information benefits robustness and accuracy of multi-atlas based anatomical segmentation. *NeuroImage*, 51(1):221, 2010.

[63] T. Heinonen, H. Eskola, P. Dastidar, P. Laarne, and J. Malmivuo. Segmentation of T1 MR scans for reconstruction of resistive head models. *Computer Methods and Programs in Biomedicine*, 54(3):173–181, Nov. 1997.

[64] S. C. Herholz, C. Lappe, and C. Pantev. Looking for a pattern: An MEG study on the abstract mismatch negativity in musicians and nonmusicians. *BMC Neuroscience*, 10(1):42, Apr. 2009. PMID: 19405970.

[65] G. N. Hounsfield. Computerized transverse axial scanning (tomography). 1. description of system. *The British Journal of Radiology*, 46(552):1016–1022, Dec. 1973. PMID: 4757352.

[66] M. Hämäläinen and R. Ilmoniemi. Interpreting magnetic fields of the brain: minimum norm estimates. *Medical and Biological Engineering and Computing*, 32(1):35–42, 1994.

[67] M. Jenkinson, M. Pechaud, and S. Smith. BET2: MR-based estimation of brain, skull and scalp surfaces. In *Eleventh annual meeting of the Organization for Human Brain Mapping*, page 12–16, 2005.

[68] M. W. Jones, J. A. Baerentzen, and M. Sramek. 3D distance fields: A survey of techniques and applications. *Visualization and Computer Graphics, IEEE Transactions on*, 12(4):581–599, 2006.

[69] J. Kaipio and E. Somersalo. Statistical inverse problems: Discretization, model reduction and inverse crimes. *Journal of Computational and Applied Mathematics*, 198(2):493–504, Jan. 2007.

[70] W. A. Kalender, W. Seissler, and P. Vock. Single-breath-hold spiral volumetric CT by continous patient translation and scanner rotation. *Radiology*, 173(P):414, 1989.

[71] T. Kapur, W. Eric, L. Grimson, R. Kikinis, W. Wells, W. Wells, A. Colchester, and S. Delp. Enhanced spatial priors for segmentation of magnetic resonance imagery. In W. M. Wells, A. Colchester, and S. Delp, editors, *Medical Image Computing and Computer-Assisted Interventation — MICCAI'98*, volume 1496 of *Lecture Notes in Computer Science*, pages 457–468. Springer Berlin / Heidelberg, 1998.

[72] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988.

[73] R. Kikinis, M. E. Shenton, G. Gerig, J. Martin, M. Anderson, D. Metcalf, C. R. G. Guttmann, R. W. McCarley, F. A. Jolesz, W. Lorensen, and H. Cline. Routine quantitative analysis of brain and cerebrospinal fluid spaces with MR imaging. *Journal of Magnetic Resonance Imaging*, 2(6):619–629, 1992.

[74] J. Kybic, M. Clerc, T. Abboud, O. Faugeras, R. Keriven, and T. Papadopoulo. A common formalism for the integral formulations of the forward EEG problem. *IEEE Transactions on Medical Imaging*, 24(1):12–18, 2005.

[75] B. Lanfer, I. Paul-Jordanov, M. Scherg, and C. H. Wolters. Influence of interior cerebrospinal fluid compartments on EEG source analysis. *Biomedical Engineering*, 57 (Suppl. 1):623–626, 2012.

[76] B. Lanfer, C. Röer, M. Scherg, S. Rampp, C. Kellinghaus, and C. Wolters. Influence of a silastic ECoG grid on EEG/ECoG based source analysis. *Brain Topography*, pages 1–17, 2012.

[77] B. Lanfer, M. Scherg, M. Dannhauer, T. Knösche, M. Burger, and C. Wolters. Influences of skull segmentation inaccuracies on EEG source analysis. *NeuroImage*, 62(1):418–431, Aug. 2012.

[78] S. Law. Thickness and resistivity variations over the upper surface of the human skull. *Brain Topography*, 6(2):99–109, 1993.

