

Developing and Validating Realistic Head Models for Forward Calculation of Electromagnetic Fields with Applications in EEG

Nielsen, Jesper Duemose

Publication date: 2022

Document Version Publisher's PDF, also known as Version of record

Link back to DTU Orbit

Citation (APA): Nielsen, J. D. (2022). *Developing and Validating Realistic Head Models for Forward Calculation of Electromagnetic Fields with Applications in EEG*. Technical University of Denmark.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

• Users may download and print one copy of any publication from the public portal for the purpose of private study or research.

- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

PhD Thesis

Developing and Validating Realistic Head Models for Forward Calculation of Electromagnetic Fields with Applications in EEG

JESPER DUEMOSE NIELSEN



Technical University of Denmark September 4, 2022

Jesper Duemose Nielsen Developing and Validating Realistic Head Hodels for Forward Calculation of Electromagnetic Fields with Applications in EEG PhD thesis September 4, 2022

Supervisors Kristoffer Hougaard Madsen, DTU, DRCMR Axel Thielscher, DTU, DRCMR Lars Kai Hansen, DTU Rong Xue, CAS

Technical University of Denmark Department of Applied Mathematics and Computer Science Section for Cognitive Systems Richard Petersens Plads, building 324 2800 Kongens Lyngby, Denmark https://www.compute.dtu.dk

Summary

Localizing the sources responsible for generating an observed electroencephalography (EEG) signal can be useful in several contexts. For example, it may help guide presurgical planning in epilepsy or improve connectivity estimates obtained from EEG data. However, EEG source analysis is an ill-posed, inverse problem and in order to solve this, we first need to solve the corresponding forward problem. The forward problem is the task of determining the measurements corresponding to individual neural sources. The accuracy with which we can solve this problem is determined (in part) by our ability to construct a valid volume conductor model. Constructing such a model includes modeling of the anatomy of the human head as well as the conductive properties of different tissue compartments. We also need to specify the positions of the electrodes relative to the head and choose a suitable resolution for our model.

In this thesis we seek to investigate how the different aspects of forward modeling affect not only the forward solution itself but also the corresponding inverse solution. We focus on the anatomical accuracy of the volume conductor model as well as the accuracy with which electrode positions are specified. As to the inverse solution, we will primarily be concerned with assessing localization accuracy (as opposed to, for example, determining the strength of a source activation).

First, we present a pipeline (available in SimNIBS) for generating reasonably realistic anatomical models of the human head with particular emphasis on reconstructing the skull. This pipeline is based on SPM12 and CAT12 where the latter is used to improve the accuracy of the brain tissue segmentation. We compare with existing tools from FSL and BrainSuite and show that the new pipeline improves skull segmentation, particularly when using both T1- and T2-weighted magnetic resonance imaging (MRI) scans. We also show that it is important to ensure that the structural images being used are of high quality (e.g., by minimizing artifacts such as fat shift) in order to facilitate good segmentation results.

Subsequently, we compare forward solutions generated by SimNIBS, MNE-Python, and FieldTrip. The major difference between these models is the extent to which they are able to capture the underlying anatomy with SimNIBS generally being more accurate than the latter two. We find increased topographic and magnitude errors of the forward solutions from MNE-Python and FieldTrip compared to SimNIBS throughout most of the brain suggesting large overall differences in the forward solutions. We also compare with a model based on a template anatomy. This too, resulted in substantial errors. In addition to comparing different pipelines, we also compare different ways of specifying the electrode positions. In particular, we compared digitizing the electrodes against using a template description of the electrode positions which is adapted to each subject. We investigated two templates; one which we created by digitizing the relevant cap on a 3D printed model of the MNI head and another which used the positions specified by the manufacturer. We found substantial topographic errors when using the manufacturer layout especially in occipital and parietal areas which was also where we found the largest errors in electrode locations (compared to the digitized positions). Our custom template performed better in these areas suggesting that the way in which the template positions have been generated can affect accuracy significantly.

In a follow-up study, we investigate the feasibility of optimizing the electrode positions obtained using our custom template based on a few measurements of distances and angles between electrodes and landmarks. We used measurements between nasion, left preauricular (LPA), right preauricular (RPA), inion, and nearby electrodes for a total of the eight measurements (four distances and four angles). We show that the result of the optimization is not particularly affected by errors in the measurements but that the effect across subjects differ substantially. Specifically, some subjects benefitted considerably whereas it did not make much of a difference for others suggesting that the procedure is good at preventing outliers.

Finally, we explore the effect of forward solution errors on source localization errors. Using the same forward models described above, we simulated data at different signal-to-noise ratio (SNR) levels and used different inverse methods (dynamic statistical parametric mapping (dSPM), standardized low resolution electromagnetic tomography (sLORETA), dipole fitting, multiple signal classification (MUSIC)) for localization. In general, we found that dipole fitting and MUSIC were slightly more sensitive to errors in the forward model compared to the minimum norm estimate (MNE)-based ones whereas all methods were sensitive to SNR level. Using a template anatomy resulted in larger errors compared to anatomical models generated in SimNIBS, MNE-Python, or FieldTrip. On the other hand, using a template anatomy with digitized electrode positions performed better than using the correct anatomical model for each subject but with the manufacturer description of electrode positions. This highlights the importance of correctly specifying electrode positions as otherwise large, spatially correlated errors may be induced.

Resumé

I mange situationer, hvor man har optaget elektroencefalografi (EEG) data, kan det være ønskværdigt at lokalisere de neurale kilder, der er ophav til det observerede signal. Det kan fx være ifm. planlægning af et kirurgisk indgreb i epilepsipatienter eller for at forbedre et estimat af forbindelsesstyrke mellem områder i hjernen. Kildelokalisation i EEG er et inverst problem og for at kunne løse dette, er det nødvendigt først at løse det tilhørende forward-problem. Forward-problemet består i at bestemme hvilke målinger, man ville observere, såfremt hver enkelt kilde var aktiv. Nøjagtigheden hvormed dette problem kan løses bestemmes (delvist) af vores evne til at konstruere realistiske modeller til bestemmelse af udbredelsen af det elektriske signal. For at kunne konstruere sådanne modeller, er det bl.a. nødvendigt at bestemme den anatomiske udformning af det enkelte individs hoved samt konduktiviteten af de forskellige vævstyper. Ydermere skal man vælge en fornuftig opløsning for sin model og det er også nødvendigt at kende positionerne af EEG-elektroderne ift. hovedet

I denne afhandling undersøges det hvordan forskellige aspekter af forwardmodellen påvirker løsningen af forward-problemet samt den efterfølgende løsning af det inverse problem. Vi fokuserer på den anatomiske nøjagtighed af hovedmodellen samt den nøjagtighed hvormed elektrodepositionerne er bestemt. Ift. det inverse problem, så fokuserer vi primært på lokalisation af kilder (i modsætning til fx at bestemme styrken af en given kilde).

I det første arbejde præsenterer vi en software pipeline (inkluderet i Sim-NIBS) til at generere forholdsvist realistiske anatomiske modeller af det menneskelige hoved med fokus på rekonstruktion af kraniet. Denne pipeline er baseret på SPM12 og CAT12 hvor sidstnævnte bruges til at forbedre nødagtigheden hvormed hjernevæv bliver segmenteret. Vi sammenligner med eksisterende værktøjer i FSL og BrainSuite og finder at den foreslåede pipeline forbedrer segmenteringen af kraniet, særligt når både et T1- og et T2-vægtet magnetisk resonans-billede er tilgængelig. Vi viser endvidere, at det, for at sikre gode segmenteringsresultater, er vigtigt at kvaliteten af de strukturelle billeder, der bliver brugt, er høj (fx at man søger at minimere artefakter såsom forskydninger i billedet pga. fedt).

Dernæst sammenligner vi forward-løsninger genereret af hhv. SimNIBS, MNEpython og FieldTrip. Den dominerende forskel mellem modeller fra disse softwarepakker er nøjagtigheden hvormed de beskriver den underliggende anatomi idet SimNIBS generelt er mere nøjagtig end de to sidstnævnte. Vi finder større topografiske og magnitude-fejl i forward-løsninger fra MNE-python og FieldTrip sammenlignet med SimNIBS i det meste af hjernen, hvilket tyder på, at der generelt er store forskelle i forward-løsningerne modellerne imellem. Endvidere sammenligner vi også med en model baseret på en standard-anatomi. Denne model resulterer også i væsentlige forskelle. Udover at sammenligne ovennævnte pipelines, så sammenligner vi også forskellige måder at angive elektrodepositioner på. Vi sammenholder elektrodedigitalisering med brug af en standardbeskrivelse af elektrodepositioner, der tilpasses hvert enkelt individ. Vi undersøger to sådanne standardbeskrivelser: én, som vi selv har genereret ved at digitalisere den relevante EEG-hætte på en 3D-printet model af MNI-hovedet og en anden, som bruger elektrodepositionerne angivet af fabrikanten af EEG-hætte. Vi finder væsentlige topografiske fejl ved brug af fabrikantens positioner, særligt occipitale og parietale områder, hvilket også er de områder hvor fejlen i elektrodepositionerer størst (sammenlignet med digitaliserede positioner). Vores egen beskrivelse af elektrodepositionerne passr bedre i disse områder, hvilket tyder på, at den måde, hvorpå sådanne standardpositioner er genereret, kan have en væsentlig indflydelse på deres nøjagtighed.

I et opfølgende studie til ovennævnte, undersøger vi hvorvidt det er muligt at optimere sådanne standard elektrodepositioner ved brug af nogle få målinger af afstande og vinkler mellem elektroder og anatomisk letidentificerbare punkter på hovedet. Vi benytter sammenlagt otte målinger (fire afstande og fire vinkler) foretaget mellem nasion, LPA, RPA, og inion samt omkringliggende elektroder. Vi viser at resultatet efter optimering ikke er nævneværdigt påvirket af fejl i målingerne hvorimod effekten på tværs af individer varierer signifikant. I særdeleshed er der nogle individer for hvem optimering resulterer i væsentlige forbedringer af elektrodepositioner, hvorimod der er andre for hvem der ikke er megen effekt. Dette antyder, at proceduren er særlig god til at modvirke "outliers".

Endeligt undersøger vi effekten af fejl i forward-løsningen på fejl i kildelokalisation. Vi bruger de samme forward-modeller som beskrevet ovenfor og simulerer data ved forskellige SNR-niveauer og bruger forskellige inverse metoder (dSPM, sLORETA, dipol-fitting, MUSIC) til at lokalisere kilder. Generelt finder vi at dipol-fitting og MUSIC er mere sensitive overfor fejl i forward-modellen sammenlignet med metoder baseret på MNE, hvorimod alle metoder er sensitive overfor SNR. Ved brug af en model baseret på standard-anatomi finder vi større fejl end ved brug af modeller genereret i SimNIBS, MNE-Python og FieldTrip. Dog viser det sig samtidig, at en model baseret på standard-anatomi men som bruger digitaliserede elektrodepositioner resulterer i mindre fejl end en model med korrekt anatomi og elektrodepositioner defineret ved brug af fabrikantens angivelser. Dette viser vigtigheden af at kunne identificere elektrodepositioner nøjagtigt da man ellers risikerer at inducere store, spatialt korrelerede fejl i modellen.

Preface

This PhD thesis has been a long time under way. The project began in October 2015. In September 2017 I sustained an injury to the head which resulted in a concussion and, as it later turned out, a neck injury as well. Since the diagnosis was initially incomplete, things got progressively worse to the point at which I was no longer able to take care of myself. When, after more than a year, things finally started getting better (due to finding the correct types of treatment) the damage had already been done, both physically and mentally, which meant that regaining normal function was very challenging. After an additional two years of rehabilitation work, I was able to join a program the purpose of which was to encourage a return back to work. In this period I gradually increased work capacity from four hours a week to near a full week. This was made possible by my primary supervisor, Kristoffer H. Madsen, who facilitated contact with relevant people at Odense Hospital (where I did the training), in particular Faisal Mahmood, something for which I am very grateful indeed. After half a year I was finally able to resume my PhD position at the Technical University of Denmark in March 2021 and now, seven years after it all started, it is finally coming to a close.

This PhD thesis was prepared between October 2015 and August 2022 at the Section for Cognitive Systems, Department of Applied Mathematics and Computer Science, Technical University of Denmark (DTU) and the Danish Research Centre for Magnetic Resonance (DRCMR), Hvidovre Hospital. It constitutes a partial fulfillment of the requirements for acquiring a PhD degree at DTU. It was funded by DTU and the Sino-Danish Center for Research and Education (SDC).

Acknowledgements

I would like to thank my supervisors, Kristoffer Hougaard Madsen and Axel Thielser, for general help and support as well as providing ideas and feedback throughout the process of working on this thesis. I thank Guilherme Bicalho Saturnino for SimNIBS related support and Oula Puonti for help with data processing, analysis, proofreading, and discussions in general. I would also like to thank my Chinese supervisors, Rong Xue and Fan Wang, for welcoming me in their department and making my stay in China possible. Additionally, I am thankful to Zuxiang Liu for providing EEG equipment and expertise; to Zhentao Zuo for help setting up and running structural and functional MRI sessions; and to Hu Yinan for general help and discussions on EEG and MEG related issues. I would like to thank my supervisor Lars Kai Hansen for contributing ideas and helping make this project possible.

Special thanks goes to my main supervisor, Kristoffer Hougaard Madsen, and head of the PhD school at DTU Compute, Kim Knudsen, for helping me resume the project after being forced to abandon it in September 2017 due to illness. Also, none of this would have been possible without the support of my family who helped me get through the most difficult time of my life. Specifically, I thank my mother for putting up with me for two and a half years.

Lastly, I would like to thank the Technical University of Denmark as well as the Sino-Danish Center for Research and Education for funding this PhD work.

List of Publications

Articles

- O. Puonti, K. Van Leemput, J. D. Nielsen, C. Bauer, H. R. Siebner, K. H. Madsen, A. Thielscher. Skull Segmentation From MR Scans Using a Higher-Order Shape Model Based on Convolutional Restricted Boltzmann Machines. SPIE Medical Imaging, 2018.
- J. D. Nielsen, K. H. Madsen, O. Puonti, H. R. Siebner, C. Bauer, C. G. Madsen, G. B. Saturnino, and A. Thielscher. Automatic Skull Segmentation From MR Images For Realistic Volume Conductor Models Of The Head: Assessment Of The State-Of-The-Art. Neuroimage, vol. 174, pp. 587–598, 2018. (Citations as of 28.9.2022 on Web of Science is 84 and 137 on Google Scholar.)
- 3. S. Farcito, O. Puonti, H. Montanaro, G. B. Saturnino, J. D. Nielsen, C. G. Madsen, H. R. Siebner, E. Neufeld, N. Kuster, B. A. Lloyd, et. al. Accurate Anatomical Head Segmentations: A Data Set For Biomedical Simulations. In: 2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), pp. 6118–6123, 2019.

Book Chapters

 G. B. Saturnino, O. Puonti, J. D. Nielsen, D. Antonenko, K. H. Madsen, and A. Thielscher. SimNIBS 2.1: A Comprehensive Pipeline For Individualized Electric Field Modelling For Transcranial Brain Stimulation. In: Brain and Human Body Modeling, pp. 3–25, Springer, Cham, 2019.

Conference Contributions

- K. H. Madsen, A. Thielscher, G. B. Saturnino, and J. D. Nielsen. Computationally Efficient and Geometrically Accurate Head Models for EEG Based on FEM. 10th FENS Forum of Neuroscience, 2016.
- J. D. Nielsen, O. Puonti, A. Thielscher, and K. H. Madsen. Source Localization Using Realistic Forward Models. 13th International Conference on Bioelectromagnetism (ICBEM), 2021.

Mathematical Notation

The following is a list of notations used throughout the thesis. Most notation is defined in the text, hence this list does not claim to be complete.

a	Scalar or function.
a	Vector.
\mathbf{A}	Matrix.
\mathbf{A}^{\top}	The transpose of matrix A .
\mathbf{A}^+	The pseudoinverse of matrix \mathbf{A} .
Ι	Identity matrix. The dimensions should be clear from context.
κ	Condition number.
$\ \cdot\ _p$	The p -norm. If p is omitted then the 2-norm is implied.
· ^r	Absolute value or cardinality. The meaning should be clear from context.
$\langle \cdot, \cdot \rangle$	Inner product.
H^1	Sobolev space.
L^1	The space of absolute integrable functions.
corr	Correlation matrix.
\det	Determinant.
\dim	Dimensionality of vector space.
diag	Operator to extract the diagonal elements of a (square) matrix.
image	Image of function or matrix.
\inf	Infimum (greatest lower bound).
$\operatorname{nullity}$	Dimension of null space.
nullspace	Null space of linear map or matrix.
rank	Rank of matrix.
\sup	Supremum (least upper bound).
trace	The trace operator, i.e., the sum of the diagonal entries of a matrix.
∇	Gradient operator.
$\nabla_{\mathbf{v}}$	Directional derivative along \mathbf{v} , i.e., $\mathbf{v} \cdot \nabla f(\mathbf{x})$.
$\nabla \cdot$	Divergence operator.
abla imes	Curl operator.
p(a)	Probability of a.
p(a b)	Conditional probability of a given b .
p(a,b)	Joint probability of a and b .

viii

Contents

Sι	ımm	ary		i	
R	esum	é		iii	
Pı	Preface				
A	Acknowledgements				
Li	st of	Public	cations	vii	
M	athe	matica	l Notation	viii	
1	Intr 1.1 1.2	roducti Resear Overvi	on ch Questions	1 3 3	
2	Bac 2.1	kgroun Electro 2.1.1 2.1.2 2.1.3 The E 2.2.1	IdDencephalographySignal OriginsLimitations of EEGData PreprocessingEG Forward ProblemProblem Formulation	5 5 6 7 13 14	
	2.3	2.2.2 2.2.3 2.2.4 The E 2.3.1 2.3.2	Analytical Solutions in Simple Geometries The Boundary Element Method The Finite Element Method EG Inverse Problem Characteristics of the Inverse Problem Stabilizing the Solution	15 15 17 27 29 33 25 $ $	
	2.42.5	2.3.3 Magne 2.4.1 2.4.2 2.4.3 2.4.4 Segme 2.5.1	EEG Inverse Solvers	$35 \\ 40 \\ 40 \\ 41 \\ 42 \\ 43 \\ 44 \\ 45$	
3	Res 3.1	earch Skull S	Contributions Segmentation From MR Images	47 47	

CONTENTS

	3.2	Evaluating EEG Forward Solutions	48			
	3.3	Optimizing Electrode Positions in EEG	49			
	3.4	Evaluating EEG Inverse Solutions	50			
		3.4.1 Methods	51			
		3.4.2 Results	57			
4	Disc	cussion	71			
	4.1	Source Localization Errors	72			
		4.1.1 Evaluation Through Simulations	74			
	4.2	Limitations and Future Research	77			
5	Conclusions 81					
A	Skull Segmentation From MR Images 8					
в	Eva	luating EEG Forward Solutions	97			
С	Opt	imizing Electrode Positions in EEG	135			
D	Additional Background Material					
	D.1	Calculating the Potential in a Multi-Layer Sphere Model	149			
	D.2 Calculating the Mapping between Cartesian and Reference Ele-					
		ments	150			
	D.3	The Geometry of $\mathbf{A}\mathbf{x} = \mathbf{b}$	151			
	D.4	Image Formation in MRI	152			
Al	Abbreviations 15					
Re	References 157					

х

Chapter 1

Introduction

The brain is the most complex organ in the human body. It is involved in everything from perception and motor control to higher cognitive functions such as memory, learning, and emotions. Understanding how the brain works is important, not only because it may provide explanations for behavioral observations but also because it may help us understand various diseases and hopefully improve treatment. Studying the brain is, however, complicated by the fact that it is enclosed within the skull thus preventing direct assess. Consequently, before the inception of modern neuroimaging technologies, studies of the brain were largely based on lesions or post-mortem analyses. Although these methods have their own merits, they also have severe limitations. Today, however, there are a host of different technologies which can be used to investigate the structure and function of the brain *in vivo* in a noninvasive way. For example, techniques such as computed tomography (CT) and magnetic resonance imaging (MRI) are crucial in understanding the structure of the brain whereas function may be investigated using for example functional magnetic resonance imaging (fMRI), positron emission tomography (PET), electroencephalography (EEG), or magnetoencephalography (MEG). Using fMRI, researchers are able to construct high resolution 3D images (on the order of a few millimeters) of the working human brain, however, the temporal resolution is low (on the order of seconds) compared to the temporal scale at which neural events take place (tens or hundreds of milliseconds). Methods based on electrophysiological measurements such as EEG and MEG have excellent temporal resolution and directly reflect the firings of neurons in the brain. Unfortunately, these methods suffer from so-called volume conduction, the effect of which is that all sensors (in principle) pick up activity from the entire brain. Drawing conclusions about the spatial origin of the observed signal is therefore not straightforward and requires the solution of a so-called inverse problem, i.e., determining the sources responsible for generating the observed data. However, in order to solve this problem, one first needs to know how the signal from individual neural sources propagate to the sensors. This is known as the forward problem. Solving this, requires knowledge about several key components. For example, the first step is to build a model which describes the geometry of human head. For this to be anatomically accurate, information from structural MRI is paramount. Additionally, we need to somehow model the neural sources in the brain and we also need information about the location of the sensors relative to the model of the head. Finally, the conductivities of the different tissues are also important. Collectively, these parameters shape the observed signal and our aim is to model them as realistically as possible. Eventually, we may attempt to recover the sources which are responsible for generating the data, however, the spatial resolution is generally expected to be low (on the order of a few centimeters) compared to other methods such as fMRI. However, it is also important to note that different modalities may be sensitive to different sources. For example, fMRI is unable to detect very transient sources whereas EEG and MEG are not sensitive to sources where the field is closed (e.g., if the sources cancel out) (Michel et al., 2004).

An obvious question, then, is why we would want to use methods such as EEG to localize activity at all given that alternatives exist which are both safe and provide higher spatial resolution? The answer, of course, lies (at least partly) in the temporal differences between these techniques as there may be situations where we are not satisfied with the temporal resolution that fMRI offers, e.g., because the events of interest are very brief. Besides, EEG is cheap and can be worn basically anywhere and for extended periods of time enabling researchers or doctors to monitor participants or patients during sleep or when looking for rare events (e.g., epileptic seizures) (van Mierlo et al., 2020).

For example, using MEG, Parkkonen et al. (2009) show how the neural generators of auditory brainstem responses can be identified by fitting dipoles to the peaks of an event-related response and they found the first cortical responses (in temporal areas) after the stimulus around 18 ms and 30 ms. This study shows how source localization of electrophysiological measurements can be used to study the evolution of the signal in the brain of early (sub)cortical responses. Similar measurements performed with fMRI, would likely have produced spatial maps showing the activity from all sources simultaneously due to the low temporal resolution of this modality. Consequently, the temporal information would have been lost. Similarly, source analysis may also be used to explore the neural generators of EEG microstates (stable topographies lasting approximately 60 ms to 120 ms) providing insight into the dynamics of global neuronal activity (Michel & Koenig, 2018).

Another situation in which reconstructing the neural sources of an observed signal may be relevant, is when we want to investigate brain connectivity. Volume conduction is known to complicate the interpretation of connectivity estimates computed on the channel level because the signal from a given source spreads to all sensors. Performing source analysis may help reduce (albeit not eliminate) such effects and also help assign more anatomically meaningful labels to the interacting entities (Schoffelen & Gross, 2009).

Finally, and perhaps most importantly, source localization based on EEG may have important clinical uses. For example, it can be used in epilepsy treatment to identify the epileptogenic zone, i.e., the brain region which must be resected in order for seizures to stop, in the presurgical stage (van Mierlo et al., 2020). Specifically, EEG can be used in a multimodal setup (including structural MRI) to assess the origin of either ictal or interictal epileptic activity. The former is preferable as it directly reflects the seizure onset zone whereas the latter likely reflects a more diffuse activity pattern due to fast propagation of the signal (Habib et al., 2016). However, ictal activity localization is complicated by low signal-to-noise ratio (SNR) as the signal is corrupted by large amounts of artifacts (e.g., muscle activity and eye movements) caused by the seizure.

Nevertheless, source localization based on EEG has shown good sensitivity and specificity (approximately 80% to 90%) in terms of identifying epileptogenic zones.

1.1 Research Questions

As suggested in the previous section, there are several scenarios in which EEG source analysis is useful. The quality of the source estimates will depend on a number of factors including the accuracy with which the corresponding forward problem was solved. However, the exact nature of this dependency is challenging to assess.

In this work, we seek to investigate how different aspects of the forward model affect source estimates in EEG. In particular, we will focus on the anatomical accuracy of the physical model of the head and the accuracy with which the position of the electrodes are known. In terms of the inverse solution, we will mostly be concerned with assessing the accuracy with which a source can be localized. Exploring these issues will help researchers answer questions such as how much effort they should put into generating accurate forward models, what aspects of forward modeling are most important, and the degree to which their results can be trusted.

The first aim was to develop a pipeline for generating reasonably realistic anatomical models of the head and assess its accuracy. This was incorporated into the open source software SimNIBS which is a tool used to simulate the fields resulting from various types of noninvasive brain stimulation. Next, we validated the simulation results obtained from SimNIBS using an improved head modeling approach¹ and compared forward solutions from SimNIBS with pipelines in wellestablished EEG analysis software packages. Finally, we continued the analysis by assessing the influence of the forward model on the inverse solution in a range of different conditions. To increase the availability of these tools, we augment the SimNIBS toolbox with the ability to export the generated solutions to a few well-known EEG software packages (specifically, MNE-Python and FieldTrip).

1.2 Overview of Thesis

The structure of this thesis is as follows. Chapter 2 presents background information and theory relevant for the work that was done. This includes a description of EEG as well as a brief overview of channel level data analysis (section 2.1). This is followed by an introduction to the EEG forward problem (section 2.2) and the inverse problem (section 2.3). We end this chapter with an overview of the basic concepts of MRI and segmentation of structural images. Chapter 3 presents the research contributions of the current thesis and includes a paper on skull segmentation and head modeling (section 3.1), a submitted manuscript on forward modeling errors due to various model imperfections (section 3.2), and a manuscript draft on optimizing electrode positions from a template layout (section 3.3). Finally, in section 3.4 we present a workin-progress on investigating the effect of forward modeling errors on the inverse

¹Since this method, termed *headreco*, was developed in 2017, a new and improved pipeline called complete head anatomy reconstruction method (CHARM) has been integrated into SimNIBS. We use this pipeline in the remaining studies.

solution. Chapter 4 contains a discussion centered around how we may expect errors in the forward solution to propagate to the inverse solution. For this, we use the results obtained in section 3.4. Finally, we present a brief recap of the most important points of this thesis in chapter 5.

Chapter 2

Background

2.1 Electroencephalography

EEG is a relatively seasoned technique dating back to 1929 where Hans Berger was the first to measure brain activity by placing electrodes on the scalp and record the potential differences between these. The result is the so-called "electroencephalogram" and it represents the collective contribution of all electrical currents in the brain due to cellular processes (ignoring other sources of noise for the moment). This means that, in principle, any transmembrane current contributes to the measured potential field (Buzsáki et al., 2012; Luck, 2005).

2.1.1 Signal Origins

The strongest transmembrane currents are generated by the fast action potentials. Although these effects are strong in the vicinity of the cell they are generally not believed to contribute substantially to the signals measured in EEG. The reason for this is that in order for cellular events to be detectable at the scalp, thousands or millions such events need to be synchronized in time (and space) as they are too weak to be detected individually. However, since axon potentials are very transient (approximately 1 ms), the probability that they are sufficiently synchronized for this condition to be met is very low (Kirschstein & Köhling, 2009).

In comparison, the temporal extent of the postsynaptic potential is much longer (tens of milliseconds). Excitatory neurotransmitters¹ released to a synapse binds to receptors on the postsynaptic cell causing influx of positively charged ions (e.g., sodium, calcium) which results in an extracellular *sink* (a point from which electrical current exits). In order for electroneutrality to be conserved, an extracellular *source* (a point at which electrical current enters), consisting of an opposing ionic flux across the membrane, needs to be formed thus restoring equilibrium. In particular, the sink forms at the apex of the dendrites where ion channels are opened (due to excitatory neurotransmitters) and the source forms where the dendrites terminate on the soma effectively creating a dipole. The resulting extracellular currents travel throughout the conductive compartments

¹Similar arguments can be made for inhibitory neurotransmitters although the direction may be inverted.

of the head. The duration of such postsynaptic potentials is substantial enough to allow them to sum up, not only within a cell (potentially generating an axon potential), but, importantly, also across thousands of synchronously active cells (Buzsáki et al., 2012).

Temporal summation of such dipolar patterns is, however, not sufficient by itself to give a detectable signal on the scalp. The summation also needs to be constructive in space as the signals would otherwise cancel each other. This is achieved when the stimulated cells are aligned in parallel. Thus, in cytoarchitecturally regular cell ensembles, as found in the cerebral cortex, the orientation of the apical dendrites allows superposition of dipoles from synchronously active cells (Buzsáki et al., 2012; Kirschstein & Köhling, 2009; Scherg, 1990).

One class of cells which fulfills the above criteria are the cortical pyramidal cells. They are the most abundant cell type in the human cortex and have long, thick dendrites allowing them to form strong dipoles when stimulated. This causes them to generate strong, extracellular fields which can be observed at considerable distance. Consequently, these cells are assumed to be the primary sources of the neural activity measured by EEG. When viewed from afar, the contribution of higher order moments to the field produced by a set of currents is neglectable. Therefore, the signal produced by an ensemble of neurons (e.g., in a patch extending 2 cm to 3 cm) can be modeled as an *equivalent current dipole* whose direction and strength is the result of spatial and temporal summation of the fields from the individual neurons (Buzsáki et al., 2012; Luck, 2005; Scherg, 1990). The primary (or source) current is that which runs between the sources and sinks described above. This creates an electric field (indicating the force acting on a test charge) which induces the flow of secondary (or return) currents throughout the conductive medium, i.e., the head, and these are what generate the potential differences observed in EEG. This phenomenon, called volume conduction, is determined by the source (location, orientation, strength) and the conductivity of the medium² and complicates the relationship between sources and observations since all sources will project to all channels. When multiple sources are active simultaneously, they will superimpose linearly further complicating matters. The relationship between the the primary sources and the observed potential differences created by volume conduction is given by (the quasi-static approximation to) Maxwell's equations (see section 2.2).

2.1.2 Limitations of EEG

Although EEG has several attractive features such as a high temporal resolution (on the order of milliseconds), being noninvasive, cheap, and extremely flexible (as opposed to methods requiring specialized environments, e.g., MRI and MEG), it also has a range of limitations.

As discussed, all active sources will contribute to the signal measured at all electrodes due to volume conduction. Additionally, certain properties of the volume conductor (i.e., the head) serve to spatially blur the signal thus exacerbating this mixing. In particular, since current tends to follow the path of least resistance, the conductivity gradient between the well conducting cerebrospinal fluid (CSF) and the low conducting skull compartment means that current will

 $^{^{2}}$ At the frequencies investigated in EEG, the medium can be considered *ohmic* meaning that it is purely resistive and thus there are no reactive or capacitive effects. This is also called the quasi-static assumption (or condition).

2.1. ELECTROENCEPHALOGRAPHY

tend to shunt through the highly conductive CSF resulting in spatial blurring of the signal. The combination of spatial blurring due to the volume conductor and the fact that high frequency variations in the signal will generally be less likely to add constructively during this averaging operation (e.g., because of noise in the measurements), renders the effective temporal resolution of EEG slightly lower than what might otherwise be expected (of course it is still much higher than methods based on vascular phenomena, e.g., fMRI).

Another important concern with EEG is that the data is corrupted by substantial amounts of noise. One source of noise is the acquisition system itself, e.g., line noise from the electrical grid (typically 50 Hz). Besides, the quality of the signal depends on having a sufficiently low impedance at the electrodes which can be difficult to achieve across all electrodes. In either case, the impedance may increase over time as the gel used to make the electrical contact between electrode and skin will tend to dry thus causing non-stationary effects of noise which may be difficult to model correctly.

The electrical signal from the brain is generally weak compared to other biological processes such as heartbeats, respiration, and eyeblinks. These may be measured and compensated for to some extent, however, this is not trivial since signal due to eyeblinks may be two orders of magnitude larger than the signal of interest. A related point is that because the brain is constantly processing information from a range of different sources, there will usually be a lot of signal due to neural processes of non-interest which is likely to obscure the signal of interest. This is true for most neuroimaging techniques, however, it is particularly severe in EEG. Consequently, long data segments will typically be required, e.g., by repeating an event of interest hundreds of times.

2.1.3 Data Preprocessing

As alluded to above, the raw data from EEG is often corrupted by substantial amounts of noise and some degree of processing will usually be required before meaningful patterns can be extracted. This section describes some of the common strategies employed to reduce noise in the data. The focus will be on preprocessing steps relevant for analysis of event-related signals. However, we will start by describing a typical approach to positioning an EEG cap on the head of a subject.

Cap Placement

The circumference of the head is measured and an appropriately sized cap is chosen. Caps are often available in standard sizes (e.g., 54 cm, 56 cm, 58 cm and 60 cm) and the one closest to the measured circumference (or slightly larger) is chosen. The point halfway between the nasion and the inion and halfway between the left preauricular (LPA) and right preauricular (RPA) points is found and the Cz electrode (in the case of a 10-20 style layout) is placed at this point. If the layout is different from the standard 10-20 system (which places electrodes at regular intervals specified by spherical angles), then the relevant electrode is likely specified by the manufacturer. The electrodes on the line from inion to nasion as well as those on the line from LPA to RPA are checked to ensure that they are on a straight line. If not, small adjustments can be made,

for example by stretching or compressing the cap (Homölle & Oostenveld, 2019). Finally, the chin strap is fastened.

Temporal Filtering

Most often some sort of filters will be applied to the data to attenuate frequencies believed not to carry information about the processes of interest. For example, a highpass filter can be used to remove low frequency drifts which may be thought to arise due to changes in impedance over time rather than neural activity. Likewise, it may be desirable to remove some of the high frequencies in the data to attenuate muscle artifacts and noise related to the equipment. Depending on the particular cutoff frequency, this may also remove line noise which may otherwise be attenuated with a notch (band-stop) filter. The goal is to increase SNR, hence the appropriateness of filter settings will depend on the time scale of the effects of interest and the amount of noise in the data. However, it is well-established that besides (hopefully) reducing noise in the data, filtering may also have unintended side-effects. Some of these (e.g., edge effects) may be reduced by applying the filters to the raw (continuous) data as opposed to the epoched data³, however, in general, filters may also distort the waveforms of the signal by changing amplitude, timing, and adding artifactual peaks (Luck, 2005) as demonstrated by Tanner et al. (2015).

These issues arise because filters are never "perfect" in the sense that the frequency cutoff is not infinitely steep at the desired frequency and the impulse response is not a delta function. Actually, the relationship between the two is such that a steep transition in frequency domain implies a longer impulse. Thus, "sharp" filters increase temporal smearing and generally result in more distortions due to increased rippling in the passband and stopband (Widmann et al., 2015).

Highpass filters remove slowly varying components of the signal meaning that they will effectively remove constant offsets in the signal as well. Consequently, they can be used in place of (and may be preferable to) baseline correction in event-related studies (Widmann et al., 2015).

Physiological Noise

One class of artifacts which can contaminate the data substantially stems from other physiological processes than those directly related to brain activity. Two such sources of noise are eye movements, which generate signals much stronger than those originating from the brain, and periodic signals due to heartbeat. An eyeblink causes a rotation of the eyeball which acts like a dipole causing a large disturbance of the signal. These effects are typically easy to identify by visual inspection. Although saccadic eye movements may also be detected this way (if they are sufficiently strong) this can be a slightly more difficult and time consuming endeavor. Heartbeats corrupt the signal throughout the entire recording session, however, as it is not likely to be time-locked to the experimental design this may not be a major issue. Whether or not to attempt to correct for such artifacts will depend on the size of the effects and how much data is available. Eyeblinks, however, generate such large signals that they will

 $^{^3\}mathrm{Epoching}$ refers to the extraction of data segments around the stimulus of interest, e.g., from 200 ms before to 500 ms after stimulus onset.

need to be dealt with somehow. Signal from eye movements and heartbeats can be measured using dedicated electrode configurations. This is known as electrooculography (EOG) and electrocardiography (ECG), respectively.

A popular method used to isolate (and remove) physiological artifacts is independent component analysis (ICA) which has proved capable of identifying and separating biologically plausible brain sources as well as artifacts such as eyeblinks (Delorme & Makeig, 2004). ICA decomposes the data into a number of statistically independent components⁴. The model is

$$\mathbf{x} = \mathbf{As} \tag{2.1}$$

where **A** is a matrix which linearly mixes the sources in **s** so as to obtain the (observed) variable **x**. Thus, the model assumes instantaneous mixing of source signals, a condition which is satisfied by the volume conduction phenomenon in EEG (Delorme & Makeig, 2004). The goal of ICA is to find an unmixing matrix, **W**, i.e., an inverse of **A**, which allows us to estimate the source components from the data. However, since the the problem of finding such a linear demixing matrix is not unique, another assumption is necessary, namely that the components are not normally distributed and consequently that they can be identified by maximizing the non-Gaussianity of each component, e.g., by maximizing a higher order moment such as kurtosis⁵ or minimizing negentropy where the latter is defined as

$$J(\mathbf{x}) = H(\mathbf{x}_{\text{Gauss}}) - H(\mathbf{x}).$$
(2.2)

Here H denotes the entropy, \mathbf{x} is a random variable, and $\mathbf{x}_{\text{Gauss}}$ is a Gaussian random variable with the same covariance as \mathbf{x} . Since the normal distribution has the highest entropy of all distributions of equal variance this is always nonnegative (Hyvärinen & Oja, 2000). The motivation for the connection between independence and non-Gaussianity comes from the central limit theorem which states that the sum of multiple random variables (irrespective of their individual distributions) tends towards a normal distribution (provided certain conditions are satisfied, e.g., that the variables have finite variance). As such, \mathbf{x} must be more Gaussian than \mathbf{s} , and any suboptimal unmixing will result in increased Gaussianity of the source estimates. Consequently, the ICA model cannot be estimated for Gaussian independent components, the reason being that any orthogonal transformation applied to \mathbf{A} would result in the same observational distribution as it has been whitened as part of the preprocessing. As such, no directions are more independent than others.

Lastly, it is worth noting that the ICA model is ambiguous wrt. the ordering of the components (as they can always be permuted) and is not able to recover the actual variance of the sources (since a scaling of \mathbf{s} can always be counteracted by a reciprocal scaling of \mathbf{A}). ICA is also not able to determine the true *number* of components and so typically a number of components corresponding to the number of sensors is returned.

⁴In contrast, principal component analysis (PCA) finds components that are uncorrelated (which is a weaker assumption) by diagonalizing the data covariance matrix and sorting them according to the amount of variance they contain. Thus, PCA is useful for data *compression*.

⁵Preprocessing in the form of whitening (and centering) is usually employed to decorrelate the components beforehand as this is implied by independence and also simplifies the subsequent problem. This can be achieved using a PCA decomposition.

ICA components may be inspected visually to identify the ones capturing physiological artifacts. Alternatively, if EOG or ECG measurements are available, correlation may also be used to identify the relevant components. These can then be discarded, however, as the separation performed by ICA is not perfect, rejecting entire components is likely to also remove some signal of interest.

To minimize such effects, one may instead attempt to filter out the signal believed to represent the noise (e.g., EOG, ECG) from the identified components (Mannan et al., 2016). This may be achieved using an adaptive filter such as recursive least squares (RLS) (He et al., 2004; He et al., 2007) which continuously updates the filter weights (as opposed to conventional linear regression). Briefly, given an input signal, \mathbf{x} , contaminated by noise, RLS solves the problem of finding an appropriate set of filter weights, \mathbf{w}_n , for each time point, n, so as to minimize the squared difference between the desired signal, \mathbf{s} , and the estimate of this signal, $\hat{\mathbf{s}}$. The loss function is defined as a weighted sum of the previous and current errors

$$L(\mathbf{w}_{n}) = \sum_{i=p}^{n} \lambda^{n-i} e_{i}^{2} = \sum_{i=p}^{n} \lambda^{n-i} \left(s_{i} - \hat{s}_{i}\right)^{2}$$
(2.3)

where e the residual, p is the length of the filter (i.e., \mathbf{w}_n is a vector of length p), and λ is the forgetting factor meaning that stationarity of the process is not strictly required. This is a weighted least squares problem with the solution

$$\mathbf{w}_n = \mathbf{C}_x^{-1}(n)\mathbf{c}_{xs}(n) \tag{2.4}$$

where \mathbf{C}_x is the weighted covariance matrix between the *p* time-lagged versions of the input signal and \mathbf{c}_{xs} is the weighted cross-covariance between the *p* timelagged versions of the input signal and the desired signal and we notice that samples close to *n* in time receive higher weights. By employing the matrix inversion lemma, this can be solved efficiently with a recursive scheme which updates \mathbf{C}_x^{-1} directly (He et al., 2004).

In EEG, one obvious issue is that we do not know the desired signal, however, if we assume that the input signal (e.g., the raw EEG signal) is a mixture of (zero-mean stationary) neural activity and noise and that we are given one (or more) reference signals (e.g., from EOG or ECG channels) which are correlated only with the noise in the input signal (but not the neural activity), then the problem becomes one of finding the set of filter weights which cancel the input signal as much as possible (in a least squares sense) using the reference signal because the neural signal has zero mean and is uncorrelated with the reference signal and noise. Thus, as the problem is stated above, \mathbf{x} is the reference signal(s), \mathbf{s} is the input signal, and $\hat{\mathbf{s}} = \mathbf{w}_n \mathbf{x}$ (He et al., 2004).

The assumptions on which this is based are clearly not valid and there will inevitably be correlations among the different components (neural activity, noise, and reference signals). By performing the regression on an appropriate set of ICA components, we hope to minimize loss of neural signal while still suppressing the noise components (Mannan et al., 2016). For example, the correlation between neural activity in a frontal electrode and an EOG signal will likely be high whereas we hope that the neural activity captured by the ICA components on which regression will be performed is less correlated with the neural activity captured by the EOG electrodes, or that, even if this is only partly true, the neural content in the relevant ICA components will be somewhat reduced.

2.1. ELECTROENCEPHALOGRAPHY

Rejecting Bad Epochs and Channels

When doing event-related analyses the raw data is segmented into epochs of a certain size relative to each event trigger. Provided that the SNR is sufficient, the data can be analyzed on an epoch-by-epoch basis but otherwise SNR can be increased by averaging multiple epochs since the signal of interest should be time-locked to the trigger whereas the noise is assumed not to be. Often, the data will also be downsampled (preferably *after* epoching to avoid jittering of the event tiggers).

As mentioned, simple visual inspection can be used to identify data segments corrupted by eye-related artifacts (and possibly others such as muscular artifacts) or bad channels. This is a relatively straightforward procedure provided that the person looking through the data has sufficient experience to reliably detect such phenomena. On the other hand, it is time consuming and a potential concern is whether this could cause a bias in the analysis.

Alternatively, once the data has been epoched, a peak-to-peak threshold can be defined and used to reject segments based on whether the threshold is surpassed or not. Setting such a threshold is, however, not trivial and the optimal value is likely to vary between subjects. As such, another strategy is to try to learn this from the data itself. One particular implementation is *Autoreject* by Jas et al. (2017). It estimates local (i.e., sensor specific) peakto-peak thresholds using cross validation (CV) and Bayesian optimization. The objective function is the average error over all folds in a K-fold CV run for a certain threshold τ . This error is computed as the Frobenius norm of the difference between the mean signal from the good trials (see below) of the kth training set, $\bar{\mathbf{X}}_{T(k)}$, and the median signal from the corresponding validation set, $\hat{\mathbf{X}}_{V(k)}$, where the split is performed over data segments (epochs)⁶, i.e.,

$$e(\rho,\kappa) = \frac{1}{K} \sum_{k=1} K \left\| \bar{\mathbf{X}}_{T(k)}^{\rho,\kappa} - \hat{\mathbf{X}}_{V(k)} \right\|_{\mathrm{F}}.$$
 (2.5)

Since the rejection is *local*, i.e., separate for each sensor, two other parameters are introduced, specifically, κ , the fraction of sensors that needs to be deemed bad in order to reject a given epoch, and ρ , the maximum number of bad sensors which will be interpolated using spherical splines. These parameters are found using grid search (Jas et al., 2017).

Noise Covariance Estimation

In EEG data analysis pipelines, it is often desirable to have a proper estimate of the noise in the data, e.g., when doing source analysis. If the amplitude of the sources are assumed to be normally distributed then so are the measured signals due to the linearity of the forward problem in EEG (see section 2.2). As such, if the data has been centered, and if we further assume additive Gaussian noise, then the source and noise covariance matrices fully characterize the signal (Engemann & Gramfort, 2015).

The noise covariance is typically estimated from data segments believed to be free from signal of interest (e.g., empty room measurements in MEG or

 $^{^{6}}$ To enable detection of globally bad channels, the data is augmented by a set of epochs where each channel is interpolated from the remaining ones.

prestimulus intervals in EEG), however, unless a sufficient amount of samples are available then the maximum likelihood (ML) estimate of the covariance, i.e., the sample covariance, may not accurately reflect the characteristics of other time intervals. Consequently, regularization is often performed, for example by decreasing the off-diagonal loading which can be achieved by adding some value to the diagonal elements. One such shrinkage model is the Ledoit-Wolf (LW) model which weighs the empirical covariance estimate, $C_{\rm ML}$, and the identity matrix, **I**, (corresponding to a prior of independence and equal variance of channels) like

$$\mathbf{C}_{\mathrm{LW}} = (1 - \alpha) \, \mathbf{C}_{\mathrm{ML}} + \alpha \mu \mathbf{I}. \tag{2.6}$$

Here, μ is the mean of the diagonal entries in \mathbf{C}_{ML} and α is a free parameter whose value can be determined in closed-form (as done by Ledoit and Wolf) or by CV by maximizing the log-likelihood on unseen data, \mathbf{Y} , given covariance matrix, \mathbf{C} ,

$$L(\mathbf{Y}|\mathbf{C}) = -\frac{1}{2t} \operatorname{trace} \left(\mathbf{Y} \mathbf{Y}^{\top} \mathbf{C}^{-1} \right) - \frac{1}{2} \log \left(\left(2\pi \right)^m \det \mathbf{C} \right)$$
(2.7)

where m is the number of channels, t is the number of samples, and \mathbf{Y} is $m \times t$. In source estimation, the noise covariance is often employed in order to whiten the data, thus the quality of the estimate can be assessed by the success with which it whitens noisy signals.

Data whitening is the operation of transforming a random vector, \mathbf{x} , with a particular covariance matrix, \mathbf{C} , into a new random vector, \mathbf{z} , of the same dimensionality but with a covariance matrix equal to the identity matrix, i.e.,

$$\mathbf{z} = \mathbf{W}\mathbf{x}.\tag{2.8}$$

Whitening can be viewed as a generalization of standardization which not only standardizes individual random variables but also decorrelates them. W is the whitener which needs to satisfy $\mathbf{W}^{\top}\mathbf{W} = \mathbf{C}^{-1}$, however, this constraint is not unique (Kessy et al., 2015). Given a covariance matrix, **C**, with eigenvalue decomposition $\mathbf{C} = \mathbf{E}\mathbf{D}\mathbf{E}^{\top}$, then we obtain the so-called PCA whitener by

$$\mathbf{W} = \mathbf{D}^{-1/2} \mathbf{E}^{\top}.$$
 (2.9)

If an average reference has been applied to the data, then the global field power (GFP) is defined as the standard deviation over channels (Murray et al., 2008). Thus, we define the (rank-adjusted) GFP of the whitened data as

$$GFP(\mathbf{z}) = \sqrt{\frac{1}{r} \mathbf{z}^{\top} \mathbf{z}} = \sqrt{\frac{1}{r} \mathbf{x}^{\top} \mathbf{W}^{\top} \mathbf{W} \mathbf{x}}$$
(2.10)

where r is the rank of the data. If successful, the whitened data should follow a standard normal distribution and therefore 95% of the data points will be expected to fall within ± 1.96 . $\mathbf{z}^{\top}\mathbf{z}$ in equation (2.10) follows a chi-squared distribution with r degrees of freedom. Hence, the GFP is expected to have a value of one. Deviation from this indicates suboptimal whitening (Engemann & Gramfort, 2015).

2.2. THE EEG FORWARD PROBLEM

Reproducibility

We have presented several (common) approaches to data preprocessing. These steps are taken with the intention of improving the SNR level of the data as many data sources in neuroimaging (including EEG) are noisy, not only because of the methodology itself but because we are interested in very specific processes embedded in an ongoing and extremely complicated system of neural signaling. For example, the magnitude of the observed changes in fMRI in response to experimental stimuli are typically only on the order of a few percent. Likewise, identifying an evoked response in a single epoch of EEG data may be very difficult as the noise is substantial (here, noise includes all brain signals which are not the focus of the particular study). As one might imagine, such issues only tend to get more severe as increasingly complex cognitive phenomena (as opposed to simple sensory or motor responses) are studied.

Because of this, preprocessing is often required—sometimes more than others depending on the quality of the data—and this may make it difficult for others to reproduce certain effects as they depend not only the experimental paradigm but also the subsequent processing steps. The problem of reproducibility can be solved by increased transparency (Jas et al., 2018), however, high levels of parameter "tweaking" so as to obtain the desired results (also known as *p*-hacking) may be more difficult to uncover.

As we would like our results to reflect effects which are actually present in the data and not be due to the particular way in which the data was processed, one may argue that the experimental effects should be visible despite modest variations in processing (e.g., small changes in peak size and location due to filtering) unless severely inappropriate (or "suspicious") steps are taken . That is, *if they are actually present in the data*. Put another way: if it takes intricate and highly tuned data analysis strategies to uncover the effects of interest, one may question the degree to which such findings can be trusted. One should always think about *why* each step is being taken and whether it has a beneficial impact on the data although this may sometimes be difficult to assess. Also, it is important to be aware that the data being submitted to a statistical testing is not the same as that which was collected during the experiment. However, if the effects are present in the data—and sufficiently strong—then the major conclusions should not be sensitive to the particular analysis strategy employed although the exact results may change somewhat.

A related problem may be encountered when testing analysis methods and statistical models. A typical step in this process is to investigate performance on synthetic data, however, by generating data according to the assumptions on which one's method is based, overly optimistic results may be obtained. A similar concept exists within the field of inverse modeling known as the *inverse crime* in which the data generating model (the forward model) used to invert the data is the same as (or very close to) to the one used to obtain the data resulting in unrealistically good performance.

2.2 The EEG Forward Problem

This section introduces the forward problem in EEG. First, we will describe the partial differential equation (PDE) that governs the relationship between electrical activity in the brain and the resulting potential differences observed on the scalp. Next, we will describe how this can be solved. We briefly present an analytical⁷ solution for a multi-layered sphere model and the boundary element method (BEM) approach before presenting the finite element method (FEM) along with different choices of source model. Finally, we will formulate the corresponding problem in transcranial electrical stimulation (TES) and show how the two are related through reciprocity.

2.2.1 Problem Formulation

The EEG forward problem is that of determining, within an appropriately defined domain (i.e., the head), the potential field resulting from a particular source current. Its solution is a matrix mapping sources (current sources in the brain) to measurements (potential differences on the scalp) often called the gain matrix or leadfield matrix.

The general equations describing how electric (and magnetic) fields are generated from charges and source currents are the Maxwell equations which are a set of coupled differential equations in space and time. In EEG, however, the derivatives wrt. time are typically ignored which is justified by the observation that no charge can be accumulated in the extracellular medium at the frequencies investigated with EEG, meaning that there are no time delays, and consequently, that the field at a particular point in time is the result of only the currently active, electric sources. As such, at a specific time point, the field is fully determined by the active sources at that time point (although the electric source distribution may vary over time). These are called the quasi-static conditions since changes in neural activity are slow compared to the propagation of the field (Hallez et al., 2007). Based on this approximation, the relationship between the electric potential, u, and the source current density, \mathbf{j}_s , can be derived as follows.

In EEG, it is common to split the total current density, \mathbf{j} , into two parts (Wolters et al., 2004),

$$\mathbf{j} = \mathbf{j}_s + \mathbf{j}_r, \tag{2.11}$$

where \mathbf{j}_s and \mathbf{j}_r are the current densities due to the source (or primary) currents and the return (or secondary) currents, respectively. That is, \mathbf{j}_s models the electrical activity in the brain whereas \mathbf{j}_r is induced by this activity (and possibly other processes) and represents the volume conduction effects. In EEG, we measure potential differences on the scalp due to \mathbf{j}_r and in source analysis we are interested in recovering \mathbf{j}_s , or alternatively, the source current integrated over a certain volume.

