

A Novel Framework for Head Imaging with Electrical Impedance Tomography

Markus Fabian Jehl



A thesis submitted for the degree of
Doctor of Philosophy
of
University College London.

Department of Medical Physics and Biomedical Engineering
Department of Mathematics
University College London

February 2016

Declaration

I, Markus Fabian Jehl, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Signed:

Abstract

Electrical Impedance Tomography (EIT) is a medical imaging technology with the potential to locate focal epilepsy, monitor patients with traumatic brain injury and diagnose stroke. EIT usually images conductivity changes in time or in frequency of the applied current and measured voltages. While it is nowadays clinically used for monitoring lung ventilation, its application in head imaging is complicated by the geometry of the head, containing tissues with strongly varying conductivities.

The aim of this thesis is to provide a novel framework for EIT head imaging by addressing the requirements for higher modelling accuracy throughout the imaging process. An introduction to EIT, its applications for head imaging and the two main components of EIT image reconstructions is given in chapter 2. A procedure for generating more accurate head models is presented in chapter 3 and is used to evaluate, whether subject specific head models are required for EIT imaging. To speed up simulations of current flow through the head and the computation of the Jacobian matrix required for image reconstructions, a fast parallel forward solver is implemented and validated in chapter 4. Stability of time-difference image reconstructions with respect to electrode modelling errors is addressed in chapter 5, followed by an evaluation of modelling error impacts on multi-frequency imaging in chapter 6. The findings of chapters 5 and 6 are finally combined in chapter 7 to recover electrode positions in multi-frequency stroke imaging, thereby reducing image artefacts and making stroke diagnosis with EIT feasible.

Acknowledgements

I would like to thank my supervisors Dr. Timo Betcke and Prof. David Holder for their support throughout my PhD research. They were always there for me when I needed advice, guidance, ideas or feedback, and also respected my views and opinions, for which I am very grateful.

The work presented in this thesis would not have been possible without the help of many present and past members of the UCL EIT group. I am particularly grateful to Dr. Kirill Aristovich, who helped me on numerous occasions with advice and code for meshing, segmentation and image reconstructions, and to Dr. James Avery, who assisted me on tank experiments and was always willing to help with technical issues. Further, the implementation of the parallel forward solver PEITS would not have been possible without the support from Prof. Andreas Dedner. I would also like to thank Dr. Tom Dowrick for his tireless workstation debugging, Dr. Nicolas Chaulet for introducing me to the electrode correction methods, Mayo Faulkner for helping me with the head segmentation and Prof. Simon Arridge for being something like a ‘third supervisor’. I am further grateful to my thesis examiners Dr. Nick Polydorides and Dr. Bangti Jin for their valuable feedback.

Throughout my PhD studies, I profited endlessly from the good atmosphere in our research group and the Department of Medical Physics and Biomedical Engineering. I would like to thank all my friends from the group and elsewhere, and the hacky sack team Nir and Mayo for the helpful distraction. This work was enabled by the full funding I was fortunate to receive, the provision of which was shared between the Department of Mathematics, the Department of Medical Physics and Biomedical Engineering and the HEAP funding scheme.

Relocating to a different country and working on an autonomous, personal project can be a daunting and isolating experience, but I was fortunate to have rarely felt that way. In contrast, I experienced it as personally and professionally enriching, thanks to the social and personal stability provided to me by my partner Lea Raible, whom I am profoundly grateful to for her understanding and encouragement. I am very thankful to my parents Ingrid Doberer and Guy Jehl, for their continuous encouragement and support, and also to Lea’s parents Thomas and Ulrike for all their kindness.