[79] A. J. Leal, V. Passão, E. Calado, J. P. Vieira, and J. P. Silva Cunha. Interictal spike EEG source analysis in hypothalamic hamartoma epilepsy. *Clinical Neurophysiology*, 113(12):1961–1969, Dec. 2002.

[80] S. Lew, C. Wolters, T. Dierkes, C. Röer, and R. MacLeod. Accuracy and run-time comparison for different potential approaches and iterative solvers in finite element method based EEG source analysis. *Applied Numerical Mathematics*, 59(8):1970–1988, 2009. http://dx.doi.org/10.1016/j.apnum.2009.02.006, NIHMSID 120338, PMCID: PMC2791331.

[81] J. Li, K. Wang, S. Zhu, and B. He. Effects of holes on EEG forward solutions using a realistic geometry head model. *Journal of Neural Engineering*, 4(3):197 – 204, 2007.

[82] S. Z. Li. *Markov random field modeling in image analysis*. Springer, Mar. 2009.

[83] W. Lorensen and H. Cline. Marching cubes: A high resolution 3d surface construction algorithm. *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 163–169, 1987.

[84] M. Lüthi, A. Lerch, T. Albrecht, Z. Krol, and T. Vetter. A hierarchical, multi-resolution approach for model-based skull-segmentation in mri volumes. Technical report, University of Basel, 2009.

[85] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. Multimodality image registration by maximization of mutualinformation. *IEEE transactions on Medical Imaging*, 16(2):187–198, 1997.

[86] R. Malladi, J. Sethian, and B. Vemuri. Shape modeling with front propagation: a level set approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(2):158 –175, Feb. 1995.

[87] T. McInerney and D. Terzopoulos. Deformable models in medical image analysis. In *Proceedings of the Workshop on Mathematical Methods in Biomedical Image Analysis, 1996*, pages 171 –180, June 1996.

[88] P. McVeigh, A. Bostan, and D. Cheyne. Study of conducting volume boundary influence on source localization using a realistic MEG phantom. *International Congress Series*, 1300:153–156, June 2007.

[89] J. W. H. Meijs, O. W. Weier, M. J. Peters, and A. van Oosterom. On the numerical accuracy of the boundary element method. *IEEE Transactions on Biomedical Engineering*, 36(10):1038–1049, 1989.

[90] C. M. Michel, M. M. Murray, G. Lantz, S. Gonzalez, L. Spinelli, and R. Grave de Peralta. EEG source imaging. *Clinical Neurophysiology*, 115(10):2195–2222, Oct. 2004.

[91] T. Moon. The expectation-maximization algorithm. *IEEE Signal Processing Magazine*, 13(6):47 –60, Nov. 1996.

[92] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on Pure and Applied Mathematics*, 42(5):577–685, 1989.

[93] S. Murakami and Y. Okada. Contributions of principal neocortical neurons to magnetoencephalography and electroencephalography signals. *The Journal of Physiology*, 575(3):925–936, Sept. 2006.

[94] J. A. Nelder and R. Mead. A simplex method for function minimization. *The Computer Journal*, 7(4):308, 1965.

[95] M. R. Nuwer, G. Comi, R. Emerson, A. Fuglsang-Frederiksen, J.-M. Guérit, H. Hinrichs, A. Ikeda, F. Jose C. Luccas, and P. Rappelsburger. IFCN standards for digital recording of clinical EEG. *Electroencephalography and Clinical Neurophysiology*, 106(3):259–261, Mar. 1998.

[96] R. Näätänen, P. Paavilainen, T. Rinne, and K. Alho. The mismatch negativity (MMN) in basic research of central auditory processing: a review. *Clinical neurophysiology: official journal of the International Federation of Clinical Neurophysiology*, 118(12):2544–2590, Dec. 2007. PMID: 17931964.

[97] T. Oostendorp, J. Delbeke, and D. Stegeman. The conductivity of the human skull: results of in vivo and in vitro measurements. *IEEE Transactions on Biomedical Engineering*, 47(11):1487–1492, 2000.

[98] R. Oostenveld and T. F. Oostendorp. Validating the boundary element method for forward and inverse EEG computations in the presence of a hole in the skull. *Human Brain Mapping*, 17(3):179–192, 2002.