Under quasi-static conditions, the curl of the electric field, **E**, is zero, i.e., $\nabla \times \mathbf{E} = 0$, in which case⁸

$$\mathbf{E} = -\nabla u. \tag{2.12}$$

In addition, Ohm's law states that

$$\mathbf{j}_r = \sigma \mathbf{E} \tag{2.13}$$

⁷Actually, the solutions are quasi-analytical in the sense that they depend on infinite polynomial series, however, for simplicity we shall simply refer to them as analytical.

⁸The electric potential is defined as the line integral of the electric field between two points. If the curl of **E** is zero then the integral does not depend on the specific path over which one integrates. In this case, **E** is said to be *conservative*.

where σ denotes the conductivity of the medium (which can be a scalar or a 3×3 tensor depending on whether the conductivity is considered to be isotropic or anisotropic). Using equations (2.11) to (2.13) and exploiting that $\nabla \cdot \mathbf{j} = 0$ under the quasi-static approximation, we get

$$-\nabla \cdot (\sigma \nabla u) = -\nabla \cdot \mathbf{j}_s = f \tag{2.14}$$

which is a PDE known as Poisson's equation relating the electric potential at a given point in space to an electric current source. When solving the EEG forward problem, f is known and we are interested in simulating the measurements, u, which we would have obtained given the specified source distribution. The abstract source term, f, is a scalar field which models the source (or sink) strength at each location in space and can be interpreted as a *monopolar* source distribution (Vorwerk, 2016). However, we are often interested in modeling the neural sources as dipoles and we will discuss how this can be achieved in section 2.2.4.

2.2.2 Analytical Solutions in Simple Geometries

The potential field at **r** generated by a current dipole with moment $\mathbf{p} = q\mathbf{d}$ (where q and \mathbf{d} denote the strength and orientation of the dipole, respectively) at position \mathbf{r}' in an infinite conductor is given by

$$u\left(\mathbf{r}\right) = \frac{\mathbf{p}^{\top}\left(\mathbf{r} - \mathbf{r}'\right)}{4\pi\sigma \|\mathbf{r} - \mathbf{r}'\|^{3}}.$$
(2.15)

If the dipole is placed at the origin and oriented along the z-axis then this simplifies to

$$u\left(\mathbf{r}\right) = \frac{q\cos\theta}{4\pi\sigma \left\|\mathbf{r}\right\|^2} \tag{2.16}$$

from which it is apparent that the field decreases with the squared distance between dipole and the point of interest. Here θ denotes the angle between **r** and the z-axis (Hallez et al., 2007).

The potential can also be computed analytically in a single sphere, however, for EEG this model is generally too simple as it disregards the low conductivity of the skull compartment. Consequently, solutions for multi-layer spheres with homogeneous conductivity have been derived (see appendix D.1). A commonly used model is the three-layer concentric sphere model where the three compartments represent brain (i.e., white matter, gray matter, and CSF), skull, and scalp. This model is fast to calculate, however, it does not allow modeling of the curvature of each tissue interface which are all assumed to be completely spherical.

2.2.3 The Boundary Element Method

The Poisson equation can only be solved analytically for very simple geometries. Therefore, if we want to model the anatomy of the volume conductor more realistically, we will need to settle for an approximate solution. BEM allows modeling of a set of realistically shaped, concentric surfaces typically corresponding to brain, skull, and scalp compartments. The integral formulation of equation (2.14) for the potential at a position \mathbf{r} on the *i*th surface in a piecewise homogeneous volume conductor (with arbitrarily shaped interfaces) due to a source current somewhere inside this conductor, is given by

$$u_{i}(\mathbf{r}) = \frac{2\sigma_{0}}{\sigma_{i}^{-} + \sigma_{i}^{+}} u_{0}(\mathbf{r}) + \frac{1}{2\pi} \sum_{j=1}^{|S|} \frac{\sigma_{j}^{-} - \sigma_{j}^{+}}{\sigma_{i}^{-} + \sigma_{i}^{+}} \int_{\mathbf{r}' \in S_{j}} u(\mathbf{r}') d\Omega(\mathbf{r}')$$
(2.17)

where σ_0 is the conductivity of the medium in which the source resides (i.e., the brain compartment), u_0 is the potential at \mathbf{r} due to this dipole in an infinite homogeneous conductor with conductivity σ_0 as defined in equation (2.15), $\sigma_j^$ and σ_j^+ are the conductivities of the inner and outer compartments wrt. the *j*th surface, denoted by S_j , and |S| is the number of surfaces in the domain. Ω denotes the solid angle subtended at \mathbf{r} by a surface element on S_j at position $\mathbf{r'}$ (Hallez et al., 2007; Hämäläinen & Sarvas, 1989). The solid angle of an object subtended from a particular point of view is a measure of how large the object appears to the observer and is equal to the area of a unit sphere covered by the object as seen from the origin. As such, it depends on the distance between the bodies and the size and orientation of the object. For example, the solid angle of an arbitrarily oriented surface element S with normal \mathbf{n} at position \mathbf{q} observed from \mathbf{p} , is given by

$$\Omega = \oint_{S} \frac{\left(\mathbf{q} - \mathbf{p}\right)^{\top} \mathbf{n}}{\left\|\mathbf{q} - \mathbf{p}\right\|^{3}} dS.$$
(2.18)

which calculates the projection of S onto a unit sphere located at **p**. Specifically, for a triangle with vertices a, b, and c whose coordinate vectors are **a**, **b**, and **c**, we have

$$\Omega = 2 \arctan \frac{\mathbf{a}^{\top} (\mathbf{b} \times \mathbf{c})}{\|\mathbf{a}\| \|\mathbf{b}\| \|\mathbf{c}\| + \mathbf{a}^{\top} \mathbf{b} \|\mathbf{c}\| + \mathbf{a}^{\top} \mathbf{c} \|\mathbf{b}\| + \mathbf{b}^{\top} \mathbf{c} \|\mathbf{a}\|}$$
(2.19)

where the numerator is the determinant of the matrix with columns \mathbf{a} , \mathbf{b} , and \mathbf{c} (de Munck, 1992).

For surfaces of arbitrary shape, integration over the whole surface cannot be done analytically and so each surface is typically discretized into a set of triangles and the integral is transformed into a summation of integrals for each triangle. However, an exact solution over the whole domain is generally not possible. The points at which the solution is required to satisfy equation (2.17) are known as collocation points and u is modeled using a set of sparse basis functions (e.g., constant or linear) placed at these points (Hallez et al., 2007). The fully discretized version of equation (2.17) can be written as (Oostendorp & van Oosterom, 1989)

$$u_{k}^{i} = \frac{2\sigma_{0}}{\sigma_{k}^{-} + \sigma_{k}^{+}} u_{0,k}^{i} + \sum_{l=1}^{|S|} \sum_{j=1}^{|S_{l}|} \frac{1}{2\pi} \frac{\sigma_{l}^{-} - \sigma_{l}^{+}}{\sigma_{k}^{-} + \sigma_{k}^{+}} \Omega_{kl}^{ij} u_{l}^{j}$$
(2.20)

which computes the potential at the *i*th point on the *k*th surface (superscripts and subscripts refer to points and surfaces, respectively). Ω_{kl}^{ij} denotes the solid angle of point *j* on surface *l* as seen from point *i* on surface *k* (which is a weighted

average of the triangles to which j belongs) and we use $\Omega_{kk}^{ii} = 0$. By letting

$$\mathbf{B}_{kl}^{ij} = \frac{1}{2\pi} \frac{\sigma_l^- - \sigma_l^+}{\sigma_k^- + \sigma_k^+} \Omega_{kl}^{ij}$$
(2.21)

$$\mathbf{g}_k^i = \frac{2\sigma_0}{\sigma_k^- + \sigma_k^+} u_{0,k}^i \tag{2.22}$$

we can write this compactly as

$$\mathbf{u} = \mathbf{B}\mathbf{u} + \mathbf{g} \tag{2.23}$$

where **u** is the potential of interest, **g** is determined by the source current, and the (typically dense) matrix **B**, which depends only on the geometry and conductivity of the compartments, describes the interactions of the system. As with equation (2.15), the solution is defined only up to a constant and thus deflation⁹ is typically performed,

$$\mathbf{C} = \mathbf{B} - \frac{1}{N} \mathbf{1} \mathbf{1}^{\top}, \qquad (2.24)$$

where N is the number of collocation points and $\mathbf{1}$ is the vector of all ones. The effect of this is to enforce an average potential of zero thus eliminating the arbitrary constant¹⁰. Finally, if we define

$$\mathbf{A} = \mathbf{I} - \mathbf{C} \tag{2.25}$$

where **I** is the identity matrix, then we get

$$\mathbf{A}\mathbf{u} = \mathbf{g} \tag{2.26}$$

which can be solved for \mathbf{u} either directly or iteratively depending on the size of \mathbf{A} . It turns out, however, that the numerical accuracy is poor when the conductivity ratio between the brain and skull compartment approaches zero (as is the case in practice). The may be alleviated by modifying the source term, \mathbf{g} , and is known as the isolated problem approach (Hallez et al., 2007; Hämäläinen & Sarvas, 1989).

2.2.4 The Finite Element Method

The finite element method is another numerical technique for obtaining approximate solutions to PDEs which uses volume elements (e.g., tetrahedra) as opposed to BEM which uses surface elements. In conforming (or continuous Galerkin) FEM, we start by defining the problem (in the strong form) along with appropriate boundary conditions. This is transformed to the weak form and Galerkin's method is used to reduce the dimensionality of the problem (i.e., discretize it). The result is a linear system of equations for which we need to assemble the relevant matrices and vectors before we are able to solve this system for the quantity of interest.

In the following, we will elaborate on each step of this process. As we presented the problem in the strong form above, we will start by defining the boundary conditions.

complement of \mathbf{v} , i.e., projecting out everything in \mathbf{A} that is aligned with \mathbf{v} . ¹⁰We can see this by $\left(\mathbf{B} - \frac{1}{N}\mathbf{1}\mathbf{1}^{\top}\right)\mathbf{u} = \mathbf{B}\mathbf{u} - \frac{1}{N}\mathbf{1}\mathbf{1}^{\top}\mathbf{u} = \mathbf{B}\mathbf{u} - \mathbf{1}\frac{1}{N}\sum_{i=1}^{N}u_i$

⁹By deflation of a matrix \mathbf{A} by a vector \mathbf{v} , we mean projection of \mathbf{A} onto the orthogonal complement of \mathbf{v} , i.e., projecting out everything in \mathbf{A} that is aligned with \mathbf{v} .

Boundary Conditions

In order to constrain the solution space one needs to specify suitable boundary conditions which the solution must satisfy. In particular, when solving the Poisson equation in EEG, the following boundary conditions are typically imposed

$$u_i = u_i \quad \text{in } \Omega \tag{2.27a}$$

$$\langle \mathbf{j}_i, \mathbf{n} \rangle = 0 \quad \text{on } \partial \Omega \tag{2.27b}$$

where i and j are two neighboring element faces and **n** is the vector normal to an element face. $\langle \mathbf{j}_i, \mathbf{n} \rangle$ is the projection of the current density on the normal direction, and Ω and $\partial \Omega$ refer to the interior of the domain and its surface, respectively. Equation (2.27a) says that the potential is continuous across the domain (i.e., no jumps between elements) and equation (2.27b) says that no current can leave the domain (i.e., be injected into the air compartment due to its very low conductivity). Equation (2.27a) is a so-called Dirichlet boundary condition (specifying a particular value at a boundary) whereas equation (2.27b)is a Neumann boundary condition (specifying a particular value of the normal derivative at the boundary).

The Weak Form

Equation (2.14) is the so-called *strong* form of Poisson's equation. Its solution is a twice differentiable function u which satisfies the boundary conditions. We can, however, relax the former requirement by restricting the solutions to a particular function space, V, where the solutions are valid only wrt. certain test functions or test vectors, $v \in V$ with $V = \{v \mid v \in H^1(\Omega)\}$ where H^1 is a so-called Sobolev space¹¹. If u is a solution to the strong form, then the equality still holds if we multiply with a test function v from V and integrate over the domain

$$-\int_{\Omega} \nabla \cdot (\sigma \nabla u) v \, dx = \int_{\Omega} f v \, dx \quad \forall v \in V.$$
(2.28)

Applying multidimensional integration by parts¹² and the divergence theorem¹³ to the left-hand side we get

$$-\int_{\Omega} \nabla \cdot (\sigma \nabla u) v \, dx = \int_{\Omega} \langle \sigma \nabla u, \nabla v \rangle \, dx - \int_{\partial \Omega} v \langle \sigma \nabla u, \mathbf{n} \rangle \, dx \tag{2.29}$$

and the surface integral vanishes due to the boundary condition in equation (2.27b)(because $\mathbf{j} = -\sigma \nabla u$) and therefore

$$\int_{\Omega} \langle \sigma \nabla u, \nabla v \rangle \, dx = \int_{\Omega} f v \, dx \quad \forall \, v \in V$$
(2.30)

18

¹¹The Sobolev space H^1 is the space of functions equipped with an L^2 norm (and hence,

an inner product) and a 1st derivative (in a weak sense). ¹²In the one-dimensional case we have $\int_{\Omega} u\nabla v = \int_{\Omega} uv - \int_{\Omega} v\nabla u$ where the idea is to flip the derivative in the original expression.

 $^{{}^{13}\}int_{\Omega} \nabla \cdot \mathbf{u} \ d\Omega = \int_{\Gamma} \langle \mathbf{u}, \mathbf{n} \rangle \ d\Gamma$, where Γ is the surface of Ω , meaning that (for a closed surface) the sum of sources and sinks within a volume is equal to the net flux across its surface.

2.2. THE EEG FORWARD PROBLEM

which is the *weak* form of Poisson's equation¹⁴. Importantly, the solution to this equation need only be once-differentiable. As such, the solutions to the weak form is not guaranteed to satisfy the strong form. It is common to write this in bilinear notation where we define

$$a(u,v) = \int_{\Omega} \langle \sigma \nabla u, \nabla v \rangle \, dx \tag{2.31}$$

$$l(v) = \int_{\Omega} f v \, dx \tag{2.32}$$

with dx being an infinitesimally small volume in the domain.

The Galerkin Method

So far, the solutions have been defined in V which is an infinite dimensional space, however, to be able to represent the problem on a computer, we shall project the solutions onto a finite dimensional subspace $V_h \subset V$. This is known as the Galerkin method and equation (2.30) in bilinear notation becomes

$$a(u_h, v_h) = l(v_h).$$
 (2.33)

If $\{\mathbf{h}_1, \mathbf{h}_2, ..., \mathbf{h}_n\}$ is a basis for V_h then we can represent the test functions and the solutions in terms of this basis, i.e., $\mathbf{v} = \sum_{i=1}^n v_i \mathbf{h}_i(x)$ and $\mathbf{u} = \sum_{i=1}^n u_i \mathbf{h}_i(x)$ and we get

$$a(\mathbf{u}, \mathbf{v}) = a\left(\sum_{i=1}^{n} u_i \mathbf{h}_i, \sum_{j=1}^{n} v_j \mathbf{h}_j\right)$$
(2.34)

$$l(\mathbf{v}) = l\left(\sum_{j=1}^{n} v_j \mathbf{h}_j\right).$$
(2.35)

The full expression becomes

$$\sum_{i=1}^{n} \sum_{j=1}^{n} u_i v_j \int_{\Omega} \left(\sigma \nabla \mathbf{h}_i \right)^\top \nabla \mathbf{h}_j \, dx = \sum_{j=1}^{n} v_j \int_{\Omega} \mathbf{f}^\top \mathbf{h}_j \, dx \tag{2.36}$$

where we may define

$$\mathbf{S}_{ij} = \int_{\Omega} \left(\sigma \nabla \mathbf{h}_i \right)^\top \nabla \mathbf{h}_j \, dx \tag{2.37a}$$

$$\mathbf{u} = [u_1, u_2, ..., u_n]^{\top}$$
 (2.37b)

$$\mathbf{v} = [v_1, v_2, ..., v_n]^{\top}$$
 (2.37c)

$$b_i = \int_{\Omega} \mathbf{f}^\top \mathbf{h}_i \, dx. \tag{2.37d}$$

This is a linear system of n equations with n unknowns (the coefficients, u_i) and may be written

$$\mathbf{v}^{\top}\mathbf{S}\mathbf{u} = \mathbf{v}^{\top}\mathbf{b} \Leftrightarrow \mathbf{S}\mathbf{u} = \mathbf{b}$$
(2.38)

¹⁴"Weak" in the sense that, compared to the initial equation, its solutions are only required to be differentiable in a weak sense and that it contains only "weaker" derivative information (first order) compared to the strong form (second order).

meaning that it suffices to use the basis as test functions¹⁵. Here, **S** is a positive definite matrix¹⁶ known as the *stiffness matrix*, **u** is the vector of coefficients (electric potentials), and **b** is the vector of current sources, often simply referred to as the right-hand side (RHS). The stiffness matrix describes the physical system and is determined by how the domain is discretized, the conductivities, and how the basis functions are defined.

In FEM, the domain is discretized into a number of non-overlapping elements (e.g., triangles in two dimensions, tetrahedra in three dimensions) defined by a set of nodes and their connectivity, collectively called a mesh. The basis functions are often defined on the nodes in such a way that the *i*th basis function takes on a value of one at the *i*th node and zero at all other nodes. Therefore, a basis coefficient, u_i , is also the function value at its associated point, \mathbf{x}_i . The basis functions are typically chosen from the polynomial family and, as mentioned, are restricted to have local support meaning that they are nonzero only over a small subset of the domain, Ω . This induces sparsity in the stiffness matrix since only functions with intersecting support will contribute nonzero entries to \mathbf{S} and is an attractive feature because the dimensions of this matrix is typically very large. The solution, \mathbf{u} , is thus piecewise polynomial.

A finite element consists of three entities: an element (e.g., a triangle or tetrahedron), the basis functions (e.g., polynomial) defined on this element, and the corresponding degrees of freedom (DOF) which is the number of support points which depends on the order of the polynomial, e.g., for a triangle in two dimensions, a polynomial of degree p will have (p+1)(p+2)/2 terms thus needing this many support points. In particular, for p = 1 the DOFs are simply the number of nodes in the domain. As such, the solution is only defined at the support points, however, we can interpolate to all positions by considering the relevant basis functions (i.e., those whose support is over the point of interest).

Assembling the Stiffness Matrix

The stiffness matrix depends not on the basis functions directly but on their gradients. Since we have discretized the domain into a number of elements, the integral over Ω may be divided into integrals over each, i.e.,

$$\mathbf{S}_{ij} = \int_{\Omega} \left(\sigma \nabla \mathbf{h}_i \right)^\top \nabla \mathbf{h}_j \ dx = \sum_e \int_{\Omega(e)} \left(\sigma \nabla \mathbf{h}_i \right)^\top \nabla \mathbf{h}_j \ dx \tag{2.39}$$

where $\Omega(e)$ is the domain of a particular element, e. Due to the limited support of each basis function, only a few of these integrals (those where i and j have intersecting support, for example in elements containing both nodes) will evaluate to something other than zero.

Assembling the stiffness matrix requires one to evaluate a lot of these integrals which can be done analytically or using numerical integration schemes where, in the latter approach, an integral is approximated as

$$\int_{\Omega} f(x) \, dx \simeq \sum_{i=1}^{n} w_i f(p_i) \tag{2.40}$$

 $^{^{15}\}mathrm{This}$ is true because the equation holds for all v.

 $^{^{16}}$ We need to impose an additional constraint on the solution in order for this to be true as the potential is defined only up to a constant.

with p_i and w_i being the integration points and weights, respectively. For polynomials, the integrals can be evaluated exactly using a suitable number of such points (e.g., for linear functions the midpoint rule gives exact results).

For convenience, the integrals over the individual elements may be transformed to an integral defined on a *reference* element. This corresponds to working in a barycentric coordinate system instead of the cartesian (or physical) coordinate system in which the elements are defined. Thus, the basis functions (and their gradients) need only be defined on the reference element from which they can be mapped to the individual elements. Once this mapping is calculated (see appendix D.2 for more details), the contribution from element e to \mathbf{S}_{ij} can be computed using integration by substitution¹⁷

$$S_{ij}^e = \int_{\Omega(e)} \sigma_e \nabla \mathbf{h}_i^e(x)^\top \nabla \mathbf{h}_j^e(x) \, dx \tag{2.41a}$$

$$= \int_{\Omega(r)} \sigma_e \left(\nabla \mathbf{h}_i^r(\lambda) \mathbf{T}_e^{-1} \right)^\top \left(\nabla \mathbf{h}_j^r(\lambda) \mathbf{T}_e^{-1} \right) |\det \mathbf{J}_{\mathbf{T}_e}| \ d\lambda$$
(2.41b)

$$\approx \sum_{l=1}^{L} w_l \sigma_e \left(\nabla \mathbf{h}_i^r(\lambda_l) \mathbf{T}_e^{-1} \right)^\top \left(\nabla \mathbf{h}_j^r(\lambda_l) \mathbf{T}_e^{-1} \right) |\det \mathbf{J}_{\mathbf{T}_e}|$$
(2.41c)

where λ_l are the barycentric coordinates associated with the *l*th integration point, **T** is the matrix associated with mapping coordinates from reference to cartesian space, and **J** denotes the Jacobian. The first integral is over the physical element, $\Omega(e)$, the second integral is over the reference element, $\Omega(r)$, and the approximation is the numerical integration. Since the gradients are evaluated on the reference element, we only have to evaluate them once for each integration point.

Assembling the Right-Hand Side

Evaluating the RHS as defined in equation (2.37d) will depend on how f is defined. This is the problem of choosing a suitable source model and is discussed in section 2.2.4.

Solving the System

Since the stiffness matrix is generally large (e.g., $n > 10^5$) and sparse, solving the system directly (e.g., using LU decomposition) is not very efficient. As such, iterative solvers which are able to exploit the particular structure of the stiffness matrix are typically employed—possibly combined with preconditioning. The condition number of **S** depends on the quality of the mesh and so it is desirable to avoid ill-shaped elements (e.g., small angles) to the extent possible.

As mentioned previously, the potential is only defined up to a constant. To obtain a unique solution, one may for example enforce that the potential is zero at a single (reference) node on the boundary (Awada et al., 1997; Vorwerk, 2016).

 $^{{}^{17}\}int_{\varphi(U)} f(\mathbf{v}) d\mathbf{v} = \int_U f(\varphi(\mathbf{u})) |\det(D\varphi)(\mathbf{u})| d\mathbf{u}$ where, in our case, \mathbf{v} and \mathbf{u} are cartesian and barycentric coordinates, respectively, φ is T, and $D\varphi$ denotes the Jacobian of φ .

Source Models

So far we have assumed that we can model the current source, abstractly termed f, but we have not specified *how*. It is common to use the *current dipole* to model the sources in EEG since it approximates reasonably well the fairly focal sources in the cortical sheet (Kirschstein & Köhling, 2009). A dipole is defined as two monopoles of opposite polarity and current flows from the negative charge (source) to the positive charge (sink) separated by a certain distance. Concretely,

$$\mathbf{p} = q\mathbf{d} = q\left[\left(\mathbf{x}_0 + \frac{\mathbf{d}}{2}\right) - \left(\mathbf{x}_0 - \frac{\mathbf{d}}{2}\right)\right]$$
(2.42)

where **p** is the dipole moment, q is the strength of the dipole, and **d** defines the displacement vector, i.e., the vector from negative to positive charge with \mathbf{x}_0 being the (center) position of the dipole. In EEG, sources are often modeled as points. Therefore, we define the ideal (or mathematical) dipole as the limit where the displacement goes to zero and the current goes to infinity (so as to keep a constant moment),

$$\mathbf{j}_s(\mathbf{x}) = \mathbf{p}\delta(\mathbf{x} - \mathbf{x}_0),\tag{2.43}$$

which corresponds to a dipole at \mathbf{x}_0 and δ is the Dirac delta distribution. Using this, we may write the RHS as $f(\mathbf{x}) = -\nabla \cdot \mathbf{j}_s(\mathbf{x}) = -\langle \mathbf{p}, \nabla \delta(\mathbf{x} - \mathbf{x}_0) \rangle$, however, this creates a singularity because it is not clear how to evaluate the gradient of the Dirac delta distribution (Vorwerk, 2016).

Traditionally, two approaches have been used to deal with this issue: direct methods and the subtraction method.

Direct Methods Direct methods approximate an ideal dipole by a discrete distribution of sources and sinks placed on the nearby nodes. Methods differ in the way these nodes are selected and how the strength is estimated. In these formulations the RHS is sparse. Here we will describe partial integration and the Saint-Venant approach.

In *partial integration* (Awada et al., 1997; Vorwerk, 2016; Yan et al., 1991), the expression for the dipole moment (i.e., equation (2.43)) is substituted into the expression for the RHS and integration by parts is applied to shift the derivative to the basis function

$$b_i = -\int_{\Omega} \left(\nabla \cdot \mathbf{j}_s(\mathbf{x}_0)\right)^{\top} \mathbf{h}_i^e(\mathbf{x}_0)$$
(2.44a)

$$= -\int_{\Omega} \mathbf{j}_s(\mathbf{x}_0)^\top \nabla \mathbf{h}_i^e(\mathbf{x}_0)$$
(2.44b)

$$= -\int_{\Omega} \mathbf{p}^{\top} \nabla \mathbf{h}_{i}^{e}(\mathbf{x}_{0}).$$
 (2.44c)

Thus, the RHS is computed as the dot product between the dipole moment and the gradient of the basis functions and will have nonzero entries only for those basis functions whose support are over \mathbf{x}_0 (rather than a single nonzero entry). That is, the dipolar source is approximated by a monopolar configuration of sources placed at the integration points of the element within which it is located.

22

2.2. THE EEG FORWARD PROBLEM

The accuracy with which a dipole can be approximated using this approach depends on the geometry of the relevant element (shape, size) as well as the position and orientation of the dipole (e.g., a dipole located close to an edge can be modeled most accurately if it is parallel to the edge but less so if it is orthogonal to it). This is particularly true when using linear basis functions (Awada et al., 1997).

In the *Saint-Venant*¹⁸ approach, the source is approximated by first finding the closest node in the domain and then placing monopoles on the neighboring nodes (Buchner et al., 1997). We start by defining

$$\mathbf{p}_k = \sum_{i=1}^n (\mathbf{x}_i - \mathbf{x}_0)^k q_i \tag{2.45}$$

such that k = 0 corresponds to the monopolar moment and k = 1 corresponds to the dipolar moment. Note, that k = 2 does *not* correspond to the quadrupolar moment (Vorwerk et al., 2019). Using $k \in \{0, 1, 2\}$, the *i*th monopole is positioned at \mathbf{x}_i and *n* is the number of monopoles to be placed, the loads of which are found as the solution to

$$\begin{pmatrix} 1 & \cdots & 1\\ (\mathbf{x}_1 - \mathbf{x}_0) & \cdots & (\mathbf{x}_n - \mathbf{x}_0)\\ (\mathbf{x}_1 - \mathbf{x}_0)^2 & \cdots & (\mathbf{x}_n - \mathbf{x}_0)^2 \end{pmatrix} \begin{pmatrix} q_1\\ \vdots\\ q_n \end{pmatrix} = \begin{pmatrix} p_0\\ \mathbf{p}_1\\ \mathbf{p}_2 \end{pmatrix}, \quad (2.46)$$

which is a system of seven equations (since \mathbf{p}_0 contains only one unique equation) and n unknowns. We can write this more succinctly as

$$\mathbf{X}\mathbf{q} = \mathbf{t} \tag{2.47}$$

where **q** is the monopolar moments (loads), **X** contains the "lever arms" (i.e., the distance weights), and **t** are the target moments. For example, to approximate a dipole with moment **p**, we set $p_0 = 0$, $\mathbf{p}_1 = \mathbf{p}$, and $\mathbf{p}_2 = \mathbf{0}$. Usually n > 7 and the monopolar source configuration is determined as the least squares solution to the problem

$$F(\mathbf{q}, \lambda) = \|\mathbf{t} - \mathbf{X}\mathbf{q}\|^2 + \lambda \|\mathbf{W}\mathbf{q}\|^2$$
(2.48)

where **W** is a weighting matrix (see Vorwerk et al., 2019, for details). In general, **t**, **X**, and **W** are scaled by a suitable reference length but this was omitted here for clarity. Again, the RHS is sparse with n nonzero entries. Also, since the source is snapped to the closest node on the domain, the resolution is limited.

In general, partial integration tends to be more accurate when the source is close to the barycenter of an element whereas Saint-Venant is more accurate the closer the source is to a node since the source is effectively snapped to the closest node on the domain. In partial integration with linear basis functions, the resolution is limited by the fact that the gradient is constant within each element rendering this approach insensitive to the intra-element positions.

¹⁸Saint-Venant's principle was originally formulated for structural mechanics where it states that the results of different load applications become insignificant when observed some distance away. In terms of electrostatics, we may use this to argue that we do not need to consider higher order moments of a multipole expansion, or, as here, that the exact way in which the dipole is modeled does not matter if we only observe the field far away.
The Subtraction Method Another approach to dealing with the issue of dipolar source singularities is the so-called *subtraction method*. The idea is to model the potential and conductivity as the sum of two parts

$$u = u_s + u_0 \tag{2.49}$$

$$\sigma = \sigma_s + \sigma_0 \tag{2.50}$$

where u_0 is the potential resulting from a point dipole in an infinitely homogeneous medium with conductivity σ_0 and u_s is a correction (or subtracted) potential. u_0 is a solution to equation (2.15) and can be computed analytically. Subtracting the strong form of Poisson's equation for u and u_0 and rearranging (see for example Awada et al., 1997) yields

$$\nabla \cdot (\sigma \nabla u_s) = -\nabla \cdot (\sigma \nabla u_0) \tag{2.51}$$

with updated nonhomogeneous Neumann boundary conditions

$$\nabla_{\mathbf{n}} u_s = -\nabla_{\mathbf{n}} u_0. \tag{2.52}$$

Thus, the left-hand side is the same as in the original formulation whereas the RHS is different. Specifically, there is no term corresponding to the original dipolar source (this is modeled by u_0 and eliminated by subtraction), however, there are sources at all interfaces between regions with different conductivities. As such, the RHS is less sparse than in the direct approaches. The subtraction method is generally more accurate than the direct approaches but also more computationally expensive due to the increased work needed to setup the RHS (Awada et al., 1997; Vorwerk, 2016). It is worth noting though, that this approach is less accurate for sources close to conductivity interfaces (tissue boundaries) (Wolters et al., 2007). Finally, the potential of interest is then calculated using equation (2.49).

FEM with Hexahedral Discretization

The approach described above should work well for a tetrahedral discretization of the domain, however, most meshing tools (e.g., NetGen, TetGen) require a set of non-intersecting surfaces which is then meshed. These tools become unstable if the surfaces touch (or are very close together) in which case they might fail (for example, when a subject is lying in an MRI scanner this might happen in the occipital area due to the brain touching the skull) (Rice et al., 2013). Recently, meshes based on hexahedra have become popular due to the ease with which they can be generated from a volume segmentation of an MRI scan. One issue with hexahedral meshes is that they are ill-suited to modeling curved geometries (unless sufficiently dense, e.g., 1 mm resolution). Consequently, in places where the skull is very thin, a node may belong to both CSF and scalp compartments effectively resulting in a hole in the skull through which current may "leak". In this case, the standard method of enforcing continuous potentials (termed continuous Galerkin) will lead to inaccurate field estimates. Discontinuous Galerkin methods, which enforce continuous *current density*, might be more appropriate in such cases where locally thin compartments with substantially different conductivities (e.g., skull and CSF) are to be modeled with limited resolution. In these approaches, the current only flows across element faces (not

through nodes) and therefore, even if, for example, CSF and skin share a node, the current will still flow through the skull compartment (Vorwerk, 2016).

Another class of FEM models are the unfitted approaches where the mesh is not fitted to the individual compartments of the domain but a (separately defined) set of interfaces are allowed to cross the elements. One such approach, CutFEM, implicitly cuts hexahedra by level-set information (i.e., anatomical information about tissue boundaries derived from, for example, MRI). These refinements do not add more degrees of freedom (i.e., more elements) *per se*; they are only used for assembling the stiffness matrix. As such, they allow for more accurate integration by cutting (or splitting) each element into several subelements each of which are assigned a conductivity value based on the level-set information before the contribution to the stiffness matrix is computed (Erdbrügger, 2021).

Solving the Forward Problem Using Reciprocity

The above formulation of the problem solves for potentials resulting from point dipoles (sources), however, due to Helmholtz' principle of reciprocity we may turn the problem around so that it is not a problem of *measurement* but rather of *stimulation*. Helmholtz' principle of reciprocity states that knowing the potential difference between two points induced by a dipole is equivalent to knowing the electric field at the position of the dipole resulting from a current running between these two points (Weinstein et al., 2000). The stimulation scenario corresponds to TES (specifically with direct current) and will be explained in the following (Weinstein et al., 2000; Ziegler et al., 2014).

In TES, we are interested in simulating the electric field resulting from applying a potential difference between two electrodes on the scalp¹⁹. The problem is very similar to the EEG forward problem described above, however, as we have no current sources in the brain, Poisson's equation reduces to Laplace's equation, i.e., the RHS is all zeros. The boundary conditions are, however, slightly different. Specifically, we may use

$$u_i = u_j \quad \text{in } \Omega \tag{2.53a}$$

$$\langle \mathbf{j}_i, \mathbf{n} \rangle = 0 \quad \text{on } \partial \Omega \setminus (\partial \Gamma_1 \cup \partial \Gamma_2)$$
 (2.53b)

$$u = 0 \quad \text{on } \partial \Gamma_1 \tag{2.53c}$$

$$u = f \quad \text{on } \partial \Gamma_2 \tag{2.53d}$$

where $\partial \Gamma_1$ and $\partial \Gamma_2$ are surfaces of the electrodes used in the stimulation (i.e., between which the potential difference is set). Γ_1 is the ground electrode. The solution is unique because of the Dirichlet boundary conditions (whereas the formulation presented above only had the Neumann boundary conditions).

Assembling the stiffness matrix is similar to the description above. Assembling the RHS is, however, different. As there are no current sources, it is initially zero. Imposing the condition in equation (2.53d), amounts to removing these rows from the stiffness matrix and the RHS and adjusting the latter accordingly (Saturnino, Madsen, & Thielscher, 2019).

 $^{^{19}}$ In TES, we are usually interested in specifying the current instead, however, this can be obtained by rescaling the result of applying a unit potential difference (Saturnino, Madsen, & Thielscher, 2019).

In TES, we are interested in the electric field (or current density). To obtain the electric field from the FEM solution, we compute the (negative) gradient in each element. For linear basis functions, this is particularly simple as it is constant over each element and therefore also not very accurate for most locations (especially at the nodes). To recover the values of the derivative function we can use a smooth estimator. Superconvergent patch recovery (SPR)²⁰ (Zienkiewicz & Zhu, 1992) is a smoothing operation which attempts to recover nodal values such that the continuous estimate of \mathbf{E} is more accurate than the initial one. Specifically, SPR uses a polynomial expansion (of the same order as the basis functions) of the derivative function in a patch of elements surrounding the node whose value is to be recovered and estimates its parameters by least squares. The patch elements are the elements to which the node of interest belongs and the choice of superconvergent points within these elements, i.e., the points which are used to fit the parameters, depends on the order of the basis function (e.g., for linear elements we use the barycenters). An electric field component at an arbitrary position, **r**, is estimated as

$$e(\mathbf{r}) = \mathbf{x}^{\top} \mathbf{b} \tag{2.54}$$

where **x** contains the appropriate polynomial terms of the coordinates of **r** (e.g., $\mathbf{x} = (1, x, y, x^2, xy, y^2)$ for a quadratic function at a point (x, y) in two dimensions) and **b** contains the associated weights which are to be estimated. This is solved by minimizing the least squares error between the values of the function, **f**, and the interpolant at the element barycenters

$$\underset{\mathbf{b}}{\operatorname{arg\,min}} \|\mathbf{f} - \mathbf{X}\mathbf{b}\|^2. \tag{2.55}$$

X is constructed by row-wise stacking of the coordinate expansions of the relevant points and we can compute the projection matrix at **r** as $\mathbf{x}^{\top} (\mathbf{X}^{\top} \mathbf{X})^{-1} \mathbf{X}^{\top}$ and apply it to each component of the electric field to interpolate the field at this location. If the positions of interest do not happen to coincide with the domain nodes, we can obtain the value of the these positions by barycentric interpolation within each element (Saturnino, Madsen, & Thielscher, 2019; Zienkiewicz & Zhu, 1992).

At present, we have a smooth estimate of the electric field at a position **r** resulting from an imposed potential difference between two electrodes, a and b, however, in EEG we are interested in the potential difference between these electrodes generated by a source current at **r**. As posited above, using Helmholtz' principle of reciprocity we are able to obtain the latter from the former, specifically, letting u_{ab}^{EEG} be the observed potential difference between electrodes a and b due to a (dipolar) current source, $\mathbf{j}^{\text{EEG}}(\mathbf{r})$, placed at position **r** and let I_{ab}^{TES} be an imposed current between electrodes a and b resulting in an electric field, $\mathbf{e}^{\text{TES}}(\mathbf{r})$, at position **r**. We get

$$u_{ab}^{\text{EEG}} I_{ab}^{\text{TES}} = -\mathbf{e}^{\text{TES}}(\mathbf{r})^{\top} \mathbf{j}^{\text{EEG}}(\mathbf{r})$$
 (2.56a)

$$u_{ab}^{\text{EEG}} = -\frac{\mathbf{e}^{\text{IES}}(\mathbf{r})^{\top}}{I_{ab}^{\text{TES}}} \mathbf{j}^{\text{EEG}}(\mathbf{r})$$
(2.56b)

$$u_{ab}^{\text{EEG}} = \mathbf{G}_{ab}(\mathbf{r})\mathbf{j}^{\text{EEG}}(\mathbf{r})$$
(2.56c)

 $^{^{20} \}rm Superconvergence refers to the ability to recover (gradient) values at a higher order of precision than that of the basis functions.$

where $\mathbf{G}_{ab}(\mathbf{r})$ is the potential field value at *a* wrt. *b* for unit dipoles at **r** aligned with each principal axis (Ruffini, 2015; Weinstein et al., 2000)²¹.

One attractive feature of this approach is that the number of electrodes is typically much less than the number of sources. A similar strategy was proposed by Wolters et al. (2004). Given the size of the stiffness matrix, it is usually not possible to invert it. Thus, in the naïve approach, one would need to solve equation (2.38) for each source yielding the potential at each DOF in the domain. However, we are only interested in the potential at a few select locations (those corresponding to the electrodes), $\mathbf{u}^{\text{EEG}} \in \mathbb{R}^m$, and these can be picked out by a so-called restriction matrix, $\mathbf{R} \in \mathbb{R}^{m \times n}$, which has a single (or a few) nonzero entry (entries) for each row corresponding to the position of each sensor. Here, m is the number of electrodes and n is the number of DOFs. If we also define a transfer matrix, \mathbf{T} ,

$$\mathbf{u}^{\text{EEG}} = \mathbf{R}\mathbf{u} \tag{2.57}$$

$$\mathbf{T} = \mathbf{R}\mathbf{S}^{-1} \tag{2.58}$$

we can express \mathbf{u}^{EEG} in terms of \mathbf{T} , like

$$\mathbf{u}^{\text{EEG}} = \mathbf{R}\mathbf{u} = \mathbf{R}\mathbf{S}^{-1}\mathbf{b} = \mathbf{T}\mathbf{b}.$$
 (2.59)

Right-multiplying equation (2.58) by **S** and using its symmetry, we get

$$\mathbf{TS} = \mathbf{RS}^{-1}\mathbf{S} \Leftrightarrow \mathbf{ST}^{\top} = \mathbf{R}^{\top}$$
(2.60)

which is a similar problem to that presented for TES with Dirichlet boundary conditions. Once the transfer matrix is computed, it can be applied to the RHSs corresponding to each dipole of interest to obtain e^{EEG} .

2.3 The EEG Inverse Problem

One potential issue in EEG is that all observations contain a mixture of all neural processes occurring at each point in time (Scherg, 1990). To localize sources of neural activity in space, we will have to undo the effects of volume conduction as best we can. This is a so-called inverse problem which, assuming that it is linear, can be written as

$$\int_{a}^{b} \text{system} \times \text{input} = \text{output}$$
(2.61)

i.e., a so-called Fredholm integral of the first kind

$$\int_{a}^{b} K(s,t)f(t) \, dt = g(s) \tag{2.62}$$

where s and t are the positions of measurement and source, respectively. We have information about the output function, g, corrupted by some amount of noise and we also know either the input, f, or the system kernel, K, and we want to determine the other. In CT, for example, the inputs are beams of

²¹Instead of a current density, we simply consider a point dipole with unit A m and no associated (or infinitesimally small) volume and thus we have dropped the δV term.

X-rays and the outputs are measurements of how these are attenuated by the tissue through which they pass and we are interested in determining the system, i.e., the particular tissue composition which gave rise to these observations. On the contrary, in EEG we are interested in determining the sources of the measurements and we acquire knowledge of the system by solving the forward problem as described in section 2.2. The integral in equation (2.62) is typically discretized into a linear system of equations,

$$\mathbf{A}\mathbf{x} = \mathbf{b}.\tag{2.63}$$

Consider **A** with dimensions $m \times n$ and full rank, i.e., rank(**A**) = min(m, n). Let us look at two different scenarios where either m > n or m < n.

If m > n then the system is *overdetermined* since we have more equations than unknowns and in general we have that $\mathbf{b} \notin \text{image}(\mathbf{A})$ meaning that we cannot find a solution, \mathbf{x} , which satisfies equation (2.63). A common strategy is to find the vector which minimizes the squared error between \mathbf{b} and $\mathbf{A}\mathbf{x}$. Since \mathbf{A} is of full rank, it has no null space and this solution is unique. Although \mathbf{A} is not invertible, $\mathbf{A}^{\top}\mathbf{A}$ is and the well known least squares solution is given by

$$\mathbf{x} = \left(\mathbf{A}^{\top}\mathbf{A}\right)^{-1}\mathbf{A}^{\top}\mathbf{b} = \mathbf{A}^{+}\mathbf{b}.$$
 (2.64)

Here, $\mathbf{A}^+ = (\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top$ is called the *pseudoinverse* and is a *left inverse* of \mathbf{A} when m > n. This is the least squares estimate of \mathbf{x} , i.e., the "solution" which minimizes the squared error.

If, on the other hand, m < n then the system is *underdetermined* since we have less equations than unknowns and $\mathbf{b} \in \text{image}(\mathbf{A})$ with infinitely many solutions satisfying equation (2.63). This time we may choose to look for the solution with minimum (squared) Euclidean norm, i.e.,

$$\begin{array}{ll} \text{minimize} & \mathbf{x}^{\top}\mathbf{x} \\ \text{subject to} & \mathbf{A}\mathbf{x} = \mathbf{b}. \end{array}$$
(2.65)

This solution is unique and can be found using Lagrange multipliers by formulating the Lagrange function

$$L(\mathbf{x},\lambda) = \mathbf{x}^{\top}\mathbf{x} + \lambda^{\top} (\mathbf{A}\mathbf{x} - \mathbf{b})$$
(2.66)

and setting

$$\frac{\partial L}{\partial \mathbf{x}} = 2\mathbf{x} + \mathbf{A}^{\top} \lambda = 0 \tag{2.67}$$

$$\frac{\partial L}{\partial \lambda} = \mathbf{A}\mathbf{x} - \mathbf{b} = 0. \tag{2.68}$$

From equation (2.67) we get

$$\mathbf{x} = -\frac{1}{2}\mathbf{A}^{\top}\lambda \tag{2.69}$$

which, when plugged into equation (2.68) gives

$$\lambda = -2 \left(\mathbf{A} \mathbf{A}^{\top} \right)^{-1} \mathbf{b} \tag{2.70}$$

and finally

$$\mathbf{x} = \mathbf{A}^{\top} \left(\mathbf{A} \mathbf{A}^{\top} \right)^{-1} \mathbf{b} = \mathbf{A}^{+} \mathbf{b}.$$
 (2.71)

Here, the pseudoinverse is $\mathbf{A}^+ = \mathbf{A}^\top (\mathbf{A}\mathbf{A}^\top)^{-1}$ and is a *right inverse* of \mathbf{A} when m < n. This is the minimum norm estimate, i.e., the solution with the smallest Euclidean norm (see appendix D.3 for more details on what applying \mathbf{A} and \mathbf{A}^+ actually do).

In summary, the pseudoinverse can be used to find the least squares solution of an overdetermined system of equations as well as choosing the solution with minimum Euclidean norm in an underdetermined system. However, it is unclear at this point whether or not this solution is actually reasonable. This depends on a number of factors, for example the amount of noise in the data and the nature of the system matrix, \mathbf{A} .

2.3.1 Characteristics of the Inverse Problem

There are several difficulties associated with solving inverse problems as they are generally ill-posed. Hadamard listed three conditions which should be satisfied for a problem to be well-posed, namely that

- 1. a solution should exists,
- 2. the solution must be unique, and
- 3. the solution should depend continuously on data and parameters (i.e., not be very sensitive to small perturbations in the initial conditions).

Otherwise a problem is ill-posed. The first issue may be dealt with by reformulating the problem such that we minimize some measure of goodness-of-fit (GOF), e.g., the l_2 -norm, instead of requiring an actual solution. The second issue may be circumvented by imposing additional requirements on the solution, e.g., that we are looking for the solution with the minimum l_2 -norm. The last condition is often violated and is related to the conditioning of the problem (Fox et al., 2010).

The condition number²² describes how much the output of a function changes due to small perturbations in the observations, i.e., how sensitive it is to noise or errors in the observations. If we consider equation (2.63), the condition number of the problem is that of the matrix \mathbf{A} , whose matrix norm, M, wrt. the norm of a certain vector, \mathbf{x} , assumed to be different from the zero vector, is given by

$$M = ||A|| \equiv \sup \frac{||\mathbf{A}\mathbf{x}||}{||\mathbf{x}||} = \sup \left\|\mathbf{A}\frac{\mathbf{x}}{||\mathbf{x}||}\right\|.$$
 (2.72)

That is, the matrix norm measures the maximum stretch of a unit vector performed by the transformation of \mathbf{A} . Similarly, we can define the minimum stretch (or shrink) of a unit vector as

$$m \equiv \inf \frac{\|\mathbf{A}\mathbf{x}\|}{\|\mathbf{x}\|} = \inf \frac{\|\mathbf{b}\|}{\|\mathbf{A}^{-1}\mathbf{b}\|} = \frac{1}{\sup \frac{\|\mathbf{A}^{-1}\mathbf{b}\|}{\|\mathbf{b}\|}} = \frac{1}{\|\mathbf{A}^{-1}\|}$$
(2.73)

where we have m = 0 for a singular (i.e., not invertible) matrix and it is apparent that the reciprocal of the minimum stretch is equal to the norm (maximum

 $^{^{22}}$ Here we are particularly interested in the condition number wrt. to inversion.

stretch) of the inverse. For inversion, the condition number is defined as the ratio of these norms,

$$\kappa(\mathbf{A}) = \frac{M}{m} = \|\mathbf{A}\| \|\mathbf{A}^{-1}\|.$$
(2.74)

For a linear system of equations perturbed by a small amount on the RHS, we have

$$\frac{\|\delta \mathbf{x}\|}{\|\mathbf{x}\|} \le \kappa(\mathbf{A}) \frac{\|\delta \mathbf{b}\|}{\|\mathbf{b}\|}$$
(2.75)

meaning that the condition number is an upper bound on the relative error magnification in the system and hence describes the sensitivity of \mathbf{x} to errors in \mathbf{b} (Fox et al., 2010). The intuition is that directions which are heavily compressed during the forward operation will need to be amplified at the inverse stage and thus errors in these directions might end up dominating the solution. A problem with low condition number is said to be well-conditioned whereas one with high condition number is said to be ill-conditioned. A condition number of infinity means that the matrix is singular.

Most real-world phenomena have the effect of lowpass filtering in the forward direction. This idea is formalized by the Riemann-Lebesgue lemma which states that the integral of the Fourier transform of K (assumed to be L^1 integrable) vanishes at high frequencies (Fox et al., 2010), i.e.,

$$\lim_{|f| \to \infty} \int_{a}^{b} K(s,t) e^{-ift} \, dx = 0.$$
(2.76)

Hence, in the process of going from sources to measurements, high frequency information tends to be lost and g will be smoother than f in equation (2.62). This is also true in EEG where the measurements are weighted averages of all sources, the primary reason for such smoothing effects being the high conductivity gradient between the CSF (high conductivity) and skull (low conductivity) compartments. As such, low frequency variations are associated with high singular values relative to high frequency variations (see figures 2.1 to 2.3). Since this difference may be substantial, it presents a clear problem for such kernels in terms of their condition number. Even if the number of measurements matched that of the sources, the "effective" number of measurements would be less because the data in nearby channels would be very similar. For example, source localization in EEG has been shown to improve substantially when the number of electrodes is increased from 25 to 100 whereas the benefits of going from 100 to 175 electrodes are much more modest (Michel et al., 2004). Obviously, the smoother the kernel, the more ill-conditioned it will be and consequently its singular values will decay faster towards zero.

In light of these observations, let us consider the "naïve" solution to equation (2.63) as given by the pseudoinverse of **A**. Let

$$\mathbf{A} = \mathbf{U}\mathbf{S}\mathbf{V}^{\top} \tag{2.77}$$

be the singular value decomposition (SVD) of the kernel and we can write the solution in terms of the singular values

$$\mathbf{x}_{\text{naive}} = \mathbf{V}\mathbf{S}^{-1}\mathbf{U}^{\top}\mathbf{b} = \sum_{i=1}^{n} \frac{\mathbf{u}_{i}^{\top}\mathbf{b}}{s_{i}}\mathbf{v}_{i}$$
(2.78)



Figure 2.1. Normalized singular values of an SVD of a forward solution (gain matrix) with fixed orientation (normal to the cortical surface) for a single subject. The faster the decay of the singular values, the more difficult the inverse problem. Using such spectra, we may compare, for example, different electrode configurations to determine which provides the most information (Molins et al., 2008).



Figure 2.2. Right singular vectors for select components of an SVD of a forward solution (gain matrix) with fixed orientation (normal to the cortical surface) for a single subject (for the corresponding singular values, see figure 2.1). These correspond to source space modes, i.e., the modes available for reconstructing the source distribution. In minimum norm estimates, the source distribution is therefore a linear combination of these spectra (Molins et al., 2008). We see how components with high indices correspond to modes of higher frequencies. The results are presented on the *fsaverage* template.



Figure 2.3. Left singular vectors for select components of an SVD of a forward solution (gain matrix) with fixed orientation (normal to the cortical surface) for a single subject (for the corresponding singular values, see figure 2.1). These are sensor space modes, i.e., the modes available for reconstructing the observations (Molins et al., 2008), and correspond to the source spaces modes presented in figure 2.2. We see how components with high indices correspond to modes of higher frequencies. If the projection of the data onto either of these components is larger than the corresponding singular values, then this suggests that the origin of this signal is not (entirely) due to neural processes but rather some sort of noise.

from which it is clear that solution coefficients associated with small singular values in the kernel will be large. If the kernel is smooth, variations associated with high frequencies will be highly compressed (small singular values) in the forward operation whereas they will be amplified correspondingly (large reciprocal singular values) during inversion. Assuming sufficient numerical accuracy and noiseless conditions, this is not an issue *per se* and the pseudoinverse will be able to recover the true parameters, however, if there is noise in our measurements (or the kernel is so smooth that the singular values approach machine precision) then the noise is added *in the compressed space* and amplified during inversion. In particular, it is problematic if the projection of the data on the left singular vectors is larger than the corresponding singular value as this suggests that there is more variation in **b** along this direction than one might expect given the properties of **A** if it was in fact due to **x**.

2.3.2 Stabilizing the Solution

An additional complication arises when we are trying to fit more parameters than observations (i.e., when the problem is underdetermined) since the data does not provide enough information to uniquely determine the parameters. This is an extreme case compared to above as the kernel is rank deficient and its condition number is infinite. Consequently, there are an infinite number of solutions because the kernel has a null space, i.e., a set of vectors which are mapped to the zero vector by the kernel and thus would not change the observations²³.

In the above discussion, a solution is a set of parameters which actually solves equation (2.63) such that the residual is zero, however, if the solutions tend to be highly unstable (e.g., very sensitive to noise), then the residual may not be a good indicator of whether or not a particular solution is "reasonable". This will often be the case for the pseudoinverse. Although it is the solution with *minimum* norm, this does not necessarily mean that the norm is small because we require the residual to be zero. That is, we are fitting all the variation in the data including the noise.