Contents

1	Overview	1
1.1	Introduction	1
1.2	Purpose	3
1.3	Statement of Originality	3
1.4	List of Publications	4
2	Literature Review	7
2.1	Electrical Impedance Tomography	7
2.1.1	Bioimpedance	7
2.1.2	Measurements	8
2.1.3	Protocols	10
2.1.4	Imaging Modalities	12
2.1.5	Experiments	14
2.1.6	Traumatic Brain Injury	15
2.1.7	Acute Stroke	16
2.1.8	Other Head EIT Applications	18
2.2	Forward Problem	19
2.2.1	Maxwell's Equations	19
2.2.2	Mathematical Formulation of the Complete Electrode Model	20
2.2.3	Domain Discretisation	21
2.2.4	Weak Formulation	25
2.2.5	Galerkin Formulation	26
2.2.6	Numerical Algorithms and Preconditioning	27
2.2.7	Derivation of the Jacobian Matrix	29
2.2.8	Electrode Movement Jacobian Matrix	31
2.2.9	Forward Solvers	32
2.3	Inverse Problem	33
2.3.1	Linear Methods Based on Singular Value Decomposition	33
2.3.2	Linear Methods With Variational Regularisation	35
2.3.3	Non-Linear Methods	37
2.3.4	Multi-Frequency Reconstruction Methods	40

Contents

3	Head Models from CT and MRI Scans	45
3.1	Overview	45
3.1.1	Introduction	45
3.1.2	Background	46
3.1.3	Purpose	46
3.2	Mesh Creation	47
3.2.1	Segmentation	47
3.2.2	Meshing	49
3.3	Methods	50
3.3.1	Voltage Simulation	50
3.3.2	Image Reconstruction	51
3.3.3	Image Quality Measures	52
3.4	Results	54
3.4.1	Analysis of the Image Quality	54
3.4.2	Analysis of the Voltage Errors	56
3.5	Discussion	57
3.5.1	Mesh Creation	57
3.5.2	Impact of Mesh Differences	58
3.5.3	Conclusion	59
4	A Fast Parallel Forward Solver (PEITS)	61
4.1	Overview	61
4.1.1	Introduction	61
4.1.2	Background	62
4.1.3	Purpose	62
4.2	Technical Background	62
4.2.1	DUNE and PETSc	62
4.2.2	Complete Electrode Model and Jacobian Matrix	63
4.2.3	Methods	64
4.2.4	Overall Structure of PEITS	65
4.3	Implementation of the Functional Steps	66
4.3.1	Parallel Substructuring	66
4.3.2	Assembly of the System Matrix	67
4.3.3	Preconditioning	69
4.3.4	Solver	70
4.3.5	Jacobian Calculation	73
4.3.6	Verification of Correct Performance	74
4.4	Performance	74
4.4.1	Total Run Times with 1st Order Elements	74
4.4.2	Comparison to EIDORS	75
4.4.3	Comparison with 2nd Order Elements	76
4.4.4	Two Applications of PEITS	77
4.5	Discussion	79

5	Electrode Model Correction in TDEIT	81
5.1	Overview	81
5.1.1	Introduction	81
5.1.2	Background	82
5.1.3	Purpose	82
5.2	Electrode Boundary Jacobian Implementation	83
5.2.1	Mathematical Formulation	83
5.2.2	Implementation	83
5.3	Electrode Boundary Jacobian Characteristics	86
5.3.1	Simulation Parameters and Analysis Methods	86
5.3.2	Voltage Dependence on Electrode Characteristics	87
5.3.3	Precision of Electrode Boundary Jacobian	88
5.4	Simulation Study	89
5.4.1	Time-Difference Image Reconstruction Algorithm	89
5.4.2	Image Error Quantification	90
5.4.3	Simulation Parameters	91
5.4.4	Electrode Position Recovery	91
5.4.5	Images	92
5.5	Experimental Validation	96
5.5.1	Experimental Setup	96
5.5.2	Images	96
5.6	Absolute Reconstructions	98
5.6.1	Absolute Reconstruction Algorithm	98
5.6.2	Images	99
5.7	Discussion	101
5.7.1	Electrode Boundary Jacobian Characteristics	101
5.7.2	Simulation Study	101
5.7.3	Experimental Validation	102
5.7.4	Absolute Reconstructions	102
5.7.5	Conclusion	102
6	Feasibility of Stroke Imaging with MFEIT	105
6.1	Overview	105
6.1.1	Introduction	105
6.1.2	Background	106
6.1.3	Purpose	106
6.2	Methods	107
6.2.1	Model and Tissue Conductivity Spectra	107
6.2.2	Image Reconstruction	108
6.2.3	Image Error Quantification	109
6.2.4	Error Simulation	110
6.3	Results	111
6.3.1	Mesh Discretisation	111
6.3.2	Erroneous Electrode Positions	113
6.3.3	Erroneous Tissue Conductivity Spectra	113
6.3.4	Erroneous Electrode Contact Impedances	117