[99] S. Osher and R. Fedkiw. *Level set methods and dynamic implicit surfaces*. Springer, 2002.

[100] S. Osher and J. A. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on hamilton-jacobi formulations. *Journal of Computational Physics*, 79(1):12–49, Nov. 1988.

[101] C. Pantev, M. Hoke, B. Lutkenhoner, and K. Lehnertz. Tonotopic organization of the auditory cortex: pitch versus frequency representation. *Science*, 246(4929):486–488, 1989.

[102] C. Pantev, B. Ross, T. Fujioka, L. J. Trainor, M. Schulte, and M. Schulz. Music and learning-induced cortical plasticity. *Annals of the New York Academy of Sciences*, 999(1):438–450, 2003.

[103] H. Park, P. Bland, and C. Meyer. Construction of an abdominal probabilistic atlas and its application in segmentation. *IEEE Transactions on Medical Imaging*, 22(4):483 –492, Apr. 2003.

[104] S. B. Petersen, R. N. Muller, and P. A. Rinck. *An Introduction to Biomedical Nuclear Magnetic Resonance*. Thieme Georg Verlag, Dec. 1985.

[105] D. L. Pham and J. L. Prince. An adaptive fuzzy c-means algorithm for image segmentation in the presence of intensity inhomogeneities. *Pattern Recognition Letters*, 20(1):57–68, Jan. 1999.

[106] R. Plonsey and D. Heppner. Considerations of quasi-stationarity in electrophysiological systems. *Bulletin of Mathematical Biology*, 29(4):657–664, 1967.

[107] J. C. Pruessner, L. M. Li, W. Serles, M. Pruessner, D. L. Collins, N. Kabani, S. Lupien, and A. C. Evans. Volumetry of hippocampus and amygdala with high-resolution MRI and three-dimensional analysis software: Minimizing the discrepancies between laboratories. *Cerebral Cortex*, 10(4):433–442, Apr. 2000.

[108] C. Pudney. Distance-ordered homotopic thinning: A skeletonization algorithm for 3D digital images. *Computer Vision and Image Understanding*, 72(3):404–413, Dec. 1998.

[109] J. Radon. Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten. *Classic papers in modern diagnostic radiology*, page 5, 2005.

[110] C. Ramon, P. Schimpf, and J. Haueisen. Influence of head models on EEG simulations and inverse source localizations. *BioMedical Engineering OnLine*, 5(1):10–10, 2006.

[111] C. Ramon, P. Schimpf, J. Haueisen, M. Holmes, and A. Ishimaru. Role of soft bone, CSF and gray matter in EEG simulations. *Brain Topography*, 16(4):245–248, 2004.

[112] M. Reiser, C. Born, F.-P. Kuhn, and J. Debus. *Radiologie : 104 Tabellen*. Thieme, Stuttgart, 2006.

[113] M. Reiser and W. Semmler. *Magnetresonanztomographie*. Springer-Verlag GmbH, 2nd edition, Apr. 1997.

[114] J. K. Rice, C. Rorden, J. S. Little, and L. C. Parra. Subject position affects EEG magnitudes. *NeuroImage*, 64:476–484, Jan. 2013.

[115] H. Rifai, I. Bloch, S. Hutchinson, J. Wiart, and L. Garnero. Segmentation of the skull in MRI volumes using deformable model and taking the partial volume effect into account. *Medical Image Analysis*, 4(3):219–233, 2000.

[116] C. Robert. *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*. Springer, Aug. 2007.

[117] N. Roche-Labarbe, A. Aarabi, G. Kongolo, C. Gondry-Jouet, M. Dümpelmann, R. Grebe, and F. Wallois. High-resolution electroencephalography and source localization in neonates. *Human Brain Mapping*, 29(2):167–176, Feb. 2008.

[118] M. Rullmann, A. Anwander, M. Dannhauer, S. Warfield, F. Duffy, and C. Wolters. EEG source analysis of epileptiform activity using a 1 mm anisotropic hexahedra finite element head model. *NeuroImage*, 44(2):399–410, Jan. 2009.