To stabilize the solution, we need to bias it in a certain direction that we believe is plausible by adding prior information to the equation system, i.e., information reflecting our expectations about the solution and which are not based on the current data. In Bayesian statistics we talk about priors, i.e., explicitly defined probability distributions over the parameters, which, when combined with the likelihood of the data (and normalized properly) gives a posterior distribution over parameters.

A similar notion is that of *regularization*, the purpose of which is to stabilize the solution. Specifically, we would like our solution to be less sensitive to noise and in order to achieve this, we sacrifice some of our ability to fit the data. Put another way: we increase the bias of our model (in a direction which we believe is sensible) so as to reduce its variance. A common strategy is known as Tikhonov regularization which penalizes the squared magnitude of the parameters, i.e.,

$$L = \|\mathbf{b} - \mathbf{A}\mathbf{x}\|^2 + \|\Gamma\mathbf{x}\|^2, \qquad (2.79)$$

 $^{^{23}}$ The null space is also commonly referred to as the "kernel", however, to avoid confusion we abstain from using this term to refer to the null space.

where Γ is known as the Tikhonov matrix. The solution to this problem is

$$\mathbf{x}_{\text{Tikhonov}} = \left(\mathbf{A}^{\top}\mathbf{A} + \Gamma^{\top}\Gamma\right)^{-1}\mathbf{A}^{\top}\mathbf{b}.$$
 (2.80)

A common choice of Γ is such that $\Gamma = \lambda \mathbf{I}$. This corresponds to ridge regression which penalizes the l_2 -norm of \mathbf{x} and we can write its solution using the SVD²⁴

$$\mathbf{x}_{\text{ridge}} = \left(\mathbf{V} \mathbf{S}^{\top} \mathbf{S} \mathbf{V}^{\top} + \lambda^2 \mathbf{I} \right)^{-1} \mathbf{V} \mathbf{S} \mathbf{U}^{\top} \mathbf{b}$$
(2.81a)

$$= \left(\mathbf{V} \left[\mathbf{S}^{\top} \mathbf{S} + \lambda^2 \mathbf{I} \right] \mathbf{V}^{\top} \right)^{-1} \mathbf{V} \mathbf{S} \mathbf{U}^{\top} \mathbf{b}$$
(2.81b)

$$= \mathbf{V} \left(\mathbf{S}^{\top} \mathbf{S} + \lambda^2 \mathbf{I} \right)^{-1} \mathbf{V}^{\top} \mathbf{V} \mathbf{S} \mathbf{U}^{\top} \mathbf{b}$$
(2.81c)

$$= \mathbf{V} \left(\mathbf{S}^{\mathsf{T}} \mathbf{S} + \lambda^2 \mathbf{I} \right)^{-1} \mathbf{S} \mathbf{U}^{\mathsf{T}} \mathbf{b}$$
(2.81d)

$$= \mathbf{V}\mathbf{D}\mathbf{S}^{-1}\mathbf{U}^{\top}\mathbf{b} \tag{2.81e}$$

$$=\sum_{i=1}^{N}\mathbf{D}_{ii}\frac{\mathbf{u}_{i}^{\top}\mathbf{b}}{s_{i}}\mathbf{v}_{i}$$
(2.81f)

$$=\sum_{i=1}^{N} \frac{s_i^2}{s_i^2 + \lambda^2} \frac{\mathbf{u}_i^{\mathsf{T}} \mathbf{b}}{s_i} \mathbf{v}_i$$
(2.81g)

$$=\sum_{i=1}^{N} \frac{\mathbf{u}_{i}^{\mathsf{T}} \mathbf{b}}{s_{i} + \lambda^{2}/s_{i}} \mathbf{v}_{i}$$
(2.81h)

where the diagonal matrix \mathbf{D} contains the "dampening factors" which are applied to the SVD components (this also makes the comparison with equation (2.78) clear).

The Tikhonov matrix, however, does not need to be diagonal. If we have information about the (full) covariance matrix of the parameters this can be used as a whitener so as to encourage a particular structure of the solution. The loss function for the generalized least squares problem with Tikhonov regularization becomes

$$L = \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_{\mathbf{C}^{-1}}^{2} + \|\mathbf{x}\|_{\mathbf{R}^{-1}}^{2}$$
(2.82)

$$= (\mathbf{b} - \mathbf{A}\mathbf{x})^{\top} \mathbf{C}^{-1} (\mathbf{b} - \mathbf{A}\mathbf{x}) + \mathbf{x}^{\top} \mathbf{R}^{-1} \mathbf{x}$$
(2.83)

where $\|\mathbf{z}\|_{\mathbf{Q}}^2 = \mathbf{z}^{\top} \mathbf{Q} \mathbf{z}$ denotes the squared norm in \mathbf{Q} . $\mathbf{R}^{-1} = \Gamma^{\top} \Gamma$ is the inverse covariance of the parameters. **C** is the covariance of the residuals. If **C** and **R** are positive definite (and hence invertible), then the solution to equation (2.82) may be written either in the least squares form

$$\mathbf{M}_{\rm LS} = \left(\mathbf{A}^{\top} \mathbf{C}^{-1} \mathbf{A} + \mathbf{R}^{-1}\right)^{-1} \mathbf{A}^{\top} \mathbf{C}^{-1}$$
(2.84)

or as the minimum norm (Mosher et al., 2003; Petersen & Pedersen, 2012)

$$\mathbf{M}_{\mathrm{MN}} = \mathbf{R}\mathbf{A}^{\top} \left(\mathbf{A}\mathbf{R}\mathbf{A}^{\top} + \mathbf{C}\right)^{-1}.$$
 (2.85)

Depending on the application, equation (2.84) or equation (2.85) may be preferable. For example, we may want to use the former in the case of an overdetermined system whereas we use the latter for an underdetermined system. In

 $^{^{24}\}text{This}$ only works because $\Gamma^{\top}\Gamma$ is a multiple of the identity matrix.

EEG, the matrix to be inverted in equation (2.84) is sources by sources whereas it is channels by channels in equation (2.85).

The equality of equation (2.84) and equation (2.85) only applies when regularization is used as otherwise either of the matrices to be inverted is singular depending on whether **A** is tall or wide (unless of course **A** is square and full rank). We see that from a numerical point of view, the purpose of the regularization is to improve the condition number of the system and, at the very least, ensure that the relevant matrix is positive definite²⁵.

The matrix **C** reflects our expectation about the covariance of the residuals. We want to avoid fitting the noise, hence, if we can estimate the pattern of the noise in the system, then we can suppress this during estimation through whitening. Likewise, we use **R** to encourage a solution with properties which we deem desirable. Common choices include penalizing the variance of the individual parameters (e.g., using $\mathbf{R}^{-1} = \lambda^2 \mathbf{I}$) or the smoothness of the parameter map (e.g., using a highpass filter such as the discrete Laplacian which penalizes differences between neighboring parameters) and the magnitude of the penalization is used to balance data fit and prior assumptions (i.e., suppressing noise) which will depend on the SNR of the data.

In EEG source analysis, penalizing the norm of the solution may be warranted based on the observation that sustaining neural activity (as reflected by the solution) requires energy which is a limited resource. Thus, being able to explain most of the data with a small amount of neural activation is believed to be more reasonable than requiring large amounts of neural activation to perfectly explain the data thus rendering solution vectors with large norms undesirable. Likewise, smoothness is justified based on the observation that neighboring neurons are more likely to synchronize their activity than distant ones generally speaking (Pascual-Marqui et al., 1994).

2.3.3 EEG Inverse Solvers

In the following, we will present several common approaches to tackling the inverse problem in EEG. We will be concerned with finding sources (or parameters), \mathbf{p} , which are somehow related to our observations, \mathbf{y} (e.g., explains our observations the best). As discussed above, the latter is related to the former through a forward operator, which in this case is a matrix that we shall denote \mathbf{G} —the gain matrix. We now assume that the data is generated in the following way

$$\mathbf{y} = \mathbf{G}\mathbf{p} + \mathbf{n} \tag{2.86}$$

where \mathbf{n} is the noise vector and our objective is to recover \mathbf{p} . If we consider a cost function which penalizes the squared error between predicted and observed data then we are implicitly assuming that the noise is normally distributed with zero mean.

If we have *m* measurements and *p* sources then the dimensions of the above quantities are $\mathbf{y} \in \mathbb{R}^m$, $\mathbf{p} \in \mathbb{R}^{3p}$, $\mathbf{n} \in \mathbb{R}^m$, and $\mathbf{G} \in \mathbb{R}^{m \times 3p}$ where 3p is because we have three orthogonal dipoles at each location. We shall use *i* to refer to a specific position in source space and $\mathbf{G}_i \in \mathbb{R}^{m \times 3}$ to refer to the three columns

²⁵A positive definite matrix is invertible. Adding a positive *definite* matrix to a positive *semidefinite* matrix yields a positive definite matrix. Since **C** and **R** are positive definite and \mathbf{ARA}^{\top} and $\mathbf{A}^{\top}\mathbf{C}^{-1}\mathbf{A}$ are positive semidefinite we know that the result is invertible.

of **G** related to location *i*. Furthermore, $\mathbf{C}(\mathbf{n}) \in \mathbb{R}^{m \times m}$ covariance of the noise, $\mathbf{C}(\mathbf{y}) \in \mathbb{R}^{m \times m}$ covariance of the data, and $\mathbf{R} \in \mathbb{R}^{3p \times 3p}$ covariance of the sources (or parameters).

Distributed Methods

In distributed approaches, the position (and possibly orientation) of each source is regarded as being fixed and the aim is to estimate all amplitudes (or dipole moments) simultaneously. This imposes minimal *a priori* assumptions on the neural activity (e.g., wrt. the number of sources), however, because the number of source locations is usually much greater than the number of measurements, the system is vastly underdetermined. As discussed previously, the solution with minimum norm (given by the pseudoinverse) generally provides a poor estimate of the parameters and regularization needs to be employed to stabilize the result.

Following equation (2.85), the weights of the standard minimum norm estimate (MNE) are given by

$$\mathbf{W} = \mathbf{R}\mathbf{G}^{\top} \left(\mathbf{G}\mathbf{R}\mathbf{G}^{\top} + \lambda^{2}\mathbf{C}(\mathbf{n})\right)^{-1}$$
(2.87)

where λ is a regularization parameter related to the SNR of the data (Lin et al., 2006). The source estimate is obtained by

$$\mathbf{p} = \mathbf{W}\mathbf{y}.\tag{2.88}$$

An estimate of the variance of the *i*th source is then given by

$$\sigma_i = \mathbf{p}_i^{\top} \mathbf{p}_i \tag{2.89}$$

where $\mathbf{p}_i \in \mathbb{R}^3$ is the vector containing the parameters associated with the *i*th source.

MNE estimators tend to suffer from depth bias because the measured signal in EEG decreases with the square of the distance between source and electrode as is evident from equation (2.16). Hence, the norm of the columns of the gain matrix associated with deep sources is generally small compared to more superficial ones. Therefore, higher source amplitudes are needed in order for a deep source to explain the data. To mitigate this, depth weighting can be included in \mathbf{R} , effectively scaling the columns of \mathbf{G} so as to normalize them in some way. For example, the depth weighting factor for location i may be calculated using

$$f_i = \operatorname{trace} \left(\mathbf{G}_i^\top \mathbf{G}_i \right)^{-\gamma} \tag{2.90}$$

where γ determines the amount of depth weighting (Lin et al., 2006). Thus, depth weighting increases the prior variance for deep sources compared to superficial ones.

Finding a depth weighting parameter which works well for all data may be difficult although sensible ranges have been suggested previously (Lin et al., 2006). Another way to achieve this is to use a noise normalized estimate for localization. Two such methods are dynamic statistical parametric mapping (dSPM) and standardized low resolution electromagnetic tomography (sLORETA). Both of these methods normalize the original MNE estimate of the source variance from equation (2.89), however, they differ in what is considered noise. dSPM normalizes by an estimate of the source variance due to noise as captured by the noise covariance matrix, i.e.,

$$F_i^{\text{dSPM}} = \frac{\mathbf{p}_i^{\top} \mathbf{p}_i}{\text{trace} \left(\mathbf{W}_i \mathbf{C}(\mathbf{n}) \mathbf{W}_i^{\top} \right)}.$$
 (2.91)

sLORETA also includes the variance of the actual sources (as estimated from the gain matrix), effectively normalizing by *ii*th block of the parameter resolution matrix,

$$F_i^{\text{sLORETA}} = \mathbf{p}_i^{\top} \left(\mathbf{W}_i \left[\mathbf{G}_i \mathbf{R}_{ii} \mathbf{G}_i^{\top} + \lambda^2 \mathbf{C}(\mathbf{n}) \right] \mathbf{W}_i^{\top} \right)^{-1} \mathbf{p}_i$$
(2.92)

$$= \mathbf{p}_{i}^{\top} \left(\mathbf{R}_{ii} \mathbf{G}_{i}^{\top} \left[\mathbf{G}_{i} \mathbf{R}_{ii} \mathbf{G}_{i}^{\top} + \lambda^{2} \mathbf{C}(\mathbf{n}) \right]^{-1} \mathbf{G}_{i} \right)^{-1} \mathbf{p}_{i}$$
(2.93)

$$=\mathbf{p}_{i}^{\top} \left(\mathbf{W}_{i} \mathbf{G}_{i}\right)^{-1} \mathbf{p}_{i}.$$
(2.94)

sLORETA has been shown to have zero localization error in noiseless conditions when the true gain matrix is known (Pascual-Marqui, 2002). Under the null hypothesis with free source orientations, $F_i^{\rm dSPM}$ and $F_i^{\rm sLORETA}$ both follow an F distribution.

In the above formulations, \mathbf{R} is usually taken to be diagonal, however, other popular approaches also exist. For example, smoothness may be imposed using the discrete Laplacian of the surface mesh (or regular grid) on which the sources are defined as is done in low resolution electromagnetic tomography (LORETA) (Pascual-Marqui et al., 1994). Smoothness can also be modeled by incorporating biophysical knowledge as in local autoregressive averages (LAURA) where the prior is based on the fact that the potential field from a dipole decreases in strength according to the squared distance (de Peralta Menendez et al., 2004) as is apparent from equation (2.16).

Another popular strategy is to penalize the l_1 -norm of the source amplitudes to encourage sparse source estimates or even mixed norms which impose sparseness in space and smoothness in time (Gramfort et al., 2012).

Dipole Fit

Another way to overcome the ill-posedness of the inverse problem is to assume that there is only a single active source responsible for generating the observations. In this case, we can use equation (2.84) to estimate the dipole moment. Without regularization, the weights for the *i*th source are given by

$$\mathbf{W}_{i} = \left(\mathbf{G}_{i}^{\top}\mathbf{C}(\mathbf{n})^{-1}\mathbf{G}_{i}\right)^{-1}\mathbf{G}_{i}^{\top}\mathbf{C}(\mathbf{n})^{-1}.$$
(2.95)

To determine the best fitting dipole, we may solve this problem on a dense grid of candidate positions by computing the generalized GOF,

$$GOF_{i} = 1 - \frac{\left(\mathbf{y} - \mathbf{G}_{i}\mathbf{W}_{i}\mathbf{y}\right)^{\top}\mathbf{C}(\mathbf{n})^{-1}\left(\mathbf{y} - \mathbf{G}_{i}\mathbf{W}_{i}\mathbf{y}\right)}{\mathbf{y}^{\top}\mathbf{C}(\mathbf{n})^{-1}\mathbf{y}}$$
(2.96)

$$=\frac{\mathbf{y}^{\top}\mathbf{W}_{i}^{\top}\mathbf{G}_{i}^{\top}\mathbf{C}(\mathbf{n})^{-1}\mathbf{G}_{i}\mathbf{W}_{i}\mathbf{y}}{\mathbf{y}^{\top}\mathbf{C}(\mathbf{n})^{-1}\mathbf{y}}$$
(2.97)

where $\mathbf{G}_i \mathbf{W}_i$ is the data resolution matrix of the model. The source which maximizes this quantity corresponds to the dipole which maximizes the likelihood of the data when the covariance of the noise is $\mathbf{C}(\mathbf{n})$ (Lütkenhöner, 1998). Put another way, it is the fraction of the variance in the whitened data which is explained by the model. Alternatively, an initial scan can be performed on a coarse grid and then, starting from the best position, a nonlinear optimization is performed over the position parameter as well.

The main issue is that we do not know the true number of dipoles responsible for generating the data. Early components (e.g., visual or somatosensory) and (inter-)ictal spikes in epilepsy are often assumed to contain only one source, however, even in these cases other activity may be present as well (Michel et al., 2004) thus questioning the validity of the fundamental assumption behind fitting a single dipole (although it is also possible to do symmetric dipole fits for example).

Beamforming

Beamformers are adaptive spatial filters which attempt to solve the problem of letting only signal from a certain location pass while blocking that of other locations as best as possible. As such, the idea is similar to a filter in the frequency domain. No explicit assumptions are made about the number of active sources, however, the time courses of spatially distinct sources are assumed to be orthogonal, i.e., uncorrelated. The ideal filter with unit gain constraint is

$$\mathbf{W}_{i}^{\top}\mathbf{G}_{j} = \begin{cases} \mathbf{I} & i = j \\ \mathbf{0} & \text{otherwise} \end{cases}$$
(2.98)

meaning that \mathbf{W}_i recovers the activity from the *i*th source and nothing else. Assuming that the leadfields of the sources are linearly independent, it is possible to perfectly suppress m/3 - 1 sources as there are *m* DOFs in a filter and each constraint uses 3 DOFs. Rather than perfectly suppressing a few select sources (a nulling beamformer), a common strategy is to minimize the overall variance of the filter output while ensuring unit gain for the source of interest²⁶. This is the linearly constrained minimum variance (LCMV) beamformer (Veen et al., 1997)

$$\begin{array}{ll} \underset{\mathbf{W}_{i}}{\operatorname{arg\,min}} & \operatorname{trace}\left(\mathbf{W}_{i}^{\top}\mathbf{C}(\mathbf{y})\mathbf{W}_{i}\right) \\ \underset{\mathbf{W}_{i}}{\operatorname{subject to}} & \mathbf{W}_{i}^{\top}\mathbf{G}_{i} = \mathbf{I} \end{array}$$
(2.99)

and the filter weights at position i are given by

$$\mathbf{W}_{i} = \left(\mathbf{G}_{i}^{\top}\mathbf{C}(\mathbf{y})^{-1}\mathbf{G}_{i}\right)^{-1}\mathbf{G}_{i}^{\top}\mathbf{C}(\mathbf{y})^{-1}$$
(2.100)

which is the generalized least squares solution where the *data* covariance represent the covariance of the residuals. It may seem counterintuitive that we are trying to suppress *all* source structure in the data, however, we rely on the unit gain constraint in equation (2.98) to ensure that the signal at position i is preserved. One caveat of this approach is that beamformers may fail to recover sources if the SNR is very high and the forward model is not completely correct (Dalal et al., 2014). This is because the gain constraint in equation (2.99) relates to the gain vectors and not the true source topography (which are unknown).

 $^{^{26}}$ Additionally, the more similar the leadfields of two sources are, the more difficult it will be to suppress one while achieving unit gain for the other which will lead to high filter weights, something may not be desirable (Veen et al., 1997).

2.3. THE EEG INVERSE PROBLEM

For localization purposes, we may use the estimated variance (or power) at each position which can be computed as

$$\sigma_i^2 = \text{trace} \left(\mathbf{W}_i^{\top} \mathbf{C}(\mathbf{y}) \mathbf{W}_i \right) = \text{trace} \left[\left(\mathbf{G}_i^{\top} \mathbf{C}(\mathbf{y})^{-1} \mathbf{G}_i \right)^{-1} \right].$$
(2.101)

This is known as the power estimate of the first kind (Sekihara, 2008). However, the spatial spectrum of the noise tends to peak at source locations far from the electrodes where the SNR is low. This means that noise is amplified for sources to which the sensors are not very sensitive (e.g., deep sources) as we are inverting the inner product of gain vectors which have very small values. Again, we can be alleviate this problem to some degree by normalizing the spatial spectrum of the data by that of the noise. This is termed the neural activity index (NAI) and is given by

$$\mathrm{NAI}_{i} = \frac{\mathrm{trace}\left(\mathbf{W}_{i}^{\top}\mathbf{C}(\mathbf{y})\mathbf{W}_{i}\right)}{\mathrm{trace}\left(\mathbf{W}_{i}^{\top}\mathbf{C}(\mathbf{n})\mathbf{W}_{i}\right)} = \frac{\mathrm{trace}\left[\left(\mathbf{G}_{i}^{\top}\mathbf{C}(\mathbf{y})^{-1}\mathbf{G}_{i}\right)^{-1}\right]}{\mathrm{trace}\left[\left(\mathbf{G}_{i}^{\top}\mathbf{C}(\mathbf{n})^{-1}\mathbf{G}_{i}\right)^{-1}\right]}.$$
 (2.102)

This is often called output SNR although, strictly speaking, the numerator also contains contributions from the noise (Sekihara, 2008). In particular, the numerator estimates the variance of the data (source plus noise) and the denominator estimates the variance of the noise (Veen et al., 1997).

We can use the last expression in equation (2.101) to understand why beamformers tend to cancel out (highly) correlated sources. As shown by Brookes et al. (2007) using the SVD, if the signal from a particular source only occurs in combination with that of another source, then the projection of its gain matrix columns on the directions of the data covariance matrix will never be high. On the other hand, if a source with a field pattern similar to some weighted combination of these two sources is present (typically somewhere in between) then its apparent variance will be higher (Veen et al., 1997).

Subspace Projection

The multiple signal classification (MUSIC) algorithm (Mosher et al., 1999) is a subspace projection method which searches through the source space to find the location where the correlation between the gain vectors of a particular source and the signal space is maximized. If the source is represented by a single vector (e.g., when \mathbf{G}_i is projected on the direction normal to the cortical sheet), then we simply find the source with maximal projection onto the signal subspace. However, if the source is represented by a set of gain vectors, \mathbf{G}_i , then we also need to determine the direction which maximizes this projection. Here, we describe the latter case.

Let the SVD of the gain vectors associated with source i be $\mathbf{G}_i = \mathbf{U}\mathbf{S}\mathbf{V}^{\top}$. The (subspace) correlation matrix between \mathbf{G}_i and $\Phi(s)$ is given by

$$\operatorname{corr}(\mathbf{G}_i, \Phi(s)) = \mathbf{U}^{\top} \Phi(s) \tag{2.103}$$

where $\Phi(s)$ is the $m \times r$ subspace projector obtained by taking the first r eigenvectors of the signal covariance matrix, i.e., the (noise) whitehed data covariance matrix.

Furthermore, the SVD of this correlation matrix is $\operatorname{corr}(\mathbf{G}_i, \Phi(s)) = \mathbf{U}_c \mathbf{S}_c \mathbf{V}_c^{\top}$. The first column of \mathbf{U}_c , denoted $\mathbf{u}_{c,1}$, is the direction which maximizes the correlation and the correlation coefficient is the associated singular value, i.e., $\mathbf{S}_{c,11}$. This is equivalent to the cosine of the smallest principal angle between these subspaces. The direction²⁷ in physical space which maximizes the subspace correlation is given by

$$\phi = \frac{\mathbf{V}\mathbf{S}^{-1}\mathbf{u}_{c,1}}{\|\mathbf{V}\mathbf{S}^{-1}\mathbf{u}_{c,1}\|} \tag{2.104}$$

and the associated gain vector is

$$\mathbf{g}(\phi) = \mathbf{G}_i \phi. \tag{2.105}$$

To find the first source, i_1 , we find the location which maximizes the subspace correlation, i.e.,

$$i_1 = \arg\max \operatorname{corr}(\mathbf{G}_i, \Phi(s))_1$$
 (2.106)

where $\operatorname{corr}(\cdot, \cdot)_1$ denotes the largest singular value of the subspace correlation matrix.

To find any subsequent sources, MUSIC may be applied recursively, a process known as recursively applied MUSIC (RAP-MUSIC) (Mosher et al., 1999). Having found the first source, we project out its estimated topography from the subspace projector as well as the gain matrix. Then we repeat the above procedure using the modified signal subspace and gain matrix. Thus, we form the quantities for the kth iteration as follows

$$\mathbf{G}_{k-1} = \begin{pmatrix} \mathbf{g}(\phi)_1 & \mathbf{g}(\phi)_2 & \dots & \mathbf{g}(\phi)_{k-1} \end{pmatrix}$$
(2.107)

$$\mathbf{P}_{k-1} = \mathbf{I} - \mathbf{G}_{k-1} \left(\mathbf{G}_{k-1}^{\top} \mathbf{G}_{k-1} \right)^{-1} \mathbf{G}_{k-1}^{\top}$$
(2.108)

$$\Phi(s)_k = \mathbf{P}_{k-1}\Phi(s) \tag{2.109}$$

$$\mathbf{G}_k = \mathbf{P}_{k-1} \mathbf{G}.\tag{2.110}$$

where \mathbf{P}_{k-1} projects onto the null space of \mathbf{G}_{k-1}^{\top} .

2.4 Magnetic Resonanace Imaging

In order to generate personalized volume conductor models, one usually relies on one or more MRI scans of a subject. These images are segmented into several different tissue types and, depending on the forward modeling approach, a volume or surface representation of the anatomy is constructed. This section describes the fundamental principles of MRI. It is based on Haacke et al. (1999) and Smith and Webb (2011).

2.4.1 Signal Origins

MRI is a medical imaging technique which relies on the magnetic properties of the tissue being studied, specifically, the interaction of nuclear spin with an externally applied magnetic field, B_0 , where the spin will align with and precess

 $^{^{27}}$ There is a sign ambiguity which can be resolved, for example, by selecting the vector whose dot product with the normal is positive.

2.4. MAGNETIC RESONANACE IMAGING

around the axis of B_0 . Precession refers to the circular movement of a rotating body about another fixed axis and the frequency with which this happens is given by

$$\omega_0 = \gamma B_0 \tag{2.111}$$

where ω_0 is the precession frequency, known as the Larmor frequency, and γ is the gyromagnetic ratio which depends on the nuclei.

In most applications of MRI, the primary nucleus contributing to the observed signal is the proton in the hydrogen atom. The abundance of hydrogen in all water and fat based tissues allows high SNR to be obtained compared to other atoms. The proton is a charged particle which rotates about its internal axis. As such, it can be thought of as a small magnet possessing a magnetic moment. Normally, these magnetic dipoles will be randomly oriented, however, upon application of an external magnetic field they will align with this. This alignment can be either parallel (low energy) or anti-parallel (high energy) to the B_0 field but because the quantum energy difference associated with these two states is much smaller than the thermal energy at room temperature, the spin excess in the parallel direction is only about one in a million. Nevertheless, due to the shear number of protons there will still be a longitudinal equilibrium magnetization, i.e., a net magnetization, M_0 , in the B_0 direction the manipulation of resulting in measurable signals. The net magnetization is given by

$$M_0 = \frac{\rho_0 \gamma^2 \hbar^2}{4kT} B_0 \tag{2.112}$$

where ρ_0 is the spin density, \hbar is Planck's constant divided by 2π , kT is the thermal energy with k being the Boltzmann constant, and T the temperature. In the following, we will assume that M_0 is aligned with the z axis of the physical coordinate system.

The receiver coils used to detect the signal are based on the principle of Faraday induction, i.e., a voltage is induced in the coils due to a *time-varying* magnetic flux. To create such a magnetic flux, the net magnetization is disturbed by a radio frequency (RF) magnetic field, B_1 , for a brief period of time the purpose of which is to rotate M_0 from the z plane to the xy plane where it will precess with the Larmor frequency thus generating a detectable magnetic flux. The angle with which the magnetization is rotated is called the flip angle, α , and it depends on the strength and duration of the applied RF pulse

$$\alpha = \gamma B_1 t. \tag{2.113}$$

2.4.2 Relaxation

Once the RF pulse has been switched off, the signal will start to diminish as the system returns to equilibrium. There are several reasons for this.

First, having been rotated (partially) into the transverse (i.e., the xy) plane, the spins will interact with their immediate surroundings causing a regrowth of the magnetic z component over time as given by the Bloch equations for an arbitrary flip angle

$$M_z(t) = M_0 \cos \alpha + (M_0 - M_0 \cos (\alpha)) \left(1 - e^{-t/T_1}\right).$$
(2.114)

The tissue-specific constant T_1 is the longitudinal relaxation time and describes the speed with which the energy obtained by the spins due to the RF pulse is dissipated to its molecular environment causing increasing vibration and rotation in its vicinity (the "lattice"). Hence, this is called spin-lattice relaxation.

Second, the transverse magnetization itself also decays. This is because in practice the frequency at which the individual protons precess vary slightly. Although aligned initially (e.g., in the y direction), dephasing will occur over time bringing M_x and M_y back to their equilibrium of zero according to (here for M_y)

$$M_{y}(t) = M_{0} \sin\left(\alpha\right) e^{-t/T_{2}}.$$
(2.115)

The transverse relaxation time, T_2 , is also a tissue-specific constant. Since this happens because of short-lived, random interactions between neighboring spins it is often termed spin-spin relaxation. Additionally, dephasing also occurs because of external inhomogeneities in the local magnetic field experienced by different spins (e.g., because of different molecular environments), however, this effect, being more stable across time, is reversible. Together these effects results in a faster relaxation time known as T_2^{\star} which is always smaller than T_2 . Longitudinal relaxation always implies transverse relaxation, however, the latter can also occur independently.

Thus, the observed contrast is determined (primarily) by how these three tissue properties are weighted: the proton density, the longitudinal relaxation time (T_1) , and the transverse relaxation time (T_2) .

2.4.3 Contrasts

The success MRI rests not only on its noninvasive nature and high level of safety but also its flexibility in terms of generating images with different contrasts. Since MRI is sensitive to a host of different properties such as proton density, relaxation times, proton diffusion, chemical shift in the resonating frequency, and field inhomogeneities it can be used to study a diverse array of phenomena. On the other hand, this great flexibility also presents a potential problem in that a given contrast is rarely "pure", hence effects due to other properties than the one(s) of interest may be viewed as noise.

Different contrasts can be achieved by the relative weighting of the different properties described above. For example, if the signal is dominated by differences in longitudinal relaxation it is said to be T_1 -weighted whereas if it is dominated by dephasing in the transverse plane it is said to be T_2 -weighted. For example, the signal intensity for a spin echo sequence using a flip angle of 90° is given by (assuming refocusing has been applied)

$$s (\mathrm{TE}) \propto \rho \left(1 - e^{-\mathrm{TR}/T_1}\right) e^{-\mathrm{TE}/T_2}$$
 (2.116)

where ρ is the spin density, repetition time (TR) is the time between RF pulses, echo time (TE) is the time from excitation to signal readout, and the corresponding contrast between tissues A and B is

$$c_{AB} = s_A (TE) - s_B (TE).$$
 (2.117)

To obtain a contrast dominated by a specific parameter we need to maximize c_{AB} by manipulating TE and TR²⁸.

42

 $^{^{28}}$ Here we ignore the flip angle for simplicity.

From equations (2.116) and (2.117) it apparent that T_1 -weighting can be achieved by having a short TR (relative to T_1) whereas if TR is long, T_1 relaxation will have little effect on the signal, since the M_z will be fully recovered irrespective of tissue type. Likewise, it is apparent that a TE approximately equal to T_2 will result in T_2 -weighting whereas a short TE will not (a long TE will also give T_2 -weighting but the signal will be compromised as dephasing is substantial at this point). In a T_1 -weighted image we emphasize the former using a short TR and a short TE whereas the opposite parameters would give an image with more T_2 -weighting. If we use a long TR and a short TE both of these effects are minimized and the resulting contrast is dominated by the spin density of the tissue.

Without refocusing and with a sufficiently long TR, the signal will be dominated by a T_2^* -weighting which is sensitive to local field inhomogeneities (susceptibility differences) and can be useful in various circumstances. For example, the blood oxygen level dependent (BOLD) contrast often used in fMRI relies on local susceptibility changes due to varying amounts of oxygenated and deoxygenated hemoglobin in blood. Since oxygenated and deoxygenated hemoglobin is diamagnetic and paramagnetic, respectively, a change in their ratio manifests as a change in the local magnetic field. In particular, as neurons need oxygen to sustain their activity, increased neural activity is accompanied by an increased blood flow to the active region. This is known as neurovascular coupling. The result is an effective increase in oxygenated hemoglobin and, being diamagnetic, this molecule interferes less with the local magnetic field than deoxygenated hemoglobin resulting in an increase in signal. FMRI data is often collected with fast sequences (e.g., echo-planar imaging) which allow whole-brain coverage in less than two seconds.

The signal as described so far does not contain any information about where the different components of the signal originates from. In order to form images, we need to be able to spatially encode the signal. This can be achieved by application of gradients to selectively excite a single slice of tissue and the signal within a slice can be encoded by means of frequency and phase (see appendix D.4)

2.4.4 Artifacts

Given its high level of flexibility and the sophisticated nature of the technique, it is perhaps not surprising that MRI is subject to a multitude of different artifacts, some of which are due to the tissue itself (e.g., chemical shifts), the subject being scanned (e.g., motion), or the technique (e.g., phase wrap-around). Here, only a select few will be described.

One potential issue in MRI is partial volume effects. If a voxel contains more than one type of tissue, its signal will be a mixture of those produced by either tissue. This is particularly problematic when trying to distinguish small or thin structures (e.g., thin layers of CSF) and can be alleviated by increasing the spatial resolution.

Images may also be corrupted by so-called chemical shift artifacts which occur because protons situated in different molecular environments experience slightly different magnetic fields. For example, the protons in fat molecules are shielded by a dense cloud of electrons whereas those in water are not since the highly electrophile oxygen atom pulls the electrons away thus exposing the protons in the hydrogen atoms. As a result, the effective magnetic field experienced by protons in fat is slightly *lower* than that experienced by protons in water and consequently they will precess at a slightly lower frequency (on the order of a few parts per million). Since frequency encoding relies on a known relationship between frequency and location, the signal from protons in fat will appear to have come from a water proton at a location corresponding to a lower field thus shifting it in the frequency encoding direction.

Finally, slow-varying intensity gradients across image, the so-called bias field, can occur due to inhomogeneity in the spatial excitation profile of the transmit coil as well as differential sensitivity of the receive coils. For example, the tissue in the center of the bore may be excited the most creating slightly higher signal from this area. The receive coils, on the other hand, will typically be most sensitive to signal in their immediate vicinity. This is mostly an issue when the field-of-view is large.

2.5 Segmentation of Structural MRI Scans

Segmentation is the process of associating a latent representation with an observed data set. Within the field of medical imaging, this typically means assigning labels to pixels or voxels. In a supervised setting, parameters are learned from a training set and applied to a new image to obtain its labeling. One attractive property of this approach is that although training may be time consuming, once this is complete, inference (i.e., generating a specific labeling) is very fast. One potential drawback of this method is that it may be sensitive to the particular image modality or specific parameters with which the data was acquired (e.g., MRI sequence) although with enough data this may not be a problem (Billot et al., 2021). To avoid this, one may choose to learn the latent representation directly from the image at hand in an unsupervised manner as this is able to adapt to the individual intensity profiles in the image (Puonti et al., 2016). Consequently, this is slower as the parameters need to be estimated every time inference is desired. A popular strategy is to assume that the intensity profile of a particular tissue class (e.g., gray matter) can be modeled using one or more normal distributions thus allowing us to model the whole image using a Gaussian mixture model (GMM). This model is generative in that, given the label (or a probability distribution over labels) of a certain voxel, we can draw values (i.e., image intensities) thus synthesizing an image of a certain modality. However, if we apply a GMM only to the image intensities, we are discarding a lot of information, namely the spatial distribution of such intensities. In particular, if we believe that a certain voxel is gray matter then we also know that neighboring voxels are likely to be gray matter as well or white matter or CSF and not, for example, skin. We can incorporate such information by using a spatially varying prior over tissue labels rather than simply assuming a uniform prior for all voxels. Such a prior over tissue probabilities is often called a probabilistic atlas and is generated by averaging over a set of (manually) labeled images from different subjects. As such, it is expressed in a standard space and needs to be registered to the image we wish to segment. Usually, an initial affine registration is performed to roughly align the atlas and the image and subsequently this is refined by a nonlinear deformation.

MRI is a particularly important medical imaging modality as it has excellent

(and tunable) soft tissue contrast, however, such scans tend to be corrupted by a spatially varying artifact (bias) which modulates intensities such that a particular tissue type may appear brighter in one part of the image than in another. This intensity nonuniformity is known as the bias field and we want our labeling to be insensitive to such variations. Since we know that this effect tends to be spatially smooth, we can model it explicitly using a small set of low frequency basis functions. On the other hand, intensity variations due to different tissue properties tend to contain high frequency information as well, e.g., at tissue borders (Ashburner & Friston, 2005).

In order for the spatial prior to meaningfully constrain the tissue classification, it needs to be registered to the image of interest. This is typically achieved by finding a set of deformation parameters such that the prior matches a particular labeling. Likewise, optimizing the parameters of the basis functions used to model the bias field, requires an initial labeling of the voxel but the labeling also depends on the current bias field correction. Thus, the procedure is inherently circular as noted by Ashburner and Friston (2005) who proposed to use an iterated conditional modes approach where each parameter set is optimized conditioned on the current values of the rest as opposed to optimizing each parameter sequentially.

Although all voxels are typically modeled as independent of each other in the likelihood function (which is clearly not plausible), spatial dependency between voxels is achieved in the conditional probabilities due to the smooth nature of the probabilistic atlas. Likewise, regularization may also be imposed on the bias field coefficients and the spatial deformations in the form of priors.

2.5.1 Unified Segmentation Using Gaussian Mixture Modeling

This section describes the problem of segmenting a structural scan from a mathematical point of view. It is based on Puonti et al. (2016) and the reader is referred to this reference for details.

We wish to model the probability of the labels, $\mathbf{y} \in \{1, ..., k\}^n$, given one or more images, $\mathbf{x} \in \mathbb{R}^{n \times m}$, where *m* is the number of MRI scans provided, *n* is the number of voxels in an image, and *k* is the number of labels in the segmentation. To do this, we need a forward model which relates labels to images. Here, we use a GMM. Below, we denote the complete set of parameters by θ . The parameters related to the GMM (which in this case includes the bias correction coefficients) and the deformation of the probabilistic atlas are denoted α and β , respectively. We have

$$p(\mathbf{y}|\mathbf{x}) = \frac{p(\mathbf{x}|\mathbf{y}) p(\mathbf{y})}{p(\mathbf{x})}$$
(2.118a)

$$\propto p\left(\mathbf{x}|\mathbf{y}\right)p\left(\mathbf{y}\right) \tag{2.118b}$$

$$\propto \int_{\theta} p(\mathbf{x}|\mathbf{y},\theta) \, p(\mathbf{y}|\theta) \, p(\theta) \, d\theta \tag{2.118c}$$

$$\propto \int_{\alpha} p(\mathbf{x}|\mathbf{y},\alpha) p(\alpha) \, d\alpha \int_{\beta} p(\mathbf{y}|\beta) \, p(\beta) \, d\beta$$
 (2.118d)

where we have made the dependence on the model parameters explicit. The data likelihood, $p(\mathbf{x}|\mathbf{y}, \alpha)$, is modeled by the GMM and the prior probability of

a certain label, $p(\mathbf{y}|\beta) p(\beta)$, is given by the values in the (deformed) probabilistic atlas. To stabilize the registration between atlas and image(s), the prior on β penalizes large deformations. The prior on the likelihood parameters is modeled with a uniform distribution and as such can be ignored as it is constant. Furthermore, we assume that the following conditional distributions factorize over voxels, *i*,

$$p(\mathbf{x}|\mathbf{y},\alpha) = \prod_{i} p(\mathbf{x}_{i}|\mathbf{y}_{i},\alpha)$$
(2.119a)

$$p(\mathbf{y}|\beta) = \prod_{i} p(\mathbf{y}_{i}|\beta)$$
(2.119b)

meaning that intensities are independent given label and model parameters, and similarly, that labels are independent given the deformation parameters. As noted above, these assumptions are clearly not correct, however, due to the circular relationship between labeling and prior deformation (and because spatial dependency is encoded in the probabilistic prior) we expect the final segmentation to be relatively smooth nonetheless (Ashburner & Friston, 2005).

One problem with equation (2.118d) is that it involves integrating over all possible models (sets of parameters) making it intractable to compute in practice. However, if we assume that the posterior of the model parameters is sufficiently peaked around its mode, then we can make the following approximation

$$p(\mathbf{y}|\mathbf{x}) \approx p\left(\mathbf{y}|\mathbf{x}, \hat{\alpha}, \hat{\beta}\right)$$
 (2.120)

where $\hat{\cdot}$ denotes the maximum a posteriori (MAP) estimate of the parameters. Thus, we have transformed the problem from one of integration to one of optimization, i.e.,

$$\underset{\alpha,\beta}{\arg\max} p\left(\alpha,\beta|\mathbf{x}\right) \tag{2.121}$$

and the posterior of the model parameters is given by

$$p(\alpha, \beta | \mathbf{x}) \propto p(\mathbf{x} | \alpha, \beta) p(\alpha) p(\beta)$$
 (2.122a)

$$\propto \sum_{y} p(\mathbf{x}|\mathbf{y},\alpha) p(\mathbf{y}|\beta) p(\alpha) p(\beta)$$
(2.122b)

$$\propto \prod_{i} \sum_{y} p(\mathbf{x}_{i} | \mathbf{y}_{i}, \alpha) p(\mathbf{y}_{i} | \beta) p(\alpha) p(\beta).$$
(2.122c)

Our GMM includes the labels \mathbf{y} , however, as these are unobserved we need to marginalize over them. In equation (2.122c), we have used the conditional independence assumption from equation (2.119) and moved the sum (over labels) inside the product (over voxels). This way, the calculation of the posterior becomes tractable (as opposed to calculating it over all possible segmentations). The MAP estimate of the parameters is obtained by iteratively updating each parameter set while keeping the other constant. For example, the GMM parameters are optimized using an expectation-maximization scheme which maximizes a lower bound on the objective function (typically the log-likelihood of the data) whereas the deformation parameters may be optimized using a generic optimizer (Ashburner & Friston, 2005; Puonti et al., 2016).

Chapter 3

Research Contributions

This chapter presents the research contributions of this thesis.

3.1 Automatic Skull Segmentation From MR Images for Realistic Volume Conductor Models of the Head: Assessment of the State-of-the-Art

The purpose of this study, presented in appendix A, is to introduce a new pipeline for generating head models for use in noninvasive brain stimulation and EEG. It uses SPM12 to segment one or more MRI scans from which surfaces are extracted and a volume mesh is created. We compare its accuracy with existing methods with particular focus on skull reconstruction. The skull plays an important role in shaping the potential fields observed on the scalp in EEG, however, identifying it on MRI scans is difficult because of very low signal emission. This is especially true when the model is based on only a T1-weighted image since CSF also emits very little signal in these images thus making accurate segmentation of these two compartments very challenging. For comparison, we use skull segmentations based on CT images.

We compare our approach, termed headreco, with two other tools, namely betsurf from FSL and skullfinder from BrainSuite. We find that headreco and betsurf generally perform better than skullfinder and that basing the segmentation on a T1- as well as a T2-weighted image helps increase performance—both in terms of average accuracy and variability—as T2-weighted images show excellent contrast between CSF and bone thus facilitating identification of the inner skull border.

We use an extended tissue prior, covering also the neck and facial area, and show how this affects field calculations in TES when an extraencephalic electrode configuration is used. We find clear stimulation of the cerebellum using an extended model as opposed to more temporal stimulation in a model restricted to the upper part of the head. On the other hand, extending the model has very little effect when both electrode pads are placed directly above cortical areas. To model the pial surface of gray matter, we expand the central surface estimate from computational anatomy toolbox 12 (CAT12) using an estimate of the cortical thickness. We show how this improves the accuracy of the gray matter compartment in our model considerably compared to the original estimate from SPM12. Again, we also show how this translates to effects on the estimated electric field in a TES stimulation setup. The SPM12-based model generally gives higher field estimates on the cortex, probably because sulci are not resolved properly (i.e., filled with gray matter) thus reducing current shunting into the CSF compartment.

In order to facilitate high quality segmentations of the human head and neck area, it is important to consider the quality of the input, i.e., the MRI scans. A T1-weighted scan has great brain tissue contrast whereas a T2-weighted scan provides good contrast between CSF and bone. Including the latter generally improves performance although T2-weighted images may suffer more from artifacts related to intensity inhomogeneity (bias) in the neck region. However, in order to achieve good tissue separation (e.g., between compact bone, spongy bone, and skin), it is also important to consider the parameters of the MRI sequences being used. For example, fat suppression or selective water excitation can be used to suppress the signal from spongy bone and avoid fat shift artifacts when the bandwidth is low.

3.2 Evaluating the Influence of Anatomical Accuracy and Electrode Positions on EEG Forward Solutions

The focus of the study in section 3.1 was on building an accurate physical representation of the anatomy of the head. In this study, presented in appendix B, we evaluate the impact of two different aspects of forward modeling on the accuracy of the final forward solution; in particular, the anatomical accuracy of the volume conductor model and the accuracy with which electrode positions are known.

We start by evaluating the numerical accuracy of using SimNIBS to solve the forward problem in EEG and show that it is comparable to existing methods.

To study the impact of anatomical accuracy, we compare different pipelines for solving the forward problem in EEG. We use FieldTrip with SPM12 and DUNEuro, MNE-Python with FreeSurfer, and SimNIBS with CHARM. We also include a model based on a template anatomy, specifically, head model constructed from the MNI152 template (see appendix B for details). This was done to model a scenario where electrode positions are known but the anatomy is not. We use manual segmentations as reference. The anatomical accuracy of CHARM was generally superior to the other pipelines and that of FieldTrip-SPM slightly better than MNE-FreeSurfer.

We found effects of anatomy on both topographic and magnitude errors in the forward solutions as measured by relative difference measure (RDM) and logarithm of the magnitude error (lnMAG), respectively. Topographic errors were particularly high for FieldTrip-SPM and MNE-FreeSurfer in deep and orbitofrontal areas whereas they were much smaller for SimNIBS-CHARM. The same was true for magnitude errors, particularly in MNE-FreeSurfer, which is a three-layer BEM model neglecting the distinction between white matter, gray matter, and CSF. The model based on a standard anatomy showed errors in between those of FieldTrip-SPM and MNE-FreeSurfer.

To study the impact of electrode accuracy, we compare models based on digitized electrode positions and positions obtained by either transforming a layout as specified by the manufacturer (i.e., spherical angles) or transforming a custom template which we created by measuring electrode positions on a 3D printed model of the MNI template head. Again, we also included a model using the digitized positions but with standard anatomy. We used the digitized positions as reference.

Our custom template generally provided a better fit than the manufacturer layout. This was particularly true in the occipital area where the mean errors of the latter reached 30 mm whereas those of the former were approximately 10 mm. We found that errors were often due to a shift along the anteriorposterior (AP) direction. Effects of electrode positions on the forward solution were seen on topographic errors only and correlated with the observations on channel level, i.e., the errors associated with using a manufacturer layout were generally increased in occipital and parietal areas. We found that using a template anatomy with digitized positions resulted in large errors on the magnitudes whereas topographic errors were similar to those using the custom template.

In this study, we only investigated errors in the forward solution, however, in EEG we mostly care about how this affects the solution of the corresponding inverse problem. There does seem to be a body of evidence suggesting that forward modeling errors do in fact impact the accuracy with which sources can be localized although the effect is likely modulated by the SNR level of the data and the choice of inverse method.

Based on our current findings we suggest to model anatomy and electrode positions as accurately as possible. As to the former, SimNIBS makes it very easy to generate realistic volume conductor models at a computational cost similar to the other pipelines. As to the latter, we suggest digitizing electrodes and, if this is not possible, use a template created by measuring electrode locations on a realistically shaped template head.

3.3 Optimizing Template-Based Channel Layouts in Electroencephalography

In section 3.2 we transformed a custom template of electrode positions measured in standard space (on a 3D printed model of the MNI head) to individual subjects. The fit was better than using the spherical angles provided by the manufacturer, however, it was still not perfect. One observation we made was that the quality of the fit varied substantially from subject to subject. In section 3.2 we also found that the accuracy with which electrode positions were known were important for determining the accuracy of the forward model. Therefore, in this work, presented in appendix C, we explore whether it is possible to optimize the electrode locations after the template has been transformed to subject space. For this to be possible, we need measurements describing the position of the cap. Here, we use four measurements of distances and angles between landmarks and nearby electrodes, i.e., eight in total. We simulate measurements at different levels of noise $(0 \text{ mm to } 10 \text{ mm for} \text{ distances}, 0^{\circ} \text{ to } 20^{\circ} \text{ for angles})$ and compare them with coordinates obtained by digitizing the electrodes.

We formulate a cost function based on inter-electrode distances and a use the measurements as constraints. The optimization procedure generally increased the accuracy of the electrode positions, particularly in the anterior-posterior direction. Performance decreased only slightly with increasing noise levels whereas performance varied substantially between subjects such that some benefitted a lot whereas it did not make much difference for others. As such, applying this procedure will help avoid outliers (i.e., subjects where the initial fit is bad) and provide modest overall improvements. Although we only used eight measurements here, the procedure can be extended to arbitrary combinations of electrode-electrode and electrode-landmarks measurements.

On a more practical note, we believe that distances are straightforward to measure, however, angles are trickier as they need to be measured wrt. some axis. Therefore, measurements should be made relative to axes defined by anatomical landmarks as in for example the Neuromag coordinate system. However, as implied, angles need not be known exactly to provide useful information. We hope to be able to corroborate our results using actual measurements in the future.

3.4 Evaluating the Influence of Anatomical Accuracy and Electrode Positions on EEG Inverse Solutions

Our aim is to investigate the extent to which different forward solutions translate to differences in source localization error. This is done for a few select inverse approaches that differ in their underlying assumptions about the data (e.g., number of sources). This study extends the work from section 3.2. Using the same forward models, we simulate sources at three SNR levels and apply four different inverse methods. We use two MNE method, dSPM and sLORETA, a least squares dipole fit, and the MUSIC subspace projection method. dSPM and sLORETA apply different noise normalizations to the standard minimum norm estimate. Thus, these represent a distributed approach where all sources are estimated simultaneously with minimal assumptions about the source configuration. On the other hand, dipole fitting assumes a single (dominant) source the position of which is found by (possibly nonlinear) scanning through the entire source space. Finally, MUSIC scans the source space to determine the source which aligns the best with an estimate of the signal subspace. As such, these techniques differ in their assumptions about the data as well as other practical aspects such as how the data covariance is determined, although, fundamentally, they are all based on a linear model (Mosher et al., 2003). This allows us to assess the influence of the forward model (modulated by anatomical accuracy and electrode position accuracy) on the inverse solution in EEG.

In the following, we present the methods and results. Discussion of the results is found in chapter 4.

3.4. EVALUATING EEG INVERSE SOLUTIONS



SMOOTH smooth using absolute value of normalized Laplacian.

Figure 3.1. Overview of the data generation and analysis procedure. Colored boxes denote different processing stages, e.g., input data and data generation . p is the number of sources (and hence the number of simulations performed). The result of the simulation using \mathbf{p}_j is stored in the *j*th column of the matrix of simulations, \mathbf{R} .

3.4.1 Methods

To explore the influence of the forward model on source localization in different settings, we use a combination of forward models, inverse methods, and SNR levels. We simulate data from approximately 20 000 sources (10 000 per hemisphere) placed on the central cortical sheet at SNR levels of 2, 4 and 8. For this, we use a reference forward model constituting the most accurate representation of anatomy and electrode layout that we are able to construct. We use four different methods to invert the data. Specifically, we use two MNE estimators, dSPM and sLORETA, a standard least squares dipole fitting approach, and MUSIC. For data inversion, we use different forward models than the one used to generate the data. Specifically, we perform two sub-studies. In the first study, the models vary in terms of how accurately they represent the underlying anatomy whereas in the second study, they vary in terms of how they represent the electrode layout. We simulate the data similarly in both studies although the noise vectors were sampled differently (see below for details).

Below we describe each of these steps as well as how localization accuracy was assessed. The procedure is illustrated in figure 3.1.