Contents

6.4	Discussion	117
6.4.1	Mesh Discretisation	117
6.4.2	Erroneous Electrode Positions	117
6.4.3	Erroneous Tissue Conductivity Spectra	118
6.4.4	Erroneous Electrode Contact Impedances	118
6.4.5	Technical Remarks	118
6.4.6	Conclusion	118
7	Electrode Model Correction in MFEIT	121
7.1	Overview	121
7.1.1	Introduction	121
7.1.2	Background	122
7.1.3	Purpose	122
7.2	Methods	122
7.2.1	Tissue Fraction Reconstruction With Electrode Position Correction . . .	122
7.2.2	Data Simulation	124
7.2.3	Image Error Quantification	125
7.3	Results	126
7.3.1	Multi-Frequency Tissue Fraction Reconstructions	126
7.3.2	Electrode Placement Correction	127
7.4	Discussion	129
7.4.1	Electrode Modelling Correction in Multi-Frequency Reconstruction . . .	129
7.4.2	Technical Remarks	129
7.4.3	Conclusion	130
8	Conclusion	131
8.1	Summary	131
8.2	Limitations	132
8.3	Outlook	133
8.4	Guidelines	135
A	PEITS User Guide	137
A.1	Downloading and Installing the Required Modules	137
A.2	Running the Solver with Different Settings	138
A.3	Making a Mesh Accessible to the Solver	140
A.4	Reading the Results into Matlab	141
A.5	Running the Solver from Matlab	141
	Bibliography	143

List of Figures

2.1	The cell modelled as an electronic circuit	7
2.2	EIT system and experimental setup	9
2.3	Electrode placement on the human head	11
2.4	Optimal current patterns for central stroke	11
2.5	Optimal current patterns for lateral stroke	12
2.6	Head shaped tanks	15
2.7	Conductivity spectra of stroke and healthy brain	17
2.8	Most common finite element shapes	22
2.9	Thorax model created from a 2D outline	23
2.10	Head model based on manual MRI segmentation	23
2.11	Meshes generated from either CT or MRI scan	24
2.12	Linear and quadratic shape functions	26
2.13	Singular values and Tikhonov filter	34
2.14	Characteristic L-curve shape	36
3.1	The Seg3D workspace after completed CT&MRI segmentation	48
3.2	3D mesh created from the segmentation shown in figure 3.1	50
3.3	Illustration of image quality measures	53
3.4	Image quality for correct reconstruction mesh vs. different head meshes	54
3.5	Exemplary reconstructions with the correct and a wrong mesh	55
3.6	Image quality dependence on mesh differences and system noise	58
4.1	Layered cut through a 5 million element head mesh	65
4.2	Partitions before and after Zoltan load-balancing	68
4.3	Strong and weak scaling of the system matrix assembly	69
4.4	Strong and weak scaling of the assembly of the AMG preconditioner ML	71
4.5	Strong and weak scaling of the CG solver with ML preconditioning	72
4.6	Strong and weak scaling of the computation of the Jacobian matrix	73
4.7	Simulated electric potential in a 2 million element mesh of the head	74
4.8	Estimated efficiency of the total run time of a realistic EIT protocol	75
4.9	Comparison of 1st and 2nd order shape functions	77
4.10	Convergence of simulated voltages by reducing element size	79
4.11	Three dimensional image of conductivity increase due to neural activity	79