[119] L. Ruthotto, H. Kugel, J. Olesch, B. Fischer, J. Modersitzki, M. Burger, and C. H. Wolters. Diffeomorphic susceptibility artifact correction of diffusion-weighted magnetic resonance images. *Physics in Medicine and Biology*, 57(18):5715, 2012.

[120] R. Sadleir and A. Argibay. Modeling skull electrical properties. *Annals of Biomedical Engineering*, 35(10):1699–1712, 2007.

[121] K. A. Salayev, N. Nakasato, M. Ishitobi, H. Shamoto, A. Kanno, and K. Iinuma. Spike orientation may predict epileptogenic side across cerebral sulci containing the estimated equivalent dipole. *Clinical neurophysiology*, 117(8):1836–1843, 2006.

[122] M. Scherg and D. Von Cramon. Evoked dipole source potentials of the human auditory cortex. *Electroencephalogr Clin Neurophysiol*, 65(5):344–60, 1986.

[123] T. H. Schiebler and W. Schmidt. *Anatomie*. Springer, Berlin, 8th edition, Sept. 2002.

[124] F. Schmidt and H.-G. Schaible, editors. *Neuro- und Sinnesphysiologie*. Springer-Lehrbuch. Springer, 2006. DOI: 10.1007/3-540-29491-0.

[125] SCIRun. A scientific computing problem solving environment, 2011. Scientific Computing and Imaging Institute (SCI).

[126] M. Sekino, H. Ohsaki, S. Yamaguchi-Sekino, N. Iriguchi, and S. Ueno. Low-frequency conductivity tensor of rat brain tissues inferred from diffusion MRI. *Bioelectromagnetics*, 30(6):489–499, 2009.

[127] J. K. Seo and E. J. Woo. Magnetic resonance electrical impedance tomography (MREIT). *SIAM Review*, 53(1):40–68, Jan. 2011.

[128] J. A. Sethian. A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences*, 93(4):1591 –1595, Feb. 1996.

[129] D. W. Shattuck, G. Prasad, M. Mirza, K. L. Narr, and A. W. Toga. Online resource for validation of brain segmentation methods. *NeuroImage*, 45(2):431–439, Apr. 2009. PMID: 19073267 PMCID: PMC2757629.

[130] SimBio Development Group. SimBio: A generic environment for bio-numerical simulations. online, https://www.mrt.uni-jena.de/simbio, accessed June 15, 2012.

[131] A. Simmons, P. S. Tofts, G. J. Barker, and S. R. Arridge. Sources of intensity nonuniformity in spin echo images at 1.5 T. *Magnetic Resonance in Medicine*, 32(1):121–128, July 1994.

[132] J. G. Sled and G. B. Pike. Standing-wave and RF penetration artifacts caused by elliptic geometry: an electrodynamic analysis of MRI. *IEEE Transactions on Medical Imaging*, 17(4):653–662, Aug. 1998.

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

[134] S. M. Smith. Fast robust automated brain extraction. *Human Brain Mapping*, 17(3):143–155, 2002.

[135] O. Steinsträter, S. Sillekens, M. Junghoefer, M. Burger, and C. H. Wolters. Sensitivity of beamformer source analysis to deficiencies in forward modeling. *Human Brain Mapping*, 31(12):1907–1927, Dec. 2010.

[136] M. Stenroos and J. Sarvas. Bioelectromagnetic forward problem: isolated source approach revis(it)ed. *Physics in Medicine and Biology*, 57(11):3517, June 2012.

[137] N. Tanaka, M. S. Hämäläinen, S. P. Ahlfors, H. Liu, J. R. Madsen, B. F. Bourgeois, J. W. Lee, B. A. Dworetzky, J. W. Belliveau, and S. M. Stufflebeam. Propagation of epileptic spikes reconstructed from spatiotemporal magnetoencephalographic and electroencephalographic source analysis. *NeuroImage*, 50(1):217–222, Mar. 2010.

[138] G. Taubin. Curve and surface smoothing without shrinkage. In *Proceedings of the Fifth International Conference on Computer Vision, 1995*, pages 852 –857, June 1995.