Data Generation

Following the convention in MNE-Python (Gramfort, 2013), we define SNR as the signal-to-noise *amplitude* ratio. Specifically, the SNR at a given point in time is defined as the rank-adjusted GFP of the whitened signal (Engemann & Gramfort, 2015; Murray et al., 2008), i.e.,

$$GFP(\mathbf{z}) = \sqrt{\frac{1}{r} \mathbf{z}^{\top} \mathbf{z}} = \sqrt{\frac{1}{r} \mathbf{x}^{\top} \mathbf{W}^{\top} \mathbf{W} \mathbf{x}}$$
(3.1)

where r is the rank of the data, $\mathbf{W} = \mathbf{C}(\mathbf{n})^{-1/2}$ is the whitener (here, we use a PCA whitener), and \mathbf{x} and \mathbf{z} are the original and whitened data vectors, respectively, where the former has an average reference. Thus, GFP, i.e., spatial standard deviation, quantifies the variability over channels at a given point in time. Once whitened, the amplitude of a signal consisting of pure Gaussian noise should fluctuate around -1.96 and 1.96 with an expected GFP of one. This would indicate proper whitening of the data. If the data contains signals of interest, this should result in increased GFP values at time points where the signal is prominent. Assuming that the noise is additive, the GFP of the data directly reflects the SNR (defined as a ratio of amplitudes) of the data at each point in time.

When evaluating the ability of inverse methods to localize sources on simulated data, one is susceptible to the so-called "inverse crime", i.e., using a model to invert the data that is "too similar" to the one used to generate the data, resulting in estimates that are generally too good (Kaipio & Somersalo, 2007). Using the same model to generate and invert the data would also not be a realistic scenario as the true forward model is never known in practice. Below we describe the procedure of simulating the data.

Let *m* be the number of electrodes and *p* be the number of source positions. We use the following model to generate a vector of observations, $\mathbf{y} \in \mathbb{R}^m$,

$$\mathbf{y} = \mathbf{G}\mathbf{p} + \mathbf{n} \tag{3.2}$$

where $\mathbf{p} \in \mathbb{R}^p$ is a vector of source activations and $\mathbf{n} \in \mathbb{R}^m$ is additive noise which we will assume to be normally distributed. The gain matrix, G, defines the "pure" source topographies. To simulate data, we use the reference forward model (see section 3.4.1 for details), however, since the true tissue conductivities are generally not known, we sample conductivity values for this model randomly using the tissue-specific distributions described in Saturnino, Thielscher, et al. (2019). This is done separately for each subject. On the other hand, when generating the forward models used to invert the data, we use the default conductivity values in the relevant pipeline (see section 3.4.1). Initially, $\mathbf{G} \in \mathbb{R}^{m \times 3p}$, i.e., each source position is associated with three gain vectors defining the gain along each of the major axes (x, y, and z). However, to create a single topography for each source, we project its gain vectors on the normal direction of the cortical sheet at the corresponding location thus resulting in $\mathbf{G} \in \mathbb{R}^{m \times p}$. To emulate smooth activations, we create \mathbf{p} from the absolute values of the graph Laplacian (i.e., the degree matrix minus the adjacency matrix) where we normalize each column vector such as to form convex combinations of neighboring sources. Consequently, each column of **Gp** is a weighted average of its own and its neighboring source topographies.

52

3.4. EVALUATING EEG INVERSE SOLUTIONS

For each subject, we sample noise vectors from a normal distribution with zero mean and a covariance matrix estimated from real EEG data. This is done in order to obtain a realistic noise estimate (see below for details). To achieve the desired SNR level, the GFP of the whitened signal topography, **WGp**, is scaled relative to whitened noise, **Wn**. Since GFP(Wn) is expected to fluctuate around one, this is particularly simple

$$\mathbf{y}_{\rm SNR} = {\rm SNR} \times \mathbf{Gp} + \mathbf{n}. \tag{3.3}$$

The data vectors obtained in this way correspond to a topography at a single point in time, e.g., at the peak of an evoked response. This is sufficient for estimating sources using certain inverse methods, e.g., MNE and dipole fitting approaches, which generally apply to a single sample¹. For these types of inverse estimation procedures, we are mostly interested in the SNR at a particular point in time. Other inverse solvers such as MUSIC rely on an estimate of the data covariance, thus requiring knowledge of the temporal evolution of the signal. In order to calculate both types of estimates, we simulated a complete evoked response by modulating the source activation patterns created above as described next.

To create dynamic source activations, we constructed a canonical time course by taking the first component of an xDAWN decomposition (Rivet et al., 2009; Rivet et al., 2011) from the first subject from the open dataset on face recognition by Wakeman and Henson (2015). This signal had a sampling frequency of 200 Hz and a duration of approximately 0.5 s resulting in 99 samples. The canonical time course, denoted c(t), was scaled such that it peaked at a value of one. Thus, we augment equation (3.3) with a temporal aspect to obtain a time-dependent data vector and, consequently, be able to simulate a full evoked response

$$\mathbf{y}(t)_{\text{SNR}} = \text{SNR} \times c(t)\mathbf{G}\mathbf{p} + \mathbf{n}(t). \tag{3.4}$$

For MNE and dipole estimators, we use the source topographies at the time point with peak SNR (corresponding to the sample at which c(t) = 1). For the MUSIC estimator, we use the data covariance calculated over the entire time window. This way we explicitly control SNR of the data to which MNE and dipole fitting were applied. However, this is not the case for MUSIC where we instead rely on a realistic temporal modulation of the source activity to obtain a reasonable SNR over time.

Noise Covariance Estimation The data used in the sub-study on electrode accuracy was collected by Madsen et al. (2019) and Karabanov et al. (2021). In addition to digitized electrode positions, this dataset also contains EEG data from an eyes open session with a duration of approximately 2 minutes for each subject. We use this data to estimate a (realistic) noise covariance matrix. For each subject, we sample noise vectors from a normal distribution with zero mean and the estimated covariance matrix. However, EEG data was only available for the subjects used in the sub-study on electrode position accuracy. Hence, for each subject in the sub-study on anatomical accuracy, we chose a random subject from the electrode position sub-study and sampled from the corresponding normal distribution. The EEG cap was the same in both studies. Below we describe how the covariance matrix was estimated.

¹Although temporal extensions exist (e.g., Scherg, 1990).

The raw EEG data was bandpass filtered between 1 and 100 Hz and a notch filter was applied at 50 and 100 Hz. The data was downsampled from 5000 Hz to 200 Hz and manually inspected for bad channels which were ignored during analysis and interpolated by the end of the preprocessing stage.

To minimize eye-related artifacts the following procedure was employed. MNE-Python was used to generate epochs associated with eyeblinks based on one of the most prefrontal electrodes and horizontal eye movements based on one of the electrodes over the temples. An ICA was performed using the Picard algorithm which has been shown to perform well on real EEG data (Ablin et al., 2018). Components were compared the with eye-related events using Pearson correlation and components of interest were identified using iterated z-scoring with a rejection threshold of three. These were removed from the data which was finally rereferenced to an average reference.

To estimate a noise covariance matrix we use the following strategy. The preprocessed data was chopped into epochs of 0.2 s and the first five epochs were discarded. We then selected 100 epochs randomly and computed the covariance matrix using a shrunk estimator (Engemann & Gramfort, 2015). This was used to simulate the noise vectors in equation (3.2). In order to avoid using the same noise covariance matrix to sample the noise and invert the data, we generated a second covariance matrix by repeating this procedure (ensuring that no epochs were selected twice) and used this as the noise covariance matrix during data inversion.

Simulation Studies

This study is a continuation of the work from section 3.2 where we investigated how differences in anatomical accuracy and electrode position accuracy affected the forward solution. Consequently, the tested forward models are the same. We will describe the models briefly here but refer to the original reference for more details.

We perform two sub-studies. One in which the anatomical model is varied and another where the position of the electrodes are varied. All forward solutions used to invert the data were corrupted by noise on the normal directions by sampling, for each source, a new vector uniformly within a cone with a width of 20 degrees. In this study, this is only relevant for the MNE estimators (see below).

Anatomy Study In the first sub-study, we construct forward models using pipelines from different software packages.

- The reference forward model (i.e., the one used to generate the data) is calculated in SimNIBS (Thielscher et al., 2015) based on manual segmentations (Farcito et al., 2019). As noted above, we sample new tissue conductivities for this model to acknowledge the fact that we never know the true conductivities in practice.
- SimNIBS-CHARM is generated using the head modeling pipeline CHARM (Puonti et al., 2020) available in SimNIBS. SimNIBS uses a first order FEM with SPR (Saturnino, Madsen, & Thielscher, 2019; Zienkiewicz & Zhu, 1992) to compute the gain matrix.

3.4. EVALUATING EEG INVERSE SOLUTIONS

- FieldTrip-SPM is created in FieldTrip (Oostenveld et al., 2011) in combination with SPM12 (https://www.fil.ion.ucl.ac.uk/spm/software/spm12/) to generate the tissue segmentation and DUNEuro (Schrader et al., 2021) to compute the gain matrix. This is also a FEM model.
- MNE-FS is a three-layer linear BEM model constructed using MNE-Python which uses FreeSurfer (Dale et al., 1999) to generate the surfaces.
- We include a model based on a template anatomy, MNI-Template, obtained by segmenting the T1-weighted image of the MNI152 template.

All models use default conductivities as defined by the different pipelines. As the source space of the template-based model was in a different space than the reference model, we morphed the template to the reference by spherical registration (Fischl et al., 1999). This was also done for SimNIBS-CHARM. FieldTrip-SPM and MNE-FSused the source space from the reference model, however, sources which were not inside the brain compartment (in FieldTrip-SPM this consisted of white matter and gray matter; in MNE-FS this was the brain compartment of the BEM model) were excluded from analysis. Hence, the simulations were restricted to these valid sources and the results smoothed when morphed to *fsaverage* space, i.e., the surface-based cortical gray matter template included in FreeSurfer.

The electrode layout was based on an EasyCap BC-TMS64-X21 transcranial magnetic stimulation (TMS) compatible cap with a modified M10 (equidistant) layout containing 64 channels (https://www.easycap.de)

Electrode Study In the second sub-study, we construct forward models using different ways of estimating the electrode positions. The cap used in this study is an EasyCap BC-TMS64-X21 TMS compatible cap with a modified M10 (equidistant) layout containing 64 channels. The head model on which the first three of the following models are based was generated using CHARM.

- The reference forward model (i.e., the one used to generate the data) uses the digitized positions available in the data set by Madsen et al. (2019) and Karabanov et al. (2021).
- Custom-Template is a model in which positions were estimated by transforming template positions defined in MNI space which were obtained by digitizing the above-mentioned cap on a 3D printed model of the MNI head. Positions were registered to subject space using nonlinear the transformation provided by CHARM.
- Man-Template uses the positions as defined by the manufacturer (i.e., spherical coordinates) which are registered to subject space by an affine transformation obtained by matching landmarks (nasion, LPA, and RPA).
- Again, we include a model based a template anatomy, MNI-Digitized, which uses the actual (digitized) electrode positions. Similar to the previous sub-study, the template-based model was morphed to the space of the reference such as to align the source spaces.

In this sub-study, all forward calculations were performed in SimNIBS.

Inverse Methods

To estimate source activations, we use two MNE estimators (dSPM and sLORETA), a standard dipole fit, and MUSIC. For MNE-based methods we use a depthweighting parameter of 0.8 (default in MNE-Python)². In each case, we use the gain matrix without orientation constraints (i.e., all three components x, y, and z), however, for the minimum norm estimators, we use an orientation prior which weighs the normal and tangential components differently (specifically, tangential components are weighted by 0.2 of that of the normal direction as is the default behavior in MNE-Python). The regularization parameter of the MNE methods was computed from the SNR level as $\lambda = \text{SNR}^{-2}$ (Lin et al., 2006).

We invert the data associated with each source and collect the results in a matrix, $\mathbf{R} \in \mathbb{R}^{p \times p}$, where each column contains the result of a simulation³. Below, *i* and *j* refer to the index of an estimated and a simulated source, respectively. Thus, \mathbf{R}_{ij} , the element of the *i*th row and *j*th column, is the source estimate at location *i* for the *j*th simulated source. Furthermore, $\mathbf{G} \in \mathbb{R}^{m \times 3p}$ is the gain matrix, $\mathbf{C}(\mathbf{n}) \in \mathbb{R}^{m \times m}$ is the estimated noise covariance matrix, $\mathbf{C}(\mathbf{y}) \in \mathbb{R}^{m \times m}$ is the estimated data covariance matrix, and $\mathbf{R} \in \mathbb{R}^{3p \times 3p}$ is the (assumed) source covariance matrix.

Since the theory behind the different inverse solvers was described in section 2.3.3, we restrict ourselves to describing how localization is performed here.

Minimum Norm Estimates We compute the source amplitudes for each orientation separately, initially resulting in a $3p \times p$ matrix. The three elements associated with the *i*th source in the *j*th simulation is obtained from

$$\mathbf{R}_{i:i+3,j} = \mathbf{W}_i \mathbf{y}_j \tag{3.5}$$

and we take the Euclidean norm of this three-element vector to obtain a single estimate of the source activation at each location

$$\mathbf{R}_{ij} = \left\| \mathbf{R}_{i:i+3,j} \right\|. \tag{3.6}$$

Dipole Fit We use the generalized GOF (Lütkenhöner, 1998) (i.e., the fraction of the variance of the whitened data which is explained by the model) associated with each element for localization purposes

$$\mathbf{R}_{ij} = 1 - \frac{\left(\mathbf{y}_j - \mathbf{G}_i \mathbf{W}_i \mathbf{y}_j\right)^{\top} \mathbf{C}(\mathbf{n})^{-1} \left(\mathbf{y}_j - \mathbf{G}_i \mathbf{W}_i \mathbf{y}_j\right)}{\mathbf{y}_j^{\top} \mathbf{C}(\mathbf{n})^{-1} \mathbf{y}_j}$$
(3.7)

$$= \frac{\mathbf{y}_{j}^{\top} \mathbf{W}_{i}^{\top} \mathbf{G}_{i}^{\top} \mathbf{C}(\mathbf{n})^{-1} \mathbf{G}_{i} \mathbf{W}_{i} \mathbf{y}_{j}}{\mathbf{y}_{j}^{\top} \mathbf{C}(\mathbf{n})^{-1} \mathbf{y}_{j}}.$$
(3.8)

Here $\mathbf{G}_i \mathbf{W}_i$ is the data resolution matrix of the model.

56

 $^{^2\}mathrm{According}$ to Lin et al. (2006), a depth-weighting parameter of 2 to 5 is appropriate for EEG analysis, however, this did not work well in our analyses. Hence, we use the default value in MNE-Python.

³For the MNE estimators, this matrix is the so-called (parameter) resolution matrix which is obtained by multiplying the filter weights with the gain matrix (Stenroos & Hauk, 2013). However, we used different gain matrices for simulating and inverting the data to simulate a more realistic setting.

Subspace Projection For localization, we use the largest singular value of the SVD of the correlation matrix between the gain vectors and the signal projector, i.e.,

$$\operatorname{corr}(\mathbf{G}_i, \Phi(s)_j) = \mathbf{USV}^{\top}$$
(3.9)

$$\mathbf{R}_{ij} = \left\| \operatorname{diag} \mathbf{S} \right\|_{\infty}, \qquad (3.10)$$

where diag extracts the diagonal elements.

Evaluation Metrics

The *j*th column of **R** contains the result of inverting the data associated with the *j*th (simulated) source and \mathbf{R}_{ij} is the response of the *i*th filter to this source. As such, we may view the columns as point spread functions (PSFs) in that the *j*th column represents how activity of the *j*th source "spreads" to other locations. Likewise, we may view the rows as cross talk functions (CTFs) since the *i*th row contains the response of the *i*th filter to all sources thus representing the degree to which activity from each source location "leaks" into the estimate of the *i*th source. Ideally, we want $\mathbf{R} \approx \mathbf{I}$ as this would indicate highly peaked PSFs and CTFs around the true source. On the contrary, if \mathbf{R} contains a lot of mass on the off-diagonal elements, then this would indicate very smooth PSFs making it difficult to localize activity accurately. Likewise, a very smooth CTF for the *i*th source means that the activity of other sources may contribute substantially to the estimate at location *i*—provided of course that these sources are active to some extent (Hauk et al., 2019). Here, we focus on the PSFs.

To evaluate the result of the simulations, we use the peak localization error (PLE) (Molins et al., 2008). PLE measures the distance between the location of the true source, j, and the source with maximal activation, \hat{i} ,

$$\hat{i} = \arg\max_{i} |\mathbf{R}_{ij}| \tag{3.11}$$

$$PLE_j = dist(\hat{i}, j) \tag{3.12}$$

where $|\cdot|$ denotes the absolute value and dist (\hat{i}, j) is the Euclidean distance between source \hat{i} and j.

3.4.2 Results

Below, we present density estimates over subjects and group level source localization errors for each of the sub-studies.

For the sub-study on anatomical accuracy, figures 3.2 and 3.3 show probability density estimates of PLE of PSF over all subjects on a linear and a logarithmic scale, respectively. The latter allows us to investigate the tails of the distributions. Figures 3.4 to 3.7 present group level average PLE of PSF when applying dSPM, sLORETA, dipole fitting, and MUSIC, respectively. Results are displayed on the *fsaverage* template surface. There is a clear effect of inverse method and SNR. As expected, performance increases with increasing SNR although the effect seems smaller for dSPM and MUSIC compared to the other two methods. dSPM shows largest errors for deep, temporal sources (figure 3.4) whereas sLORETA (figure 3.5) and dipole fitting (figure 3.6) show larger errors in orbitofrontal areas and the temporal pole. MUSIC (figure 3.7)



Figure 3.2. Sub-study on anatomical accuracy. Density estimates (over all subjects) of source localization error quantified using PLE of PSF with free orientation. The y-axis was restricted to the range 0 to 1.

shows very good performance, perhaps suggesting some kind of ceiling effect. In terms of forward models, MNI-Template consistently seems worse than the others, particularly for superficial sources and as SNR increases. MUSIC and dipole fitting seem more affected by the type of forward model compared to the estimate based on MNE.

As to the sub-study on electrode positions, figures 3.8 and 3.9 show probability density estimates of PLE of PSF over all subjects on a linear and a logarithmic scale, respectively. Figures 3.10 to 3.13 present group level average PLE of PSF of when applying dSPM, sLORETA, dipole fitting, and MUSIC, respectively. Results are displayed on the *fsaverage* template surface. Again, there is a clear effect of the inverse method and SNR similar to that described above. However, forward model effects seem more pronounced here. In particular, we see increased errors for Man-Template in occipital areas which is also where the largest errors in electrode positions were found in the study presented in section 3.2. This effect seems fairly consistent across SNR levels except for dSPM (figure 3.10) where the effect seems to disappear at high SNR. In general, using the (true) digitized positions and correct anatomy performs better than using our custom template (Custom-Template) or a template anatomy with digitized positions (MNI-Digitized). These, in turn, are better than using the manufacturer template as is done in Man-Template.



Figure 3.3. Sub-study on anatomical accuracy. Density estimates (over all subjects) of source localization error quantified using PLE of PSF with free orientation. This is the same data as in figure 3.2 but with a logarithmic scale on the y-axis emphasizing the tail of the distributions.


Figure 3.4. Sub-study on anatomical accuracy. Group average source localization error quantified using PLE of PSF using dSPM with free orientation (although in this case, an orientation prior was used as described in the main text). Results are presented on the *fsaverage* template.



Figure 3.5. Sub-study on anatomical accuracy. Group average source localization error quantified using PLE of PSF using sLORETA with free orientation (although in this case, an orientation prior was used as described in the main text). Results are presented on the *fsaverage* template.



Figure 3.6. Sub-study on anatomical accuracy. Group average source localization error quantified using PLE of PSF using dipole fitting with free orientation. Results are presented on the *fsaverage* template.



Figure 3.7. Sub-study on anatomical accuracy. Group average source localization error quantified using PLE of PSF using MUSIC with free orientation. Results are presented on the *fsaverage* template.

Figure 3.8. Sub-study on electrode accuracy. Density estimates (over all subjects) of source localization error quantified using PLE of PSF with free orientation. The y-axis was restricted to the range 0 to 1.

Figure 3.9. Sub-study on anatomical accuracy. Density estimates (over all subjects) of source localization error quantified using PLE of PSF with free orientation. This is the same data as in figure 3.8 but with a logarithmic scale on the y-axis emphasizing the tail of the distributions.

Figure 3.10. Sub-study on electrode accuracy. Group average source localization error quantified using PLE of PSF using dSPM with free orientation (although in this case, an orientation prior was used as described in the main text). Results are presented on the *fsaverage* template.

Figure 3.11. Sub-study on electrode accuracy. Group average source localization error quantified using PLE of PSF using sLORETA with free orientation (although in this case, an orientation prior was used as described in the main text). Results are presented on the *fsaverage* template.

Figure 3.12. Sub-study on electrode accuracy. Group average source localization error quantified using PLE of PSF using dipole fitting with free orientation. Results are presented on the *fsaverage* template.

Figure 3.13. Sub-study on electrode accuracy. Group average source localization error quantified using PLE of PSF using MUSIC with free orientation. Results are presented on the *fsaverage* template.

70

Chapter 4

Discussion

In section 3.1, we proposed a method for generating head models. This method performs reasonably well in terms of reconstructing the skull compartment which is known to be important in EEG forward modeling (Lanfer et al., 2012; Vorwerk et al., 2014). We also made the point that the construction of high quality head models starts at the stage of acquiring the structural MRI scans upon which these models are based. Ensuring that scans are of good quality (e.g., reducing artifacts and ensuring good contrast) is crucial in order to facilitate the success of the subsequent steps in such a pipeline.

In section 3.2, we showed how the anatomical accuracy of the volume conductor model impacts the EEG forward solution. We compared head models constructed using different software packages. In particular, we used SimNIBS, MNE-Python, and FieldTrip. The former generates reasonably realistic representations of the anatomy whereas the latter two model the anatomy more coarsely. We showed how the pipelines in MNE-Python and FieldTrip generally resulted in high magnitude errors and moderate topography errors throughout the brain.

We also investigated the forward solution errors resulting from using template descriptions of electrode positions compared to using digitized locations for each individual subject. We found that by using a template layout measured on a realistic head model, we were able to achieve smaller errors compared to using the layout as specified by the manufacturer. The latter, in particular, resulted in large errors in the forward solution.

In spite of the reasonable accuracy of our custom template, we still found systematic inaccuracies in electrode position (and subsequently in the forward solution). Therefore, in section 3.3 we proposed a simple method for optimizing the electrode layout after it had been transformed to subject space. The main power of the method is that it helps reduce outliers substantially which is crucial in terms of getting an accurate forward solution (as evidenced by the relatively large errors incurred by using the manufacturer layout in the previous study).

In these studies we mostly focused on evaluating effects on the forward solution. However, it is still an open question to what extent the observed differences due to forward modeling translate to differences in inverse solutions.

Figure 4.1. A toy illustration of the relationship between RDM and dipole localization error based on the results of Dannhauer et al. (2011) who used scatter plots to explore this connection. The gray area shows the region with highest density of points.

4.1 Source Localization Errors

As discussed in section 3.2, several studies have explored the connection between forward and inverse modeling errors. An interesting attempt at making this explicit was made by Dannhauer et al. (2011) who used scatter plots to visualize the connection between RDM values of forward models and the corresponding dipole localization errors. This showed that small RDM values were associated with a wide range of localization errors (both small and large) whereas high RDM values were associated with medium localization errors (see figure 4.1for an illustration). This is not explored further in the article. However, one might wonder why this connection does not appear to be stronger. Due to the specific manipulations performed by Dannhauer et al. (2011), high and low RDM values were observed primarily for superficial and deep sources, respectively. However, accurate localization of deep sources is intrinsically difficult due to low sensitivity to these sources (even when the forward model is correct). Consider, for example, two sources separated by 1 cm. In general, it is probably reasonable to expect the topographic differences between such sources to be larger if they are superficial compared to deep meaning that the difference in sensor space (resulting from a difference in source space) is modulated by the positions (as well as other factors) of the sources. Consequently, RDM only tells part of the story; we also need to consider the sensitivity towards a given source as well as how smooth the forward operator is for such sources as this says something about how difficult it is to distinguish between neighboring sources.

It is also important to consider the SNR of the data. First of all, if the amount of noise in the data is high, then accurate localization of sources will be difficult (even with the true forward model) and the error due to noisy measurements will likely dominate over errors incurred by using a flawed forward model (see Montes-Restrepo et al., 2014, for an example). Besides, minimum norm estimators are usually regularized based on the estimated SNR of the data. As regularization is increased, solutions are essentially forced towards zero regardless of the forward model on which they are based. This is illustrated in figure 4.2 where we have applied ridge regression according to equation (2.81).

However, noiseless simulations are often employed in the literature (e.g., Dannhauer et al., 2011; Fiederer et al., 2016; Lanfer et al., 2012; McCann & Beltrachini, 2022) corresponding to an infinite SNR, thus likely maximizing

Figure 4.2. Effect of regularization on the similarity of two different MNE filters (estimated using equation (2.81)). One set of filters is based on a FEM model generated from a manual segmentation whereas the other is a BEM model constructed using MNE-Python. These correspond to the forward models used to investigate the effects of anatomical accuracy in the study presented in section 3.2. (The reference model was also used to generate figures 2.1 to 2.3.) Here we use RDM to assess the similarity of the filter weights. The regularization parameter takes on the values $\lambda^2 = \{2^i \mid i \in \{1, 2, \ldots, 30\}\}$. Left. Probability density of filter weight similarity as a function of the regularization parameter λ . Right. The dampening factor (from equation (2.81)) applied to the *i*th component of the SVD of the reference forward model. (We expect the plot for the BEM model to look very similar.)

the influence of the forward model. Other authors add Gaussian white noise to the measurements (e.g., Neugebauer et al., 2017) which may be reasonable for example to simulate electrical noise in the acquisition system. However, in general we might expect spatial correlations in the noise covariance estimates obtained from the data (e.g., because of noise due to neuronal processes of noninterest or artifacts). Interestingly, very high SNR levels may actually degrade the performance of beamformers when the true forward model is not known (Brookes et al., 2007; Dalal et al., 2014; Neugebauer et al., 2017). This is due to a mismatch between the topography of the gain vector(s) (to which the unit gain constraint is related) at a certain location and a source at that location.

The specific method employed when inverting the data is also very important in shaping the final results and may also modulate the degree to which forward modeling inaccuracies are expected to affect localization errors. For example, methods related to MNE solve for all sources simultaneously and generally make few assumptions about the data. This serve to stabilize the estimates with regard to (local) forward modeling errors, as shown by Stenroos and Hauk (2013) for errors in skull conductivity. This is because the responsibility of explaining the data is distributed among all possible sources. As mentioned above, increasing the amount of regularization will also stabilize the solution wrt. forward modeling errors (which, of course, does not necessarily mean that the solution will be good as it implies low SNR).

The scanning approaches (e.g., dipole fitting, MUSIC, beamformers) solve local problems. The source estimate at a certain location only depends on the gain vector(s) associated with this location. Thus, local imperfections in the forward model will directly propagate to the source estimate. In dipole fitting, we are examining each source to see how well it can explain the data. That is, we project the data vector onto the gain vectors to thus computing the similarity between the space spanned by the gain vectors for a particular source and the data vector (Sekihara, 2008). Modifying the gain vectors changes this relationship. Similarly, in MUSIC we are looking at the similarity of the spaces spanned by the signal and the gain vectors at a certain location although the signal space is defined differently than in the dipole scanning approach. Likewise, in the LCMV beamformer, changing the gain vectors for a certain position, directly impacts the filter weights due to the unit gain constraint.

4.1.1 Evaluation Through Simulations

In section 3.4 we presented selected results of localizing sources using different forward models, inverse methods, and SNR levels. Unfortunately, a comprehensive evaluation of these effects was outside the scope of the current project. Hence, we focused on source localization errors on group level.

As expected, we found decreasing localization error with increasing SNR although this effect was not as pronounced for MUSIC as the other methods. In particular, MUSIC generally performed well for all noise levels. This suggests some kind of ceiling effect perhaps due to the noise whitener being "too good" in this case. We used different covariance matrices for sampling the noise vectors and whitening the data. However, in practice these matrices were likely very similar. We ensured that we achieved the desired SNR at the time of peak activation and thus did not explicitly control the SNR over the entire time window. We see that dipole scanning is much more sensitive to the SNR level compared to MUSIC. Since the key assumptions of both methods were satisfied in our simulations (for dipole scanning that the data was generated by a single, active source; and, for MUSIC, that source time courses are not correlated), we do not expect this to explain these differences. However, in dipole scanning, the signal is represented by the (whitened) data vector at the time of interest whereas in MUSIC we use the (whitened) data covariance (or, rather, a subspace hereof) to model the signal subspace. The data covariance matrix is estimated from the entire time window (and hence is shaped by noise vectors from the entire time window) whereas the signal vector at a single time point is corrupted by a single noise vector drawn from a normal distribution. Thus, it is probably fair to assume that the whitening procedure will work better when applied to the data covariance matrix as opposed to applying it to a *single* data vector. Lastly, we see that sLORETA and dipole scanning are particularly sensitive to SNR in areas of low electrode coverage (orbitofrontal area and temporal pole). This is also somewhat true for MUSIC.

Comparing the MNE-based approaches, dSPM and sLORETA, we find that the latter performs better for sources to which the sensitivity is high whereas it performs worse for most deep sources suggesting that the normalization performed by sLORETA performs better for some sources than others. Since sLORETA uses the gain vectors to perform normalization (to estimate the source variance), it seems reasonable that the normalization will be less stable in areas with low sensitivity.

In general, the effects of the forward model on the accuracy of the inverse solution were modest. The model based on a template anatomy seems slightly worse than the others, particularly for superficial sources. We also see larger effects of the forward model on MUSIC and dipole fitting estimates compared to minimum norm-based estimates which is in line with the discussion presented above. In the sub-study where we varied electrode positions, effects of forward model were more pronounced. In particular, the solutions based on Man-Template are worse in the occipital area although effects are generally smaller than those of SNR except for MUSIC. Dipole scanning also shows this effect across SNR levels as is the case for sLORETA. This is in line with the increased errors in channel locations observed in the study presented in appendix B which were found to be largest over the occipital and parietal areas. Thus, the fact that the errors induced by using such a template configuration are highly spatially correlated seem to be very important. In practice, the coregistration between EEG and MRI coordinate systems also has the potential to cause substantial source localization errors (Shirazi & Huang, 2019) as the induced errors in the forward model will also be spatially correlated. On the other hand, if this spatial correlation is absent, then effects may be neglectable (Y. Wang & Gotman, 2001). Our results on electrode position accuracy agree with those of Homölle and Oostenveld (2019) who also found a manufacturer layout to be worse than a custom template, however, they only explored dipole scanning in noiseless conditions.

In terms of the impact of anatomical accuracy on the inverse solution, we did not find substantial effects of including the CSF compartment (e.g., comparing MNE-FS with SimNIBS-CHARM and FieldTrip-SPM). However, as we did not do a systematic evaluation of particular model properties (but rather, a practical comparison between different pipelines) this is perhaps not unreasonable. For example, the FieldTrip-SPM model includes CSF whereas MNE-FS does not, however, the former was not found to perform better than the latter. On the other hand, it is clear that the skull and CSF compartments are not particularly well modeled by FieldTrip-SPM. SimNIBS-CHARM and achieved better anatomical accuracy (and smaller source localization errors) although this was more general and not restricted to the skull and CSF compartments. Thus, these results seem to suggest that including CSF is indeed valuable provided that it can be identified with a certain accuracy.

In most studies where the effect of CSF is investigated, CSF is either omitted or it is modeled without errors (i.e., exactly as in the reference model) (e.g., Antonakakis et al., 2019; Cho et al., 2015; Conte & Richards, 2021; Neugebauer et al., 2017; Ramon et al., 2006). However, in general, it is more complicated. If we choose to increase the model complexity, we need to be confident that we are able to represent the new aspect(s) of the model with reasonable accuracy—both in terms of anatomical accuracy and conductivity. On the other hand, we have also previously made the point that increasing the complexity of the model so as to distinguish between compact and spongy bone may be beneficial because the skull compartment is otherwise very heterogeneous which may complicate its accurate segmentation (Nielsen et al., 2018). In this sense we may argue that, given our ability to generate an anatomical model, there is an optimal level of complexity and that more complicated is not always better. This is the problem of balancing bias and variance of our model.

In general, although we found substantial differences between the different forward solutions, these differences are typically reduced in the process of going from the forward to the inverse solution. This is not surprising because in order to solve the latter problem, we need to bias our solution towards what we believe is plausible since, as discussed in section 2.3, the data by itself it not sufficient to ensure a stable solution.

The current investigation is also subject to multiple limitations. We took several steps in order to avoid committing the inverse crime (e.g., adding noise, smoothing source activations, using different conductivities). However, we used the same source positions and the same model resolution for simulating and inverting the data. A possible way of dealing with the former is to shift the source positions slightly, e.g., 0.5 mm as done by Lanfer et al. (2012). Compared to the errors we observed, this modification seems very small indeed. As to the model resolution, it is possible to use a very high density model for simulating the data. However, in our study in section 3.2 we found that the difference in forward solution between models of resolution 0.5 nodes/mm^2 and 1.0 nodes/mm^2 were small compared to errors incurred by, for example, modeling the anatomy incorrectly. Therefore, we opted to use the same resolution in the reference forward model and those used to invert the data in section 3.4.

Another limitation is that we only simulated one source at a time. In practice, this is unlikely to be true due to rapid parallel activations occurring even at very early latencies (Michel et al., 2004). Besides, this constitutes a scenario where, in general, we expect it to be reasonably easy to properly localize sources since assumptions about number of sources (e.g., dipole scanning) and their correlation (e.g., MUSIC, LCMV beamformer) are not violated.

Finally, it is important to note that we only reported group level results. As such, we have not evaluated the extent to which errors may vary across individual subjects. In section 3.2 we found substantial variability in how well some of the pipelines represented the anatomy of different subjects. For example, the template anatomy could fit some areas of some subjects well whereas it would be very inaccurate in other areas. The same can be said of the MNE-Python- and FieldTrip-based models, particularly wrt. the skull compartment (the CHARM model seemed, in general, to capture the anatomy well). Consequently, it is possible that some models could result in outliers which exhibit large errors although in most cases the model fit will be reasonable. The same argument can be made when fitting electrode positions using a template as we also observed variation in how well the templates were able to fit each subject. With this in mind, particularly if one is interested in making statements at the level of individual subjects, there might be benefits to ensuring that the forward model is as accurate as possible even if on a group level such benefits are modest.

Figure 3.3 shows the density plots for the sub-study on anatomical accuracy on a log scale allowing us to assess the tails of the distributions. It is clear that the minimum norm-based estimates are robust to errors in these models whereas for dipole fitting and MUSIC we see that FieldTrip-SPM, in particular, have more dense tails but only when SNR is high. This seems to fit with the reasoning presented above, i.e., that local imperfections in the forward model are more likely to propagate to source estimates when using scanning techniques as opposed to distributed methods. Similarly, figure 3.9 shows the density plots for the sub-study on electrode accuracy on a log scale. We see largely the same pattern in terms of distributed methods compared to scanning methods. Apart from performing worse than the other models in general, Man-Template is also associated with heavier tails.

In practice, though, it may be difficult to know how accurate a forward model actually is. However, if one or more MRI scans are available, then the anatomical accuracy of the model can actually be evaluated qualitatively (at least to some extent) by comparing with the scans used to generate the head model. Other aspects of the forward may be difficult to assess, e.g., whether default conductivities are appropriate or how accurate the electrode positions are known.

4.2 Limitations and Future Research

Validating forward models and assessing their accuracy is challenging as the true forward model is never known in practice (aside from very simple, idealized models such as a sphere). Although experimental validation is difficult, one potential technique is magnetic resonance current density imaging which allow imaging of the magnetic field resulting from weak current injections on the scalp. This technique is, however, still new and currently only allows modeling of simple setups (Göksu et al., 2018).

Thus, when investigating source localization accuracy, studies usually rely on a reference model to simulate data. This model is typically taken to be the most sophisticated model that one is able to generate.

The influence of forward modeling errors on the inverse solution has often been investigated using simulations (e.g., Acar & Makeig, 2013; Dannhauer et al., 2011; McCann & Beltrachini, 2022; Montes-Restrepo et al., 2014) and this is also what we did in section 3.4. The obvious problem with this approach is that making the simulations *realistic* is very challenging because it is difficult to assess how (un)realistic they are. Thus, one is often susceptible to committing the "inverse crime", i.e., generating and reconstructing sources with the same (or a very similar) model (Kaipio & Somersalo, 2007). For example, since the true forward model is never known, one will have to generate a model with which to simulate the data and this forward model will often tend to look very similar to (a subset of) the ones used to invert the data (often, the latter models are constructed by introducing one or more specific imperfections, i.e., those which one is interested in assessing the effect of).

Additionally, the way the sources are simulated are often largely compatible with assumptions of the inverse solver used (e.g., using a dipole scanning approach when simulating a single source). Making simulations that are "equally fair" to all inverse solvers is very difficult thus making it difficult to make statements about the (absolute) errors that one might expect.

When using simulations, most studies seem to simulate a single source at a time (as we also did in section 3.4). This is convenient but may not be realistic in most circumstances as brain activity is usually not limited to a single location (Michel et al., 2004). Looking at more complex source patterns (e.g., multiple active, possibly correlated, sources) is therefore also highly relevant, not only in terms of making the simulations more realistic but also to explore the sensitivity of different inverse methods to violations of their assumptions. Of course, this type of analysis is usually performed when validating a certain method (Mosher et al., 1999; Veen et al., 1997), however, the degree to which the importance of the forward model is modulated by, for example, number of sources and their relative locations, is less clear.

A way to (at least partially) circumvent these issues may be to use real data instead of simulations. However, this presents its own challenges since the true location of the sources are rarely known in this case. In addition, one is limited to investigating effects specific to the source configuration present in the data. One way to try to alleviate the former issue is to compare with other modalities such as MEG or fMRI. Although also affected by volume conduction (Stenroos et al., 2014; Vorwerk et al., 2014), MEG usually has more sensors than an EEG system as well as decent SNR. The source of the signal is similar to that of EEG (i.e., apical dendrites of pyramidal cells oriented perpendicular

to the cortex), however, the sensitivity profile is slightly different (for example, MEG is known to be insensitive to radial sources) and the inverse problem is still ill-posed (although perhaps slightly less so). This has to be kept in mind when comparing EEG and MEG.

FMRI has excellent spatial resolution, however, using this as "ground truth" is also not unproblematic since mismatch between EEG and fMRI signals can happen due to several reasons. For example, fMRI activation will be apparent even in areas where the field is closed, i.e., where the neuronal activity is invisible to EEG. On the other hand, EEG is sensitive to very transient signals which may not survive thresholding in statistical maps based fMRI data if they are sufficiently brief. Thus, care has to be taken when using fMRI activation maps as source prior or as "ground truth" (Michel et al., 2004).

Finally, there are of course other aspects of the forward which are important but which have been ignored in this thesis. This includes, for example, the number of EEG electrodes. Previous results have shown substantial improvements of using (approximately) 60 electrodes compared to 30 electrodes whereas the improvement seems less significant going from 60 to 120 electrodes (Acar & Makeig, 2013; Michel et al., 2004). The positions of the electrodes are also important and generally more uniform coverage is preferred (Acar & Makeig, 2013; Song et al., 2015).

Another property of the forward model which has received a lot of attention is the conductivity profile of different tissues, particularly the skull compartment. For example, the conductivity ratio between brain and skull tissue as well as the effect of modeling the details of the skull composition have been investigated and this has also been shown to impact source localization results in EEG (e.g., Acar & Makeig, 2013; Chen et al., 2010; Dannhauer et al., 2011; Marin et al., 1998; Vallaghe & Clerc, 2009; G. Wang & Ren, 2013). Likewise, white matter anisotropy has also been explored (Wolters et al., 2006). However, as alluded to above, it is always important to weigh potential benefits against the added modeling complexity (Vorwerk et al., 2014).

Our analysis in section 3.4 also showed the importance of high SNR measurements. EEG data is often very noisy and a certain degree of preprocessing is important. However, as discussed in section 2.1.3, the quality of the data cannot be expected to be ensured by preprocessing alone as other factors are equally important (e.g., proper conductivity between skin and electrodes, minimization of artifacts during recording, for example, by ensuring that the subject is comfortable). If the quality of the EEG data is low, we cannot hope to recover the underlying sources no matter the accuracy with which the forward problem is solved.

Lastly, we would like to emphasize that we will to make the code for exporting gain matrices created with SimNIBS openly available. Although there are many challenges when doing source analysis, we believe that this is an important contribution to the field which has the potential to help improve the accuracy of EEG source localization across many studies. Initially, we plan on supporting MNE-Python and FieldTrip. However, the gain matrices created with SimNIBS may be subject to a few limitations compared to those generated directly with either software. For example, the dipole fitting procedures will not work because they need to be able to continuously sample new gain vectors in order for the nonlinear part (i.e., the optimization of the position) to work¹. Technically, however, it should be relatively straightforward to implement a similar functionality for gain matrices from SimNIBS as it simply requires barycentric interpolation.

¹In FieldTrip, the nonlinear step can be skipped

80

Chapter 5 Conclusions

In this work we have investigated several aspects of the EEG forward model and explored how it affects the forward solution as well as inverse solution.

Initially, we presented a pipeline (available in SimNIBS) for generating reasonably realistic anatomical models of the human head with particular emphasis on reconstructing the skull. We pointed out that it is important to ensure high quality of the structural images being used (e.g., by minimizing artifacts such as fat shift). We also found that including a T2-weighted MRI scan was very beneficial for delineating the border between CSF and skull and we highly recommend using this if possible. Finally, using an extended volume conductor model (which also includes the neck region) is beneficial for certain TES configurations as it allows a more realistic placement of the return electrode.

Subsequently, we compared forward solutions generated by SimNIBS, MNE-Python, and FieldTrip. The major difference between these models was the extent to which they were able to capture the underlying anatomy with SimNIBS generally being more accurate than the latter two. We found increased topographic and magnitude errors of MNE-Python and FieldTrip compared to SimNIBS throughout most of the brain suggesting large overall differences in the forward solutions. Thus, careful modeling of the anatomy seems to be important and this is easily achieved using SimNIBS. In addition to comparing different pipelines, we also compared different ways of specifying the electrode positions. In particular, we compared digitizing the electrodes against using either a custom template or a manufacturer template description of the electrode positions which was adapted to each subject. Using the manufacturer template induced large, spatially correlated errors as the electrode positions were poorly modeled in occipital and parietal areas whereas our custom template performed better. Thus, if it is not possible to directly measure electrode positions then we suggest using a template generated by measuring positions on a realistically shaped object (as opposed to the manufacturer positions which are defined on a sphere).

In a follow-up study, we investigated the feasibility of optimizing the electrode positions obtained using our custom template based on a few measurements of distances and angles between electrodes and landmarks. We used a total of the eight measurements (four distances and four angles). We showed that the results of the optimization is relatively robust against measurement errors but that the effect across subjects differ substantially. Specifically, some subjects benefitted considerably whereas it did not make much of a difference for others suggesting that the procedure is good at preventing outliers. Additional measurements are likely to further improve performance, however, we believe that electrode digitization is still superior.

Finally, we explored the effect of forward solution errors on source localization errors. Using the same forward models described above, we simulated data at different SNR levels and used different inverse methods for localization. Effects of anatomical accuracy on source localization error were modest although using template anatomy generally degraded the results. However, we only inspected the results on group level and not the level of individual subjects. Given the variability in the anatomical accuracy of the models from MNE-Python and FieldTrip, we find it reasonable to anticipate somewhat larger effects when looking at the solution for an individual subject. Effects of electrode locations were generally larger with the manufacturer layout performing most poorly. Interestingly, using template anatomy with digitized electrode positions performed better than using the correct anatomical model combined with the manufacturer description of electrode positions. This highlights the importance of correctly specifying electrode positions as otherwise large, spatially correlated errors may be introduced.

In order to facilitate the use of forward models with high anatomical accuracy in EEG data analysis, we will include code in SimNIBS for exporting gain matrices to various EEG analysis software packages. Initially, we intend to support MNE-Python and FieldTrip. By making these tools easily assessible, we hope to encourage their use by other researchers as well. The tools are easy to use and we would like to mention that although we did not systematically assess the computational demands of each pipeline, our testing suggests that running a SimNIBS pipeline is not more expensive than running the other pipelines. Appendix A

Automatic Skull Segmentation From MR Images for Realistic Volume Conductor Models of the Head: Assessment of the State-of-the-Art

This article was published in NeuroImage, March 2018.

NeuroImage 174 (2018) 587-598

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/neuroimage

Automatic skull segmentation from MR images for realistic volume conductor models of the head: Assessment of the state-of-the-art

Jesper D. Nielsen ^{a,b,e}, Kristoffer H. Madsen ^{a,b,e}, Oula Puonti ^{a,c}, Hartwig R. Siebner ^{a,d}, Christian Bauer ^a, Camilla Gøbel Madsen ^a, Guilherme B. Saturnino ^{a,c}, Axel Thielscher ^{a,c,*}

^a Danish Research Centre for Magnetic Resonance, Centre for Functional and Diagnostic Imaging and Research, Copenhagen University Hospital Hvidovre, Hvidovre, Denmark

^b Department of Applied Mathematics and Computer Science, Technical University of Denmark, Kgs. Lyngby, Denmark

^c Center for Magnetic Resonance, Department of Electrical Engineering, Technical University of Denmark, Kgs Lyngby, Denmark

^d Department of Neurology, Copenhagen University Hospital Bispebjerg, Copenhagen, Denmark

^e Sino-Danish College, University of Chinese Academy of Sciences, Beijing, 100190, PR China

ARTICLE INFO

Keywords: Skull segmentation Volume conductor model Forward modeling Transcranial brain stimulation Electroencephalography

ABSTRACT

Anatomically realistic volume conductor models of the human head are important for accurate forward modeling of the electric field during transcranial brain stimulation (TBS), electro- (EEG) and magnetoencephalography (MEG). In particular, the skull compartment exerts a strong influence on the field distribution due to its low conductivity, suggesting the need to represent its geometry accurately. However, automatic skull reconstruction from structural magnetic resonance (MR) images is difficult, as compact bone has a very low signal in magnetic resonance imaging (MRI). Here, we evaluate three methods for skull segmentation, namely FSL BET2, the unified segmentation routine of SPM12 with extended spatial tissue priors, and the skullfinder tool of BrainSuite. To our knowledge, this study is the first to rigorously assess the accuracy of these state-of-the-art tools by comparison with CT-based skull segmentations on a group of ten subjects. We demonstrate several key factors that improve the segmentation quality, including the use of multi-contrast MRI data, the optimization of the MR sequences and the adaptation of the parameters of the segmentation methods. We conclude that FSL and SPM12 achieve better skull segmentations than BrainSuite. The former methods obtain reasonable results for the upper part of the skull when a combination of T1- and T2-weighted images is used as input. The SPM12-based results can be improved slightly further by means of simple morphological operations to fix local defects. In contrast to FSL BET2, the SPM12-based segmentation with extended spatial tissue priors and the BrainSuite-based segmentation provide coarse reconstructions of the vertebrae, enabling the construction of volume conductor models that include the neck. We exemplarily demonstrate that the extended models enable a more accurate estimation of the electric field distribution during transcranial direct current stimulation (tDCS) for montages that involve extraencephalic electrodes. The methods provided by FSL and SPM12 are integrated into pipelines for the automatic generation of realistic head models based on tetrahedral meshes, which are distributed as part of the open-source software package SimNIBS for field calculations for transcranial brain stimulation.

Introduction

Volume conductor models of the head are key components of several neuroscientific methods such as electric field simulations for transcranial brain stimulation (TBS) and source localization in electro- (EEG) and magnetoencephalography (MEG). The anatomical accuracy of the head models has a strong influence on the accuracy of the calculated field distributions (Cho et al., 2015; Dannhauer et al., 2011; Eichelbaum et al., 2014; Lanfer et al., 2012; Montes-Restrepo et al., 2014; Wolters et al., 2006) and attempts to use individualized models based on structural magnetic resonance (MR) images are gaining momentum (Vorwerk et al., 2014). Recently available open-source software, including FSL (Smith et al., 2004), BrainSuite (Shattuck and Leahy, 2002), and SPM12 (http://www.fil.ion.ucl.ac.uk/spm/), facilitates the adoption of this approach by

E-mail address: axelt@drcmr.dk (A. Thielscher).

https://doi.org/10.1016/j.neuroimage.2018.03.001

^{*} Corresponding author. Danish Research Centre for Magnetic Resonance, Centre for Functional and Diagnostic Imaging and Research, Copenhagen University Hospital Hvidovre, Section 714, Kettegaard Allé 30, 2650 Hvidovre, Denmark.

Received 17 August 2017; Received in revised form 27 February 2018; Accepted 1 March 2018 Available online 12 March 2018 1053-8119/© 2018 Elsevier Inc. All rights reserved.

offering automatic segmentation procedures for the head. These tools have been integrated into software pipelines for the forward modeling of electric fields for TBS (e.g., SimNIBS; Thielscher et al., 2015) and EEG/MEG (e.g., FieldTrip; Oostenveld et al., 2011 and Brainstorm; Tadel et al., 2011). Accurate modeling of the skull compartment is an important aspect of individualized head models as the skull strongly shapes the forward solution due to its low conductivity (Dannhauer et al., 2011; Indahlastari et al., 2016; Lanfer et al., 2012; Montes-Restrepo et al., 2014; Stenroos et al., 2014). However, its automatic segmentation is still a major challenge, as the compact bone parts have a very low signal in conventional magnetic resonance imaging (MRI) sequences.

While the performance of most software packages in segmenting the brain have been thoroughly validated, similar tests are scarce for the skull. Thus, in this study we investigate the performance of three widely used neuroimaging software packages, FSL (Smith et al., 2004), Brain-Suite (Shattuck and Leahy, 2002), and SPM12 (http://www.fil.ion.ucl. ac.uk/spm/). Specifically, we assess FSL BET2 which includes the BET and betsurf tools (Pechaud et al., 2006), BrainSuite skullfinder (Dogdas et al., 2005), and the unified segmentation routine (Ashburner and Friston, 2005) implemented in SPM12. The latter was tested with spatially extended tissue priors in order to avoid clipping of the lower parts of the head (Huang et al., 2013). In contrast to BrainSuite, FSL and SPM12 support the use of multi-spectral MRI for segmentation. We therefore also compare the results when basing the segmentations on a single, high-resolution T1-weighted structural MR image, as often acquired in neuroimaging studies and used in clinical standard of care, versus a combination of high-resolution T1- and T2-weighted MR images. In addition, for the SPM12-based segmentations, we assess to which extent the results can be improved when applying morphological operations to "clean up" the raw segmentations. We test the quality of the segmentations by systematic comparisons against skull segmentations from computed tomography (CT) scans of the same subjects. To the best of our knowledge, this study is the first to rigorously assess the performance of these tools on skull segmentation and thus serves as important evaluation of the state-of-the-art on this topic.

Whereas the main focus of the paper is on skull segmentation, we further compare the accuracy of the reconstructed brain surfaces derived from SPM12-based segmentations with surfaces obtained using Free-Surfer 5.3.0 (Dale et al., 1999; Fischl et al., 1999). Finally, we exemplarily demonstrate the importance of selecting adequate MRI sequence parameters and adjusting the parameters of the SPM12 segmentation routine to the properties of the MR images in order to achieve robust and accurate results, particularly in non-brain regions. As such, our study gives useful guidelines for the adoption of individualized volume conductor models in neuroscientific research.

Material and methods

Subjects

Ten healthy subjects (five Caucasians [three males], five East Asians [two males], 20–50 years old; 28.9 ± 9.3 [mean age \pm SD]) were included in this study. They had no previous history of neurological or psychiatric disorders and were screened for contraindications to MRI and CT. In addition, the structural MR images were checked by a radiologist. Written informed consent was obtained from all participants prior to the scans. The study was approved by the Ethical Committee of the Capital Region of Denmark.

Data acquisition

The structural MR images were acquired on a 3.0 T Philips Achieva MRI scanner using a 32-channel head coil. A high-resolution T1-weighted scan (T1w; repetition time = 6.0 ms; echo time = 2.7 ms; flip angle = 8° ; 245 sagittal slices; matrix = 288×288 ; field of view = $245 \times 245 \times 208 \text{ mm}^3$; voxel size = 0.85 mm^3 ; bandwidth = 299.3 Hz; selective water excitation;

SENSE factor 2 along AP direction) and a high-resolution T2-weighted scan (T2w; repetition time = 2500 ms; echo time = 272 ms; flip angle = 90° ; 224 sagittal slices; matrix = 288×288 ; field of view = $245 \times 245 \times 190 \text{ mm}^3$; voxel size = 0.85 mm^3 ; bandwidth = 880.6 Hz, SENSE factor 2 along AP and 1.8 along RL) were performed. These sequence parameters were chosen to give good results for the skull segmentation using FSL BET2 and the FreeSurfer-based brain segmentation, and are based on our prior experience with the tools (Windhoff et al., 2013). Specifically, the readout bandwidth of the T1w image was chosen low enough to give a good signal-to-noise in the brain region. In addition, the parameters were selected to give a good contrast between gray matter (GM) and white matter (WM). Selective water excitation was chosen to ensure that most of the signal of the fatty spongy bone of the skull was suppressed. Otherwise, the signal of the spongy bone would have been merged with that of the pial surface of the brain at some positions due to the increased fat shift at low bandwidth, rendering an accurate segmentation of the inner skull boundary difficult. The fat shift refers to a systematic displacement of the spatial position of fat along the frequency encoding direction in the MR image that is caused by slightly different resonance frequencies of water and fat (King, 2004). The T2w image had a sufficiently high readout bandwidth to ensure a small fat shift that allowed for a good separation between the scalp, the compact and spongy bone, and the cerebro-spinal fluid (CSF). It is reasonable to assume that the above strategy to ensure a good spatial separation and a good contrast between the tissue classes was at least not detrimental for the SPM12-and BrainSuite-based segmentations even though we did not specifically optimize the MRI parameters for these.