List of Figures

5.1	Surface coordinate system on a 800k mesh with one refined electrode	84
5.2	Generation of the surface coordinate system	84
5.3	Ratio of voltage change over EBJ prediction for adjacent injection	87
5.4	Same than figure 5.3 for polar current injection	88
5.5	Slices through the coarse hexahedral mesh used for reconstruction	91
5.6	Recovered electrode movement	93
5.7	Reconstructions when electrodes have not moved	93
5.8	Electrodes connected by the 31 unique current injections	94
5.9	Reconstructions when one electrode moved by 5 mm	94
5.10	Reconstructions when all electrodes moved around 1 mm	95
5.11	Image errors for three different perturbation locations	95
5.12	Reconstructions of a simulated sponge in the back of the head	95
5.13	Image errors for reconstructions of a sponge	96
5.14	The KHU Mark 2.5 system and the 3D printed head shaped tank with skull . . .	97
5.15	Experimental images in the head tank with realistic skull	97
5.16	Image error quantification results for experimental images	98
5.17	Absolute reconstruction and absolute electrode position recovery	100
5.18	Absolute reconstructions with 1 mm electrode movement	100
6.1	Conductivity spectra and simulated stroke positions	107
6.2	Effect of mesh discretisation on image reconstructions	112
6.3	Quantification of image errors caused by mesh discretisation	112
6.4	Electrode areas on the fine and the coarse mesh	113
6.5	Image reconstructions with inaccurately modelled electrode positions	114
6.6	Quantification of image errors caused by inaccurate electrode modelling	114
6.7	Image reconstructions with erroneous tissue spectra	115
6.8	Quantification of image errors caused by spectral errors	115
6.9	Conductivity differences and noise with respect to the lowest frequency	116
6.10	Image reconstructions with inaccurately modelled contact impedances	116
6.11	Quantification of image errors caused by contact impedance errors	117
7.1	Image reconstructions without electrode modelling errors	127
7.2	Traditional MF fraction reconstructions with electrode modelling errors	127
7.3	Novel MF fraction reconstructions correcting for electrode modelling errors . .	128
7.4	Recovered electrode position errors for different stroke types and positions . .	129

List of Tables

2.1	Typical head tissue conductivities	16
3.1	Dimensions of the four different heads used for meshing	50
3.2	Stroke detection statistics for reconstruction and simulation meshes	55
3.3	Stroke detection statistics for noise levels and positions	56
4.1	Total time required for partitioning and loading the meshes	67
4.2	Total time required to assemble the system matrix	69
4.3	Times for the assembly of the preconditioner and subsequent CG solutions . . .	70
4.4	Time for one forward solution for CG with ML preconditioning	71
4.5	Iterations of the CG solver on different mesh sizes	71
4.6	Performance of the MUMPS direct solver	72
4.7	Number of forward solutions for MUMPS to be faster than ML	72
4.8	Time taken for the computation of one row of the Jacobian matrix	73
4.9	Total estimated run times for a protocol with 1000 lines	75
4.10	Comparison of EIDORS/SuperSolver and PEITS	75
4.11	Run time comparison 1st vs. 2nd order elements	78
5.1	Accuracy and amplitudes of the EBJ	88
5.2	2-norm of the electrode movement recovery error	92
5.3	2-norm of the electrode position recovery error	99
7.1	Accuracy of electrode position recovery	128