[139] D. M. Thomasson, D. E. Purdy, J. R. Moore, and P. Finn. Phase-modulated binomial Spatial-Spectral pulse design for Spin-Echo applications. In *Proceedings of the Society of Magnetic Resonance, third scientific meeting and exhibition, and the European Society for Magnetic Resonance in Medicine and Biology, twelfth annual meeting and exhibition, Nice, France*, volume 1, page 561, 1995.

[140] D. S. Tuch, V. J. Wedeen, A. M. Dale, J. S. George, and J. W. Belliveau. Conductivity tensor mapping of the human brain using diffusion tensor MRI. *Proceedings of the National Academy of Sciences*, 98(20):11697–11701, 2001.

[141] S. P. van den Broek, F. Reinders, M. Donderwinkel, and M. J. Peters. Volume conduction effects in EEG and MEG. *Electroencephalography and Clinical Neurophysiology*, 106(6):522–534, June 1998.

[142] B. Vanrumste, G. V. Hoey, R. V. de Walle, M. D'Havé, I. Lemahieu, and P. Boon. Dipole location errors in electroencephalogram source analysis due to volume conductor model errors. *Medical and Biological Engineering and Computing*, 38(5):528–534, 2000.

[143] L. Vese and T. Chan. A multiphase level set framework for image segmentation using the mumford and shah model. *International Journal of Computer Vision*, 50(3):271–293, 2002.

[144] P. Vock, H. Jung, and W. A. Kalender. Single-breathhold spiral volumetric CT of the lung. *Radiology*, 173(P):400, 1989.

[145] N. von Ellenrieder, C. H. Muravchik, and A. Nehorai. Effects of geometric head model perturbations on the EEG forward and inverse problems. *IEEE Transactions on Biomedical Engineering*, 53(3):421–429, Mar. 2006.

[146] J. Vorwerk, M. Clerc, M. Burger, and C. H. Wolters. Comparison of boundary element and finite element approaches to the EEG forward problem. *Biomedizinische Technik*, 57:1, 2012.

[147] D. Wang, L. Shi, W. C. Chu, J. C. Cheng, and P. A. Heng. Segmentation of human skull in MRI using statistical shape information from CT data. *Journal of Magnetic Resonance Imaging*, 30(3):490–498, 2009.

[148] G. Wang, G. Worrell, L. Yang, C. Wilke, and B. He. Interictal spike analysis of high-density EEG in patients with partial epilepsy. *Clinical Neurophysiology*, 122(6):1098–1105, June 2011.

[149] J. C. Weinreb, L. Brateman, E. E. Babcock, K. R. Maravilla, J. M. Cohen, and S. D. Horner. Chemical shift artifact in clinical magnetic resonance images at 0.35 T. *AJR. American Journal of Roentgenology*, 145(1):183–185, July 1985. PMID: 3873844.

[150] D. Weinstein, L. Zhukov, and C. Johnson. Lead-field bases for electroencephalography source imaging. *Annals of Biomedical Engineering*, 28(9):1059–1065, Sept. 2000.

[151] W. M. Wells, W. E. L. Grimson, R. Kikinis, and F. A. Jolesz. Adaptive segmentation of MRI data. *IEEE Transactions on Medical Imaging*, 15(4):429–442, 1996.

[152] W. M. Wells, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis. Multi-modal volume registration by maximization of mutual information. *Medical image analysis*, 1(1):35–51, 1996.

[153] K. Wendel, N. Narra, M. Hannula, P. Kauppinen, and J. Malmivuo. The influence of CSF on EEG sensitivity distributions of multilayered head models. *IEEE Transactions on Biomedical Engineering*, 55(4):1454–1456, 2008.

[154] K. Wendel, O. Väisänen, J. Malmivuo, N. Gencer, B. Vanrumste, P. Durka, R. Magjarevic, S. Supek, M. Pascu, H. Fontenelle, et al. EEG/MEG source imaging: methods, challenges, and open issues. *Computational intelligence and neuroscience*, 2009:13, 2009.