The low-dose CT scans were acquired on a Siemens Biograph mCT (PET-CT) with axial slices having a voxel size of 0.42×0.42 mm² and a field of view of 215×215 mm². The resolution along the Z-direction was 0.60 mm³. The extent in Z-direction was individually adjusted to cover the complete neck while minimizing the radiation dose, and was on average 236 mm³. The tube current-time product (number of electrons per helical rotation) was 115 mAs, the tube potential (intensity) was 80 KeV, and the maximum effective dose was below 0.35 mSv for all participants.

Segmentation and preparation of CT images

An initial segmentation of the CT images was provided by thresholding the image at an intensity of 450 Hounsfield units (HU), a value, which seemed a reasonable compromise between retaining signal from the bone parts and suppressing noise. In order to achieve a robust segmentation of the bone regions while keeping the details as good as possible, the Chan-Vese level set method (Chan and Vese, 2001) was used, as implemented in the 2D/3D image segmentation toolbox (Zhang et al., 2008) for MATLAB. This method minimizes an energy function, which combines an image term penalizing large intensity variations within the segmented area and a curvature term penalizing the curve length resulting in smooth segmentations. As a final step, a few morphological operations were performed to fix minor inaccuracies in the segmentations (e.g., remove small holes). Visual inspection of the resulting bone masks confirmed the very good quality of the final segmentations, suggesting that they were suitable as "ground truth" for validating the MRI-based segmentations.

Segmentation of MR images

Three different methods for obtaining skull segmentations from MR images were tested. Their choice was based on own prior work (Windhoff et al., 2013) and results reported in related studies (Perdue and Diamond, 2014; Huang et al., 2013). The first method is based on tools provided by FSL 5.0.9 (https://fsl.fmrib.ox.ac.uk/fsl/fslwiki). Specifically, *BET* was employed to generate a brain mask, which was then used by *betsurf* as initialization to detect the inner and outer skull boundaries (Pechaud et al., 2006). *Betsurf* uses local intensity profiles along vectors perpendicular to the brain surface to identify the boundaries, based on the

assumption that the intensity variations along this normal follow a specific pattern, i.e. low (compact bone), high (spongy bone), low (compact bone). This approach recovers the skull regions directly surrounding the brain but fails to segment the skull parts belonging to the facial region and the vertebrae. In case both T1w and T2w images were used for segmentation, the T2w image was initially coregistered to the T1w image using the *FLIRT* registration tool (Jenkinson et al., 2002; Jenkinson and Smith, 2001). As the tools were integrated into our pipeline for the automatic construction of volume head meshes (details are stated below), an additional post-processing step was used to ensure a minimal distance between the inner and outer skull boundaries of 1 mm in case *betsurf* had failed to accurately detect the boundaries. We also ensured that there is at least a one voxel thick layer of CSF between the skull and gray matter.

The second method, *skullfinder* from BrainSuite 17a (http:// brainsuite.org/), consists of a series of morphological operations (Dogdas et al., 2005) following an initial skull-stripping step (Shattuck et al., 2001). *Skullfinder* automatically computes intensity thresholds used for initial skull and scalp segmentation, which is a starting point for the morphological operations. However, as these thresholds yielded poor skull segmentations on our dataset, we manually optimized the thresholds using the T1w images of three subjects and then applied those thresholds to all data.

The third method for skull segmentation was based on the unified segmentation routine (Ashburner and Friston, 2005) implemented in SPM12 (http://www.fil.ion.ucl.ac.uk/spm/) for MATLAB which combines spatial normalization to MNI space, intensity inhomogeneity correction, and tissue segmentation into one model. It uses a Gaussian mixture model for modeling tissue intensities and so fits one (or several) Gaussians to the intensity histogram of each tissue class. This is guided by a spatial prior on the probability of different tissues at a given position. Finally, SPM12 employs light Markov random field post-processing (default setting) to clean up the segmentations. The default spatial prior in SPM12 does not include the neck region. Thus, in order to avoid clipping of the lower parts of the head, we employed the spatially extended tissue priors introduced by (Huang et al., 2013). SPM12 outputs probability maps for each tissue, which were binarized by assigning the tissue type with the largest probability to each voxel. Since several local defects were typically still present in these binarized segmentations, we applied additional post-processing steps to correct the masks. The post-processing consists of simple morphological operations, similar to the strategy suggested in (Huang et al., 2013), including closing (dilation followed by erosion) of the bone structure to remove small holes, eroding the skin compartment, keeping only the largest connected component, and dilating it again to recover its original size (see the Appendix for details). These steps are automatically applied in our pipeline and help to improve the skull reconstruction, mostly by correcting the false labeling of spongy bone as skin. We assessed the quality of the segmentations both before and after these operations, and further tested the approach with only a T1w image or both T1w and T2w images as input.

In addition to testing the accuracy of the skull segmentations, further aspects of the SPM12-based segmentations were assessed:

• Binarized GM masks were extracted from the SPM12-based segmentation of the T1w images and compared with GM reconstructions obtained via FreeSurfer 5.3.0. It was also tested how much using the Computational Anatomy Toolbox 12 (CAT12) toolbox (http://www.neuro.uni-jena.de/cat) improved the SPM12-based GM segmentations. FreeSurfer has been shown to achieve reliable reconstructions of the GM and WM surfaces (Eggert et al., 2012; Han et al., 2006). We therefore chose it as suitable tool to assess the quality of the SPM12-and CAT12-based segmentations, but do not claim that the FreeSurfer-based segmentations should be considered as "ground truth". In order to allow for a comparison of the results, binarized GM masks were obtained from the surface reconstructions provided by FreeSurfer by filling in the region between the GM and WM surfaces, with voxel size defined by the MR image, using custom-written

software. A similar strategy was employed for the results of CAT12, which provides a surface that delineates the middle of the GM sheet and additionally an estimate of the local GM thickness as output. Custom-written software implemented in Python 2.7 was used to reconstruct the GM and WM surfaces from the CAT12 results and then filling in the regions between the two surfaces. The WM surface was constructed from the WM segmentation provided by CAT12 using the marching cubes algorithm (Lewiner et al., 2003; van der Walt et al., 2014). The GM surface was constructed by moving the nodes of the central GM estimate outwards by half of the GM thickness to generate the pial GM surface. In the latter case, the movement was locally stopped whenever necessary to prevent self-intersections.

- MR images suffer from inhomogeneous intensity profiles due to imperfections in the radio-frequency (RF) excitation and receive coils so that a given tissue type will have varying intensities across the image which can cause errors in the resulting segmentations if not accounted for (Van Leemput et al., 1999). SPM12 deals with this problem by estimating the MRI-related intensity inhomogeneities, typically referred to as "the bias field" (Wells et al., 1996), during the segmentation process. The estimated bias field is then used to correct the intensity profile of the MR scan. Estimation of the bias field can, however, be problematic when performed across the whole head rather than only the brain as is usually done in neuroimaging. We found that careful adjustment of the regularization parameter, which controls how quickly the estimated bias field is allowed to change across space, helped to improve the segmentation in particular of the lower parts of the head. Given the practical relevance of this topic, we exemplarily demonstrate the impact of the regularization parameter on the segmentation results.
- The sequence parameters were chosen based on prior experience (Windhoff et al., 2013), and the rationale behind this choice is outlined in the discussion. To demonstrate the importance of careful selection and testing of the employed MRI sequence, we show an example using a T1w image from a publicly accessible dataset (obtained from the OpenfMRI database, accession number ds000171).

Validation and comparison of the MRI-based segmentations

We validated the skull reconstructions of all three methods against those obtained from CT scans. Images were downsampled to $1\,\mathrm{mm}^3$ voxels and the CT scans were coregistered to the T1w images using FLIRT. Subsequently, an additional, nonlinear coregistration step was performed using the Elastix software package (Klein et al., 2010) to properly coregister the neck. After applying FLIRT and Elastix, careful visual inspection was performed to ensure the quality of the coregistrations. In addition, we compared the SPM12-based GM segmentations (both with and without CAT12) against FreeSurfer-based GM segmentations. The full set of analyses is listed in Table 1. For the skull analysis, neither changing the downsampling parameter in SPM12 from threefold to none, nor using CAT12 affected the result of the SPM12-based segmentations (both with and without post-processing using morphological operations); hence, we report results only for the default settings. Also for the GM analysis, changing the downsampling parameter in SPM12 did not affect the result, so that we report the results only for the default settings. The binarized posteriors were used for all comparisons.

For validating the bone reconstructions against the CT scans, we used the Dice coefficient and the modified Hausdorff distance (Dubuisson and Jain, 1994) to provide summary scores of the overall quality of the segmentations. The Dice coefficient measures the similarity of sets, *A* and *B*, by their overlap and is defined as

$$D = \frac{2|A \cap B|}{|A| + |B|}$$

where |X| denotes the number of points in set X. Thus, $0 \le D \le 1$ with perfect overlap resulting in D = 1. The modified Hausdorff distance H(A,B) measures the average of two directed Hausdorff distances $h_m(A,B)$ and $h_m(B,A)$ between the borders of sets A and B, and is defined as

$$H(A,B) = \operatorname{mean}\{h_m(A,B), h_m(B,A)\}$$

with the directed Hausdorff distance

$$h_m(A,B) = \frac{1}{|A|} \sum_{a \in A} \min_{b \in B} \{d(a,b)\}$$

being the average minimum distance from set A to set B for some suitable distance measure d (Dubuisson and Jain, 1994). In our case, d denotes the Euclidean distance and the boundary voxels (surface) were obtained by subtraction with the eroded segmentation. The reason for reporting both the Dice score and the modified Hausdorff distance is that the Dice score is dominated by large clusters of voxels and is not very sensitive to differences along the surfaces of the segmentations (which are often the most difficult parts to properly model).

In addition to providing the summary scores, we also produced spatial maps of false positive and negative rates of the correspondence between the MRI and CT segmentations to document how the segmentation accuracy varied between different parts of the head. These maps were obtained by labeling a voxel that was incorrectly classified as bone in the MRI-based segmentations as a false positive and a missed bone voxel correspondingly as a false negative. Finally, the subject-specific false positive and negative maps were transformed to the MNI space and averaged, resulting in false positive and negative rates across subjects for each anatomical location. We produced similar maps for the accuracy of the SPM12-and CAT12-based GM segmentations by comparing them to the FreeSurfer GM segmentations.

For the SPM12- and BrainSuite-based skull segmentations, we report the Dice coefficients and modified Hausdorff distances for the complete head, including the vertebrae, and in addition for the upper part of the head only. We included the latter results to allow for a better comparison with the FSL-based skull segmentations, which do not include the vertebrae.

Generation of head volume meshes and TDCS field simulations

The FSL- and SPM12-based methods for skull segmentation were integrated into software pipelines for the automatic generation of individualized volume conductor models of the head. These head models can be employed in our open-source software SimNIBS 2 (www.simnibs.org; Thielscher et al., 2015) that uses the finite element method (FEM) for the calculation of the electric fields generated by transcranial magnetic stimulation (TMS) and transcranial direct current stimulation (TDCS). Their usefulness has further been demonstrated in field calculations for electro- and magnetoencephalography (EEG and MEG) based on the boundary element method (Stenroos and Nummenmaa, 2016), as well as for FEM-based forward models for EEG (Ziegler et al., 2014).

The software pipeline *mri2mesh* has already been released as part of SimNIBS and uses FSL *BET2* for skull segmentation and FreeSurfer 5.3.0

Table 1

Summary of the tested software configurations.

Segmentation	Postprocessing of segmentations		
Skull segmentations (compared against CT; input: T1w only, or T1w&T2w)			
FSL BET2	ensure 1 mm thickness		
SPM12 unified segmentation	MRF		
SPM12 unified segmentation	MRF + morph		
GM segmentations (compared against FreeSurfer; input: T1w only)			
SPM12 unified segmentation	MRF		
SPM12&CAT12	_		

MRF denotes the Markov random field procedure integrated in SPM12, "morph" refers to a number of morphological operations to clean up the binarized skull segmentations

for brain segmentation. The methods used by mri2mesh to convert the voxel segmentations into a volume mesh based on tetrahedral elements are described in detail in (Windhoff et al., 2013). Briefly, voxel segmentations of the inner and outer skull boundaries (and the scalp) are converted into surfaces based on triangle elements. These surfaces and the GM and WM surfaces reconstructed by FreeSurfer then undergo cleaning steps to ensure a good triangle quality and to avoid mutual intersections. Finally, volume meshes are constructed by filling the spaces between the surfaces by tetrahedra using Gmsh (Geuzaine and Remacle, 2009). The meshes contain around 3-4 million tetrahedra and distinguish between five tissue types, namely WM, GM, CSF, skull and scalp. Due to the chosen meshing approach, the tissue surfaces need to be nested and thus we ensure that there is at least one voxel layer of CSF between skull and gray matter as described in section Segmentation of MR images. This artificial restriction might not be exactly anatomically accurate, and is a limitation enforced solely by the current meshing approach.

Our new software pipeline, *headreco*, uses the voxel segmentations generated by SPM12 cleaned using morphological operations and, if available, the reconstruction of the middle of the GM sheet supplied by CAT12 to build individualized head models. It will be added to the next release of SimNIBS 2. Similar methods as described in (Windhoff et al., 2013) are used to create triangle surfaces from the voxel segmentations, which are then cleaned and finally used as input for the generation of tetrahedral volume meshes. As part of this process, the GM surface was constructed by moving the surface representing the middle of the GM sheet outwards as described above. The final head meshes produced by *headreco* extend the head meshes from *mri2mesh* by also including the vitreous bodies of the eyes and, in addition, modeling the paranasal sinuses, parts of the mouth and the throat as air cavities.

The numerical accuracy and stability of the FEM calculations relies on the geometrical quality of the tetrahedra (i.e., their "well-shapedness"). We therefore compared the tetrahedral quality of the head meshes generated with the new pipeline with those reported in (Windhoff et al., 2013) for *mri2mesh*, using three metrics, ρ , η , and γ which relate to the shape of the tetrahedra (Liu and Joe, 1994; Zhang et al., 2005)

$$\rho = C_1 \frac{\min_i e_i}{\max_i e_i}$$
$$\eta = C_2 \frac{V^{2/3}}{\sum_i e_i^2}$$
$$\gamma = C_3 \frac{V}{\max_i \{e_i\} \sum_i A_i}$$

where e_i is the length of the *i*th edge, V is the volume of a particular tetrahedron and A_i is the area of the *i*th face of a particular tetrahedron. Constants C1, C2 and C3 were chosen so that all of these metrics are bounded between 0 (worst quality) and 1 (best quality). For example, ρ is the ratio of the lengths of the shortest and longest edges and thus penalizes long, thin tetrahedra. The metric η assesses the ratio between the tetrahedral volume and the edge lengths, again favoring tetrahedra, which achieve a given volume V with on average shorter edges compared to tetrahedra requiring on average longer edges. Metric γ works similarly by relating the tetrahedral volume to the surface area. The metrics show different sensitivities to various possible deviations from an ideal symmetric pyramid shape, rendering an assessment of all three metrics useful. All metrics are calculated per tetrahedron, and then average statistics are extracted. We report the minimal values occurring per head mesh (averaged across the 10 subjects), the mean values, and the percentage of tetrahedra having a value lower than 0.1. The latter value of 0.1 was selected as lower threshold for a "good" tetrahedral quality based on our subjective experience with the stability of the FEM results.

In order to illustrate the relevance of using extended head models

J.D. Nielsen et al.

that include the neck region when calculating the electric field distribution for TDCS, we exemplarily show the calculated fields from two tDCS montages; a standard montage for motor cortex stimulation with a return electrode on the contralateral supraorbital region and a montage with an extraencephalic return electrode. Both montages were simulated using a head model restricted to the upper part of the head (as generated by mri2mesh) and an extended head model (generated by headreco). In case of the restricted head model, the extraencephalic return electrode was positioned on the bottom surface to mimic the correct position "as good as possible". We give two further examples to show how a misestimation of the skull thickness or a coarse segmentation of the gyral folding pattern can affect the electric field in the brain, both times based on the standard montage. All tissues, including WM, were modelled as isotropic conductors, and the assigned conductivity values were 0.126 S/m (WM), 0.275 S/m (GM), 1.79 S/m (CSF), 0.01 S/m (bone), 0.50 S/m (vitreous bodies) and 0.465 S/m (scalp) (Baumann et al., 1997; Saturnino et al., 2015). The electrodes were modelled as rubber pads $(40 \times 40 \text{ mm}^2, 1 \text{ mm} \text{ thickness}, 0.1 \text{ S/m})$ embedded in saline-soaked sponges $(70 \times 50 \text{ mm}^2, 5 \text{ mm} \text{ thickness},$ 1 S/m), with the connectors placed centrally on the rubber pads. A current strength of 1 mA was simulated.

Results

MRI-based skull reconstructions

Comparison of the MRI- with the CT-based segmentations reveals better results for the segmentations based on combined T1w- and T2wimages versus those using only a T1w image (Fig. 1), consistently for FSL BET2 and the unified segmentation routine of SPM12. Inclusion of a T2w image generally serves to improve the segmentations and stabilizes the results (i.e., decrease the variance across subjects). Importantly, outliers with very bad segmentations are mostly prevented. This is likely due to the better contrast between skull and CSF, which is mostly absent in the T1w images. Comparison between SPM12 and FSL BET2, using the reduced field-of-view (FoV) and a T1w image as input, shows that both perform comparably in terms of Dice score but that SPM12 obtains a lower Hausdorff distance, which is likely due to a more accurate segmentation of the boundary of the skull. When both T1w and T2w images are given as input, SPM12 performs better in both metrics compared to FSL BET2. Finally, BrainSuite consistently performed worse than the other methods irrespective of FoV.

Post-processing the SPM12-based voxel segmentations using simple morphological operations improves the results moderately, but consistently (bold vs. patterned boxes of same colors in Fig. 1). In addition, we note improved accuracy of the SPM12 segmentations in the upper part of the head (covering the skull) compared to the lower part (covering the

Dice Coefficient

neck) as indicated by better results for the reduced vs. full FoVs.

The false positive and negative maps shown in Fig. 2A allow for a more fine-grained evaluation of the segmentation accuracy for the different parts of the skull and the vertebrae. Generally, most errors occur at the boundaries, particularly at the inner skull boundary if a T2w image is not used. While FSL BET2 both over- and underestimates the inner skull boundary in this case, BrainSuite and SPM12 systematically underestimate the skull thickness (the results after applying the morphological cleanup operations are shown for the SPM12 segmentations). Using also a T2w image, FSL BET2 slightly, but consistently, underestimates the skull thickness, while the SPM12-based results are quite accurate in the upper part of the skull. In general, air cavities such as the frontal and paranasal sinuses tend to be underestimated by all methods and be represented as bone (FSL BET2 covers only the region of the frontal sinuses). FSL BET2 and BrainSuite oversimplify the shape of the skull base. SPM12 recovers this part better in case both T1w and T2w images are given as input, but generally underestimates the vertebrae.

MRI-based gray matter reconstructions

Fig. 2B shows the map of false positive and negative rates across the subjects for the GM segmentations. Comparing the SPM12-based GM segmentations with those of FreeSurfer (upper row in Fig. 2B and first row in Table 2) reveals that SPM12 tends to miss CSF in thin and deep sulci, which are instead labeled as GM (visible as yellow false positives and higher average Hausdorff distance compared to CAT12 in Table 2). In contrast, CAT12 generates a better segmentation of these sulci, resulting in lower rates of false positives and negatives when compared to FreeSurfer (lower row in Fig. 2B and second row in Table 2). This effect is also clearly visible in pial surfaces reconstructed via the three different methods, as exemplarily demonstrated on a single subject in Fig. 2C.

Effect of regularization of bias field correction and of MR image parameters

While helpful for a robust segmentation of the CSF-bone border, using a T2w image in addition to the T1w image in the SPM12-based segmentation sometimes had detrimental effects on the segmentation of the spinal cord in several subjects. This was caused by an incorrect estimate of the bias field, which could be resolved by adjusting the regularization parameter that controls the flexibility of the estimated bias field. Fig. 3A shows an example of the effect of the regularization parameter on the estimated probability map of white matter (which is binarized to obtain the segmentation). Note the spinal cord which is missed when using too large (0.1) or too low (0.001) values of the regularization parameter.

In addition, the quality of the segmentation depends clearly on the chosen MRI sequence parameters. In Fig. 3B, this is exemplarily demonstrated for the SPM12-based skull segmentation using a T1w

Fig. 1. Global accuracy measures for the skull segmentation. A) Modified Hausdorff distance. B) Dice coefficient. The black lines indicate the medians across the 10 subjects. The boxes indicate the 25 and 75 percentiles, and the error bars indicate the most extreme datum within 150% of this range. Data points outside the error bars are treated as outliers and are shown as small circles. Analyses using a reduced FoV restricted to the upper part of the head are marked by (R) in the legend. Analyses including the full FoV including the neck and vertebrae are marked by (F). Solid blue: FSL BET2 results assessed in the reduced FoV; Solid red: SPM12-reduced FoV; Striped red: SPM12 + morphological operations reduced FoV; Solid green: SPM12-full FoV; Striped green: SPM12 + morphological operations - full FoV; Solid yellow: BrainSuite - reduced FoV; Solid magenta: BrainSuite - full FoV. To aid visualization in A, outliers at 4.25 mm and 4.42 mm for SPM12 (T1+T2, full FoV) and BrainSuite (full FoV), respectively, are not shown.

J.D. Nielsen et al.

Fig. 2. A) Local accuracy of bone reconstruction assessed by false positive and false negative rates across subjects overlaid on the bone prior of the extended atlas. Red denotes false positives with a ratio of one corresponding to a false positive labeling of non-bone tissue as bone in all subjects. Blue denotes false negatives with a one indicating a labeling of bone as non-bone in all subjects. Shown are exemplary coronal and sagittal slices. The gray box in the upper row indicates the neck region that is ignored by FSL BET2. B) Local comparison of the gray matter reconstruction of SPM12 (upper row) and CAT12 (lower row) with FreeSurfer assessed by false positive and false negative rates across subjects overlaid on the MNI T1 template. Red and yellow colors indicate false positives, blue denotes false negatives. C) Representative gray matter surface reconstruction on a single subject based on the GM segmentations of SPM12 (green) and CAT12 (red), and the pial surface reconstructed by FreeSurfer (yellow).

image from an openly available dataset as input. This T1w image was acquired without fat suppression, resulting in strong signal from subdural fat belonging to the scalp and the neighboring spongy parts of the bone. As a result, large parts of the skull were incorrectly segmented as scalp and the skull was strongly underestimated. Interestingly, facial bones were modelled in the defaced region based on the skull priors.

Mesh quality

The quality of the head meshes created by mri2mesh (which employs

FSL BET2 and FreeSurfer for segmentation) and headreco (which uses SPM12 and CAT12) is comparable (Tables 3-5). This is expected as similar methods are used for generation and optimization of the triangle surfaces and for the subsequent volume meshing. A more accurate modeling of fine sulcal structures slightly increases the number of lowquality tetrahedra, as revealed by a comparison of the results for mri2mesh and headreco with CAT12 versus headreco using only SPM12 in Tables 3-5 However, a good mesh quality is reached in all cases. The number of low-quality tetrahedra is lower than originally reported in Table 2 in (Windhoff et al., 2013). This is most likely due to improvements in the volume meshing routines in the newer version of Gmsh

Table 2

Similarity of the gray matter segmentation to the FreeSurfer result for default SPM12 and when using CAT12. We report the mean and standard deviation of the modified Hausdorff distance for the 10 subjects.

Method/Score	Modified Hausdorff distance in mm (mean/std)
SPM12	
T1	0.7168/0.057
T1+T2	0.7242/0.061
SPM12&CAT12	
T1	0.3740/0.047
T1+T2	0.3604/0.029

employed here.

Electric field simulations for TDCS

Exemplary volume conductor models, as generated by the different methods, can be seen in Fig. 4A&B. Compared to *mri2mesh, headreco* offers an extended spatial coverage including a (coarse) modeling of the vertebrae, a finer modeling of the skull base and of the air cavities, and the modeling of the vitreous bodies of the eyes as a separate class. Interestingly, these differences only seem to play a minor role for a standard TDCS montage targeting the motor cortex, which places both sponge electrodes on the upper part of the head (electric field images in Fig. 4A&B). In contrast, very clear differences are seen when modeling an extraencephalic return electrode (Fig. 4C&D). While the field distribution underneath the orange stimulation electrode is similar for the two head models (image not shown), very clear differences occur in the

Table 3

Mesh quality as quantified by η . Summary statistics include minimum quality tetrahedron, percentage of tetrahedra below $\eta = 0.1$, and mean quality of tetrahedra. Values reflect mean (standard deviation) across all subjects.

	η		
	min (std)	% < 0.1 (std)	mean (std)
mri2mesh			
T1	0.034 (0.023)	7.5e-04 (3.8e-04)	0.768 (0.004)
T1+T2	0.033 (0.020)	5.5e-04 (2.6e-04)	0.763 (0.005)
headreco			
T1	0.086 (0.035)	2.3e-05 (2.5e-05)	0.784 (0.005)
T1+T2	0.082 (0.047)	1.2e-04 (2.4e-04)	0.785 (0.005)
T1 (CAT)	0.025 (0.014)	7.0e-04 (4.0e-04)	0.780 (0.004)
T1+T2 (CAT)	0.013 (0.007)	1.1e-03 (4.8e-04)	0.780 (0.005)

cerebellum and the temporal lobe of the right hemisphere.

The effects of a gross underestimation of the average skull thickness on the electric field in the brain are exemplarily shown in Fig. 5A. The inaccurate skull segmentation (similar in quality to the one shown in Fig. 3B) results in less blurring of the electric field on the brain surface and the occurrence of localized field peaks. Interestingly, also a coarse segmentation of the gyral folding pattern clearly affects the field estimates in the brain (Fig. 5B). Not modeling the CSF in thin and deep sulci increases the field strength in large parts of gray matter, presumably due to reduced shunting of currents through well-conductive CSF pathways.

Fig. 3. A) Probability map of white matter from SPM12 for different strengths of the regularization of the bias field correction for the T2 image. The regularization strength of 0.001 used in the right column is the default value in SPM12. The results have been overlaid on the bias corrected T2 image. B) Exemplary demonstration of the impact of the MRI parameter settings on the results of the skull segmentations. A T1w image from an open-source dataset was used as input. Please note that the images of this dataset were defaced to protect the privacy of the subjects. The left column shows sagittal and coronal slices through the T1w image; the middle column shows the posterior probability given by the SPM12 unified segmentation routine; the right column shows the final skull segmentation after the morphological operations.

Table 4

Mesh quality as quantified by γ . Summary statistics include minimum quality tetrahedron, percentage of tetrahedra below $\gamma = 0.1$, and mean quality of tetrahedra. Values reflect mean (standard deviation) across all subjects.

	<u>γ</u>		
	min (std)	% < 0.1 (std)	mean (std)
mri2mesh			
T1	0.031 (0.023)	1.0e-03 (4.5e-04)	0.630 (0.004)
T1+T2	0.032 (0.020)	8.0e-04 (4.0e-04)	0.625 (0.005)
headreco			
T1	0.067 (0.034)	4.2e-05 (4.2e-05)	0.646 (0.005)
T1+T2	0.068 (0.039)	1.6e-04 (2.9e-04)	0.647 (0.005)
T1 (CAT)	0.023 (0.014)	1.0e-03 (4.8e-04)	0.642 (0.004)
T1+T2 (CAT)	0.011 (0.007)	1.4e-03 (6.0e-04)	0.643 (0.005)

Table 5

Mesh quality as quantified by ρ . Summary statistics include minimum quality tetrahedron, percentage of tetrahedra below $\rho = 0.1$, and mean quality of tetrahedra. Values reflect mean (standard deviation) across all subjects.

	ρ		
	min (std)	% < 0.1 (std)	mean (std)
mri2mesh			
T1	0.043 (0.011)	1.4e-03 (6.8e-04)	0.548 (0.005)
T1+T2	0.057 (0.010)	1.1e-03 (4.2e-04)	0.543 (0.005)
headreco			
T1	0.104 (0.015)	3.2e-05 (4.5e-05)	0.565 (0.005)
T1+T2	0.102 (0.022)	3.9e-05 (1.0e-04)	0.566 (0.005)
T1 (CAT)	0.066 (0.019)	4.4e-04 (3.1e-04)	0.561 (0.004)
T1+T2 (CAT)	0.057 (0.016)	6.8e-04 (3.4e-04)	0.562 (0.005)

Discussion

We have validated the accuracy of skull segmentations obtained by three methods (FSL *BET2*, BrainSuite *skullfinder*, and the SPM12 *unified segmentation* routine) by comparing against CT-based skull segmentations in ten subjects. Both FSL and SPM12 give reasonable results for the upper part of the skull, in particular when both a T1w and T2w image are used as input. The results of BrainSuite are less accurate. For FSL and SPM12, including a T2w image strongly reduces the variability of the segmentation accuracy across subjects, with the SPM12-based being moderately more accurate than FSL *BET2* in that case. Post-processing of the SPM12based segmentation using simple morphological operations increased the accuracy slightly. Visual inspection of the local accuracy maps demonstrated the remaining specific weaknesses of all methods – even when using both T1w and T2w images.

BrainSuite *skullfinder* (which uses only a T1w image) was constructed primarily with the purpose of segmenting the upper part of the skull (Dogdas et al., 2005) explaining its relatively low accuracy in segmenting the jaw and neck areas. We see, however, that it generally underestimates the inner skull border, i.e., the resulting skull segmentation is thinner compared to the CT-based ground truth. The underlying reason seems to be the use of too large brain masks that are used as starting points of the inner skull estimates. Consequently, the manual optimization of the threshold parameters (as we did in this study) helps to improve the outer but not inner skull estimate. The reconstructed vertebrae seem to be spatially misplaced, resulting in a mixture of false positives and false negatives in the lower part of the segmentations.

FSL *BET2* relatively coarsely approximates the shape of the skull base, resulting in both false positives and negatives in this region. This approximation is caused by *betsurf*, which assumes a particular pattern of intensity variations along vectors from the brain to the outer scalp and identifies the skull boundaries as minima along this intensity profile. This strategy works reasonably well for the upper part of the skull, but not for the lower parts with more complicated bone structures.

In contrast to the results of FSL BET2, the SPM12-based segmentation

NeuroImage 174 (2018) 587-598

Fig. 4. A) Extended volume conductor model including the neck of an exemplary subject, created by the headreco script using SPM12 and CAT12. A "standard" montage for left motor cortex stimulation is simulated, with one electrode positioned above the left motor cortex and the return electrode placed above the right supraorbital region (the black bars were added to camouflage the facial features). The middle column shows a sagittal slice through the volume mesh. The right column shows the induced field over the target area. B) Volume conductor model with limited coverage of the same subject, created by mri2mesh using FSL BET2 and FreeSurfer. The injected electric field in GM is very similar to the one calculated with the extended head model. C) Norm of the electric field for a montage targeting the left motor cortex, with one electrode placed above the motor cortex and an extraencephalic return electrode on the right side of the neck. The extended volume conductor model was used. D) Norm of the electric field for the same montage as in C, but adapted for the volume conductor model covering only the upper part of the head. The position of the extraencephalic return electrode was mimicked as good as possible by placing it on the right bottom surface. The estimated field clearly deviates from that in C. While the field strength in the cerebellum is underestimated, the field in the right temporal lobe is overestimated. The rightmost view in subfigures C and D is from below the brain.

also covers facial bone structures and the vertebrae. Remaining inaccuracies compared to the CT-based segmentations manifest themselves mainly as false positives around the air cavities and false negatives around the vertebrae. Both air and compact bone do not give signal in conventional MR images, so that the *unified segmentation* routine of SPM12 has to rely on the spatial tissue priors in order to distinguish both. The employed spatial priors did not enforce the shapes of the air cavities strongly enough, resulting in inaccurate segmentations. The vertebrae were consistently underestimated, which was likely a result of

Fig. 5. A) Demonstration of the effect of inaccurate (top row) and accurate (bottom row) skull segmentation on the simulated electric fields. The inaccurate skull model is the result of an automatic segmentation of a T1w image using FSL BET2. The accurate model used SPM12 with standard tissue probability maps for the initial segmentation of the same image, followed by manual corrections. The norm of the electric field is shown on the FreeSurfer cortical (pial) surface, where the electric field strength was read out on the mid layer between the white matter and pial surfaces. **B)** Demonstration of the effect of SPM12 (left) and SPM12&CAT12 (right) gray matter segmentations on the simulated electric fields. For both models, the electric field strength was read out on the mid layer of the gray matter sheet, as reconstructed by CAT12. In both subfigures, a "standard" montage for left motor cortex and the return electrode placed above the right supraorbital region.

insufficient contrast between different tissue types in the MR images together with smooth spatial priors missing the details of the vertebrae structures. For example, in the T2w images, the spinal cord, the spongy bone of the vertebrae, and subdural fat show similar intensities and occur in close proximity, so that recovering the vertebrae structures would necessarily require strong prior information. Furthermore, the spatial priors do not account for abnormalities, such as tumors, stroke or surgical intervention, as occurring in patients. Large deviations from normal anatomy can lead to significant degradation of the resulting head models (Birot et al., 2014; Minjoli et al., 2017) and need to be accounted for explicitly to ensure good quality of the head models in these cases. This can be done either by manual editing (Datta et al., 2011; Minjoli et al., 2017) or by tailored automated segmentation tools (Kamnitsas et al., 2017; Menze et al., 2015).

Commercial software packages were excluded from this study. We are aware of two tools: BESA MRI (http://www.besa.de/products/besa-mri/ besa-mri-overview/) and EGI's GeoSource 3 (https://www.egi.com/ research-division/electrical-source-imaging/geosource), which both provide automated head model generation, including skull segmentation, from MR scans. The skull segmentation approach in BESA MRI combines tissue probability atlases with a spatial Markov random field model (Lanfer et al., 2014), which on a general level is closely related to what SPM12 does, whereas GeoSource 3 uses landmark-based nonlinear registration to warp a skull template to the target MR scan to produce a skull segmentation (Li et al., 2016). However, quantitative comparison of these methods was not possible, as we do not have access to either of the software packages.

While our results indicate that relatively good skull segmentations can be obtained from automatic procedures, the achieved quality depended on the MRI sequence parameters and (for SPM12 and Brain-Suite) also the input parameters and thresholds for the segmentation routines. FSL BET2 relies on clearly detectable intensity variations between brain, compact and spongy bone and scalp along the sampled vectors, which is usually well achieved in T2w, but not in T1w images. In the latter, CSF is dark, rendering it similar to compact bone. In addition, the fat-water shift due to a rather low readout bandwidth can camouflage the boundaries between scalp, spongy bone and brain. During prior testing, this led us to combine a T1w image with selective water excitation (in which compact and, to a large extent, spongy bone gives homogenously low intensity) with a non-fat suppressed T2w image. While we did not separately optimize the MR image parameters for use with SPM12 or BrainSuite, our results indicate that they are also a reasonable choice for these methods. Testing BrainSuite with a T1w image without fat suppression, we observed a similar pattern as in our previous tests, namely that the outer skull border seemed to be well delineated whereas the inner skull border was recovered less successfully. Testing SPM12 on this data gave far worse results than obtained for our chosen MRI parameters, even when compared to the results obtained without an additional T2w image. To confirm our prior experience, we also tested FSL BET2 on this data, which resulted in a segmentation where large parts of the spongy bone were mislabeled as skin. This highlights the importance of testing the segmentation performance on some pilot data sets prior to the start of a study. We do not claim that our chosen parameters are the only possible or optimal choice, but hope that they are a helpful guidance for the reader. Some of the problems related to bone modeling could be resolved by using a combination of CT and MR scans (Eichelbaum et al., 2014). The good bone contrast from CT allows to distinguish small details, such as small openings due to blood vessels, and would likely help with modeling the vertebrae. However, acquiring a CT scan in addition to the MRI requires access to additional equipment, and results in an increased workload and exposure to radiation. Moving to ultra high-field MRI (7T) could facilitate the modeling of small blood vessels or openings in the skull (Fiederer et al., 2016), but access to high-field scanners is often limited in practice.

We tested the impact of changes of the image downsampling factor (the default setting in SPM12 is downsampling the resolution of the input image by a factor of three) and the regularization parameter for the bias field correction in SPM12 on the segmentation accuracy. While the segmentations were robust to changes in the downsampling factor, optimization of the regularization parameter helped to ensure better segmentations of the spinal cord. This was only necessary for the T2w images as they suffered more from intensity inhomogeneities in the lower part than the T1w image. The high flip angles employed in T2w images likely render this image type more susceptible to spatially varying flip angles at the lower boundary of the RF excitation coil. Therefore, the incorrectly estimated bias field when using the standard regularization parameter setting caused the intensity profiles of the tissues in the bottom part (e.g., the spinal cord) to be different from that in the upper part (e.g., cortical white matter), so that spinal cord voxels were assigned a low probability of being WM, and consequently were classified as scalp. However, as the same MRI hardware and image parameters were used for all subjects, the selection of optimized regularization parameter had to be done only once by testing a few subjects.

Our results are broadly in line with those presented in (Perdue and Diamond, 2014) for the upper part of the head. In that study, a fair segmentation performance of FSL and SPM8 for the skull was

documented when using a T1w image as input. We further replicated their finding that BrainSuite offers a lower segmentation performance, which is in contrast to the original results for BrainSuite presented in (Dogdas et al., 2005). We also demonstrated the relevance of using a combination of T1w and T2w images to improve the segmentation accuracy. In addition, we based our validations on individual CT-based skull segmentations as ground truth, tested a new version of SPM and included spatially extended tissue priors that covered the neck.

Whereas automatic skull segmentation was the focus of this study, we further demonstrated that the CAT12 toolbox in SPM12 achieves GM segmentations similar to the ones produced by FreeSurfer 5.3.0, thus confirming the results of Dahnke et al., 2013. As a result, the segmentations based on SPM12 combined with CAT12 enabled automatic construction of volume conductor models that had good anatomical accuracy in the upper part of the head (in particular when a T2w image was used), and included the coarse features of the neck and vertebrae. Specifically, accurately modeling the skull and the CSF-GM and GM-WM boundaries have been suggested to be important for the realistic estimation of volume conduction effects in the head (Vorwerk et al., 2014). In addition, for EEG and MEG, the reconstructed cortical surface is often considered important for determining the orientation of the neural dipole sources and for spanning the source space. In practice, including Figs. 4 and 5, we define the cortical surface to be the mid layer between the white matter and pial surfaces. The placement of this surface, as long as it is not right at the white matter or pial boundary, should not have a large effect on the resulting field strength as the electric field is discontinuous only right at the tissue boundaries. We exemplarily demonstrated the importance of a spatially extended head model for electric field estimations of tDCS montages that involve extraencephalic electrodes. Similarly, for EEG source modeling, cutting the model right below the cerebellum (and artificially closing the bone compartment) has been shown to result in non-negligible errors in particular for deep sources (Indahlastari et al., 2016; Lanfer et al., 2012). Thus, extended volume conductor models including the neck should help to improve the accuracy of the calculated fields and sources.

Despite these advantages, the SPM12-based segmentations still have clear limitations. Large parts of the air cavities are wrongly segmented as bone. However, as both bone and air have a low conductivity, this might not have a strong impact on the accuracy of the resulting field calculations unless the erroneously labeled voxels are very close to the brain area of interest (Indahlastari et al., 2016; Lanfer et al., 2012). While modeling of the spinal cord and the vertebrae likely allows for a more accurate estimation of the current pathways in the cerebellum that is positioned mostly superior to these structures, underestimation of the vertebrae renders the calculated fields in the spinal cord itself inaccurate. None of the segmentation methods tested here currently distinguish between spongy and compact skull. Given their different conductivities, distinguishing these tissues might be beneficial to improve the accuracy of the field estimates (Dannhauer et al., 2011), though using an optimized bone conductivity may also suffice (Cho et al., 2015; Dannhauer et al., 2011; Huang et al., 2017; Vorwerk et al., 2014). Regarding the SPM12-based segmentations, the explicit modeling of spongy bone, including a separate spatial prior, might also help to improve the robustness of the skull segmentation by reducing the extent to which spongy bone is mislabeled as, for example, scalp. Finally, scalp is currently a lumped representation of skin, subdural fat and muscle tissue, rendering the reliable assignment of a single conductivity value to this tissue class difficult. In particular, this is the case for the segmented scalp

in the neck region, in which fat and muscle tissue are spatially intermingled. In contrast, the scalp in the upper part of the head mostly consists of a regularly layered structure of skin and subdural fat with only little muscle tissue. This again suggests that the field estimates in the upper part of the head model will likely be more accurate than in the neck region when a spatially extended head model is used. However, estimations of the size of the error caused by this simplification would require a systematic sensitivity analysis. It is worth noting in this respect that the anatomical accuracy of volume conductor models is only one, albeit important, aspect to obtain accurate forward models of the electric field distributions. In addition, the applied conductivity values are average values taken from a rather sparse literature, in which the reported values often vary considerably between studies (e.g., Gabriel et al., 1996; Geddes and Baker, 1967). This limitation adds to the uncertainty of the field estimates, and the inter-subject variabilities in the tissue conductivities which have been observed (Dabek et al., 2016; Huang et al., 2017) suggest the need for subject-specific calibrations, either through dedicated measurements (Dabek et al., 2016) or optimization procedures (Aydin et al., 2014; Huang et al., 2017).

Conclusion

In summary, our study demonstrates the current state-of-the-art of automatic skull segmentation from MR images, including the identification of remaining shortcomings, and introduces a novel, easily accessible and validated open-source tool for the automatic creation of volume meshes of the complete head. We have compared three methods (FSL BET2, BrainSuite skullfinder, and the unified segmentation routine of SPM12 with extended spatial tissue priors) to automatically segment the human skull. We conclude that FSL and SPM12 achieve reasonable results for the upper part of the skull when a combination of T1w and T2w images is used as input and suited MRI sequence parameters are chosen. In that case, the skull segmentation of SPM12 seems to be moderately better than that of FSL BET2, in particular when combined with a postprocessing step based on simple morphological operations to clean up local defects. In contrast to FSL BET2, the SPM12-based segmentation also covers the lower part of the head including the neck and recovers details of the facial bones. For EEG and TDCS, the skull has a strong influence on the electric field distribution occurring in the brain and on the scalp surface due to its low conductivity. Employing an extended volume conductor model that builds upon the SPM12-based segmentation can thus help to improve the accuracy of the forward models. However, the reconstructed vertebrae suffer from a low anatomical accuracy so that the volume conductor models are still too coarse to allow for reliable estimation of the electric field in the spinal cord. The validated tools (FSL and SPM12) are integrated into our open-source pipeline SimNIBS (www.simnibs.org) that allows for the automatic generation of highquality tetrahedral head meshes and their subsequent use in field calculations for brain stimulation. The released version supports head meshes using the FSL-based skull segmentation, while the SPM12-based segmentation will be added in the upcoming version.

Acknowledgements

This study was supported by the Lundbeck foundation (grant R118-A11308 to AT and grant R59-A5399 - PI Hartwig Siebner), the Novonordisk foundation (grant no. NNF14OC0011413) and a PhD stipend of the Sino-Danish Center to JDN.

APPENDIX

Table A1

List of the main post-processing operations applied to the different tissues. *Dilation* flips background voxels to foreground voxels if any of the neighboring voxels, defined by the kernel, is a foreground voxel. *Erosion* removes foreground voxels if any of the neighboring voxels, defined by the kernel, is background voxel. *Closing* is defined as a dilation followed by an erosion (Gonzalez and Woods, 2002).

Tissue	Operation	Kernel	Explanation
Bone	Closing	Six closest neighbors (6-connectivity)	Remove holes from the bone structure by closing.
Skin	Erosion	Six closest neighbors (6-connectivity)	Erode to remove small erroneously segmented skin areas.
Skin	Keep largest component	Six closest neighbors (6-connectivity)	Find the largest connected component, and discard the rest, which are considered noise.
Skin	Dilation	Six closest neighbors (6-connectivity)	Dilate the largest component to account for the initial erosion.
All tissues	Iterative Gaussian smoothing (sigma $=$ 1) and assignment of voxels.	N/A	After the post-processing steps, some of the voxels might be unassigned. Iteratively smooth the binary masks, and assign unassigned voxels.

References

- Ashburner, J., Friston, K.J., 2005. Unified segmentation. Neuroimage 26, 839–851. https://doi.org/10.1016/j.neuroimage.2005.02.018.
- Aydin, Ü., Vorwerk, J., Küpper, P., Heers, M., Kugel, H., Galka, A., Hamid, L., Wellmer, J., Kellinghaus, C., Rampp, S., Wolters, C.H., 2014. Combining EEG and MEG for the reconstruction of epileptic activity using a calibrated realistic volume conductor model. PLoS One 9, e93154. https://doi.org/10.1371/journal.pone.0093154.
 Baumann, S.B., Wozny, D.R., Kelly, S.K., Meno, F.M., 1997. The electrical conductivity of
- Baumann, S.B., Wozny, D.R., Kelly, S.K., Meno, F.M., 1997. The electrical conductivity of human cerebrospinal fluid at body temperature. IEEE Trans. Biomed. Eng. 44, 220–223. https://doi.org/10.1109/10.554770.
- Birot, G., Spinelli, L., Vulliémoz, S., Mégevand, P., Brunet, D., Seeck, M., Michel, C.M., 2014. Head model and electrical source imaging: a study of 38 epileptic patients. NeuroImage Clin. 5, 77–83. https://doi.org/10.1016/j.nicl.2014.06.005.
- Chan, T.F., Vese, L.A., 2001. Active contours without edges. IEEE Trans. Image Process 10, 266–277. https://doi.org/10.1109/83.902291.
- Cho, J.H., Vorwerk, J., Wolters, C.H., Knösche, T.R., 2015. Influence of the head model on EEG and MEG source connectivity analyses. Neuroimage 110, 60–77. https:// doi.org/10.1016/j.neuroimage.2015.01.043.
- Dabek, J., Kalogianni, K., Rotgans, E., van der Helm, F.C.T., Kwakkel, G., van Wegen, E.E.H., Daffertshofer, A., de Munck, J.C., 2016. Determination of head conductivity frequency response in vivo with optimized EIT-EEG. Neuroimage 127, 484-495. https://doi.org/10.1016/j.neuroimage.2015.11.023.
- Dahnke, R., Yotter, R.A., Gaser, C., 2013. Cortical thickness and central surface estimation. Neuroimage 65, 336–348. https://doi.org/10.1016/ j.neuroimage.2012.09.050.
- Dale, A.M., Fischl, B., Sereno, M.I., 1999. Cortical surface-based analysis: I. Segmentation and surface reconstruction. Neuroimage 9, 179–194. https://doi.org/10.1006/ nimg.1998.0395.
- Dannhauer, M., Lanfer, B., Wolters, C.H., Knösche, T.R., 2011. Modeling of the human skull in EEG source analysis. Hum. Brain Mapp. https://doi.org/10.1002/ hbm.21114.
- Datta, A., Baker, J.M., Bikson, M., Fridriksson, J., 2011. Individualized model predicts brain current flow during transcranial direct-current stimulation treatment in responsive stroke patient. Brain Stimul. 4, 169–174. https://doi.org/10.1016/ j.brs.2010.11.001.
- Dogdas, B., Shattuck, D.W., Leahy, R.M., 2005. Segmentation of skull and scalp in 3-D human MRI using mathematical morphology. Hum. Brain Mapp. 26, 273–285. https://doi.org/10.1002/hbm.20159.
- Dubuisson, M.-P., Jain, A.K., 1994. A modified Hausdorff distance for object matching. Proc. 12th Int. Conf. Pattern Recognit. 1, 566–568. https://doi.org/10.1109/ ICPR.1994.576361.
- Eggert, L.D., Sommer, J., Jansen, A., Kircher, T., Konrad, C., 2012. Accuracy and reliability of automated gray matter segmentation pathways on real and simulated structural magnetic resonance images of the human brain. PLoS One 7. https:// doi.org/10.1371/journal.pone.0045081.
- Eichelbaum, S., Dannhauer, M., Hlawitschka, M., Brooks, D., Knösche, T.R., Scheuermann, G., 2014. Visualizing simulated electrical fields from electroencephalography and transcranial electric brain stimulation: a comparative evaluation. Neuroimage 101, 513–530. https://doi.org/10.1016/ j.neuroimage.2014.04.085.
- Fiederer, L.D.J., Vorwerk, J., Lucka, F., Dannhauer, M., Yang, S., Dümpelmann, M., Schulze-Bonhage, A., Aertsen, A., Speck, O., Wolters, C.H., Ball, T., 2016. The role of blood vessels in high-resolution volume conductor head modeling of EEG. Neuroimage 128, 193–208. https://doi.org/10.1016/j.neuroimage.2015.12.041.
- Fischl, B., Sereno, M.I., Dale, A.M., 1999. Cortical surface-based analysis: II: inflation, flattening, and a surface-based coordinate system. Neuroimage 9, 195–207. https:// doi.org/10.1006/nimg.1998.0396.
- Gabriel, C., Gabriel, C., Gabriel, S., Gabriel, S., Corthout, E., Corthout, E., 1996. The dielectric properties of biological tissues: I. Literature survey. Phys. Med. Biol. 41, 2231–2249. https://doi.org/10.1088/0031-9155/41/11/001.

- Geddes, L.A., Baker, I.E., 1967. The specific resistance of biological material a compendium of data for the biomedical engineer and physiologist. Med. Biol. Eng. 5, 271–293.
- Geuzaine, C., Remacle, J.-F., 2009. Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities. Int. J. Numer. Methods Eng. 79, 1309–1331.
- Gonzalez, R., Woods, R., 2002. Digital Image Processing. Prentice Hall. https://doi.org/ 10.1016/0734-189X(90)90171-Q.
- Han, X., Jovicich, J., Salat, D., van der Kouwe, A., Quinn, B., Czanner, S., Busa, E., Pacheco, J., Albert, M., Killiany, R., Maguire, P., Rosas, D., Makris, N., Dale, A., Dickerson, B., Fischl, B., 2006. Reliability of MRI-derived measurements of human cerebral cortical thickness: the effects of field strength, scanner upgrade and manufacturer. Neuroimage 32, 180–194. https://doi.org/10.1016/ j.neuroimage.2006.02.051.
- Huang, Y., Dmochowski, J.P., Su, Y., Datta, A., Rorden, C., Parra, L.C., 2013. Automated MRI segmentation for individualized modeling of current flow in the human head. J. Neural Eng. 10 https://doi.org/10.1088/1741-2560/10/6/066004, 66004.
- Huang, Y., Liu, A.A., Lafon, B., Friedman, D., Dayan, M., Wang, X., Bikson, M., Doyle, W.K., Devinsky, O., Parra, L.C., 2017. Measurements and models of electric fields in the in vivo human brain during transcranial electric stimulation. Elife 6. https://doi.org/10.7554/eLife.18834.
- Indahlastari, A., Chauhan, M., Schwartz, B., Sadleir, R.J., 2016. Changing head model extent affects finite element predictions of transcranial direct current stimulation distributions. J. Neural Eng. 13 https://doi.org/10.1088/1741-2560/13/6/066006, 66006.
- Jenkinson, M., Bannister, P., Brady, M., Smith, S., 2002. Improved optimisation for the robust and accurate linear registration and motion correction of brain images. Neuroimage 17, 825–841. https://doi.org/10.1016/S1053-8119(02)91132-8.
- Jenkinson, M., Smith, S., 2001. A global optimisation method for robust affine registration of brain images. Med. Image Anal. 5, 143–156. https://doi.org/10.1016/ \$1361-8415(01)00036-6.
- Kamnitsas, K., Ledig, C., Newcombe, V.F.J., Simpson, J.P., Kane, A.D., Menon, D.K., Rueckert, D., Glocker, B., 2017. Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. Med. Image Anal. 36, 61–78. https:// doi.org/10.1016/j.media.2016.10.004.
- King, K.F., 2004. Operator parameter selection. In: BERNSTEIN, M.A., KING, K.F., ZHOU, X.J. (Eds.), Handbook of MRI Pulse Sequences. Elsevier Academic Press, Amsterdam, pp. 374–377.
- Klein, S., Staring, M., Murphy, K., Viergever, M.A., Pluim, J.P.W., 2010. Elastix: a toolbox for intensity-based medical image registration. IEEE Trans. Med. Imaging 29, 196–205. https://doi.org/10.1109/TMI.2009.2035616.
- Lanfer, B., Paul-Jordanov, I., Wolters, C., 2014. Individual FEM pipeline for EEG source analysis requiring minimal user interventio. In: Proceedings of the Annual Meetin of the Organization for Human Brain Mapping. Hamburg, Germany, p. 1657.
- Lanfer, B., Scherg, M., Dannhauer, M., Knösche, T.R., Burger, M., Wolters, C.H., 2012. Influences of skull segmentation inaccuracies on EEG source analysis. Neuroimage 62, 418–431. https://doi.org/10.1016/j.neuroimage.2012.05.006.
- Lewiner, T., Lopes, H., Vieira, A.W., Tavares, G., 2003. Efficient implementation of marching cubes' cases with topological guarantees. J. Graph. Tools 8, 1–15. https:// doi.org/10.1080/10867651.2003.10487582.
 Li, K., Papademetris, X., Tucker, D.M., 2016. BrainK for structural image processing:
- Li, K., Papademetris, X., Tucker, D.M., 2016. BrainK for structural image processing: creating electrical models of the human head. Comput. Intell. Neurosci. 2016 https:// doi.org/10.1155/2016/1349851, 1349851.
- Liu, A., Joe, B., 1994. Relationship between tetrahedron shape measures. BIT 34, 268–287. https://doi.org/10.1007/BF01955874.
- Menze, B.H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., Lanczi, L., Gerstner, E., Weber, M.A., Arbel, T., Avants, B.B., Ayache, N., Buendia, P., Collins, D.L., Cordier, N., Corso, J.J., Criminisi, A., Das, T., Delingette, H., Demiralp, Ç., Durst, C.R., Dojat, M., Doyle, S., Festa, J., Forbes, F., Geremia, E., Glocker, B., Golland, P., Guo, X., Hamamci, A., Iftekharuddin, K.M., Jena, R., John, N.M., Konukoglu, E., Lashkari, D., Mariz, J.A., Meier, R., Pereira, S., Precup, D., Price, S.J., Raviv, T.R., Reza, S.M.S., Ryan, M.,

J.D. Nielsen et al.