Chapter 1

Overview

1.1 Introduction

Electrical Impedance Tomography (EIT) is a non-invasive imaging modality in which the three dimensional conductivity distribution in an object is imaged. It is actively used in industrial settings and a broad range of medical EIT applications are envisaged. Fundamentally, EIT measures the electrical properties of tissues in the body by applying Ohm's law in a defined series of measurements. Small, imperceptible currents are injected into the body in different directions and the resulting electrical potentials are measured on the surface. Even though the name suggests that tomographic images are obtained, this is not precisely true. While the x-rays in Computed Tomography (CT) pass through the body in straight lines, the current in EIT spreads through the whole body. As a result, all measurements contain information about the conductivity throughout the imaged body and have to be considered as a whole, in order to produce an image. This makes EIT a very interesting field for mathematicians, since the quality of acquired images depends strongly on the modelling of the experimental setup and on the used image reconstruction algorithm.

EIT will never compete with imaging techniques like CT or magnetic resonance imaging (MRI) in terms of image quality. The spatial resolution is limited due to the incomplete electrode covering of the body surface and experimental error dependence of the reconstruction methods. However, EIT has a very high temporal resolution and the required equipment is small, portable and relatively inexpensive. These benefits have lead to a wide interest in clinical research, resulting in successful applications for the monitoring of lung ventilation (Frerichs, 2000) and gastric emptying (Brown, 2003). The EIT group of University College London focuses on applications of EIT to detect brain pathology and neural activity (Holder and Tidswell, 2004).

A typical EIT system consists of three main components: a current source applying an

1. Overview

alternating current, a voltmeter measuring the resulting voltages and a switching network, to connect the current source and voltmeter to different electrode pairs. Data collection is then guided by the current and measurement protocols, which define the consecutive electrode pairs used for injection and measurement. EIT imaging can be categorised into three modalities. In *absolute* imaging, a conductivity distribution explaining the measurements is computed directly from data collected at one point in time. *Time-difference* imaging is agnostic to the absolute conductivity distribution, and aims to reconstruct changes in conductivity between two measurements at different times. *Frequency-difference* (or multi-frequency) imaging uses measurements at different modulation frequencies to differentiate biological tissues on the basis of their characteristic conductivity spectra.

The image reconstruction relies on a good knowledge of the current path through the body, the modelling of which is called the *forward problem*. The forward problem is commonly solved with the finite element method (FEM), in which the body is modelled as the combination (called *mesh*) of many small elements with piecewise static conductivities. The current flow and resulting electric potential can be computed with FEM, by formulating the partial differential equation given by Maxwell's equations on this discrete mesh. However, the modelling never perfectly matches reality and discrepancies of real and modelled body geometry, tissue conductivities, electrode placement, size, shape and contact impedance all include additional noise to the inherent instrumentation noise. The three categories of EIT imaging mentioned in the previous paragraph rely on an accurate forward model to a different extent. While the resulting voltages in the body are linearly dependent on the applied current (modelled in the forward problem), their dependence on the conductivity is non-linear (inverted in the inverse problem).

The mathematical problem of the conductivity image reconstruction (the *inverse problem*) in EIT was first formulated by Calderón (2006) in 1980, and was proven to be unique if the whole boundary is known. In practice this is not the case, since current is applied and voltages are measured through a discrete set of electrodes. Furthermore, the inverse problem is severely *ill-posed*, meaning that small measurement or modelling errors can lead to large artefacts in the reconstructed image. As a consequence, the actual problem is unstable and not unique, and prior information has to be considered in order to reconstruct meaningful images.

These mentioned difficulties explain, why satisfactory absolute images have only been achieved in controlled measurement situations on simple geometries. As soon as the modelling does not accurately reflect the measurement, strong artefacts appear in the reconstructed image. In realistic EIT applications, the inverse problem is commonly solved by linearising it around the forward model (with the *Jacobian* matrix). Small and localised changes can then be recovered with time-difference image reconstructions. Taking time-differential measurements has the advantage that most modelling errors are cancelled out, because they are present in both the reference and the data measurement.