[155] C. H. Wolters. *Influence of tissue conductivity inhomogeneity and anisotropy on EEG/MEG based source localization in the human brain*. PhD thesis, 2003.

[156] C. H. Wolters, A. Anwander, X. Tricoche, D. Weinstein, M. A. Koch, and R. S. MacLeod. Influence of tissue conductivity anisotropy on EEG/MEG field and return current computation in a realistic head model: a simulation and visualization study using high-resolution finite element modeling. *NeuroImage*, 30(3):813–826, 2006.

[157] C. H. Wolters, L. Grasedyck, and W. Hackbusch. Efficient computation of lead field bases and influence matrix for the FEM-based EEG and MEG inverse problem. *Inverse problems*, 20:1099, 2004.

[158] C. H. Wolters, H. Köstler, C. Möller, J. Härdtlein, L. Grasedyck, and W. Hackbusch. Numerical mathematics of the subtraction method for the modeling of a current dipole in EEG source reconstruction using finite element head models. *SIAM Journal on Scientific Computing*, 30(1):24–45, 2007.

[159] C. F. J. Wu. On the convergence properties of the EM algorithm. *The Annals of Statistics*, 11(1):95–103, Mar. 1983. Mathematical Reviews number (MathSciNet): MR684867; Zentralblatt MATH identifier: 0517.62035.

[160] Y. Yan, P. L. Nunez, and R. T. Hart. Finite-element model of the human head: scalp potentials due to dipole sources. *Medical and Biological Engineering and Computing*, 29(5):475 – 481, 1991.

[161] P. A. Yushkevich, J. Piven, H. C. Hazlett, R. G. Smith, S. Ho, J. C. Gee, and G. Gerig. User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *NeuroImage*, 31(3):1116–1128, July 2006.

[162] B. Yvert, O. Bertrand, M. Thévenet, J. F. Echallier, and J. Pernier. A systematic evaluation of the spherical model accuracy in EEG dipole localization. *Electroencephalography and Clinical Neurophysiology*, 102(5):452–459, May 1997.

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

[164] H.-K. Zhao, T. Chan, B. Merriman, and S. Osher. A variational level set approach to multiphase motion. *Journal of Computational Physics*, 127(1):179–195, Aug. 1996.

[165] X. Zhou, T. Kitagawa, K. Okuo, T. Hara, H. Fujita, R. Yokoyama, M. Kanematsu, and H. Hoshi. Construction of a probabilistic atlas for automated liver segmentation in non-contrast torso CT images. *International Congress Series*, 1281(0):1169–1174, May 2005.

[166] S. C. Zhu and A. Yuille. Region competition: Unifying snakes, region growing, and Bayes/MDL for multiband image segmentation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 18(9):884–900, 1996.

List of Figures

List of Tables

Lebenslauf

Name:	Benjamin Lanfer	
Geburtsdatum:	08.06.1983 in Vreden	
Familienstand:	verheiratet	
Eltern:	Norbert Lanfer; Gabriele Wilhelmine Lanfer, geb. Fischer	
Schulbildung:	Grundschule:	Von-Galen-Grundschule, Gescher, von August 1989 bis Juni 1993
	Gymnasium:	Gymnasium Nepomucenum, Coesfeld, von August 1993 bis Juni 2002
Hochschulreife (Abitur):	am 03. Juni 2002 in Coesfeld	
Studium:	Physik, Westfälische Wilhelms-Universität Münster, von Oktober 2002 bis September 2007	
Promotionsstudiengang:	Mathematik, Westfälische Wilhelms-Universität Münster, von Oktober 2007 bis März 2014	
Prüfungen:	Diplom im Fach Physik am 25. September 2007, Westfälische Wilhelms-Universität Münster	
Tätigkeiten:	Studentische Hilfskraft, Institut für Biomagnetismus und Biosignalanalyse, WWU Münster, von August 2005 bis Mai 2006	
	Software-Entwickler, BESA GmbH, Gräfelfing, von Oktober 2007 bis heute	
Begin der Dissertation:	Oktober 2007; Betreuer: Priv.-Doz. Dr. Carsten H. Wolters; Mitbetreuer: Prof. Dr. Martin Burger	