Sarikaya, D., Schwartz, L., Shin, H.C., Shotton, J., Silva, C.A., Sousa, N., Subbanna, N.K., Szekely, G., Taylor, T.J., Thomas, O.M., Tustison, N.J., Unal, G., Vasseur, F., Wintermark, M., Ye, D.H., Zhao, L., Zhao, B., Zikic, D., Prastawa, M., Reyes, M., Van Leemput, K., 2015. The multimodal brain tumor image segmentation benchmark (BRATS). IEEE Trans. Med. Imaging 34, 1993–2024. https://doi.org/ 10.1109/TMI.2014.2377694.

Minjoli, S., Saturnino, G.B., Blicher, J.U., Stagg, C.J., Siebner, H.R., Antunes, A., Thielscher, A., 2017. The impact of large structural brain changes in chronic stroke patients on the electric field caused by transcranial brain stimulation. NeuroImage Clin. 15, 106–117. https://doi.org/10.1016/j.nicl.2017.04.014.

- Montes-Restrepo, V., Van Mierlo, P., Strobbe, G., Staelens, S., Vandenberghe, S., Hallez, H., 2014. Influence of skull modeling approaches on EEG source localization. Brain Topogr. 27, 95–111. https://doi.org/10.1007/s10548-013-0313-y.
- Oostenveld, R., Fries, P., Maris, E., Schoffelen, J.M., 2011. FieldTrip: open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. Comput. Intell. Neurosci. https://doi.org/10.1155/2011/156869, 2011.
- Pechaud, M., Jenkinson, M., Smith, S., 2006. BET2 MRI-Based Estimation of Brain, Skull and Scalp Surfaces FMRIB Technical Report TR06MP1. Report XXXIII, pp. 81–87.
- Perdue, K.L., Diamond, S.G., 2014. T1 magnetic resonance imaging head segmentation for diffuse optical tomography and electroencephalography. J. Biomed. Opt. 19 https:// doi.org/10.1117/1.JBO.19.2.026011, 26011.
- Saturnino, G.B., Antunes, A., Thielscher, A., 2015. On the importance of electrode parameters for shaping electric field patterns generated by tDCS. Neuroimage 120, 25–35. https://doi.org/10.1016/j.neuroimage.2015.06.067.
- Shattuck, D.W., Leahy, R.M., 2002. Brainsuite: an automated cortical surface identification tool. Med. Image Anal. 6, 129–142. https://doi.org/10.1016/S1361-8415(02)00054-3.
- Shattuck, D.W., Sandor-Leahy, S.R., Schaper, K.A., Rottenberg, D.A., Leahy, R.M., 2001. Magnetic resonance image tissue classification using a partial volume model. Neuroimage 13. 856–876. https://doi.org/10.1006/nimg.2000.0730.
- Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E.J., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobnjak, I., Flitney, D.E., Niazy, R.K., Saunders, J., Vickers, J., Zhang, Y., De Stefano, N., Brady, J.M., Matthews, P.M., 2004. Advances in functional and structural MR image analysis and implementation as FSL. In: NeuroImage. https://doi.org/10.1016/j.neuroimage.2004.07.051.
- Stenroos, M., Hunold, A., Haueisen, J., 2014. Comparison of three-shell and simplified volume conductor models in magnetoencephalography. Neuroimage 94, 337–348. https://doi.org/10.1016/j.neuroimage.2014.01.006.
- Stenroos, M., Nummenmaa, A., 2016. Incorporating and compensating cerebrospinal fluid in surface-based forward models of magneto- and electroencephalography. PLoS One 11. https://doi.org/10.1371/journal.pone.0159595 e0159595.

- Tadel, F., Baillet, S., Mosher, J.C., Pantazis, D., Leahy, R.M., 2011. Brainstorm: a userfriendly application for MEG/EEG analysis. Comput. Intell. Neurosci. 1–13. https:// doi.org/10.1155/2011/879716, 2011.
- Thielscher, A., Antunes, A., Saturnino, G.B., 2015. Field modeling for transcranial magnetic stimulation: a useful tool to understand the physiological effects of TMS? Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. EMBS 222–225. https://doi.org/ 10.1109/EMBC.2015.7318340.
- van der Walt, S., Schönberger, J.L., Nunez-Iglesias, J., Boulogne, F., Warner, J.D., Yager, N., Gouillart, E., Yu, T., 2014. scikit-image: image processing in Python. PeerJ 2. https://doi.org/10.7717/peerj.453 e453.
- Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P., 1999. Automated model-based bias field correction of MR images of the brain. IEEE Trans. Med. Imaging 18, 885–896. https://doi.org/10.1109/42.811268.

Vorwerk, J., Cho, J.H., Rampp, S., Hamer, H., Knösche, T.R., Wolters, C.H., 2014. A guideline for head volume conductor modeling in EEG and MEG. Neuroimage 100, 590–607. https://doi.org/10.1016/j.neuroimage.2014.06.040.

- Wells, W.M., Grimson, W.L., Kikinis, R., Jolesz, F. a, 1996. Adaptive segmentation of MRI data. IEEE Trans. Med. Imaging 15, 429–442. https://doi.org/10.1109/42.511747.
- Windhoff, M., Opitz, A., Thielscher, A., 2013. Electric field calculations in brain stimulation based on finite elements: an optimized processing pipeline for the generation and usage of accurate individual head models. Hum. Brain Mapp. 34, 923–935. https://doi.org/10.1002/hbm.21479.
- Wolters, C.H., Anwander, A., Tricoche, X., Weinstein, D., Koch, M.A., MacLeod, R.S., 2006. 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, 813–826. https:// doi.org/10.1016/j.neuroimage.2005.10.014.
- Zhang, Y., Bajaj, C., Sohn, B.-S., 2005. 3D finite element meshing from imaging data. Comput. Methods Appl. Mech. Eng. 194, 5083–5106. https://doi.org/10.1016/ j.cma.2004.11.026.
- Zhang, Y., Matuszewski, B.J., Shark, L.-K., Moore, C.J., 2008. Medical image segmentation using new hybrid level-set method. In: 2008 Fifth International Conference BioMedical Visualization: Information Visualization in Medical and Biomedical Informatics, pp. 71–76. https://doi.org/10.1109/MediVis.2008.12.
- Ziegler, E., Chellappa, S.L., Gaggioni, G., Ly, J.Q.M., Vandewalle, G., André, E., Geuzaine, C., Phillips, C., 2014. A finite-element reciprocity solution for EEG forward modeling with realistic individual head models. Neuroimage 103, 542–551. https:// doi.org/10.1016/j.neuroimage.2014.08.056.
Appendix B

Evaluating the Influence of Anatomical Accuracy and Electrode Positions on EEG Forward Solutions

This manuscript was published on bioRxiv on September 4, 2022, and submitted to NeuroImage for review on September 4, 2022.

Evaluating the Influence of Anatomical Accuracy and Electrode Positions on **EEG** Forward Solutions

Jesper Duemose Nielsen^{a,b,c,*}, Oula Puonti^b, Rong Xue^{g,e,f}, Axel Thielscher^{d,b}, Kristoffer Hougaard Madsen^{a,b}

^aDepartment of Applied Mathematics and Computer Science, Technical University of Denmark, Denmark

^bDanish Research Centre for Magnetic Resonance, Centre for Functional and Diagnostic Imaging and Research, Copenhagen University Hospital Amager and Hvidovre, Denmark

^cSino-Danish Centre for Education and Research, Aarhus, Denmark ^dDepartment of Health Technology, Technical University of Denmark, Kgs. Lyngby, Denmark

^eState Key Laboratory of Brain and Cognitive Science, Beijing MRI Center for Brain Research, Institute of Biophysics,

^fBeijing Institute for Brain Disorders, Beijing, China

^gUniversity of Chinese Academic of Sciences, Beijing, China

Abstract

Generating realistic volume conductor models for forward calculations in electroencephalography (EEG) is not trivial and several factors contribute to the accuracy of such models, two of which are its anatomical accuracy and the accuracy with which electrode positions are known. Here, we investigate effects of anatomical accuracy by comparing forward solutions from SimNIBS, a tool which allows state-of-the-art anatomical modeling, with well-established pipelines in MNE-Python and FieldTrip. We also compare different ways of specifying electrode locations when digitized positions are not available such as transformation of measured positions from standard space and transformation of a manufacturer layout.

Substantial effects of anatomical accuracy were seen throughout the entire brain both in terms of field topography and magnitude with SimNIBS generally being more accurate than the pipelines in MNE-Python and FieldTrip. Topographic and magnitude effects were particularly pronounced for MNE-Python which uses a three-layer boundary element method (BEM) model. We attribute these mainly to differences in skull and cerebrospinal fluid (CSF) modeling. Effects of electrode specification method were evident in occipital and posterior areas when using a transformed manufacturer layout whereas transforming measured positions from standard space generally resulted in smaller errors.

We suggest modeling the anatomy of the volume conductor as accurately possible and we hope to facilitate this by making it easy to export simulations from SimNIBS to MNE-Python and FieldTrip for further analysis. Likewise, if digitized electrode positions are not available, a set of measured positions on a standard head template may be preferable to those specified by the manufacturer.

Keywords: electroencephalography, forward model, head model, source localization

1. Introduction

Electroencephalography (EEG) can be used to analyze brain activity with high temporal resolution. The measurements consist of potential differences observed on the scalp which directly reflect the synchronous activity of a large body of pyramidal cells oriented perpendicular to the cortical surface (Kirschstein and Köhling, 2009). However, the relationship between measurements and brain activity is complicated by the fact that the former is a spatially low-pass filtered representation of the latter where the filter consists of

Chinese Academy of Sciences, Beijing, China

^{*}Corresponding author

the tissues separating the neural sources from the sensors (Buzsáki et al., 2012). Due to volume conduction, all sources in the brain affect the potential at all electrodes simultaneously.

In some situations, it may be beneficial to analyze the data in the source domain rather than the sensor domain, e.g., when investigating the connectivity between different brain regions (Schoffelen and Gross, 2009; Mahjoory et al., 2017; Nguyen-Danse et al., 2021) or when localizing spike activity in epilepsy (Kaiboriboon et al., 2012; van Mierlo et al., 2020). Reconstructing the neural generators of an observed EEG signal is an inverse problem. In order to solve this, one first needs to solve the corresponding forward problem which consists of estimating the potential distribution on the scalp due to a neural source placed arbitrarily in the brain.

To solve the forward problem in EEG, one typically starts by constructing a physical model of the head, the volume conductor model, which enables simulation of volume conduction effects due to activity in the brain. Naturally, one would expect the accuracy of this volume conductor model to impact the accuracy with which such effects can be modeled (Vorwerk et al., 2014).

In this work, we investigate some of the elements which affect the accuracy of the forward solution, specifically, the anatomical accuracy of the head model and the accuracy with which electrode positions are known. We compare existing methods with SimNIBS which is able to generate state-of-the-art head models automatically based on magnetic resonance imaging (MRI) scans. Our aim is to explore the extent to which more accurate modeling leads to improvements in the forward solution and perhaps also in source localization accuracy (although the latter is not investigated here).

Numerous methods have been employed to solve the EEG forward problem. In simple geometries, e.g., nested sphere models, (quasi-)analytical solutions can be derived, however, for realistically shaped models numerical methods are needed. In the boundary element method (BEM), it is common to use a three-layer surface model (consisting of inner skull, outer skull, and skin) of relatively low resolution due to the dense nature of the problem¹ whereas high-resolution, multi-compartment models are often seen in the finite element method (FEM). However, the sparsity of the FEM system means that it is generally inferior to BEM in terms of numerical accuracy for geometries where BEM is applicable (Hallez et al., 2007; Vorwerk et al., 2012).

One important aspect of such a model is its anatomical accuracy. In a systematic evaluation of the effect of distinguishing different tissue compartments, Vorwerk et al. (2014) concluded that modeling cerebrospinal fluid (CSF) is important as well as distinguishing white matter and gray matter, however, explicit modeling of spongy bone as well as the anisotropic conductivity of white matter were found to be less important. Using a similar strategy, Azizollahi et al. (2016) corroborated these results in neonates.

The skull compartment has received special attention in the literature due to its importance in shaping the observed fields in EEG (Hämäläinen et al., 1993). Lanfer et al. (2012) investigated the effect of several geometrical simplifications and errors on the forward solution, concluding that localized modeling errors (e.g., skull holes, erroneous thickness) result in forward errors mostly in the vicinity of such geometrical inaccuracies whereas simplifications of a more general nature (e.g., cutting the model at the base of the skull or approximating the base of the skull with constant thickness) show increased errors for large array of positions. Similar findings were reported by Chauveau et al. (2004) and Li et al. (2007). Forward modeling errors have also been observed when not modeling skull openings (Fiederer et al., 2016). Errors in skull thickness were also investigated by Chauveau et al. (2004), who found sources close to the skull to be most affected, however, their results also demonstrated that errors were seen predominantly on magnitudes and not topography. Several of these studies also report increased dipole localization errors in the vicinity of such forward modeling errors (Chauveau et al., 2004; Fiederer et al., 2016; Lanfer et al., 2012) suggesting that these effects indeed translate to errors in source localization in EEG.

To construct anatomically accurate volume conductor models, images from one or more structural sequences (e.g., T1- and T2-weighted MRI or in rare cases even computed tomography (CT)) are used. However, they may not always be available. To reconstruct sources without such anatomical information, one will have to rely on some kind of average anatomy. This was investigated by Acar and Makeig (2013)

¹Efficient ways of solving the system, e.g., the fast multipole method, exist allowing for high resolution BEM models as well.

who compared individualized head models to models based on template anatomy and found median source localization errors of about 5 mm, particularly in the inferior part of the brain, since the template model was cut at the base of the skull.

Taken together, it seems that the anatomical accuracy of the forward model plays an important role in shaping not only the forward solution but likely also the final source reconstruction. FEM allows us to specify geometrically complex models, however, realistic modeling of the head geometry is difficult and has so far not been easily accessible to the EEG community. SimNIBS (Thielscher et al., 2015) is a tool for simulating electrical fields in the brain due to noninvasive brain stimulation which can construct realistic anatomical volume conductor models of the human head with reasonable accuracy from MRI scans as validated against manual segmentations based on MRI and CT (Nielsen et al., 2018; Puonti et al., 2020). Given the intimate relationship between the EEG forward problem and the problem of simulating the effect of transcranial electrical stimulation (TES)—related by reciprocity (Wolters et al., 2004; Ruffini, 2015)—SimNIBS can also be used to solve the EEG forward problem and we explore this in further detail in the current work.

Another important aspect of the forward model is how electrode positions are estimated and coregistered to the anatomical model. Different systems to acquire electrode positions exist (e.g., the Polhemus Fastrak system, https://polhemus.com), however, if such systems are not available, one may have to use a template description of the electrode positions. To this end, EEG cap manufacturers typically provide spherical angles of the different electrodes which can then be mapped onto a sphere of a radius corresponding to the size of the subject head.

Dalal et al. (2014) found that localization accuracy and output signal-to-noise ratio (SNR) of a linearly constrained minimum variance beamformer was impacted by electrode digitization technique and coregistration method. Another study by Wang and Gotman (2001) found only minimal effect of electrode position errors while simulating displacements individually for each electrode. This may be a realistic model of pure measurement errors in a scenario where positions are digitized. However, errors due to coregistration between MRI scans and digitized electrode positions or the use of template positions, might be expected to be more spatially correlated, thus introducing a general bias in electrode positions (e.g., due to tilting, stretching etc.) (Acar and Makeig, 2013; Homölle and Oostenveld, 2019). Homölle and Oostenveld (2019) showed that digitizing electrodes using a structured-light 3D scanner resulted in slightly smaller errors compared to transformation of a custom template which again performed better than simply using the manufacturer template positions. Errors were also evident in the forward solution and a subsequent dipole fit. The errors of the latter reached 30 mm over the parietal cortex for the manufacturer template coinciding with the largest spatial bias in electrode positions. Acar and Makeig (2013) simulated coregistration errors by tilting electrode positions and found increased source localization errors (on average 5 mm) affecting predominantly superficial sources. Thus, we expect the accuracy with which electrode positions are determined to affect not only the forward model but also subsequent source estimates.

Here we explicitly consider how these two aspects of the forward model, namely its anatomical accuracy as well as how electrode positions are estimated, affect the forward calculations. The geometrical specification of the model will, to a large extent, be determined by the particular choice of forward modeling pipeline. Consequently, we first compare pipelines from different software packages for generating such models. This will result in aggregated effects from various sources but also provide practically relevant insights into the performance of the tested pipelines. Specification of electrode locations will depend on the available equipment and resources. Since it may not always be possible to obtain individual estimates during the experimental session, one may have to resort to a set of template positions. Thus, we compare different template layouts to determine their effect on the final solution in the second part of the study.

Finally, although not the focus here, we would like to point out that several other parameters of the forward model are important as well, such as the influence of the conductivity profile of different tissues. The conductivity of some tissues in the human head are fairly well known (e.g., that of CSF) (Baumann et al., 1997) whereas that of others (e.g., compact and spongy bone) (Ümit Aydin et al., 2014) generally show large variability across studies. Even though this is known to affect modeling accuracy (Bangera et al., 2010; Dannhauer et al., 2011; Marin et al., 1998; McCann and Beltrachini, 2022; Vallaghe and Clerc, 2009; Vorwerk et al., 2019), one will typically have to resort to standard values from the literature although calibration paradigms have also been suggested with (Ümit Aydin et al., 2014) or without (Acar et al., 2016)

the use of magnetoencephalography. Finally, model resolution is also important (Kaipio and Somersalo, 2007). In most cases, though, researchers will probably rely on the default settings of a particular pipeline (as determined from validation experiments).

We start by evaluating the numerical accuracy of SimNIBS in a spherical model and a realistic three-layer head model. Next, we compare forward solutions from SimNIBS with standard pipelines in MNE-Python and FieldTrip. We also include a model based on standard anatomy representing a scenario where electrode positions are known but the anatomy is not. As we will see, the pipelines differ substantially in terms of anatomical accuracy. Finally, we compare different ways of specifying electrode positions when digitized positions are not available such as transformation of positions from standard space and transformation of manufacturer layout positions. As in the first sub-study, we also include a model based on standard anatomy.

2. Methods

2.1. Calculating EEG Forward Solutions in SimNIBS

Starting from the quasi-static approximation of the Maxwell equations, the EEG forward problem takes the form

$$-\nabla \cdot (\sigma \nabla u) = -\nabla \cdot \mathbf{j}_s = f \tag{1}$$

subject to boundary conditions

$$u_i = u_j \, \ln \, \Omega \tag{2}$$

$$\langle \mathbf{j}_i, \mathbf{n} \rangle = 0 \quad \text{on } \partial \Omega$$
(3)

where *i* and *j* are two neighboring elements (e.g., tetrahedra) and **n** is the vector normal to their interface. $\langle \mathbf{j}, \mathbf{n} \rangle$ is the projection of the current density, **j**, in the normal direction, and Ω and $\partial \Omega$ refer to the interior of the domain and its surface, respectively. Equation (2) ensures that the solution is continuous across interfaces of neighboring elements whereas equation (3) states that no current can leave the domain. We are interested in finding *u*, the potential distribution, for a given *f* (assumed to be known), which can be interpreted as a monopolar source configuration, i.e., the source of the current density, **j**_s (Hallez et al., 2007).

Equation (1) states that the divergence of the return current density (left-hand side) is equal to the divergence of the source current density. This is important for EEG source analysis as we measure potential differences on the scalp due to the return currents induced by the primary currents (i.e., the neural sources) along with other effects of non-interest (e.g., respiration, heartbeat).

Due to the principle of reciprocity, the EEG forward problem is related to the estimation of electric fields induced by noninvasive brain stimulation. Specifically, the potential between two points, a and b, due to a dipolar source at a certain position is related to the electric field at that position resulting from running a current between a and b (Weinstein et al., 2000). SimNIBS is a software package for solving the latter problem. To generate so-called leadfields for EEG, we exploit the intimate relationship between the EEG forward problem (estimating scalp potentials from current sources in the brain) and TES (estimating the electric field resulting from applying a potential difference between two electrodes or, equivalently, applying a current at one electrode and removing it at another). SimNIBS returns a smooth estimate of the electric field from which the gain (or leadfield) matrix can be computed. We describe how this is achieved below.

In SimNIBS, the volume conductor model is constructed using tetrahedral elements (potentially with anisotropic conductivities) and linear basis functions are employed to model the potential. The electric field is estimated by differentiating the solution and nodal values are recovered separately for each tissue compartment using superconvergent patch recovery (SPR) (Zienkiewicz and Zhu, 1992). In EEG, the sources are located in the gray matter and we therefore interpolate the field on a surface representing the center of the gray matter sheet.

Having solved the equation system for each electrode wrt. the reference, we start by computing the (negative) gradient of the solution to obtain the electric field, \mathbf{e}_i , in each element, i,

$$\mathbf{e}_i = -\mathbf{D}_i \mathbf{u}_i. \tag{4}$$

where \mathbf{D}_i is the gradient operator and \mathbf{u}_i contains the potentials at the nodes of *i*. As we use linear basis functions, this field estimate is not continuous and therefore very inaccurate for locations other than the element barycenters. To recover nodal field values, we use SPR to construct a smooth, interpolating function for each node by fitting a linear model to the values in an element patch (i.e., the collection of all elements associated with a given node) around each node such that a field component at an arbitrary position, *r*, can be estimated as

$$e(r) = \mathbf{x}^{\top} \mathbf{b} = \begin{pmatrix} 1 & x & y & z \end{pmatrix} \mathbf{b}$$
(5)

where (x, y, z) are the coordinates at position r and **b** is a vector of parameters to be estimated. This is solved by minimizing the least squares error between the values of the function **f** which we are trying to fit (here the values of the electric field) and the interpolant at the barycenters

$$\underset{\mathbf{b}}{\arg\min} \|\mathbf{f} - \mathbf{X}\mathbf{b}\|^2.$$
(6)

Here, **X** is constructed by row-wise stacking of components \mathbf{x}_i^{\top} corresponding to the barycenter of each element in the patch (augmented by a column of ones as above). We can compute the projection matrix at r as $\mathbf{x}^{\top} (\mathbf{X}^{\top} \mathbf{X})^{-1} \mathbf{X}^{\top}$ and apply it to each component of the electric field to interpolate the field at this location (which, in this case, is the coordinates of the node defining the element patch). Finally, the electric field at any location can be estimated using barycentric interpolation (see supplementary material of Saturnino et al., 2019, for details).

Following Ruffini $(2015)^2$, if we consider a particular location in the brain, r, and a pair of electrodes, a and b, on the scalp, we have

$$u_{ab}^{\text{EEG}} I_{ab}^{\text{TES}} = -\mathbf{e}^{\text{TES}}(r)^{\top} \mathbf{j}^{\text{EEG}}(r)$$
(7)

$$u_{ab}^{\text{EEG}} = -\frac{\mathbf{e}^{\text{TES}}(r)^{+}}{I_{ab}^{\text{TES}}} \mathbf{j}^{\text{EEG}}(r)$$
(8)

$$u_{ab}^{\text{EEG}} = \mathbf{G}_{ab}(r)\mathbf{j}^{\text{EEG}}(r) \tag{9}$$

where $\mathbf{e}(r)$ is the electric field induced by TES at position, r, using current I, and u is the potential difference at a wrt. b due to a dipole source $\mathbf{j}(r)$. $\mathbf{G}_{ab}(r)$ is the potential field value at a wrt. b for unit dipoles at r aligned with each principal axis. Collecting these for all locations and electrode pairs and augmenting it with the reference electrode (effectively, adding zeros corresponding to the reference electrode, and applying an average reference), we arrive at the final matrix of leadfields, i.e., a matrix describing potential fields induced by (ideal) dipolar sources in the brain. Efficient implementations based on reciprocity have also been presented by others (Weinstein et al., 2000; Wolters et al., 2004).

2.2. Validation of SimNIBS for EEG

We validate the leadfields generated using SimNIBS in two ways. First, we compare with an analytical solution in a simple geometry consisting of nested spheres. Second, using a realistically shaped three-compartment model, we compare with existing BEM implementations as BEM generally has high numerical accuracy compared to first order FEM.

2.2.1. Spherical Model

We constructed a model consisting of three concentric spheres representing the inner skull, outer skull, and outer skin surfaces with radii 80, 86, and 92 mm and conductivities 0.3, 0.006, $0.3 \,\mathrm{S}\,\mathrm{m}^{-1}$, respectively. We generated such models with different node densities, 0.065, 0.125, 0.25, 0.5, and 1.0 nodes/mm² (see table 1 for information on number of nodes and elements in the volume meshes), to investigate the effect of resolution on the numerical accuracy and used Gmsh (Geuzaine and Remacle, 2009) to generate a volume

²Instead of a current density, we simply consider a point dipole with unit A m and no associated (or infinitesimally small) volume and thus we have dropped the δV term.

mesh representation of each of the spherical models. Following Vorwerk et al. (2012), we placed point sources aligned with each axis (x, y, and z) in steps of 1 mm ranging from 2 to 77 mm to the sphere origin. This was done for 1,000 random directions. Thus, all sources resided within the innermost compartment of the model. The outermost sources were at a distance of 15 mm to the skin surface which is similar to the closest sources in a realistic head model (Lu et al., 2019). Finally, we placed 100 electrodes equidistantly on the outermost surface. In the simulations, these were modeled as rings with 10 mm radius and 4 mm thickness and were meshed onto the spherical models whereas in the analytical solutions they were treated as point electrodes³. The conductivity of the electrodes was set to 29.4 S m^{-1} . Analytical potentials were computed at the location of the electrodes using equations 2 and 2a from Ary et al. (1981).

2.2.2. Realistic Geometry

To validate the models in a realistically shaped geometry, we extracted surfaces corresponding to inner skull, outer skull, and skin of a head model generated using the *headreco* pipeline in SimNIBS. We down-sampled each surface to approximately 5,000 nodes and used these in the simulations. Additionally, we upsampled these surfaces by uniform remeshing using MeshFix (Attene, 2010) to a resolution close to what we would normally use in our FEM calculations (approximately 80,000 nodes). Running this model in BEM was not feasible with the implementations used here due to the excessive computational resources needed. Conductivities were the same as in the three-layer sphere model. Electrodes were modeled as described above and were placed by creating a standard montage of the 64 channel BioSemi cap (https://www.biosemi.com) included in MNE-Python. For comparison, we used BEM implementations from FieldTrip (specifically, the *dipoli* implementation) and MNE-Python, the latter of which was used as a reference.

We also show numerical differences between FEM models with different surface densities (specifically, 0.5 and 1.0 nodes/mm^2) and refined (achieved by splitting the tetrahedra) versions of these models to illustrate the effect of geometrical accuracy compared to that of resolution. See table 1 for information on number of nodes and elements in the volume meshes.

Lastly, we show results of a comparison between a five-compartment FEM model and a three-layer BEM model. This allows us to compare the numerical differences between FEM and BEM (from the previous simulation) with the differences resulting from modeling the anatomy at different levels of geometrical accuracy and detail. Here, the FEM model was used as reference.

2.3. Impact of Anatomical Accuracy on the Forward Solution

The dataset used in this study consists of 20 subjects each of which had a T1- and a T2-weighted MRI scan (for details on this dataset, please see Farcito et al., 2019). Importantly, manual segmentations of 16 tissue classes were available for all subjects. However, in this study we only distinguished 9 tissues⁴ (10 including air pockets).

2.3.1. Volume Conductor Models

To study the effect of anatomical accuracy we generated four head models using different approaches, specifically, (1) SimNIBS with complete head anatomy reconstruction method (CHARM) (a new segmentation and meshing pipeline which is available in SimNIBS 4.0) (Puonti et al., 2020), (2) MNE-Python (Gramfort, 2013) which is based on FreeSurfer (Dale et al., 1999), (3) FieldTrip (Oostenveld et al., 2011) which uses SPM12 (https://www.fil.ion.ucl.ac.uk/spm/software/spm12/) for segmentation and DUNEuro (Schrader et al., 2021) for FEM modeling, and (4) SimNIBS with CHARM using standard anatomy. As a reference, we used the manual segmentations which we meshed using SimNIBS⁵. Here we briefly describe each of the pipelines.

 $^{^{3}}$ Using point electrodes (as is common in EEG forward calculations) did not have a strong influence and so here we present the results for meshed electrodes which is the default approach in SimNIBS.

⁴White matter, gray matter, CSF, spongy bone, compact bone, skin, eyes, veins, muscle.

 $^{{}^{5}}$ Upon visual inspection, many sulci were closed in the manual segmentations, hence these were "corrected" using FreeSurfer by (1) enforcing gray matter within the cortical gray matter mask from FreeSurfer, (2) relabeling any cortical gray matter voxels outside of this mask which overlapped with the white matter mask from FreeSurfer to white matter, and (3) relabeling any unlabeled voxels to CSF. This seemed to work well for opening sulci, however, at the price of slightly reduced accuracy around cerebellum.

Table 1.	Resolution of	f surface	meshes a	and m	umber	of nodes	and	elements	in the	corresponding	volume	meshes	used	in t	he
validation	experiments.	The num	ber and	"C" ii	n paren	thesis re	efers t	to the nur	mber of	f compartments	s in the r	nodel.			

surface	volume						
$nodes/mm^2$	nodes	tetrahedra					
spherical models (3C)							
0.065	24880	126844					
0.125	61154	322872					
0.25	133952	710780					
0.5	356742	1934983					
1.0	903074	4977163					
realistic models (3C)							
low	22969	120130					
high	697793	3861891					
realistic models $(5C)$							
0.5	638247	3530439					
0.5 (refined)	4880956	28243512					
1.0	1955900	11171938					
1.0 (refined)	15228577	89375504					

SimNIBS-CHARM. The CHARM pipeline uses the SAMSEG framework (Puonti et al., 2016) to segment one or more MRI scans into fifteen different tissue classes. Subsequently, the segmentation is upsampled to 0.5 mm^3 isotropic resolution and post-processed using morphological operations to smooth the segmentation and to ensure that the brain is surrounded by at least one voxel (0.5 mm) thick layer of CSF. The central cortical surface is reconstructed using adapted versions of the corresponding CAT12 (Dahnke et al., 2013) functionalities. The pial surface is then estimated from the central cortical surface and used to improve the sulci representations in the original segmentation. Finally, a volume mesh representation is generated using GGAL (https://www.cgal.org) by directly meshing the updated segmentation. The size of the tetrahedral elements are locally adapted using sizing fields such that elements that are far from tissue borders are larger and elements that are close to tissue borders are smaller. The initial tetrahedral mesh is post-processed to improve the quality of the tissue surfaces. This is done by reassigning the tissue type of tetrahedra representing "spikes" in the surfaces. In addition, the surface nodes are smoothed using a Taubin approach (Taubin, 1995) which is adapted to ensure good quality of the tetrahedra connected to the surfaces. These steps are implemented as standard steps of the CHARM segmentation and meshing pipeline.

Please note that the dataset described in section 2.3 was used to create the tissue priors used in CHARM (Puonti et al., 2020), however, when running CHARM, we used priors from a four-fold cross-validation split such that a subject was never included in the prior used when creating its segmentation to circumvent biased accuracy estimates due to overfitting. The FEM simulations were run using standard settings except that the Pardiso FEM solver included in the Intel Math Kernel Library (version 2022.0.1) was used to calculate multiple FEM solutions efficiently.

MNE-FS. MNE-Python integrates with FreeSurfer and relies on its *recon-all* pipeline. Here we created the BEM surfaces (inner skull, outer skull, and skin) based only on a T1-weighted image⁶. This procedure uses the watershed algorithm in FreeSurfer. The final node count in each surface was 5,120. We note that since the inner and outer skull surfaces sometimes overlapped, we had to fix this in order to be able to run the BEM calculations. This was done using MeshFix by pushing the inner skull inwards to ensure a minimum

 $^{^{6}}$ MNE-Python also provides a tool to create surfaces from a T1-weighted image and a sequence of fast low angle shot (FLASH) images, however, since FLASH was not available for this dataset we decided to only use the T1-weighted image (we tried using the T2-weighted image instead but with disappointing results).

distance of 2 mm between the inner and outer skull since the problem seemed to be that the inner skull surface extended too far outwards. Consequently, we view this as an improvement of the model. We used MNE-Python 0.24 and FreeSurfer 7.1.0.

FieldTrip-SPM. FieldTrip⁷ uses SPM12 to segment a T1-weighted image into five tissue classes (gray matter, white matter, CSF, bone, and skin). This segmentation is post-processed using morphological operations (e.g., smoothing) to, for example, ensure that the brain is surrounded by CSF which again is surrounded by bone. Based on this segmentation, a hexahedral mesh was created with a shift of 0.3 and a downsampling factor of two. We used the DUNEuro solver with default settings.

MNI-Template. This model is created simply by running CHARM on the MNI152 T1-weighted image. We padded the image such that the lower part of the face and neck was filled from the prior. Otherwise the procedure is exactly as described for SimNIBS-CHARM.

The purpose of this work is not to provide a thorough evaluation of the anatomical accuracy of each of these models (we refer to Nielsen et al. (2018) and Puonti et al. (2020)), however, visual inspection revealed that the SimNIBS-CHARM models were generally more anatomically accurate than the other models (see figures 6 and 7 for two examples). To summarize, SimNIBS-CHARM is a first order FEM model with generally high anatomical accuracy; FieldTrip-SPM is a FEM model with decent accuracy of brain tissue but simplified non-brain tissue (e.g., bone); MNE-FS is a BEM model with simplified anatomy (no distinction between white matter, gray matter, and CSF) and limited accuracy; MNI-Template is a FEM model based on average anatomy which, in general, is quite dissimilar to the individual anatomy but may estimate CSF and bone reasonably well.

In each case, we used the default conductivities specified by the particular pipeline⁸.

2.3.2. Electrode Positions

To model electrode positions we used an EasyCap BC-TMS64-X21 transcranial magnetic stimulation (TMS) compatible cap with a modified M10 (equidistant) layout (https://www.easycap.de). See section 2.4.1 for details on how these positions were obtained.

We placed the electrodes on the individual reference models (i.e., the models based on the manual segmentations) by registering the template positions (in MNI space) to subject MRI space by a seven parameter affine registration (one uniform scaling parameter) based on coregistration of landmarks (nasion, left preauricular (LPA), right preauricular (RPA)). For SimNIBS-CHARM, MNE-FS, and FieldTrip-SPM, the electrodes were then projected onto the skin surface of the respective model. For MNI-Template, the mesh was deformed nonlinearly using thin-plate splines to head points sampled around the electrodes (four points per electrode resulting in approximately 200 points in total) similar to the procedure used in Brainstorm (Darvas et al., 2006). Given that the skin surface was well modeled in all pipelines, we do not anticipate major effects due to different electrode positions.

2.3.3. Source Space

In the MNE-FS and FieldTrip-SPM models, we defined the source space as the central gray matter surface extracted from the reference model using CAT12 and downsampled to approximately 10,000 nodes per hemisphere. In the event that a source was not inside the brain (defined as the brain compartment in MNE-FS and white matter or gray matter in FieldTrip-SPM) this was excluded from analysis⁹. In SimNIBS-CHARM and MNI-Template, we used the the central gray matter surface created during the respective runs. To compare the gain matrices (leadfields) we mapped them to *fsaverage* before computing the evaluation metrics.

⁷We used the master branch from the GitHub repository (https://github.com/fieldtrip) at commit 4b50f71105f6bb250018b815bb628e6f3e9de2b6 (from 22.11.2021).

⁸For FieldTrip-SPM, we specified conductivities as defined in the tutorial https://www.fieldtriptoolbox.org/tutorial/ headmodeleegfem.

 $^{^{9}}$ When sources were excluded, smoothing was employed to ensure that all nodes in *fsaverage* space were covered. To avoid unnecessary smoothing, this was only done for the nodes where it was required.

2.4. Impact of Electrode Positions on the Forward Solution

The dataset used in this study is a subset of the dataset described in Madsen et al. (2019) and Karabanov et al. (2021) in which noninvasive brain stimulation was performed conditional on the brain state as determined by EEG measurements. Of particular relevance to the current study, structural MRI scans were acquired and electrode positions digitized using an infrared optic stereo tracking system from Localite (https://localite.de). For this study, we extracted 32 subjects for which T1- and T2-weighted images as well as digitized electrode positions were available. During our analysis, we found that some digitization runs contained electrode positions which were clearly the result of some kind of equipment malfunction (e.g., a situation where an electrode was placed in the center of the head). To formalize the identification of such errors, and possibly "correct" them, we first transformed all electrode positions to standard space by nonlinear deformation. We then defined a "spurious digitization" as a position which was more than five centimeters away from the group median and we decided to correct a maximum of one electrode per subject. Corrections were performed by replacing the erroneous coordinates with the group median. Spurious digitization errors were detected in four of 32 subjects. In each of these subjects there was only a single error which we then corrected.

2.4.1. EEG Custom Template Generation

We generated a custom template in standard space of the EasyCap BC-TMS64-X21 EEG cap using the following procedure. We reconstructed the skin surface of an augmented MNI152 template (which included the jaw and neck as well) from the final segmentation of the MNI-Template described in section 2.3.1 using the marching cubes implementation from *scikit-image* (van der Walt et al., 2014) and smoothed the result using PyVista (Sullivan and Kaszynski, 2019). Since the MNI152 model is quite large (approximately 60 cm circumference of the head), we scaled it to a circumference of 57 cm which allowed us to comfortably fit a 58 cm cap on it. This scaled model was 3D printed on an Ultimaker 2 Extended+, the aforementioned cap was placed on it, and electrode positions were digitized using a Polaris Spectra infrared camera from NDI with custom software (https://www.ndigital.com). To securely fasten the sensor responsible for tracking the head position, we augmented and flattened the nasal bridge such that it could be easily attached here. The final template was scaled back to the original space of the MNI template.

To facilitate coregistration between the coordinate system of the digitizer and MNI space, we added 13 small dents (which served as "landmarks") in the surface (nasion, LPA, RPA, and ten points spread across the top of the head). We placed the EEG cap according to standard guidelines (e.g., Homölle and Oostenveld, 2019) and digitized the electrode positions as well as the landmarks. We repeated this three times, registered the electrode positions to MNI space and used the median position for the final template. In the final template we noticed a few irregularities. First, two parietal channels, 50 and 52, were not exactly equidistant to the surrounding electrodes. Upon inspection of the different cap sizes, we found that these channels were placed on one side of a stitching in caps smaller than 56 cm whereas they were on the opposite side in caps larger than or equal to 58 cm. Since this effect was evident mostly in the lateral direction, channels 50 and 52 were corrected by replacing the x coordinate with the mean of electrodes (49, 51) and (51, 53), respectively. Likewise, we saw variations in channel 22 between caps. Channel 22 was corrected by replacing the x and y coordinate with the means from 21 and 23. Finally, all channels were (re)projected onto the MNI skin surface. These steps were taken to improve the average fit of the template to subjects of different head sizes.

2.4.2. Volume Conductor Models

To study the effect of different ways of estimating electrode positions, we first constructed a model of the head for each subject using the CHARM pipeline in SimNIBS. To assess the errors in electrode positions when using template layouts, we transformed the layout of the EEG cap as defined by the manufacturer¹⁰

 $^{^{10}}$ The cap used here has an equidistant M10 layout and includes two additional channels in the temporal area. Here we use the manufacturer layout for this cap shipped with MNE-Python augmented by these two channels. To adhere to the equidistant properties of the layout, we defined the angles (theta, phi) of these electrodes as (-138, -22) and (138, 22).

as well as our custom positions from standard space to subject space. The former was transformed using an affine transformation obtained by coregistering fiducials (nasion, LPA, and RPA) between subject space and the space of the manufacturer layout (after mapping the positions to a sphere). The latter was transformed in three different ways using (1) the nonlinear deformation field calculated by CHARM as part of the segmentation, (2) the affine registration between subject space and MNI space estimated by CHARM before performing the segmentation, and (3) the affine transformation obtained by coregistration of fiducials in subject space and MNI space. Subject specific fiducials were identified on the T1-weighted MRI of each subject. As reference, we used the digitized positions available in the data set.

We also included a model based on standard anatomy (similar to the one described in section 2.3.1). As in the reference model, we used the digitized electrode positions from the data set. This was done to model a scenario where electrode positions have been accurately defined but structural information is unavailable.

As we did not find any appreciable differences in the accuracy of the channel positions between different ways of transforming the custom template (see section 3.3), we chose to only use the nonlinearly deformed positions in the forward simulations as this does not require identification of landmark positions for each individual. Thus, we ran simulations using the models with the digitized positions, nonlinear deformation of custom template (Custom-Template), with affine registration of the manufacturer layout (Man-Template), and with standard anatomy (MNI-Digitized). All simulations were carried out in SimNIBS similarly to what was described for SimNIBS-CHARM in section 2.3.1.

2.4.3. Source Space

As source space we used the central gray extracted using the CAT12 functionalities embedded in CHARM and downsampled to approximately 10,000 vertices per hemisphere, thus the source space was the same for models based on individual subject anatomy (the reference using the digitized positions from the data set, Custom-Template, Man-Template) but different for the one based on standard anatomy (MNI-Digitized). To compare the gain matrices we mapped them to *fsaverage* before computing the evaluation metrics.

2.5. Evaluation Metrics

To evaluate differences in electrode position, we used the Euclidean distances to the reference positions. We show the full distribution of these errors and also how the mean errors are distributed over channels. Additionally, we also show the spatial distribution of the absolute error along each axis (x, y, z) to explore the direction of the errors. For evaluation, we always used channel positions after projecting onto the skin surface.

All channel positions were evaluated in the head coordinate system used by MNE-Python. Here, the x-axis goes through the LPA to the RPA, the y-axis is normal to the x-axis and passes through the nasion, and the z-axis is normal to the xy-plane and points upwards according to the right-hand rule. Thus, this constitutes a right-anterior-superior (RAS) oriented coordinate system.

To evaluate the differences between forward models we use the relative difference measure (RDM)

$$RDM = \left\| \frac{u}{\|u\|} - \frac{u_{ref}}{\|u_{ref}\|} \right\|$$
(10)

and the logarithm of the magnitude error (lnMAG)

$$\ln MAG = \ln \frac{\|u\|}{\|u_{ref}\|} \tag{11}$$

which computes the relevant measure of u with respect to u_{ref} and where $\|\cdot\|$ is the 2-norm (Meijs et al., 1989; Vorwerk et al., 2014). RDM measures the topographic error and is bounded between 0 and 2 where 0 means no difference whereas lnMAG measures the magnitude errors with a value of 0 corresponding to no difference.

All forward solutions were re-referenced to an average reference and converted to a fixed orientation (normal to the cortical surface). The final gain matrix consisted of approximately 20,000 columns. Before



Figure 1. CRF of RDM and lnMAG for different model densities (nodes/mm²) in a three-layer sphere model.

computing the evaluation metrics, gain matrices were morphed to the fifth subdivision of the *fsaverage* template ("*fsaverage5*") which consists of 20,484 nodes in total (10,242 per hemisphere).

We show the cumulative relative frequency (CRF) of the error metrics using the data from all subjects and also the spatial distribution of the average of each metric. All surface-based plots were created using $PyVista \ 0.35.1$ (Sullivan and Kaszynski, 2019). Remaining plots were created with *Matplotlib* 3.5.1 (Hunter, 2007).

3. Results

3.1. Validation of SimNIBS for EEG

3.1.1. Spherical Model

Figure 1 shows the CRF of RDM and lnMAG for each model and figure 2 shows RDM and lnMAG as a function of source eccentricity. Both figures illustrate that increasing the resolution of the model is beneficial for numerical accuracy. Errors get larger with increasing source eccentricity. However, for the high density models, the RDM stays below 0.02 and the lnMAG between ± 0.01 . From figure 2, we see a small bias in the magnitude in that it is smaller in the SimNIBS simulations.

3.1.2. Realistic Geometry

Figure 3 shows CRF of RDM and lnMAG in a realistic three-layer model using the BEM solution from MNE-Python as reference. We see that errors of the FEM model with a resolution used in practice, i.e., the high resolution model, are generally low (below 0.05) except for a few outliers (heavy tail). The numerical accuracy generally deteriorates as sources get closer to the sensors (see section S1 and fig. S1.1 in the supplementary material). The BEM implementations are very similar in terms of topography but less so for magnitude.

Figure 4 shows CRF of RDM and lnMAG for different model comparisons. We see that purely numerical errors (as evidenced by the comparison of a model with its refined counterpart, e.g., FEM 0.5 vs FEM 0.5r) are lower than errors due to modeling the anatomy in greater detail (i.e., FEM 0.5 vs. FEM 1.0) which again is much lower than the errors incurred by simplifications of the model (i.e., the full FEM models vs. three-layer BEM). An example of the effect of resolution and refinement on the geometry of the FEM models are illustrated in figure 5. Errors due to model simplifications are about ten times those of the errors



Figure 2. RDM and lnMAG as a function of source eccentricity for different model densities in a three-layer sphere model. Shaded area denote the interval from the 0.5th to 99.5th percentile.

due to numerical inaccuracies. The pattern is the same for both RDM and lnMAG. Contrary to the simple three-layer model evaluated above, source eccentricity does not seem to be as good a predictor of the overall errors (numerical and otherwise) in more complicated geometries (see section S1 and figs. S1.2 and S1.3 in the supplementary material).

3.2. Impact of Anatomical Accuracy on the Forward Solution

Figures 6 and 7 show the surfaces extracted using each pipeline overlaid on the T1-weighted image for two subjects. It is not our intention to do a systematic assessment of the anatomical accuracy of the models as this has been done previously (Nielsen et al., 2018; Puonti et al., 2020). However, we believe this to be representative of the differences between the pipelines we have seen in the data.

Generally, we found high anatomical accuracy of SimNIBS-CHARM whereas substantial simplifications of the modeling of the CSF and bone compartments are apparent in FieldTrip-SPM. As a result, FieldTrip-SPM is not particularly accurate for non-brain tissues. MNE-FS collapses white matter, gray matter, and CSF into one brain compartment. Additionally, we found that the BEM surfaces were not very accurate (as evidenced by the need for us to correct intersections between these surfaces). MNI-Template does not capture the individual structure of the brain tissues very well. However, in several cases it actually provided a decent fit of the skull compartment, whereas in others it was not accurate at all. We provide similar figures to figures 6 and 7 for more subjects in section S2 figs. S2.4 to S2.7 in the supplementary material.

Figure 8 shows CRF of the error metrics, RDM and lnMAG, using the forward solution associated with the manual segmentations as reference. In terms of topographical errors, SimNIBS-CHARM shows best performance followed by FieldTrip-SPM, MNI-Template, and finally MNE-FS. As to the magnitude errors, SimNIBS-CHARM and MNI-Template are both centered around zero with the latter having a larger standard deviation than the former. FieldTrip-SPM and MNE-FS generally show decreased and increased sensitivity, respectively. The effect is most severe for MNE-FS both in terms of the mean and standard deviation.

Figures 9 and 10 present the spatial distribution of the mean RDM and lnMAG on the *fsaverage* surface. From figure 9 we see slightly elevated errors in the orbitofrontal area for SimNIBS-CHARM but otherwise low errors. The spatial pattern for FieldTrip-SPM is similar, however, the effects are more pronounced and also extends into the temporal cortex. MNE-FS shows high errors on the gyral crowns both for superior and deep sources. For MNI-Template we see mostly errors comparable to those of FieldTrip-SPM, but with



Figure 3. CRF of RDM and lnMAG for a high and low resolution FEM model and a BEM model in a realistically shaped three-compartment model. The BEM solution from MNE-Python was used as a reference.



Figure 4. CRF of RDM and lnMAG for FEM comparing five-compartment FEM models of different resolution (0.5 and 1.0 nodes/mm^2), refined (i.e., upsampled) versions of these models, and a three-layer BEM model. The model on the right-hand side of the greater-than sign is the reference.



Figure 5. Zoomed view of a coronal slice showing the effect of resolution and refinement on the head models. The corresponding effects on the forward solution can be found in figure 4.

a few hot spots and also slightly increased errors in parts of the parietal cortex. From figure 10 we that SimNIBS-CHARM generally has the lowest overall errors. FieldTrip-SPM shows decreased sensitivity to deep, occipital, and temporal sources whereas MNE-FS shows increased sensitivity everywhere, particularly for shallow sources on the gyral crowns.

We expect the magnitudes to be sensitive to the ohmic tissue conductivities. Since we used the default values in each pipeline—and because the conductivities are the same in SimNIBS-CHARM and MNI-Template as in the reference model—it is not surprising that these are similar in terms of sensitivity. On the other hand, differences in conductivity is a potential confounder for the comparison of the FieldTrip-SPM and MNE-FS models to the reference. In an attempt to disentangle such effects, we created another reference model with conductivities similar to those of FieldTrip-SPM¹¹ and computed RDM and lnMAG. From figure 8, we see that most of the differences remain even after this correction, and so we believe that our statements above are indeed valid. Due to the vast geometrical simplifications in the MNE-FS model (collapsing brain tissue and CSF into one compartment) and their known effects on the forward solution (Vorwerk et al., 2014), we did not do a similar analysis with the MNE-FS model as we feel the geometrical adjustments to the reference model would be too large.

3.3. Impact of Electrode Positions on the Forward Solution

We first report errors on the channel positions before describing the resulting effects on the forward solution.

Figure 11 is a density plot of the Euclidean distance between electrode positions estimated using the different methods and the digitized positions across all subjects. It is evident that Man-Template generally performs worse than Custom-Template and also that the particular transformation method of Custom-Template does not have a strong influence on the accuracy of the electrode positions.

Figure 12 shows the spatial distribution of the mean error over channels. We see that the errors in the posterior areas (mostly occipital but also parietal and temporal) are large for Man-Template whereas

¹¹The conductivity of compact and spongy bone were both set to that of bone in the FieldTrip-SPM model and thus a minor geometrical simplification of the reference was introduced.



Figure 6. Tissue compartments as estimated by each of the pipelines in the anatomy study for subject one. It is apparant that MNE-FS does not capture the inner skull very well for this subject. The skull thickness varies substantially but this does not seem to agree with the structural MRI scan. Also, the lower part of the head is mostly modeled as skull as the outer skull surface is obtained by eroding the skin surface. Furthermore, the skull thickness in FieldTrip-SPM is almost constant throughout the volume (e.g., thickness is overestimated around the occipital cortex). MNI-Template captures the skull surprisingly well in this subject but the brain tissues are obviously not accurate.



Figure 7. Tissue compartments as estimated by each of the pipelines in the anatomy study for subject two. MNE-FS captures the inner skull better in this subject compared to figure 6, however, the skull thickness is not well modeled. FieldTrip-SPM captures inner skull and thickness well except for the prefrontal part. MNI-Template displaces the skull outwards in the zoomed-in view whereas other parts are modeled more accurately. The spongy bone, however, is not.



Figure 8. CRF of RDM and lnMAG for each model. FieldTrip-SPM (Matched) is a comparison of the FieldTrip-SPM model with a reference with matched conductivities (see section 4.2.2) allowing us to disentangle effects of geometry and conductivity.



 ${\bf Figure}~{\bf 9.}~{\rm Spatial}$ distribution of mean RDM for each model.



Figure 10. Spatial distribution of mean $\ln MAG$ for each model.



Figure 11. Density of errors in electrode position as measured by the Euclidean distance between each model and the digitized positions. The parenthesis denotes the registration method where "LM" means that the registration was based on landmarks.

this is much less severe for the Custom-Template. Likewise, the standard deviation is also larger for Man-Template, again in posterior but also anterior areas. Interestingly, although the mean error is similar across the different transformation methods for Custom-Template, the standard deviation is not. The MRI-based coregistrations show larger variability in the anterior parts of the head (suggesting that other parts of the image are driving the registration). The landmark-based registration exhibits larger variability in the posterior regions (perhaps unsurprisingly given that the landmarks used here are the nasion, LPA, and RPA).

To determine potential spatial bias in the template registrations, we show the absolute errors along each axis in figure figure 13. In general, it seems that errors in anterior and posterior areas are largest in the z direction and errors in the superior areas are largest in the y direction. This suggests that the errors here are mostly due to an anterior-posterior (AP) shift (i.e., a rotation around the x-axis). Interestingly, errors in the most inferior electrodes in the temporal areas are primarily due to shifts along the AP direction and not up or down. Finally, please note the different scales of Custom-Template and Man-Template and that errors may be exaggerated close to the edges of the head contour due to extrapolation.

For both Custom-Template and Man-Template, we found that channel errors in occipital and frontal regions were mostly along the z-axis and along the y-axis for superior areas. To investigate the extent to which errors were due to bias (i.e., a systematic difference between the template and the average positions across subjects) or variability in the reference positions (i.e., the standard deviation of each position, e.g., due to inter-individual differences in cap placement), we transformed the reference positions for all subjects to standard space and computed bias and standard deviation for each position. The standard deviations (figure 14) showed the same pattern as that observed in figure 13, suggesting more variability in AP alignment of the cap compared to aligning it laterally. On the other hand, we also found a higher bias at some positions than others (figure 15). For example, the digitized positions of the electrodes between the ears and the eyes were, generally speaking, more anterior compared to Custom-Template.

We also show all digitized electrode positions in standard space on the MNI surface template (figure 16) from which it is apparent that the positions towards the back are lower in the left side compared to the right side. This points to some kind of misalignment, perhaps due to the way the chin strap was fastened being different between subjects and the MNI template head. As such, it is not clear whether the observed bias is due to suboptimal (biased) cap placement in the data set used here or because the template fails to capture how a cap fits on an actual human head.

Figure 17 shows CRF of the error metrics, RDM and lnMAG, for the forward solutions using the solution associated with the digitized positions as reference. It is evident that the RDM errors are generally higher



Figure 12. Spatial distribution of errors (mean and standard deviation) in electrode position as measured by the Euclidean distance between each model and the digitized positions. The parenthesis denotes the registration method where "LM" means that the registration was based on landmarks.

for Man-Template than for Custom-Template whereas the errors for MNI-Digitized fall between the two. As to lnMAG, Custom-Template is again better than Man-Template but here MNI-Digitized is clearly worse than the other two as evidenced by the heavy tails.

Figures 18 and 19 present the spatial distribution of the mean RDM and lnMAG on the *fsaverage* surface. From figure 18, we see slightly elevated topographical errors for Custom-Template in orbitofrontal and occipital errors. However, this is more pronounced in Man-Template which also shows increased errors in parietal and temporal regions. The errors of MNI-Digitized are generally similar to those of Custom-Template. However, we see several hot spots of increased error, often on the gyral crowns. From figure 19, we see slightly increased magnitudes for Custom-Template in frontal and right temporal regions. Generally, Man-Template shows increased and decreased magnitudes for deep and shallow sources, respectively. Again the magnitudes are elevated in the right temporal region. MNI-Digitized shows large magnitude errors with increased and decreased sensitivity to sources on the gyral crowns and sulcal walls, respectively. Decreased sensitivity is also observed for orbitofrontal and occipital sources.

4. Discussion

4.1. Validation of SimNIBS for EEG

The simulations in spherical models, assuming sufficient model resolution of about 0.5 or 1.0 nodes/mm², showed that the numerical errors in topography and magnitude are low and similar to those of existing FEM implementations for EEG (Vorwerk et al., 2012). The same was true for simulations using a realistic geometrical models. In addition, by comparing numerical errors of our FEM to errors incurred by not distinguishing different tissue compartments in the BEM head models (see figures 3 and 4), we show that the latter are about ten times those of the former. Hence, we suggest that in general the decrease in numerical accuracy associated with FEM is offset by its flexibility when it comes to modeling complex anatomical structures.

Close inspection of figure 2 reveals that there is a small bias in the magnitude in the SimNIBS simulations. Decreasing the thickness of the electrodes (thus minimizing potential drop over the electrode itself) and matching the conductivity to that of skin (since high conducting electrodes effectively makes the skin compartment a better conductor) removed this effect. Thus, we attribute this to the electrodes being



Figure 13. Spatial distribution of mean errors in electrode position as measured by the absolute difference between each model and the digitized positions. Please note the different scale of the rightmost column. The parenthesis denotes the registration method where "LM" means that the registration was based on landmarks. Positive z direction is out the page.



Figure 14. Standard deviations of the digitized positions in the electrode study calculated using the overall Euclidean distance or distances along each major axis. This illustrates the variability of the reference positions.



Figure 15. Distance between the average of the digitized positions in the electrode study and the positions of the custom template calculated using the overall Euclidean distance or distances along each major axis. This illustrates the bias of the template compared to the reference positions.



Figure 16. Our custom template positions (white spheres) and the digitized positions for each subject (colored spheres) in MNI coordinates.



Figure 17. CRF of RDM and lnMAG for each model.



Figure 18. Spatial distribution of mean RDM for each model.



Figure 19. Spatial distribution of mean lnMAG for each model.

meshed in SimNIBS whereas the analytical solution assumes point sensors. One could argue that the correct approach is to model the electrodes. However, the effects are very small in practice.

As to the source model used in SimNIBS, calculating the electric field in each element is similar to the partial integration approach (Vorwerk, 2016) with dipoles along each axis. However, to improve the accuracy of the estimate, we use a postprocessing procedure (SPR) which consists in fitting a smooth function such that the value of a node is given by a linear combination of its neighboring elements. Although different from the St. Venant approach (Vorwerk, 2016), we found the numerical accuracy to be similar.

4.2. Impact of Anatomical Accuracy on the Forward Solution

4.2.1. Topographical Errors

Interpreting the topographical errors reflected by RDM can be difficult since the quality of the anatomical models may vary from subject to subject. However, we will summarize some general observations about the different models.

Both FieldTrip-SPM and MNE-FS showed elevated topographical errors in orbitofrontal and deep temporal areas. Given the close proximity of such sources to the facial area and the fact that neither of these models attempt to characterize the complex anatomy in this region, we suspect that such simplifications may explain the increased RDM values here. Since the anatomy is simpler around the occipital cortex, the modeling decision of FieldTrip-SPM to simply enclose the brain with a certain amount of bone (and CSF) seems more appropriate here. The same holds for MNE-FS which also showed large RDM values for superficial sources in the superior part of the brain. Similar observations were made by Vorwerk et al. (2014) when distinguishing CSF and brain tissues. The inaccuracy of the inner and outer skull surfaces in MNE-FS likely exacerbates this issue (Lanfer et al., 2012).

In line with the better anatomical accuracy, the topographical errors of SimNIBS-CHARM are generally small compared to the other models (see figures 6 and 7). MNI-Template seems slightly better than MNE-FS but worse than the others. This seems to highlight the importance of accurately modeling CSF since the brain tissue compartments in MNI-Template are generally very different from the individual anatomy.

4.2.2. Magnitude Errors

As with the topographical errors, the magnitude errors of SimNIBS-CHARM are low whereas they are slightly higher for MNI-Template (heavier tails in the CRF curve). The errors of MNE-FS and FieldTrip-SPM are larger still and biased towards higher and lower sensitivity, respectively.

The magnitude errors in FieldTrip-SPM suggest that sensitivity is generally decreased, particularly for deep sources. From figures 6 and 7 it is evident that the models tend to contain too much CSF, particularly in the deep and temporal areas where bone is also too thick, which will increase shunting effects. MNE-FS shows the opposite effect, i.e., an overall increase in magnitude, most prominently for sources close to the sensors. MNE-FS does not model the CSF compartment explicitly but rather collapses white matter, gray matter, and CSF into a single compartment with an adjusted conductivity. Thus, the conductivity gradient between the brain and bone compartments is smaller and hence we expect less current shunting. These results are in line with Vorwerk et al. (2014) in terms of the relevance of distinguishing CSF from brain tissue as well as white matter from gray matter.

4.3. Impact of Electrode Positions on the Forward Solution

Similar to Homöle and Oostenveld (2019), we also found that channel errors were generally much larger for Man-Template than Custom-Template in occipital and parietal areas suggesting that approximating the head shape with a sphere is problematic when the sphere is aligned using frontal and temporal landmarks. Including the inion in the coregistration could potentially help, although we did not test this. On average, we did not find similar effects in posterior regions when transforming Custom-Template using landmarks (although standard deviations were higher), suggesting that the misalignment of Man-Template is in fact due to this layout being a poor fit of the actual positions in these areas.

We also investigated the extent to which the errors observed in Custom-Template were due to bias (i.e., systematic differences to the reference positions) or variability (in the reference positions). We found that

the errors were partly due to a bias of the template but also that some electrodes were less consistently placed on the individual subjects. We did not do a similar analysis for Man-Template.

Given this bias and variability in the electrode positions, it is not reasonable to expect a perfect fit of the template to all subjects. We noticed that the fit was particularly bad in a few cases where the cap was clearly not aligned well (e.g., rotated around the x or z axis).

4.3.1. Topographical errors

In line with the mean error observed on the channel level, topographical errors in the forward solutions, as estimated by RDM were more severe for Man-Template than Custom-Template. For the former, errors were high in occipital and posterior areas. Errors were higher on the gyral crowns (i.e., where the sensitivity is high) than the sulcal walls and valleys, something which is most clearly seen for Man-Template.

Topographical errors for MNI-Digitized are slightly higher than Custom-Template but smaller than Man-Template suggesting that electrode positions are important in determining such errors and that template electrode positions as supplied by the manufacturer result in higher RDM values compared to a template anatomy with digitized positions. This is in line with the results of Homölle and Oostenveld (2019).

We also see a few RDM hot spots in MNI-Digitized which, as they stand out on a group level, are likely due to local inaccuracies in this model. We were not able to identify any obvious errors upon visual inspection of the final segmentation and mesh though.

4.3.2. Magnitude errors

We observed that Man-Template showed increased and decreased magnitudes for deep and superior sources, respectively. Visual inspection of the final positions showed that electrodes were generally placed further down on the head model, thus giving increased sensitivity to deep (or ventral) sources, and, conversely, decreased sensitivity to superior sources (most notably parietal, occipital, and prefrontal) as the montage was spread more thinly across the head. Increased sensitivity to temporal sources was also observed as these are generally not very well covered by electrodes.

Magnitude errors were much larger for MNI-Digitized. In particular, we observed increased sensitivity to sources on the gyral crowns and decreased sensitivity to sources on the sulcal walls. Since the MNI152 T1-weighted template represents a standard anatomy, it is generally quite blurred. As a result, fine details such as gray matter curvature and sulcal structure are lost. Visual inspection (see figure 20) showed that the gray matter compartment of the final head model was much smoother than for individual subjects, hence sources on sulcal walls and in sulcal valleys were surrounded almost exclusively by gray matter and not highly conducting CSF as was the case in the models based on the individual anatomy. Additionally, we also found that the CSF layer was much thinner compared to the individualized models resulting in less current shunting, particularly for sources on the gyral crowns to which sensitivity was increased (see figure 20).

4.4. Effect on Inverse Solution

Despite clear differences in the forward solutions obtained from the different pipelines investigated in this study, the impact on the inverse solution still needs to be established. There is, however, already a body of literature suggesting that the forward model is indeed an important factor that influences source reconstructions. Below, we discuss this in relation to the effects observed in the current study.

As mentioned in section 1, the modeling of the skull compartment has received much attention in the EEG community and from figures 6 and 7 it is apparent that there are clear differences in how the pipelines model this. The thickness is not accurately captured in FieldTrip-SPM and MNE-FS whereas SimNIBS-CHARM is able to delineate the skull with decent accuracy (Puonti et al., 2020). However, previous studies report only minor source localization errors (less than 5 mm) due to such erroneous thickness estimation (Chauveau et al., 2004; Lanfer et al., 2012). It is also apparent that FieldTrip-SPM and MNE-FS simplifies the modeling of the inferior part of the skull quite substantially. This was shown by Lanfer et al. (2012) to result in large errors (more than 1 cm) for sources close to the base of the skull. A related point has been made by studies showing that failure to model skull openings (foramina) (Fiederer et al., 2016) and



Figure 20. The top row illustrates the reduced amount of CSF between the cortex and the inner skull in MNI-Digitized compared to the SimNIBS-CHARM for a particular subject (the MNI-Digitized model is almost identical for all subjects except for the nonlinear deformation). The bottom row shows the corresponding gray matter compartments.

sutures (McCann and Beltrachini, 2022) may decrease localization accuracy particularly for sources in the vicinity of such errors. Finally, Montes-Restrepo et al. (2014) showed that low quality skull segmentations (generated from T1-weighted images compared to CT scans) resulted in increased source localization errors of up to 1 cm, but only for sufficiently low noise levels.

The human skull, however, is not a homogeneous compartment since it is comprised of both spongy (soft) and compact (hard) bone. SimNIBS-CHARM distinguishes between these, which is not the case for FieldTrip-SPM and MNE-FS. Dannhauer et al. (2011) suggests that modeling such inhomogeneity is important and that failing to do so may result in increased localization errors up to 1 cm. On the other hand, using anisotropic modeling of the skull (different conductivities in the radial and tangential directions) to account for the layered structure of the bone may not lead to any noticeable improvements and may even deteriorate performance in areas of complex geometry (Dannhauer et al., 2011; Montes-Restrepo et al., 2014). We would like to point out though—irrespective of whether or not the particular distinction is important for source localization *per se*—that within the Bayesian segmentation framework used by CHARM, modeling compact and spongy bone separately may be beneficial for the overall segmentation result because spongy bone and skin often look similar on a T1-weighted image. Thus, this might improve the segmentation quality in general, particularly in subjects where spongy bone is very prominent or where the image quality is low (e.g., corrupted by fat shift) (Nielsen et al., 2018; Puonti et al., 2020).

Another important aspect of forward modeling is whether or not to distinguish between different brain tissues and CSF or simply model it as one compartment. The latter strategy is employed by MNE-FS (and most other pipelines using BEM) whereas both FieldTrip-SPM and SimNIBS-CHARM includes white matter, gray matter and CSF. Ramon et al. (2006) concluded that the CSF compartment is important for shaping the observed potential distribution on the scalp and neglecting this increases average source localization errors by a few millimeters. Likewise, Conte and Richards (2021) found that including CSF was the most important addition to a model consisting initially of brain, skull, and scalp and that when doing so, it performed similarly to more complex models distinguishing even more tissue types. The latter suggests that the increased complexity of the SimNIBS-CHARM model (including for example muscle tissue and eyes) may not matter much in terms of source localization accuracy. Similarly, differentiating between white and gray matter was found to be less important compared to CSF (Ramon et al., 2006).

In terms of anatomical accuracy, there is substantial variation as to how well the template model fits each subject. Generally, the brain tissues and CSF do not match at all. However, the inner and outer skull borders are a more reasonable fit. Acar and Makeig (2013) observed median dipole localization errors of about 5 mm when using a realistic four compartment template model warped to the head shape of each subject (as was also done in this study). They found high errors towards the base of the skull since their template was cut just below the nasal area. However, the template used here has similar coverage to that of the individualized models so we expect this to be less of an issue (Lanfer et al., 2012).

In the study investigating the impact of the accuracy of the electrode positions, we found spatially correlated errors concentrated around occipital and parietal areas. This was true for both Custom-Template and Man-Template, but was more prominent for the latter. Similar results were reported by Homölle and Oostenveld (2019), who found average dipole localization errors of 8.5 and 11 mm for custom and manufacturer templates, respectively. Again, particularly high errors for the latter occurred over the above mentioned areas. In an attempt to simulate coregistration errors, Acar and Makeig (2013) tilted electrode positions 5 degrees and found mean localization errors about 6 mm along the direction in which the tilt was performed (e.g., left-right tilt caused larger errors in temporal areas). Thus, spatially correlated errors in electrode locations seem to propagate to source localization errors. On the other hand, Wang and Gotman (2001) found only minimal effect of electrode position on source localization accuracy. however, in this study the simulated errors were uncorrelated between neighboring sensors.

Studies have also suggested that the choice of digitization method may impact inverse solutions. For example, Dalal et al. (2014) reported decreased beamformer sensitivity and slightly increased localization errors when electrode positions were recorded using a suboptimal digitization procedure and Shirazi and Huang (2019) found that digitization errors of about 1 cm could potentially result in source reconstruction errors of up to 2 cm. In the latter study, they report degraded localization accuracy using a template of electrode positions. However, it is not clear from the study exactly how these positions were obtained, and, as suggested by our results and those of Homölle and Oostenveld (2019), this may significantly impact the accuracy.

While the literature reviewed here suggests that errors in the head model and electrode positions used for forward modeling affect the accuracy of the inverse solution, this influence may be modulated by other factors such as the SNR of the measurements and the inverse modeling approach used. For example, Montes-Restrepo et al. (2014) show that source localization accuracy benefits from accurate skull modeling based on CT, however, they also found that differences between CT and MRI based skull models disappeared as noise levels increased (e.g., going from SNR of 10 dB to SNR < 5 dB). Ramon et al. (2006) investigated the effect of distinguishing different tissue compartments (e.g., white matter from gray matter or merging CSF, white matter, and gray matter) on source localization accuracy. They found that differences between models only started to emerge above an SNR of approximately 5 dB. Thus, a sufficient SNR seems to be required to benefit from more complex forward modeling. Both of these studies used dipole fitting for localizing sources. On the other hand, distributed inverse methods are generally expected to be more robust to forward modeling errors at the expense of being less focal (even in ideal conditions) as well as suffering from depth bias (Stenroos and Hauk, 2013).

4.5. Limitations

The comparison of different pipelines is likely biased in favor of SimNIBS-CHARM since the tissue priors used by the CHARM pipeline was built from this dataset. Although we used alternative priors from a four-fold split such that a particular subject never contributed to the prior used for generating its own segmentation, we might still expect some bias. On the other hand, it is clear from simple visual inspection of the segmentations that SimNIBS-CHARM is more accurate than FieldTrip-SPM and MNE-FS. As such, we believe the observed differences are robust and will generalize to other MRI datasets as well. Similar to SimNIBS-CHARM, the results of the model based on a template anatomy might also be biased given that it contains the same tissue classes and was constructed from the same pipeline. To minimize bias, we again used priors from a four-fold split. The reference models used in the anatomy study, although based on manual segmentations, were still not perfect and certain simplifications were made. For example, we failed to model brain conductivity anisotropy. However, assuming that the effect is similar across different pipelines, we do not expect it to strongly affect relative differences. Besides, white matter anisotropy seems to mostly affect certain deep sources and effects are generally small compared to other modeling aspects (Vorwerk et al., 2014).

In the electrode study, we used a dataset collected at our department as part of a previous project. Since electrode positions were digitized in this study, the experimenter may have been less careful with cap placement than if it was known beforehand that a template was to be used to model the electrode positions. As such, they might not constitute the best possible case when evaluating the fit of an electrode template. For example, visual inspection (as well as the standard deviation some of the electrodes across the dataset) suggested that in some subjects the cap was clearly misaligned in the anterior-posterior direction or twisted left or right. On the other hand, we believe that this is likely to happen in an actual experimental setting and as such our results may give a realistic impression of the errors that might be incurred using such methods.

5. Conclusion

In summary, it seems that topographical errors, as measured by RDM, are sensitive to errors in the modeled anatomy and the sensor positions, whereas magnitude errors, as measured by lnMAG, are sensitive to anatomical errors only. Consequently, magnitudes from a template model or simplified models cannot be trusted. Based on the results presented here, we suggest to use a forward model as realistic as possible. The availability of pipelines able to generate realistic geometrical descriptions of the human head have so far been limited. SimNIBS makes it easy to automatically generate models and the computational cost of doing so is not greater than running a FreeSurfer pipeline (as required when using MNE-Python) or the DUNEuro solver (used in the Fieldtrip pipeline). See Appendix A for an example.

If digitized electrode positions are unavailable, we suggest to use a template created by measuring the positions on a realistically shaped head template. Alternatively, average positions from a previous study using the same cap could be used. The template described here will be included in SimNIBS and we plan to add other caps in the future. Perhaps somewhat surprisingly, using a geometrical model based on template anatomy did not result in larger errors than coarse anatomical modeling of the individual anatomy. This suggests that if individual anatomical information is lacking then it might still be possible to generate a usable forward model based on standard anatomy. Even though we did not evaluate the impact of the observed effects on source localization accuracy, previous studies suggest that our main findings will translate to source localization although this association will likely be modulated by the noise level in the data and the particular choice of inverse solver.

Acknowledgements

JDN is supported by a PhD stipend from Sino-Danish Center for Education and Research. OP was supported by the Lundbeck foundation (grant R360-2021-395). AT was supported by the Lundbeck foundation (grants R244-2017-196 and R313-2019-622) and the NIH (grant 1RF1MH117428-01A1). The authors would like to thank Fan Wang (Chinese Academy of Sciences) for helpful discussions regarding electrophysiological data acquisition and modeling.

Data Availability

The data used in this study has been presented previously and we refer to the original articles for statements about data availability. The data used to compare the different forward modeling pipelines was presented by Farcito et al. (2019). The data used to compare electrode specification strategies was presented by Madsen et al. (2019) and Karabanov et al. (2021).

```
# Create head model based on T1- and (optionally) T2-weighted MRI scan. This creates
# /m2m_subid (directory)
charm subid t1.nii.gz t2.nii.gz
# Create montage file. This creates
# eeg_montage.csv
prepare_eeg_montage mne eeg_montage.csv info.fif trans.fif
# Run TDCS leadfield calculations. This creates
# /fem_subid/subid_leadfield_eeg_montage.hdf5
prepare_tdcs_leadfield to forward solution in MNE-Python format. This creates
# /fem_subid/subid_leadfield_eeg_montage_subsampling-10000-fwd.fif
# /fem_subid/subid_leadfield_eeg_montage_subsampling-10000-src.fif
prepare_eeg_forward mne subid fem_subid/subid_leadfield_eeg_montage_subsampling-10000-src.fif
prepare_eeg_forward mne subid fem_subid/subid_leadfield_eeg_montage_subsampling-10000-src.fif
```

Listing 1. Generating an EEG forward solution for MNE-Python with SimNIBS. info.fif can be a Raw, Epochs, Evoked or Info object. trans.fif is a Trans object containing an affine mapping between head coordinates and subject MRI coordinates. -fwd.fif, -morph.fif, and -src.fif are standard MNE-Python objects containing the forward solution, the sparse matrix to morph to *fsaverage* space, and the source space definition, respectively.

Appendix A. Generating EEG Forward Solutions With SimNIBS

Here we show how to generate EEG forward solutions with SimNIBS which can be used with MNE-Python (listing 1) or FieldTrip (listing 2). In both examples, the central gray matter surface of each hemisphere is subsampled to 10,000 nodes per hemisphere and a mapping to *fsaverage5* (which contains 10,242 nodes per hemisphere) is constructed. In both cases, the procedure is the same but the output format is different: when exporting to MNE-Python format (listing 1), the forward solution is in the head coordinate system used by MNE-Python whereas when exporting to FieldTrip (listing 2), the forward solution is in subject MRI coordinate system.

References

- Acar, Z.A., Acar, C.E., Makeig, S., 2016. Simultaneous head tissue conductivity and EEG source location estimation. NeuroImage 124, 168–180. doi:10.1016/j.neuroimage.2015.08.032.
- Acar, Z.A., Makeig, S., 2013. Effects of forward model errors on EEG source localization. Brain Topogr 26, 378-396. doi:10. 1007/s10548-012-0274-6.
- Ary, J.P., Klein, S.A., Fender, D.H., 1981. Location of sources of evoked scalp potentials: Corrections for skull and scalp thicknesses. IEEE Transactions on Biomedical Engineering BME-28, 447–452. doi:10.1109/tbme.1981.324817.
- Attene, M., 2010. A lightweight approach to repairing digitized polygon meshes. The Visual Computer 26, 1393–1406. doi:10.1007/s00371-010-0416-3.
- Ümit Aydin, Vorwerk, J., Küpper, P., Heers, M., Kugel, H., Galka, A., Hamid, L., Wellmer, J., Kellinghaus, C., Rampp, S., Wolters, C.H., 2014. Combining EEG and MEG for the reconstruction of epileptic activity using a calibrated realistic volume conductor model. PLoS ONE 9, 1–17. doi:10.1371/journal.pone.0093154.
- Azizollahi, H., Aarabi, A., Wallois, F., 2016. Effects of uncertainty in head tissue conductivity and complexity on EEG forward modeling in neonates. Human Brain Mapping 37, 3604–3622. doi:10.1002/hbm.23263.
- Bangera, N.B., Schomer, D.L., Dehghani, N., Ulbert, I., Cash, S., Papavasiliou, S., Eisenberg, S.R., Dale, A.M., Halgren, E., 2010. Experimental validation of the influence of white matter anisotropy on the intracranial eeg forward solution. Journal of Computational Neuroscience 29, 371–387. doi:10.1007/s10827-009-0205-z.
- Baumann, S.B., Wozny, D.R., Kelly, S.K., Meno, F.M., 1997. The electrical conductivity of human cerebrospinal fluid at body temperature. IEEE Transactions on Biomedical Engineering 44, 220–223. doi:10.1109/10.554770.
- Buzsáki, G., Anastassiou, C.A., Koch, C., 2012. The origin of extracellular fields and currents EEG, ECoG, LFP and spikes. Nature Reviews Neuroscience 13, 407–420. doi:10.1038/nrn3241.
- Chauveau, N., Franceries, X., Doyon, B., Rigaud, B., Morucci, J.P., Celsis, P., 2004. Effects of skull thickness, anisotropy, and inhomogeneity on forward EEG/ERP computations using a spherical three-dimensional resistor mesh model. Human Brain Mapping 21, 86–97. doi:10.1002/hbm.10152.

```
# Create head model based on T1- and (optionally) T2-weighted MRI scan. This creates
# /m2m_subid (directory)
charm subid t1.nii.gz t2.nii.gz
# Create montage file. This creates
#
    eeg_montage.csv
prepare_eeg_montage fieldtrip eeg_montage.csv info.mat trans.mat
# Run TDCS leadfield calculations. This creates
   /fem_subid/subid_leadfield_eeg_montage.hdf5
prepare_tdcs_leadfield subid eeg_montage.csv -s 10000 -o fem_subid
# Convert TDCS leadfield to forward solution in FieldTrip format. This creates
    /fem_subid/subid_leadfield_eeg_montage_subsampling-10000-fwd.mat
#
    /fem_subid/subid_leadfield_eeg_montage_subsampling-10000-morph.mat
    /fem_subid/subid_leadfield_eeg_montage_subsampling-10000-src.mat
#
prepare_eeg_forward fieldtrip subid fem_subid/subid_leadfield_eeg_montage.hdf5 --fsaverage 5
```

Listing 2. Generating an EEG forward solution for FieldTrip with SimNIBS. info.mat contains information about the electrodes and trans.mat contains an affine mapping between the coordinate system of the electrodes and MRI space. -fwd.mat contains the forward solution, -morph.mat is a sparse matrix which can be used to "morph" from subject space to *fsaverage* space, and -src.mat contains the source space definition.

- Conte, S., Richards, J.E., 2021. The influence of the head model conductor on the source localization of auditory evoked potentials. Brain Topography 34, 793–812. doi:10.1007/s10548-021-00871-z.
- Dahnke, R., Yotter, R.A., Gaser, C., 2013. Cortical thickness and central surface estimation. NeuroImage 65, 336-348. doi:10.1016/j.neuroimage.2012.09.050.
- Dalal, S.S., Rampp, S., Willomitzer, F., Ettl, S., 2014. Consequences of EEG electrode position error on ultimate beamformer source reconstruction performance. Frontiers in Neuroscience 8. doi:10.3389/fnins.2014.00042.
- Dale, A.M., Fischl, B., Sereno, M.I., 1999. Cortical surface-based analysis. NeuroImage 9, 179–194. doi:10.1006/nimg.1998.0395.
- Dannhauer, M., Lanfer, B., Wolters, C.H., Knösche, T.R., 2011. Modeling of the human skull in eeg source analysis. Human Brain Mapping 32, 1383–1399. doi:10.1002/hbm.21114.
- Darvas, F., Ermer, J.J., Mosher, J.C., Leahy, R.M., 2006. Generic head models for atlas-based EEG source analysis. Human Brain Mapping 27, 129–143. doi:10.1002/hbm.20171.
- Farcito, S., Puonti, O., Montanaro, H., Saturnino, G.B., Nielsen, J.D., Madsen, C.G., Siebner, H.R., Neufeld, E., Kuster, N., Lloyd, B.A., Thielscher, A., 2019. Accurate anatomical head segmentations: a data set for biomedical simulations, in: 2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), IEEE. pp. 193–208. doi:10.1109/embc.2019.8857041.
- Fiederer, L.D.J., Vorwerk, J., Lucka, F., Dannhauer, M., Yang, S., Dümpelmann, M., Schulze-Bonhage, A., Aertsen, A., Speck, O., Wolters, C.H., Ball, T., 2016. The role of blood vessels in high-resolution volume conductor head modeling of eeg. NeuroImage 128, 193–208. doi:10.1016/j.neuroimage.2015.12.041.
- Geuzaine, C., Remacle, J.F., 2009. Gmsh: A 3-d finite element mesh generator with built-in pre- and post-processing facilities. International Journal for Numerical Methods in Engineering 79, 1309–1331. doi:10.1002/nme.2579.
- Gramfort, A., 2013. MEG and EEG data analysis with MNE-python. Frontiers in Neuroscience 7. doi:10.3389/fnins.2013. 00267.
- Hallez, H., Vanrumste, B., Grech, R., Muscat, J., De Clercq, W., Vergult, A., D'Asseler, Y., Camilleri, K.P., Fabri, S.G., Van Huffel, S., 2007. Review on solving the forward problem in eeg source analysis. J NeuroEngineering Rehabil 4, 46. doi:10.1186/1743-0003-4-46.
- Hämäläinen, M., Hari, R., Ilmoniemi, R.J., Knuutila, J., Lounasmaa, O.V., 1993. Magnetoencephalographytextemdashtheory, instrumentation, and applications to noninvasive studies of the working human brain. Reviews of Modern Physics 65, 413–497. doi:10.1103/revmodphys.65.413.
- Homölle, S., Oostenveld, R., 2019. Using a structured-light 3d scanner to improve EEG source modeling with more accurate electrode positions. Journal of Neuroscience Methods 326, 108378. doi:10.1016/j.jneumeth.2019.108378.
- Hunter, J.D., 2007. Matplotlib: A 2d graphics environment. Computing in Science & Engineering 9, 90–95. doi:10.1109/mcse. 2007.55.
- Kaiboriboon, K., Lüders, H.O., Hamaneh, M., Turnbull, J., Lhatoo, S.D., 2012. EEG source imaging in epilepsy—practicalities and pitfalls. Nature Reviews Neurology 8, 498–507. doi:10.1038/nrneurol.2012.150.
- Kaipio, J., Somersalo, E., 2007. Statistical inverse problems: Discretization, model reduction and inverse crimes. Journal of Computational and Applied Mathematics 198, 493–504. doi:10.1016/j.cam.2005.09.027.
 Karabanov, A.N., Madsen, K.H., Krohne, L.G., Siebner, H.R., 2021. Does pericentral mu-rhythm "power" corticomotor
- Karabanov, A.N., Madsen, K.H., Krohne, L.G., Siebner, H.R., 2021. Does pericentral mu-rhythm "power" corticomotor excitability? – a matter of EEG perspective. Brain Stimulation 14, 713–722. doi:10.1016/j.brs.2021.03.017.

Kirschstein, T., Köhling, R., 2009. What is the source of the EEG? Clinical EEG and Neuroscience 40, 146–149. doi:10.1177/155005940904000305.

- Lanfer, B., Scherg, M., Dannhauer, M., Knösche, T.R., Burger, M., Wolters, C.H., 2012. Influences of skull segmentation inaccuracies on eeg source analysis. NeuroImage 62, 418–431. doi:10.1016/j.neuroimage.2012.05.006.
- Li, J., Wang, K., Zhu, S., He, B., 2007. Effects of holes on EEG forward solutions using a realistic geometry head model. Journal of Neural Engineering 4, 197–204. doi:10.1088/1741-2560/4/3/004.
- Lu, H., Lam, L.C.W., Ning, Y., 2019. Scalp-to-cortex distance of left primary motor cortex and its computational head model: Implications for personalized neuromodulation. CNS Neuroscience & Therapeutics 25, 1270–1276. doi:10.1111/cns.13204.
- Madsen, K.H., Karabanov, A.N., Krohne, L.G., Safeldt, M.G., Tomasevic, L., Siebner, H.R., 2019. No trace of phase: Corticomotor excitability is not tuned by phase of pericentral mu-rhythm. Brain Stimulation 12, 1261–1270. doi:10.1016/j.brs. 2019.05.005.
- Mahjoory, K., Nikulin, V.V., Botrel, L., Linkenkaer-Hansen, K., Fato, M.M., Haufe, S., 2017. Consistency of EEG source localization and connectivity estimates. NeuroImage 152, 590–601. doi:10.1016/j.neuroimage.2017.02.076.
- Marin, G., Guerin, C., Baillet, S., Garnero, L., Meunier, G., 1998. Influence of skull anisotropy for the forward and inverse problem in EEG: Simulation studies using FEM on realistic head models. Human Brain Mapping 6, 250-269. doi:10.1002/ (sici)1097-0193(1998)6:4<250::aid-hbm5>3.0.co;2-2.
- McCann, H., Beltrachini, L., 2022. Impact of skull sutures, spongiform bone distribution, and aging skull conductivities on the EEG forward and inverse problems. Journal of Neural Engineering 19, 016014. doi:10.1088/1741-2552/ac43f7.
- Meijs, J.W.H., Weier, O.W., Peters, M.J., Van Oosterom, A., 1989. On the numerical accuracy of the boundary element method (eeg application). IEEE transactions on biomedical engineering 36, 1038–1049. doi:10.1109/10.40805.
- van Mierlo, P., Vorderwülbecke, B.J., Staljanssens, W., Seeck, M., Vulliémoz, S., 2020. Ictal EEG source localization in focal epilepsy: Review and future perspectives. Clinical Neurophysiology 131, 2600–2616. doi:10.1016/j.clinph.2020.08.001.
- Montes-Restrepo, V., van Mierlo, P., Strobbe, G., Staelens, S., Vandenberghe, S., Hallez, H., 2014. Influence of skull modeling approaches on eeg source localization. Brain topography 27, 95–111. doi:10.1007/s10548-013-0313-y.
- Nguyen-Danse, D.A., Singaravelu, S., Chauvigné, L.A.S., Mottaz, A., Allaman, L., Guggisberg, A.G., 2021. Feasibility of reconstructing source functional connectivity with low-density EEG. Brain Topography 34, 709–719. doi:10.1007/s10548-021-00866-w.
- Nielsen, J.D., Madsen, K.H., Puonti, O., Siebner, H.R., Bauer, C., Madsen, C.G., Saturnino, G.B., Thielscher, A., 2018. Automatic skull segmentation from MR images for realistic volume conductor models of the head: Assessment of the stateof-the-art. NeuroImage 174, 587-598. doi:10.1016/j.neuroimage.2018.03.001.
- Oostenveld, R., Fries, P., Maris, E., Schoffelen, J.M., 2011. FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. Computational Intelligence and Neuroscience 2011, 1–9. doi:10.1155/2011/156869.
- Puonti, O., Iglesias, J.E., Leemput, K.V., 2016. Fast and sequence-adaptive whole-brain segmentation using parametric bayesian modeling. NeuroImage 143, 235–249. doi:10.1016/j.neuroimage.2016.09.011.
- Puonti, O., Leemput, K.V., Saturnino, G.B., Siebner, H.R., Madsen, K.H., Thielscher, A., 2020. Accurate and robust wholehead segmentation from magnetic resonance images for individualized head modeling. NeuroImage 219, 117044. doi:10. 1016/j.neuroImage.2020.117044.
- Ramon, C., Schimpf, P.H., Haueisen, J., 2006. Influence of head models on EEG simulations and inverse source localizations. BioMedical Engineering OnLine 5. doi:10.1186/1475-925x-5-10.
- Ruffini, G., 2015. Application of the reciprocity theorem to eeg inversion and optimization of eeg-driven transcranial current stimulation (tcs, including tdcs, tacs, trns). doi:10.48550/ARXIV.1506.04835.
- Saturnino, G.B., Madsen, K.H., Thielscher, A., 2019. Electric field simulations for transcranial brain stimulation using FEM: an efficient implementation and error analysis. Journal of Neural Engineering 16, 066032. doi:10.1088/1741-2552/ab41ba.
- Schoffelen, J.M., Gross, J., 2009. Source connectivity analysis with meg and eeg. Human Brain Mapping 30, 1857–1865. doi:10.1002/hbm.20745.
- Schrader, S., Westhoff, A., Piastra, M.C., Miinalainen, T., Pursiainen, S., Vorwerk, J., Brinck, H., Wolters, C.H., Engwer, C., 2021. DUNEuro—a software toolbox for forward modeling in bioelectromagnetism. PLOS ONE 16, e0252431. doi:10.1371/ journal.pone.0252431.
- Shirazi, S.Y., Huang, H.J., 2019. More reliable EEG electrode digitizing methods can reduce source estimation uncertainty, but current methods already accurately identify brodmann areas. Frontiers in Neuroscience 13. doi:10.3389/fnins.2019.01159.
- Stenroos, M., Hauk, O., 2013. Minimum-norm cortical source estimation in layered head models is robust against skull conductivity error. NeuroImage 81, 265–272. doi:10.1016/j.neuroimage.2013.04.086.
- Sullivan, C., Kaszynski, A., 2019. PyVista: 3d plotting and mesh analysis through a streamlined interface for the visualization toolkit (VTK). Journal of Open Source Software 4, 1450. doi:10.21105/joss.01450.
- Taubin, G., 1995. Curve and surface smoothing without shrinkage, in: Proceedings of IEEE International Conference on Computer Vision, IEEE Comput. Soc. Press. pp. 852–857. doi:10.1109/iccv.1995.466848.
- Thielscher, A., Antunes, A., Saturnino, G.B., 2015. Field modeling for transcranial magnetic stimulation: A useful tool to understand the physiological effects of tms?, in: 2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), Institute of Electrical & Electronics Engineers (IEEE). pp. 222–225. doi:10.1109/ embc.2015.7318340.
- Vallaghe, S., Clerc, M., 2009. A global sensitivity analysis of three- and four-layer EEG conductivity models. IEEE Transactions on Biomedical Engineering 56, 988–995. doi:10.1109/tbme.2008.2009315.
- Vorwerk, J., 2016. New Finite Element Methods to Solve the EEG/MEG Forward Problem. phdthesis. University of Münster. Vorwerk, J., Ümit Aydin, Wolters, C.H., Butson, C.R., 2019. Influence of head tissue conductivity uncertainties on EEG dipole
- reconstruction. Frontiers in Neuroscience 13. doi:10.3389/fnins.2019.00531.

Vorwerk, J., Cho, J.H., Rampp, S., Hamer, H., Knösche, T.R., Wolters, C.H., 2014. A guideline for head volume conductor modeling in eeg and meg. NeuroImage 100, 590–607. doi:10.1016/j.neuroimage.2014.06.040.

Vorwerk, J., Clerc, M., Burger, M., Wolters, C.H., 2012. Comparison of boundary element and finite element approaches to the EEG forward problem. Biomedical Engineering / Biomedizinische Technik 57. doi:10.1515/bmt-2012-4152.

van der Walt, S., Schönberger, J.L., Nunez-Iglesias, J., Boulogne, F., Warner, J.D., Yager, N., Gouillart, E., Yu, T., 2014. scikit-image: image processing in python. PeerJ 2, e453. doi:10.7717/peerj.453.

Wang, Y., Gotman, J., 2001. The influence of electrode location errors on EEG dipole source localization with a realistic head model. Clinical Neurophysiology 112, 1777–1780. doi:10.1016/s1388-2457(01)00594-6.

Weinstein, D., Zhukov, L., Johnson, C., 2000. Lead-field bases for electroencephalography source imaging. Annals of Biomedical Engineering 28, 1059–1065. doi:10.1114/1.1310220.

Wolters, C.H., Grasedyck, L., Hackbusch, W., 2004. Efficient computation of lead field bases and influence matrix for the FEM-based EEG and MEG inverse problem. Inverse Problems 20, 1099–1116. doi:10.1088/0266-5611/20/4/007.

Zienkiewicz, O.C., Zhu, J.Z., 1992. The superconvergent patch recovery and posteriori error estimates. part 1: The recovery technique. International Journal for Numerical Methods in Engineering 33, 1331–1364. doi:10.1002/nme.1620330702.

Supplementary Material for "Evaluating the Influence of Anatomical Accuracy and Electrode Positions on EEG Forward Solutions"

S1 SimNIBS Validation

Figure S1.1 shows that the numerical accuracy of finite element method (FEM) generally deteriorates as sources get closer to the sensors in a realistically shaped three-compartment model. The boundary element method (BEM) solution from MNE-Python was used as reference. We also compare with a BEM solution from FieldTrip to show how this is more numerically accurate for high-eccentricity sources (at these resolutions). Figures S1.2 and S1.3 show error as a function of eccentricity in a realistically shaped five-compartment model, however, in this case the resulting impact of eccentricity is less clear. A high resolution FEM model from SimNIBS was used as reference.

S2 Additional Illustrations of Volume Conductor Geometry

To support the claims in the main text about the general differences between the volume conductor models from the different pipelines in the anatomy study, we show a few more examples in figures S2.4 to S2.7.



Figure S1.1. Average error (relative difference measure (RDM) or logarithm of the magnitude error (lnMAG)) per source location as a function of eccentricity for a high and low resolution FEM model and a BEM model in a realistically shaped three-compartment model. The BEM solution from MNE-Python was used as a reference.



Figure S1.2. Average error (RDM or lnMAG) per source location as a function of eccentricity for the FEM models.



Figure S1.3. Average error (RDM or lnMAG) per source location as a function of eccentricity for the BEM model.



Figure S2.4. Tissue compartments as estimated by each of the pipelines in the anatomy study for subject 3.


Figure S2.5. Tissue compartments as estimated by each of the pipelines in the anatomy study for subject 4.



Figure S2.6. Tissue compartments as estimated by each of the pipelines in the anatomy study for subject 5.



Figure S2.7. Tissue compartments as estimated by each of the pipelines in the anatomy study for subject 6.

Appendix C

Optimizing Template-Based Channel Layouts in Electroencephalography

This is a paper draft.

Optimizing Template-Based Channel Layouts in Electroencephalography

Jesper Duemose Nielsen^{1,2,3}, Oula Puonti^{2,4}, Axel Thielscher^{2,4}, and Kristoffer Hougaard Madsen^{1,2}

¹Department of Applied Mathematics and Computer Science, Technical University of Denmark, Denmark

²Danish Research Centre for Magnetic Resonance, Centre for Functional and Diagnostic Imaging and Research, Copenhagen University Hospital Amager and Hvidovre, Denmark ³Sino-Danish Centre for Education and Research, Aarhus, Denmark

⁴Department of Health Technology, Technical University of Denmark, Kgs. Lyngby, Denmark

Abstract

Accurate knowledge of electrode positions in electroencephalography (EEG) is important for constructing forward models which are needed when performing source analysis. Ideally, electrodes are accurately digitized using a dedicated system during data acquisition, however, this may not always be possible. In such situations, using a set of template positions corresponding to the EEG cap being used is the only option (e.g., as provided by the manufacturer or as measured by placing the cap on a reference model). This, however, decreases accuracy of the forward model and impairs source analysis. In this work we show that it is possible to optimize the template positions once they have been transformed to subject space. This is achieved by performing a few distance and angle measurements according to anatomical landmarks and selected electrodes at the time of the experiment. The template positions are then deformed so as to conform with these measurements. The procedure improved the overall accuracy of the layout and we observed the largest improvements in situations where large discrepancies between the initial and the target positions exist. On the other hand, it makes little difference in cases where the initial registration is reasonably accurate.

1 Introduction

In situations where we need to know the locations of the electrodes in electroencephalography (EEG) (e.g., when performing source analysis), the best option is to measure these positions directly using a dedicated system (Dalal et al., 2014; Homölle and Oostenveld, 2019). If this is not possible, we may transform a template representation of the given montage to adapt it to the head shape of a particular subject. However, upon placing an EEG cap on the human head, the layout tends to deform nonlinearly (e.g., due to head shape and size as well as any adjustments made to the cap by the experimenter). Hence, the accuracy with which a transformed template is able to fit the actual positions of the electrodes depends on the degree to which such template positions captures the specific deformations present in a particular layout. If the transformation is obtained by matching landmark positions or head shape, then any deviations from the idealized or average positions as defined by the template positions are unlikely to be captured. Here we describe how—given a few measurements of distances and angles between channels and landmarks—electrode positions from a template layout can be optimized to better fit the "actual" positions for a particular subject. The hypothesis is that if such measurements are available, it will allow us to adapt the initial estimate of the channel positions by encouraging that they agree with said measurements.

In a previous study, Munck, Vijn, and Spekreijse (1991) used a dense set of distance measurements between electrodes and reference points (landmarks) to locate electrodes on a sphere. However, this procedure is very time consuming since all landmark-electrode distances have to be measured, e.g., 192 measurements for a 64 channel layout and three landmarks. Using a similar strategy, Le et al. (1998) suggested a method for locating electrode positions in a 10-10 layout based on measurements of only 14 electrode-electrode distances and 9 electrode-landmark distances. The remaining channel positions were estimated by interpolation. Both of these methods have the attractive property that they require no equipment other than a pair of calipers to measure distancee. However, the method of Munck, Vijn, and Spekreijse (1991) is very time-consuming and that of Le et al. (1998) is only applicable to caps with a 10-10 (or similar) layout.

Another approach was suggested by He and Estepp (2013) who estimated a high-density array of electrode positions based on digitizing only a subset (19 to be specific) of the electrodes and then interpolated the remaining positions from a dense measurement of the same layout on a reference head. This has the advantage that it is generalizable to any cap layout, however, it requires the availability of a set of reference positions as well as a digitization system. The method proposed here also requires a set of reference positions by only uses a simple set of measurements of distances and angles avoiding the need for a digitization system.

Here, we use simulations to test the hypothesis that optimizing electrode positions based on a few measurements of distances and angles is feasible and can improve their accuracy. Specifically, we optimize electrode positions obtained by transforming a set of template positions (measured on a 3D model of the MNI head) and we use actual digitized positions as reference when evaluating the procedure.

2 Methods

We use data from Madsen et al. (2019) and Karabanov et al. (2021) which contain 32 subjects for which EEG electrode positions were digitized using an infrared optic stereo tracking system from Localite (https://localite.de). As part of the digitization procedure, these coordinates were transformed to subject magnetic resonance imaging (MRI) space by matching of fiducial landmarks. The montage used in these studies was an EasyCap BC-TMS64-X21 transcranial magnetic stimulation (TMS) compatible cap with a modified M10 (equidistant) layout containing 64 electrodes (https://www.easycap.de).



Figure 1. An example of how to perform angle measurements. Angles are measured relative to the major axes as defined in the Neuromag coordinate system (see main text). The yellow plane is the xy plane, the red plane is the yz plane. Green points denote landmarks, blue points are the initial positions of the template, and red points are the target positions (i.e., the digitized electrode positions). The angles (and distances) are measured using the red points during an experiment. The polar axis insets define how angles are measured. For each landmark, we use the horizontal axis and set 0° to the direction in which this axis points (e.g., for RPA, this would be positive y). Positive angles always in the direction of the positive z axis. Black arcs and magenta lines denote angle and distance measurements, respectively.

The template positions, which we obtained by measuring the same EEG cap on a 3D printed model of the MNI head, were transformed from MNI space to subject MRI space using the spatial normalization parameters estimated by complete head anatomy reconstruction method (CHARM) (Puonti et al., 2020), the head modeling pipeline of SimNIBS 4. The transformation is nonlinear and is estimated during segmentation of a T1- and a T2-weighted MRI scan from each subject.

Inter-electrode and electrode-landmark distances are fairly straightforward to measure using either a pair of calipers or perhaps even a ruler. As our implementation works on Euclidean distances, the former is expected to be preferable. Angles are trickier to measure as they have to be measured wrt. some axis. In this work, we only use simulated measurements and could therefore simply use the axes from the subject MRI space. However, in a real-world setting, we suggest measuring angles relative a set of major axes which are defined by landmarks on the head, e.g., the head coordinate system used by the Neuromag system which is constructed from the nasion, left preauricular (LPA), and right preauricular (RPA) (see figure 1 for an example). In the current implementation, points are projected on either the xz- or yz-plane depending on the landmark-electrode pair such that measurements related to LPA and RPA are done in the yz-plane and measurements related to nasion and inion are done in the *xz*-plane. A consequence of this is that although the procedure is very flexible (in principle allowing any combination of measurements of distances and angles irrespective of whether these relate electrode-landmark pairs or pairs of electrodes) some angles may be difficult to estimate whereas this is not the case for distances.

It is likely that the accuracy of the procedure scales with the number of measurements, as each measurement serves to constrain the final layout. In this study, we simulate measurements of distances and angles between four electrodelandmarks pairs (see figure 2). These were chosen because they describe the



Figure 2. Illustration of the electrodes (red) used for generating the measurements. Each electrode was paired with its closest landmark, i.e., (43, nasion), (64, inion), (70, LPA), and (71, RPA), and their distance an angle were measured. The lines show what electrodes are considered neighbors.

positioning of the cap both left to right and front to back but also because we believe them to be relatively easy to measure. To investigate the sensitivity of the method to noise, we run the experiment using different amounts of error in the distance and angle measurements. Specifically, we draw errors from a normal distribution with mean zero and a standard deviation defined such that 95 % of the errors are within a certain range (e.g., ± 6 mm for distance and $\pm 10^{\circ}$ for angles). For distances, we use 0 mm to 10 mm in steps of 2 mm and for angles we use 0° to 20° in steps of 4°. We sample each combination 10 times.

Let $D = \{(i_1, j_1), \ldots, (i_m, j_m)\}$ and $A = \{(i_1, j_1), \ldots, (i_n, j_n)\}$ be the sets of points (electrodes and/or landmarks) for which distance and angle measurements have been made, then we formulate the following optimization problem for finding a set of the parameters (associated with a certain parameterization of the electrode positions, see below) such that the final electrode positions agree with the measurements,

minimize
$$\sum_{i \in E} \sum_{j \in N(i)} ||d_{ij} - d_{ij}^0||^2$$
subject to $\begin{vmatrix} d_{ij} - d_{ij}^{\text{measured}} \\ a_{ij} - a_{ij}^{\text{measured}} \end{vmatrix} \leq t_d, \quad i, j \in D$ (1)

Here, E is the set of all electrodes and N(i) is the set of electrodes in the neighborhood of i (as defined by the lines in figure 2), and t_d and t_a are the tolerances associated with distances and angles (in this work we used 2.5 mm and 5 degrees), respectively. These tolerances are related to the expected accuracy of the measurements. d_{ij} denotes the Euclidean distance between i and j, d_{ij}^0 is the initial distance between i and j (i.e., before optimization), and d_{ij}^{measured} is the measured distance between i and j.

To (approximately) constrain the electrodes to the skin surface during the optimization, we express them in spherical coordinates by fitting a sphere to each electrode and its neighbors (see below). This is necessary in order to avoid excessive warping of the montage. We fit each sphere following the procedure described by Dorst (2014) in which the problem of minimizing the least squared distances between a sphere and a set of points is expressed in a special form (using conformal geometric algebra) which allows an approximate solution to be derived in closed form as it is transformed to an eigenvalue problem. To stabilize the fit, we use a minimum of five neighbors per electrode for a total of six points. If an electrode has fewer than five neighbors (see figure 2) then we add the required number of nearest electrodes. During the subsequent optimization, we keep the radius of the estimated spheres fixed and optimize the polar and azimuthal angles.

Although the human head as a whole is not spherical, we believe that a locally fitted sphere (here, one per electrode) is able to capture the local curvature reasonably well. This is true even in the temporal regions where the spheres are generally very large thus resembling a plane. Besides being able to approximate the local curvature of the skin surface, another key assumption of this reparameterization is that the initial positions are reasonable and thus will not change considerably during the optimization as each sphere only approximates the curvature in a certain subset of space.

To optimize equation (1), we use a trust region solver with constraints as implemented in SciPy (Virtanen et al., 2020) which we run for 100 iterations. We use the squared error between distances in the cost function to discourage outliers as this encourages deformation of the layout as a whole as opposed to only locally, i.e., to avoid moving only the electrodes subjected to constraints.

3 Results

Figure 3 shows the result of optimizing electrode positions for two examples. Figure 4 shows density estimates of channel position errors before and after optimization across all subjects and noise levels. It is clear that there is an overall improvement in the positions. However, although we see an improvement on average, it is not clear how consistent this effect is. For example, one could imagine that the position of some electrodes may actually get more inaccurate after optimization. To investigate this, we look at the difference in error between the initial and optimized layout, i.e., $e_{\text{difference}} = e_{\text{optimized}} - e_{\text{initial}}$. Here, positive and negative values correspond to smaller errors in the initial and optimized layout, respectively.

Figure 6 shows probability densities of $e_{\text{difference}}$ as a function of noise level and subject. Effects of noise are relatively modest whereas effects of subject are much more pronounced suggesting that some subjects benefit substantially from the optimization whereas others do not. See also figure 5 which shows densities of the errors before and after optimization. Improvement seem to be consistent over all subjects (expect perhaps for subject 21) although the magnitude varies substantially. Hence, the cost function seems to be reasonable at ensuring a sensible layout.

It is also interesting to explore which electrodes benefit the most (or the least) from optimization. Figure 7 show the spatial distribution of mean of $e_{\text{difference}}$. Unsurprisingly, the errors are generally low for the electrodes to which constraints are attached (i.e., where the measurements were made). Additionally, the optimizations performed here seem to be good at correcting AP misalignments which were the largest source of error we observed in the initial



Figure 3. The result of optimizing electrode positions in two subjects *without* measurement noise. The top row corresponds to subject 0 in figure 5 and top left in figures 8 to 10. This is a case where optimization does not make much difference. The bottom row corresponds to subject 4 in figure 5 and fifth from the left in the first row of figures 8 to 10. This is a case where optimization makes a big difference by successfully realigning the template with the digitized positions in the anterior-posterior (AP) direction. Red, blue, and green are target, initial, and optimized positions, respectively.



Figure 4. Distribution of channel location errors before and after optimization (across all noise levels and subjects) wrt. the digitized positions. Densities were estimated using a Gaussian kernel density estimator.



Figure 5. Errors in electrode positions for the initial and optimized layouts for each subject. The densities are over all noise levels.



Figure 6. Probability densities of $e_{\text{difference}}$. Negative values on the x-axis mean that the error is lower after optimization whereas positive values mean that the error increased after optimization (i.e., the initial estimate was better). Left. The curves correspond to different noise levels (as coded by the inserted matrix) averaged over subjects. It is apparent that the final electrode positions are fairly insensitive to (these levels of) noise in the measurements. Right. Each curve correspond to individual subjects averaged over noise levels. Subjects were sorted by mean $e_{\text{difference}}$ and colored from purple (low) to yellow (high).

positions. Since this particular correction relies primarily on knowing whether a selected electrode was below or above the inion (and how much) it is fairly insensitive to the angular errors introduced here as the error before optimization can be as large as 180°. Using a similar argument, the angle is also not very critical for the frontal electrode in terms of AP alignment. Measurements between LPA and RPA and temporal electrodes will generally be more sensitive to angular errors. On the other hand, all landmark-electrode measurements are expected to be equally sensitive to distance errors.

Figures 8 to 10 show the spatial distribution of $e_{\text{difference}}$ as a function of each subject. As was clear from figure 6 there is large variability between subjects in how much benefit (harm) the optimization provides (does) and this variability is much larger than the decrease in performance due to noise in the measurements. In particular, subjects where the cap was poorly aligned in the AP direction benefit the most (and consistently) from this procedure.

4 Discussion and Conclusions

Our method shares features with some of the previously mentioned methods. As in He and Estepp (2013), we require a template layout to be available, however, we dispense with the requirement of having a digitization system and use simple measurements of distances and angles instead. Similar to Le et al. (1998), our method requires only a few measurements (here we use eight) making it feasible in an experimental setting whereas Le et al. (1998) requires 23 distance measurements and is only applicable to 10-10 (or similar) layouts.

Our method generally performs worse than the previous studies, however, it is important to keep in mind that we used considerably fewer measurements here. In principle, given enough measurements (distances would suffice), it can perform arbitrarily well.

One obvious limitation of this small study is that we only used simulated



Figure 7. Spatial distribution of the mean of $e_{\text{difference}}$ for each noise level. Highlighted electrodes correspond to statistically significant differences (improvements) between the initial and the optimized layouts at $\alpha = 0.05$ as assessed by a paired *t*-test using 9999 permutations. Correction for multiple comparisons was achieved by forming a permutation-based empirical null distribution of the maximum test statistic (over electrodes) (Nichols and Holmes, 2001).

Distance Error 0.00 mm; Angular Error 0.00°



Figure 8. Spatial distribution of $e_{\text{difference}}$ for each subject at 0 mm distance error and 0° angular error. A blue (red) star denotes statistically significant improvement (worsening) for this subject after optimizing the layout. The tests were performed as described in the text to figure 7 except that the distribution of maximal test statistic was over subjects in this case.



Figure 9. Spatial distribution of $e_{\text{difference}}$ for each subject at 4 mm distance error and 8° angular error. A blue (red) star denotes statistically significant improvement (worsening) for this subject after optimizing the layout. The tests were performed as described in the text to figure 7 except that the distribution of maximal test statistic was over subjects in this case.

Distance Error 10.00 mm; Angular Error 20.00°



Figure 10. Spatial distribution of $e_{\text{difference}}$ for each subject at 10 mm distance error and 20° angular error. A blue (red) star denotes statistically significant improvement (worsening) for this subject after optimizing the layout. The tests were performed as described in the text to figure 7 except that the distribution of maximal test statistic was over subjects in this case.

data in the sense that the measurements were constructed by using the actual positions of the electrodes with some additive noise. However, unless the level of errors in real-world measurements are substantially larger than the ones simulated here, we expect performance to be comparable.

Accuracy is important but so is the ease with which measurements can be made. Distance measurements are straightforward to make. By using a caliper we also expect them to be reasonably accurate. Angle measurements are trickier but can be made using a protractor. Additionally, the measurements will have to be typed into a computer which may be cumbersome. However, preliminary experimentation indicates that an accuracy well within the simulated bounds can easily be archived in less than two minutes.

The choice of running the optimization for 100 iterations is somewhat arbitrary, although we found this to give satisfactory results during our testing. However, it is important to note that the purpose of the cost function as it is formulated here is simply to facilitate a reasonable deformation of the initial layout. Since it is unlikely that the initial distances are identical to the actual distances, we are not necessarily interested in finding a minimum of the cost function *per se*. The primary objective is to satisfy the constraints (within the given tolerances) and it is difficult to say whether optimizing the cost function beyond a certain point after this has been achieved will lead to improvements in the fit or not. This is likely to vary from subject to subject.

Taken together we have shown that optimization of electrodes position based on a few measurements of distances and/or angles can be used to adapt an initial template fit to individual subjects. Performance vary substantially across subjects but is quite insensitive to moderate levels of noise in distance and, in particular, angle measurements.

References

- Dalal, Sarang S. et al. (2014). "Consequences of EEG electrode position error on ultimate beamformer source reconstruction performance." Frontiers in Neuroscience 8. DOI: 10.3389/fnins.2014.00042.
- Dorst, Leo (2014). "Total Least Squares Fitting of k-Spheres in n-D Euclidean Space Using an (n+2)-D Isometric Representation." Journal of Mathematical Imaging and Vision 50.3, pp. 214–234. DOI: 10.1007/s10851-014-0495-2.
- He, Ping and Justin R. Estepp (2013). "A practical method for quickly determining electrode positions in high-density EEG studies." *Neuroscience Letters* 541, pp. 73–76. DOI: 10.1016/j.neulet.2013.02.028.
- Homölle, Simon and Robert Oostenveld (2019). "Using a structured-light 3D scanner to improve EEG source modeling with more accurate electrode positions." *Journal of Neuroscience Methods* 326, p. 108378. DOI: 10.1016/j. jneumeth.2019.108378.
- Karabanov, Anke Ninija et al. (2021). "Does pericentral mu-rhythm "power" corticomotor excitability? – A matter of EEG perspective." Brain Stimulation 14.3, pp. 713–722. DOI: 10.1016/j.brs.2021.03.017.
- Le, Jian et al. (1998). "A rapid method for determining standard 10/10 electrode positions for high resolution EEG studies." *Electroencephalography* and Clinical Neurophysiology 106.6, pp. 554–558. DOI: 10.1016/s0013-4694(98)00004-2.
- Madsen, Kristoffer Hougaard et al. (2019). "No trace of phase: Corticomotor excitability is not tuned by phase of pericentral mu-rhythm." *Brain Stimulation* 12.5, pp. 1261–1270. DOI: 10.1016/j.brs.2019.05.005.
- Munck, J. C. De, P. C. M. Vijn, and H. Spekreijse (1991). "A practical method for determining electrode positions on the head." *Electroencephalography and Clinical Neurophysiology* 78.1, pp. 85–87. DOI: 10.1016/0013-4694(91) 90023-w.
- Nichols, Thomas E. and Andrew P. Holmes (2001). "Nonparametric permutation tests for functional neuroimaging: A primer with examples." *Human Brain Mapping* 15.1, pp. 1–25. DOI: 10.1002/hbm.1058.
- Puonti, Oula et al. (2020). "Accurate and robust whole-head segmentation from magnetic resonance images for individualized head modeling." *NeuroImage* 219, p. 117044. DOI: 10.1016/j.neuroimage.2020.117044.
- Virtanen, Pauli et al. (2020). "SciPy 1.0: fundamental algorithms for scientific computing in Python." Nature Methods 17.3, pp. 261–272. DOI: 10.1038/ s41592-019-0686-2.

148 APPENDIX C. OPTIMIZING ELECTRODE POSITIONS IN EEG

Appendix D

Additional Background Material

D.1 Calculating the Potential in a Multi-Layer Sphere Model

The potential at \mathbf{r} on the outer surface of a three-layer concentric spherical model generated by a dipole at a certain position in the innermost compartment is given by

$$u\left(\mathbf{r}\right) = \frac{1}{4\pi\sigma} \sum_{n=1}^{\infty} \frac{2n+1}{n} \left(\frac{\|\mathbf{r}\|}{r_3}\right)^{n-1} \cdot \left[nm_r P_n\left(\cos\alpha\right) + m_t P_n^1\left(\cos\alpha\right)\cos\beta\right] \left[\frac{\lambda\left(2n+1\right)^2}{\left(n+1\right)d_n}\right]$$
(D.1)

where the conductivity of the inner- and outermost compartments are given by σ_{13} (modeling brain and scalp) and that of the central compartment by σ_2 (modeling skull). P_n and P_n^1 are Legendre and associated Legendre polynomials of degree n, respectively, and

$$\lambda = \frac{\sigma_2}{\sigma_{13}} \tag{D.2}$$

$$d_{n} = [\lambda (n+1) + n] \left[\frac{\lambda n}{n+1} \right]$$

$$(1-\lambda) [\lambda (n+1) + n] \left(f_{1}^{2n+1} - f_{2}^{2n+1} \right)$$

$$+ n (\lambda - 1)^{2} \left(\frac{f_{1}}{f_{2}} \right)^{2n+1}$$
(D.3)

$$f_1 = \frac{r_1}{r_3} \tag{D.4}$$

$$f_2 = \frac{r_2}{r_3}$$
(D.5)

where r_i indicates the radius of the *i*th sphere and m_r and m_t are the radial and tangential components of the dipole moment, respectively. α and β are the spherical angles (polar and azimuth, respectively) of **r** relative to the position vector of the dipole (Ary et al., 1981; Hallez et al., 2007). This model is fast to calculate and, for sufficiently high n, very accurate.

D.2 Calculating the Mapping between Cartesian and Reference Elements

We start by defining the linear mapping from a cartesian element to the reference element. Specifically, we shall define this for a tetrahedron in three dimensions. The cartesian coordinates of a point $\mathbf{r} = [x, y, z]^{\top}$ may be expressed in terms of the reference (barycentric) coordinates $\lambda = [\lambda_1, \lambda_2, \lambda_3, \lambda_4]^{\top}$ like so

$$x = \lambda_1 x_1 + \lambda_2 x_2 + \lambda_3 x_3 + \lambda_4 x_4 \tag{D.6a}$$

$$y = \lambda_1 y_1 + \lambda_2 y_2 + \lambda_3 y_3 + \lambda_4 y_4 \tag{D.6b}$$

$$z = \lambda_1 z_1 + \lambda_2 z_2 + \lambda_3 z_3 + \lambda_4 z_4 \tag{D.6c}$$

where x_i , y_i , and z_i refer to the coordinates of the *i*th node on the element. Exploiting that $\lambda_4 = 1 - \lambda_1 - \lambda_2 - \lambda_3$ we get

$$x = \lambda_1(x_1 - x_4) + \lambda_2(x_2 - x_4) + \lambda_3(x_3 - x_4) + x_4$$
 (D.7a)

$$y = \lambda_1(y_1 - y_4) + \lambda_2(y_2 - y_4) + \lambda_3(x_3 - y_4) + y_4$$
 (D.7b)

$$z = \lambda_1(z_1 - z_4) + \lambda_2(z_2 - z_4) + \lambda_3(x_3 - z_4) + z_4$$
 (D.7c)

which in matrix form is

$$\mathbf{r} = \mathbf{T}\lambda + \mathbf{r}_4 \tag{D.8}$$

where

$$\mathbf{T} = \begin{pmatrix} x_1 - x_4 & x_2 - x_4 & x_3 - x_4 \\ y_1 - y_4 & y_2 - y_4 & y_3 - y_4 \\ z_1 - z_4 & z_2 - z_4 & z_3 - z_4 \end{pmatrix}$$
(D.9)

implements the affine map $T(\lambda) = \mathbf{T}\lambda$ from barycentric coordinates to cartesian coordinates. Likewise, we have

$$\lambda = \mathbf{T}^{-1} \left(\mathbf{r} - \mathbf{r}_4 \right) \tag{D.10}$$

where

$$\mathbf{T}^{-1} = \frac{1}{\det \mathbf{T}} \begin{bmatrix} (\mathbf{t}_1 \times \mathbf{t}_2)^\top \\ (\mathbf{t}_2 \times \mathbf{t}_0)^\top \\ (\mathbf{t}_0 \times \mathbf{t}_1)^\top \end{bmatrix}$$
(D.11)

with \mathbf{t}_i being the columns of \mathbf{T} .

To map the integrals from reference space to cartesian space, we will need the determinant of the Jacobian (i.e., the deformation gradient) of this operator. The Jacobian is given by

$$\mathbf{J}_{T} = \nabla T(\lambda) = \begin{pmatrix} \frac{\partial T_{x}}{\partial \lambda_{1}} & \frac{\partial T_{x}}{\partial \lambda_{2}} & \frac{\partial T_{x}}{\partial \lambda_{3}} \\ \frac{\partial T_{y}}{\partial \lambda_{1}} & \frac{\partial T_{y}}{\partial \lambda_{2}} & \frac{\partial T_{y}}{\partial \lambda_{3}} \\ \frac{\partial T_{z}}{\partial \lambda_{1}} & \frac{\partial T_{z}}{\partial \lambda_{2}} & \frac{\partial T_{z}}{\partial \lambda_{3}} \end{pmatrix} = \mathbf{T}$$
(D.12)

150

and its determinant by the triple product, i.e.,

$$\det \mathbf{J}_T = \det \mathbf{T} = \mathbf{t}_0 \cdot (\mathbf{t}_1 \times \mathbf{t}_2) \tag{D.13a}$$

which is the signed volume of the parallelepiped spanned by the columns of \mathbf{T} (which, in this case, is equals six times the volume of the corresponding tetrahedron) representing the scaling factor of the mapping between reference and cartesian element. Evaluating the *i*th basis function on the physical element, *e*, can be achieved by

$$\mathbf{h}_{i}^{e}(\mathbf{x}) = \mathbf{h}_{i}^{r}\left(T_{e}^{-1}(\mathbf{x})\right) \tag{D.14a}$$

$$=\mathbf{h}_{i}^{r}(\lambda) \tag{D.14b}$$

and the gradient can be evaluated using the chain rule

$$\nabla \mathbf{h}_{i}^{e}(\mathbf{x}) = \nabla \mathbf{h}_{i}^{r} \left(T_{e}^{-1}(\mathbf{x}) \right) \nabla T_{e}^{-1}(\mathbf{x})$$
(D.15a)

$$= \nabla \mathbf{h}_i^r(\lambda) \mathbf{T}_e^{-1} \tag{D.15b}$$

where \mathbf{x} and λ are cartesian and reference coordinates, respectively, and \mathbf{h}_i^e and \mathbf{h}_i^r refer to the *i*th basis function evaluated on *e* (the physical element) and *r* (the reference element).

D.3 The Geometry of Ax = b

This section explores the vector spaces in which ${\bf x}$ and ${\bf b}$ exist in the following equation system

$$\mathbf{A}\mathbf{x} = \mathbf{b}.\tag{D.16}$$

The column space or image of \mathbf{A} is the space spanned by the columns of \mathbf{A} . The rank of \mathbf{A} is the dimension of the column space of \mathbf{A} , i.e., the number of linearly independent columns of \mathbf{A} , and is also equal to the dimension of the row space of \mathbf{A} (i.e., the column space of \mathbf{A}^{\top}). The null space of \mathbf{A} is the set of all vectors which are mapped to the zero vector upon application of \mathbf{A} .

Let **A** be an $m \times n$ matrix with rank $r \leq \min(m, n)$ which maps between vector spaces $U \in \mathbb{R}^n$ and $V \in \mathbb{R}^m$. We have

$$\operatorname{rank}(\mathbf{A}) + \operatorname{nullity}(\mathbf{A}) = \dim V = m$$
 (D.17)

where

$$\operatorname{rank}(\mathbf{A}) = \operatorname{dim}\operatorname{image}(\mathbf{A})$$
 (D.18)

$$nullity(\mathbf{A}) = \dim nullspace(\mathbf{A}). \tag{D.19}$$

image(\mathbf{A}) is a subspace of \mathbb{R}^m with dimensionality r spanned by the columns of \mathbf{A} . image(\mathbf{A}^\top) is a subspace in \mathbb{R}^n with dimensionality r spanned by the columns of \mathbf{A}^\top . nullspace(\mathbf{A}) is a subspace of \mathbb{R}^n with dimensionality n-r and is equal to image(\mathbf{A}^\top)^{\perp}, i.e., the image of \mathbf{A}^\top is the orthogonal complement of the null space of \mathbf{A} . nullspace(\mathbf{A}^\top) is a subspace in \mathbb{R}^m with dimensionality m-rand is equal to image(\mathbf{A})^{\perp}, i.e., the image of \mathbf{A} is the orthogonal complement of the null space of \mathbf{A}^\top . We can explore what \mathbf{A} and its pseudoinverse, \mathbf{A}^+ , do using the SVD. Specifically, let

$$\mathbf{A} = \mathbf{U}\mathbf{S}\mathbf{V}^{\top} \tag{D.20}$$

$$\mathbf{A} = \mathbf{U}_r \mathbf{S}_r \mathbf{V}_r^{\top} \tag{D.21}$$

be the full and compact SVD, respectively, where the latter is restricted to rank(**A**). The columns of **U** form an orthonormal basis for the vector space which **A** maps to (the codomain of **A**) and \mathbf{U}_r is an orthonormal basis for image(**A**). Similarly, the columns of **V** form an orthonormal basis for the vector space which **A** maps from (the domain) and \mathbf{V}_r is an orthonormal basis for image (\mathbf{A}^{\top}). Applying **A** amounts to projecting a vector **x** onto image (\mathbf{A}^{\top}), scaling it, and rotating to image (**A**). As to the pseudoinverse, we have

$$\mathbf{A}^+ = \mathbf{V}\mathbf{S}^+\mathbf{U}^\top \tag{D.22}$$

$$\mathbf{A}^{+} = \mathbf{V}_{r} \mathbf{S}_{r}^{-1} \mathbf{U}_{r}^{\top} \tag{D.23}$$

where the pseudoinverse of \mathbf{S} is the reciprocal of every non-zero entry (transposed). Applying \mathbf{A}^+ amounts to projecting a vector \mathbf{b} onto image (\mathbf{A}), scaling it, and rotating to image (\mathbf{A}^\top). Thus, any component in the null space of \mathbf{A} and \mathbf{A}^\top is effectively ignored when applying \mathbf{A} and \mathbf{A}^+ , respectively. For example, if equation (D.16) is underdetermined, then the solution with minimum norm is a vector which is orthogonal to nullspace(\mathbf{A}) as it would otherwise contain a component to which \mathbf{A} is insensitive thus giving the same \mathbf{b} but with a larger norm.

D.4 Image Formation in MRI

Spatial encoding can be achieved by applying gradients along each of the major $\operatorname{axes}^1(x, y, z)$ resulting in a gradient on the overall magnetic field. Concurrent application of a gradient and a frequency-selective RF pulse enables excitation of a particular slice along this axis (e.g., z). This is called slice selection and slice thickness is controlled by the strength of the gradient. A side effect of this is that different amounts of phase will be accumulated throughout the selected volume (resulting in dephasing), however, this can be reverse by a applying a negative gradient of the half duration after excitation.

Spatial encoding within a slice is achieved by phase and frequency encoding. Brief application of a gradient along an axis perpendicular to the slice excitation (here for example y) results in varying amounts of phase accumulation for spins along this axis. During readout, i.e., when the receiver coil is turned on and signal is being acquired, a third gradient is applied orthogonally to the previous ones (here x) resulting in different precession frequencies depending on location of the spins. Thus, the induced voltages in the receiver coils are a mixture of signal from spins of different precession frequencies at different phase lags.

Frequency encoding can be achieved by an initial dephasing and subsequent rephasing thus generating an echo. In the beginning of the readout, the spins are highly different in terms of frequency. At time TE, there are no differences and at the end of the readout, there is again a high (but opposite to the initial)

152

¹Gradients may also be combined such that slices can be of arbitrary orientation.

difference in frequency of spins and consequently the signal has been encoded in different spatial frequencies.

Since the sequence of phase encodings is equivalent to a frequency encoding², the signal can be considered a function of two spatial frequency components denoted k_x and k_y ,

$$k_x = \frac{\gamma}{2\pi} G_x t_x \tag{D.24}$$

$$k_y = \frac{\gamma}{2\pi} G_y t_y \tag{D.25}$$

where the gradient, G, is assumed to be constant throughout the entire time interval t_x or t_y for the frequency and phase encoding, respectively. The signal from a single excitation is given by

$$s(k_x, k_y) \propto \int_x \int_y \rho(x, y) e^{-i2\pi(k_x x + k_y y)} dx dy$$
 (D.26)

which is the Fourier transform of the spin density over the excited slice. To retrieve spin density information as a function of position within a slice, this has to be inverted which can be achieved using the two-dimensional inverse Fourier transform,

$$\rho(x,y) \propto \int_{k_x} \int_{k_y} s(k_x,k_y) e^{i2\pi(k_x x + k_y y)} dk_x dk_y.$$
(D.27)

The result is a complex image and typically the magnitude is what is used in most imaging applications. To estimate ρ , the space in which the spatial frequencies reside, termed k-space, needs to be sampled sufficiently. All else being equal, image space field of view (FOV) can be manipulated by the density with which k-space is sampled (related to how much the gradient changes between successive measurements) whereas image space resolution is determined by k-space coverage (related to the strength of the gradients) since the edges correspond to high frequency content in the image.

²At each application of the phase encoding gradient (with varying strength) the transverse magnetization will have traveled shorter or further. Thus, if viewed as a function of time, the encoding will effectively be one of frequency.

Abbreviations

AP	anterior-posterior
BEM	boundary element method
BOLD	blood oxygen level dependent
CAT12	computational anatomy toolbox 12
CHARM	complete head anatomy reconstruction method
CSF	cerebrospinal fluid
CT	computed tomography
CTF	cross talk function
CV	cross validation
DOF	degrees of freedom
dSPM	dynamic statistical parametric mapping
ECG	electrocardiography
EEG	electroencephalography
EOG	electrooculography
FEM	finite element method
fMRI	functional magnetic resonance imaging
FOV	field of view
GFP	global field power
GMM	Gaussian mixture model
GOF	goodness-of-fit
ICA	independent component analysis
LAURA	local autoregressive averages
LCMV	linearly constrained minimum variance
lnMAG	logarithm of the magnitude error
LORETA	low resolution electromagnetic tomography
LPA	left preauricular
LW	Ledoit-Wolf
MAP	maximum a posteriori
MEG	magnetoencephalography

ML MNE MRI MUSIC	maximum likelihood minimum norm estimate magnetic resonance imaging multiple signal classification
NAI	neural activity index
PCA PDE	principal component analysis partial differential equation
PET	positron emission tomography
PLE	peak localization error
PSF	point spread function
RAP-MUSIC	recursively applied MUSIC
RDM	relative difference measure
RF	radio frequency
RHS	right-hand side
RLS	recursive least squares
RPA	right preauricular
sLORETA	standardized low resolution electromagnetic tomography
SNR	signal-to-noise ratio
SPR	superconvergent patch recovery
SVD	singular value decomposition
TE	echo time
TES	transcranial electrical stimulation
TMS	transcranial magnetic stimulation
TR	repetition time

156

References

- Ablin, P., Cardoso, J.-F., & Gramfort, A. (2018). Faster independent component analysis by preconditioning with hessian approximations. *IEEE Trans*actions on Signal Processing, 66(15), 4040–4049. https://doi.org/10. 1109/tsp.2018.2844203
- Acar, Z. A., & Makeig, S. (2013). Effects of forward model errors on EEG source localization. Brain Topogr, 26(3), 378–396. https://doi.org/10.1007/ s10548-012-0274-6
- Antonakakis, M., Schrader, S., Wollbrink, A., Oostenveld, R., Rampp, S., Haueisen, J., & Wolters, C. H. (2019). The effect of stimulation type, head modeling, and combined EEG and MEG on the source reconstruction of the somatosensory p20/n20 component. *Human Brain Mapping*, 40(17), 5011–5028. https://doi.org/10.1002/hbm.24754
- Ary, J. P., Klein, S. A., & Fender, D. H. (1981). Location of sources of evoked scalp potentials: Corrections for skull and scalp thicknesses. *IEEE Transactions on Biomedical Engineering*, *BME-28*(6), 447–452. https: //doi.org/10.1109/tbme.1981.324817
- Ashburner, J., & Friston, K. J. (2005). Unified segmentation. NeuroImage, 26(3), 839–851. https://doi.org/10.1016/j.neuroimage.2005.02.018
- Awada, K. A., Jackson, D. R., Williams, J. T., Wilton, D. R., Baumann, S. B., & Papanicolaou, A. C. (1997). Computational aspects of finite element modeling in EEG source localization. *IEEE Transactions on Biomedical Engineering*, 44 (8), 736–752. https://doi.org/10.1109/10.605431
- Billot, B., Greve, D. N., Puonti, O., Thielscher, A., Van Leemput, K., Fischl, B., Dalca, A. V., & Iglesias, J. E. (2021). Synthese: Domain randomisation for segmentation of brain scans of any contrast and resolution. https: //doi.org/10.48550/ARXIV.2107.09559
- Brookes, M. J., Stevenson, C. M., Barnes, G. R., Hillebrand, A., Simpson, M. I. G., Francis, S. T., & Morris, P. G. (2007). Beamformer reconstruction of correlated sources using a modified source model. *NeuroImage*, 34(4), 1454–1465. https://doi.org/10.1016/j.neuroimage.2006.11.012
- Buchner, H., Knoll, G., Fuchs, M., Rienäcker, A., Beckmann, R., Wagner, M., Silny, J., & Pesch, J. (1997). Inverse localization of electric dipole current sources in finite element models of the human head. *Electroen*cephalography and Clinical Neurophysiology, 102(4), 267–278. https: //doi.org/10.1016/s0013-4694(96)95698-9
- Buzsáki, G., Anastassiou, C. A., & Koch, C. (2012). The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, 13(6), 407–420. https://doi.org/10.1038/nrn3241

- Chen, F., Hallez, H., & Staelens, S. (2010). Influence of skull conductivity perturbations on EEG dipole source analysis. *Medical Physics*, 37(8), 4475– 4484. https://doi.org/10.1118/1.3466831
- Cho, J.-H., Vorwerk, J., Wolters, C. H., & Knösche, T. R. (2015). Influence of the head model on eeg and meg source connectivity analyses. *NeuroImage*, 110, 60–77. https://doi.org/10.1016/j.neuroimage.2015.01.043
- Conte, S., & Richards, J. E. (2021). The influence of the head model conductor on the source localization of auditory evoked potentials. *Brain Topography*, 34(6), 793–812. https://doi.org/10.1007/s10548-021-00871-z
- Dalal, S. S., Rampp, S., Willomitzer, F., & Ettl, S. (2014). Consequences of EEG electrode position error on ultimate beamformer source reconstruction performance. *Frontiers in Neuroscience*, 8. https://doi.org/10.3389/ fnins.2014.00042
- Dale, A. M., Fischl, B., & Sereno, M. I. (1999). Cortical surface-based analysis. *NeuroImage*, 9(2), 179–194. https://doi.org/10.1006/nimg.1998.0395
- Dannhauer, M., Lanfer, B., Wolters, C. H., & Knösche, T. R. (2011). Modeling of the human skull in eeg source analysis. *Human Brain Mapping*, 32(9), 1383–1399. https://doi.org/10.1002/hbm.21114
- Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. Journal of Neuroscience Methods, 134(1), 9–21. https://doi.org/ 10.1016/j.jneumeth.2003.10.009
- de Munck, J. C. (1992). A linear discretization of the volume conductor boundary integral equation using analytically integrated elements (electrophysiology application). *IEEE Transactions on Biomedical Engineering*, 39(9), 986–990. https://doi.org/10.1109/10.256433
- de Peralta Menendez, R. G., Murray, M. M., Michel, C. M., Martuzzi, R., & Andino, S. L. G. (2004). Electrical neuroimaging based on biophysical constraints. *NeuroImage*, 21(2), 527–539. https://doi.org/10.1016/j. neuroimage.2003.09.051
- Engemann, D. A., & Gramfort, A. (2015). Automated model selection in covariance estimation and spatial whitening of meg and eeg signals. *NeuroIm*age, 108, 328–342. https://doi.org/10.1016/j.neuroimage.2014.12.040
- Erdbrügger, T.-R. (2021). Cutfem forward modeling for geometries with touching surfaces in bioelectromagnetism [Master's thesis, University of Münster].
- Farcito, S., Puonti, O., Montanaro, H., Saturnino, G. B., Nielsen, J. D., Madsen, C. G., Siebner, H. R., Neufeld, E., Kuster, N., Lloyd, B. A., & Thielscher, A. (2019). Accurate anatomical head segmentations: A data set for biomedical simulations. 2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), 193–208. https://doi.org/10.1109/embc.2019.8857041
- Fiederer, L. D. J., Vorwerk, J., Lucka, F., Dannhauer, M., Yang, S., Dümpelmann, M., Schulze-Bonhage, A., Aertsen, A., Speck, O., Wolters, C. H., & Ball, T. (2016). The role of blood vessels in high-resolution volume conductor head modeling of eeg. *NeuroImage*, 128, 193–208. https://doi.org/10.1016/j.neuroimage.2015.12.041
- Fischl, B., Sereno, M. I., Tootell, R. B., & Dale, A. M. (1999). High-resolution intersubject averaging and a coordinate system for the cortical surface.

Human brain mapping, 8, 272–284. https://doi.org/10.1002/(sici)1097-0193(1999)8:4<272::aid-hbm10>3.0.co;2-4

- Fox, C., Nicholls, G. K., & Tan, S. M. (2010). An introduction to inverse problems.
- Göksu, C., Hanson, L. G., Siebner, H. R., Ehses, P., Scheffler, K., & Thielscher, A. (2018). Human in-vivo brain magnetic resonance current density imaging (MRCDI). *NeuroImage*, 171, 26–39. https://doi.org/10.1016/ j.neuroimage.2017.12.075
- Gramfort, A. (2013). MEG and EEG data analysis with MNE-python. Frontiers in Neuroscience, 7. https://doi.org/10.3389/fnins.2013.00267
- Gramfort, A., Kowalski, M., & Hämäläinen, M. (2012). Mixed-norm estimates for the m/eeg inverse problem using accelerated gradient methods. *Physics in medicine and biology*, 57, 1937–1961. https://doi.org/10. 1088/0031-9155/57/7/1937
- Haacke, E. M., Brown, R. W., Thompson, M. R., & Venkatesan, R. (1999). Magnetic resonance imaging: Physical principles and sequence design. Wiley-Liss.
- Habib, M. A., Ibrahim, F., Mohktar, M. S., Kamaruzzaman, S. B., Rahmat, K., & Lim, K. S. (2016). Ictal EEG source imaging for presurgical evaluation of refractory focal epilepsy. World Neurosurgery, 88, 576–585. https: //doi.org/10.1016/j.wneu.2015.10.096
- Hallez, H., Vanrumste, B., Grech, R., Muscat, J., De Clercq, W., Vergult, A., D'Asseler, Y., Camilleri, K. P., Fabri, S. G., & Van Huffel, S. (2007). Review on solving the forward problem in eeg source analysis. J Neuro-Engineering Rehabil, 4(1), 46. https://doi.org/10.1186/1743-0003-4-46
- Hämäläinen, M. S., & Sarvas, J. (1989). Realistic conductivity geometry model of the human head for interpretation of neuromagnetic data. *IEEE Transactions on Biomedical Engineering*, 36(2), 165–171. https://doi. org/10.1109/10.16463
- Hauk, O., Stenroos, M., & Treder, M. (2019). EEG/MEG source estimation and spatial filtering: The linear toolkit. In *Magnetoencephalography* (pp. 167–203). Springer International Publishing. https://doi.org/10. 1007/978-3-030-00087-5 85
- He, P., Wilson, G., & Russell, C. (2004). Removal of ocular artifacts from electroencephalogram by adaptive filtering. *Medical & biological engineering & computing*, 42, 407–412.
- He, P., Wilson, G., Russell, C., & Gerschutz, M. (2007). Removal of ocular artifacts from the EEG: A comparison between time-domain regression method and adaptive filtering method using simulated data. *Medical & Biological Engineering & Computing*, 45(5), 495–503. https://doi.org/ 10.1007/s11517-007-0179-9
- Homölle, S., & Oostenveld, R. (2019). Using a structured-light 3d scanner to improve EEG source modeling with more accurate electrode positions. *Journal of Neuroscience Methods*, 326, 108378. https://doi.org/10. 1016/j.jneumeth.2019.108378
- Hyvärinen, A., & Oja, E. (2000). Independent component analysis: Algorithms and applications. *Neural Networks*, 13(4-5), 411–430. https://doi.org/ 10.1016/s0893-6080(00)00026-5
- Jas, M., Engemann, D. A., Bekhti, Y., Raimondo, F., & Gramfort, A. (2017). Autoreject: Automated artifact rejection for MEG and EEG data. *Neu*-

roImage, 159, 417–429. https://doi.org/10.1016/j.neuroimage.2017.06.030

- Jas, M., Larson, E., Engemann, D. A., Leppäkangas, J., Taulu, S., Hämäläinen, M., & Gramfort, A. (2018). A reproducible MEG/EEG group study with the MNE software: Recommendations, quality assessments, and good practices. Frontiers in Neuroscience, 12. https://doi.org/10.3389/ fnins.2018.00530
- Kaipio, J., & Somersalo, E. (2007). Statistical inverse problems: Discretization, model reduction and inverse crimes. Journal of Computational and Applied Mathematics, 198(2), 493–504. https://doi.org/10.1016/j.cam. 2005.09.027
- Karabanov, A. N., Madsen, K. H., Krohne, L. G., & Siebner, H. R. (2021). Does pericentral mu-rhythm "power" corticomotor excitability? – a matter of EEG perspective. Brain Stimulation, 14(3), 713–722. https://doi.org/ 10.1016/j.brs.2021.03.017
- Kessy, A., Lewin, A., & Strimmer, K. (2015). Optimal whitening and decorrelation. The American Statistician. https://doi.org/10.1080/00031305. 2016.1277159
- Kirschstein, T., & Köhling, R. (2009). What is the source of the EEG? Clinical EEG and Neuroscience, 40(3), 146–149. https://doi.org/10.1177/155005940904000305
- Lanfer, B., Scherg, M., Dannhauer, M., Knösche, T. R., Burger, M., & Wolters, C. H. (2012). Influences of skull segmentation inaccuracies on eeg source analysis. *NeuroImage*, 62(1), 418–431. https://doi.org/10.1016/j. neuroimage.2012.05.006
- Lin, F.-H., Witzel, T., Ahlfors, S. P., Stufflebeam, S. M., Belliveau, J. W., & Hämäläinen, M. S. (2006). Assessing and improving the spatial accuracy in MEG source localization by depth-weighted minimum-norm estimates. *NeuroImage*, 31(1), 160–171. https://doi.org/10.1016/j. neuroimage.2005.11.054
- Luck, S. J. (2005). An introduction to the event-related potential technique (cognitive neuroscience). The MIT Press.
- Lütkenhöner, B. (1998). Dipole source localization by means of maximum likelihood estimation. i. theory and simulations. *Electroencephalography and Clinical Neurophysiology*, 106(4), 314–321. https://doi.org/10.1016/s0013-4694(97)00140-5
- Madsen, K. H., Karabanov, A. N., Krohne, L. G., Safeldt, M. G., Tomasevic, L., & Siebner, H. R. (2019). No trace of phase: Corticomotor excitability is not tuned by phase of pericentral mu-rhythm. *Brain Stimulation*, 12(5), 1261–1270. https://doi.org/10.1016/j.brs.2019.05.005
- Mannan, M. M. N., Jeong, M. Y., & Kamran, M. A. (2016). Hybrid ICA—regression: Automatic identification and removal of ocular artifacts from electroencephalographic signals. *Frontiers in Human Neuroscience*, 10. https://doi.org/10.3389/fnhum.2016.00193
- Marin, G., Guerin, C., Baillet, S., Garnero, L., & Meunier, G. (1998). Influence of skull anisotropy for the forward and inverse problem in EEG: Simulation studies using FEM on realistic head models. *Human Brain Mapping*, 6(4), 250–269. https://doi.org/10.1002/(sici)1097-0193(1998)6:4<250:: aid-hbm5>3.0.co;2-2

- McCann, H., & Beltrachini, L. (2022). Impact of skull sutures, spongiform bone distribution, and aging skull conductivities on the EEG forward and inverse problems. *Journal of Neural Engineering*, 19(1), 016014. https: //doi.org/10.1088/1741-2552/ac43f7
- Michel, C. M., & Koenig, T. (2018). EEG microstates as a tool for studying the temporal dynamics of whole-brain neuronal networks: A review. *NeuroImage*, 180, 577–593. https://doi.org/10.1016/j.neuroimage. 2017.11.062
- Michel, C. M., Murray, M. M., Lantz, G., Gonzalez, S., Spinelli, L., & de Peralta, R. G. (2004). Eeg source imaging. *Clinical Neurophysiology*, 115(10), 2195–2222. https://doi.org/10.1016/j.clinph.2004.06.001
- Molins, A., Stufflebeam, S. M., Brown, E. N., & Hämäläinen, M. S. (2008). Quantification of the benefit from integrating MEG and EEG data in minimum l2-norm estimation. *NeuroImage*, 42(3), 1069–1077. https: //doi.org/10.1016/j.neuroimage.2008.05.064
- Montes-Restrepo, V., van Mierlo, P., Strobbe, G., Staelens, S., Vandenberghe, S., & Hallez, H. (2014). Influence of skull modeling approaches on eeg source localization. *Brain topography*, 27(1), 95–111. https://doi.org/ 10.1007/s10548-013-0313-y
- Mosher, J. C., Baillet, S., & Leahy, R. M. (2003). Equivalence of linear approaches in bioelectromagnetic inverse solutions. *IEEE Workshop on Statistical Signal Processing*, 2003. https://doi.org/10.1109/ssp.2003. 1289402
- Mosher, J. C., Baillet, S., & Leahy, R. M. (1999). Eeg source localization and imaging using multiple signal classification approaches. *Journal of Clinical Neurophysiology*, 16(3), 225–238. https://doi.org/10.1097/ 00004691-199905000-00004
- Murray, M. M., Brunet, D., & Michel, C. M. (2008). Topographic ERP analyses: A step-by-step tutorial review. Brain Topography, 20(4), 249–264. https://doi.org/10.1007/s10548-008-0054-5
- Neugebauer, F., Möddel, G., Rampp, S., Burger, M., & Wolters, C. H. (2017). The effect of head model simplification on beamformer source localization. Frontiers in Neuroscience, 11. https://doi.org/10.3389/fnins. 2017.00625
- Nielsen, J. D., Madsen, K. H., Puonti, O., Siebner, H. R., Bauer, C., Madsen, C. G., Saturnino, G. B., & Thielscher, A. (2018). Automatic skull segmentation from MR images for realistic volume conductor models of the head: Assessment of the state-of-the-art. *NeuroImage*, 174, 587–598. https://doi.org/10.1016/j.neuroimage.2018.03.001
- Oostendorp, T. F., & van Oosterom, A. (1989). Source parameter estimation in inhomogeneous volume conductors of arbitrary shape. *IEEE Transac*tions on Biomedical Engineering, 36(3), 382–391. https://doi.org/10. 1109/10.19859
- Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J.-M. (2011). FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Computational Intelligence and Neuroscience*, 2011, 1–9. https://doi.org/10.1155/2011/156869
- Parkkonen, L., Fujiki, N., & Mäkelä, J. P. (2009). Sources of auditory brainstem responses revisited: Contribution by magnetoencephalography. *Human Brain Mapping*, 30(6), 1772–1782. https://doi.org/10.1002/hbm.20788

- Pascual-Marqui, R. D. (2002). Standardized low-resolution brain electromagnetic tomography (sloreta): Technical details. Methods and findings in experimental and clinical pharmacology, 24 Suppl D, 5–12.
- Pascual-Marqui, R. D., Michel, C. M., & Lehmann, D. (1994). Low resolution electromagnetic tomography: A new method for localizing electrical activity in the brain. *International Journal of Psychophysiology*, 18(1), 49–65. https://doi.org/10.1016/0167-8760(84)90014-x
- Petersen, K. B., & Pedersen, M. S. (2012). The matrix cookbook [Version 20121115].
- Puonti, O., Iglesias, J. E., & Leemput, K. V. (2016). Fast and sequence-adaptive whole-brain segmentation using parametric bayesian modeling. *NeuroImage*, 143, 235–249. https://doi.org/10.1016/j.neuroimage.2016.09. 011
- Puonti, O., Leemput, K. V., Saturnino, G. B., Siebner, H. R., Madsen, K. H., & Thielscher, A. (2020). Accurate and robust whole-head segmentation from magnetic resonance images for individualized head modeling. *NeuroImage*, 219, 117044. https://doi.org/10.1016/j.neuroimage.2020. 117044
- Ramon, C., Schimpf, P. H., & Haueisen, J. (2006). Influence of head models on EEG simulations and inverse source localizations. *BioMedical Engineering OnLine*, 5(1). https://doi.org/10.1186/1475-925x-5-10
- Rice, J. K., Rorden, C., Little, J. S., & Parra, L. C. (2013). Subject position affects eeg magnitudes. *NeuroImage*, 64, 476–484. https://doi.org/10. 1016/j.neuroimage.2012.09.041
- Rivet, B., Souloumiac, A., Attina, V., & Gibert, G. (2009). xDAWN algorithm to enhance evoked potentials: Application to brain computer interface. *IEEE Transactions on Biomedical Engineering*, 56(8), 2035–2043. https://doi.org/10.1109/tbme.2009.2012869
- Rivet, B., Cecotti, H., Souloumiac, A., Maby, E., & Mattout, J. (2011). Theoretical analysis of xDAWN algorithm: application to an efficient sensor selection in a P300 BCI. 19th European Signal Processing Conference (EUSIPCO 2011), 1382–1386.
- Ruffini, G. (2015). Application of the reciprocity theorem to eeg inversion and optimization of eeg-driven transcranial current stimulation (tcs, including tdcs, tacs, trns). https://doi.org/10.48550/ARXIV.1506.04835
- Saturnino, G. B., Madsen, K. H., & Thielscher, A. (2019). Electric field simulations for transcranial brain stimulation using FEM: An efficient implementation and error analysis. *Journal of Neural Engineering*, 16(6), 066032. https://doi.org/10.1088/1741-2552/ab41ba
- Saturnino, G. B., Thielscher, A., Madsen, K. H., Knösche, T. R., & Weise, K. (2019). A principled approach to conductivity uncertainty analysis in electric field calculations. *NeuroImage*, 188, 821–834. https://doi.org/ 10.1016/j.neuroimage.2018.12.053
- Scherg, M. (1990). Fundamentals of dipole source potential analysis. In F. Grandori, M. Hoke, & G. L. Romani (Eds.), Auditory evoked magnetic fields and electric potentials (pp. 40–69, Vol. 6). Advances in Audiology.
- Schoffelen, J.-M., & Gross, J. (2009). Source connectivity analysis with meg and eeg. Human Brain Mapping, 30(6), 1857–1865. https://doi.org/10. 1002/hbm.20745

- Schrader, S., Westhoff, A., Piastra, M. C., Miinalainen, T., Pursiainen, S., Vorwerk, J., Brinck, H., Wolters, C. H., & Engwer, C. (2021). DUNEuro—a software toolbox for forward modeling in bioelectromagnetism (C. Papadelis, Ed.). *PLOS ONE*, 16(6), e0252431. https://doi.org/10.1371/journal.pone.0252431
- Sekihara, K. (2008). Adaptive spatial filters for electromagnetic brain imaging. Springer.
- Shirazi, S. Y., & Huang, H. J. (2019). Influence of mismarking fiducial locations on EEG source estimation. 2019 9th International IEEE/EMBS Conference on Neural Engineering (NER). https://doi.org/10.1109/ner. 2019.8717065
- Smith, N. B., & Webb, A. (2011). Introduction to medical imaging: Physics, engineering, and clinical applications. Cambridge University Press.
- Song, J., Davey, C., Poulsen, C., Luu, P., Turovets, S., Anderson, E., Li, K., & Tucker, D. (2015). EEG source localization: Sensor density and head surface coverage. *Journal of Neuroscience Methods*, 256, 9–21. https: //doi.org/10.1016/j.jneumeth.2015.08.015
- Stenroos, M., & Hauk, O. (2013). Minimum-norm cortical source estimation in layered head models is robust against skull conductivity error. *NeuroIm*age, 81, 265–272. https://doi.org/10.1016/j.neuroimage.2013.04.086
- Stenroos, M., Hunold, A., & Haueisen, J. (2014). Comparison of three-shell and simplified volume conductor models in magnetoencephalography. *NeuroImage*, 94, 337–348. https://doi.org/10.1016/j.neuroimage.2014. 01.006
- Tanner, D., Morgan-Short, K., & Luck, S. J. (2015). How inappropriate highpass filters can produce artifactual effects and incorrect conclusions in ERP studies of language and cognition. *Psychophysiology*, 52(8), 997– 1009. https://doi.org/10.1111/psyp.12437
- Thielscher, A., Antunes, A., & Saturnino, G. B. (2015). Field modeling for transcranial magnetic stimulation: A useful tool to understand the physiological effects of tms? 2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), 222–225. https://doi.org/10.1109/embc.2015.7318340
- Vallaghe, S., & Clerc, M. (2009). A global sensitivity analysis of three- and four-layer EEG conductivity models. *IEEE Transactions on Biomedical Engineering*, 56(4), 988–995. https://doi.org/10.1109/tbme.2008. 2009315
- van Mierlo, P., Vorderwülbecke, B. J., Staljanssens, W., Seeck, M., & Vulliémoz, S. (2020). Ictal EEG source localization in focal epilepsy: Review and future perspectives. *Clinical Neurophysiology*, 131(11), 2600–2616. https://doi.org/10.1016/j.clinph.2020.08.001
- Veen, B. D. V., Drongelen, W. V., Yuchtman, M., & Suzuki, A. (1997). Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. *IEEE Transactions on Biomedical Engineer*ing, 44(9), 867–880. https://doi.org/10.1109/10.623056
- Vorwerk, J. (2016). New finite element methods to solve the EEG/MEG forward problem [Doctoral dissertation, University of Münster].
- Vorwerk, J., Cho, J.-H., Rampp, S., Hamer, H., Knösche, T. R., & Wolters, C. H. (2014). A guideline for head volume conductor modeling in eeg and meg.

NeuroImage, 100, 590-607. https://doi.org/10.1016/j.neuroimage.2014. 06.040

- Vorwerk, J., Hanrath, A., Wolters, C. H., & Grasedyck, L. (2019). The multipole approach for EEG forward modeling using the finite element method. *NeuroImage*, 201, 116039. https://doi.org/10.1016/j.neuroimage.2019. 116039
- Wakeman, D. G., & Henson, R. N. (2015). A multi-subject, multi-modal human neuroimaging dataset. *Scientific Data*, 2, 150001. https://doi.org/10. 1038/sdata.2015.1
- Wang, G., & Ren, D. (2013). Effect of brain-to-skull conductivity ratio on eeg source localization accuracy. *BioMed Research International*, 2013, 1– 10. https://doi.org/10.1155/2013/459346
- Wang, Y., & Gotman, J. (2001). The influence of electrode location errors on EEG dipole source localization with a realistic head model. *Clinical Neurophysiology*, 112(9), 1777–1780. https://doi.org/10.1016/s1388-2457(01)00594-6
- Weinstein, D., Zhukov, L., & Johnson, C. (2000). Lead-field bases for electroencephalography source imaging. Annals of Biomedical Engineering, 28(9), 1059–1065. https://doi.org/10.1114/1.1310220
- Widmann, A., Schröger, E., & Maess, B. (2015). Digital filter design for electrophysiological data – a practical approach. Journal of Neuroscience Methods, 250, 34–46. https://doi.org/10.1016/j.jneumeth.2014.08.002
- Wolters, C. H., Anwander, A., Tricoche, X., Weinstein, D., Koch, M. A., & MacLeod, R. S. (2006). 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. https://doi.org/10.1016/ j.neuroimage.2005.10.014
- Wolters, C. H., Grasedyck, L., & Hackbusch, W. (2004). Efficient computation of lead field bases and influence matrix for the FEM-based EEG and MEG inverse problem. *Inverse Problems*, 20(4), 1099–1116. https:// doi.org/10.1088/0266-5611/20/4/007
- Wolters, C. H., Köstler, H., Möller, C., Härdtlein, J., & Anwander, A. (2007). Numerical approaches for dipole modeling in finite element method based source analysis. *International Congress Series*, 1300, 189–192. https://doi.org/10.1016/j.ics.2007.02.014
- Yan, Y., Nunez, P. L., & Hart, R. T. (1991). Finite-element model of the human head: Scalp potentials due to dipole sources. *Medical & Biological Engineering & Computing*, 29(5), 475–481. https://doi.org/10.1007/ bf02442317
- Ziegler, E., Chellappa, S. L., Gaggioni, G., Ly, J. Q. M., Vandewalle, G., André, E., Geuzaine, C., & Phillips, C. (2014). A finite-element reciprocity solution for eeg forward modeling with realistic individual head models. *NeuroImage*, 103, 542–551. https://doi.org/10.1016/j.neuroimage.2014. 08.056
- Zienkiewicz, O. C., & Zhu, J. Z. (1992). The superconvergent patch recovery anda posteriori error estimates. part 1: The recovery technique. International Journal for Numerical Methods in Engineering, 33(7), 1331– 1364. https://doi.org/10.1002/nme.1620330702