To use EIT for imaging acute stroke (Holder and Tidswell, 2004) or breast cancer (Zou and

Guo, 2003), a baseline image of the healthy patient is not available, and frequency-difference measurements have been proposed to overcome this issue (Brown et al., 1995). For homogeneous backgrounds, frequency-difference images can be reconstructed similarly to time-difference by linearising the inverse problem. For more realistic applications, non-linear reconstruction algorithms are required. In terms of dependence on model accuracy, these non-linear frequency-difference algorithms occupy the middle ground between time-difference and absolute imaging.

1.2 Purpose

The purpose of the work in this thesis is to improve the mathematical and algorithmic methods in head EIT. The goal is to provide a framework for accurate modelling and stable image reconstruction in both time-difference and frequency-difference imaging, with a focus on the detection of stroke during monitoring of traumatic brain injury patients and the imaging of acute ischaemic and haemorrhagic stroke. The connection of the forward and the inverse problem is studied, and improvements on the former are implemented to increase performance of the latter. The main contributions to the field are the following:

1. The proposal of a novel method for creating accurate head meshes and the analysis, whether subject specific head meshes are required for time-difference stroke imaging.
2. The implementation of a fast, parallel forward solver, to improve the speed and the memory requirement of boundary voltage and Jacobian matrix computations.
3. The study of voltage errors introduced by inaccurate modelling of electrode parameters and subsequent correction for electrode movement artefacts in time-difference and absolute imaging on the human head in simulation and experiments.
4. The comparison of modelling errors commonly occurring in non-linear multi-frequency imaging and iterative recovery of inaccurately modelled electrode positions on the human head.

1.3 Statement of Originality

Electrical Impedance Tomography is a very interdisciplinary field, as reflected in the presence of computer scientists, engineers, mathematicians, medical doctors, neuroscientists and physicists in the UCL EIT group. Therefore, most EIT research is collaborative and mine is no exception. Any problem I faced during the PhD studies, could immediately be discussed with an expert in the respective field, and consequently this thesis contains contributions of most current and former group members.

1. Overview

My supervisors Dr. Timo Betcke and Prof. David Holder, as well as co-authors on the publications listed in the next section — Dr. James Avery, Dr. Kirill Aristovich, Dr. Emma Malone —, have helped me with study design, presentation, language and structuring of the thesis.

The meshing procedure used in chapter 3 was originally devised by Dr. Kirill Aristovich for generating rat brain models, and was adapted by me for human head segmentations and to refine elements near the electrodes. One of the four CT and MRI segmentations used in this chapter was done by Mayo Faulkner under my supervision.

The work presented in chapter 4 would not have been possible without the support of Prof. Andreas Dedner, who has helped me throughout the implementation of PEITS with advice based on his in-depth knowledge of the DUNE-FEM package. The image of neural activity measured on a rat brain and the convergence study on the rat brain mesh were provided by Dr. Kirill Aristovich.

The 3D printed head tank used for the experimental validation in chapter 5, was designed and printed by Dr. James Avery, who also instructed me on how to use the KHU Mark 2.5 system for the tank experiments. The absolute images shown in this chapter were reconstructed by Dr. Emma Malone, based on my implementation of the electrode movement correction.

The feasibility study on multi-frequency stroke imaging in chapter 6 was jointly designed and written with Dr. Emma Malone. All forward simulations, image reconstructions and image quality measures were computed by me. For the image reconstructions I used an adapted implementation of Emma's fraction reconstruction algorithm.

All work on the correction of electrode modelling in MFEIT in chapter 7 was done by myself. The used code was adapted from the implementation of the fraction reconstruction algorithm by Dr. Emma Malone.

1.4 List of Publications

The work presented in this thesis has been published or accepted for publication in the following peer-reviewed journal papers:

Chapter 3 M. Jehl, K. Aristovich, M. Faulkner, D. Holder. '*Are patient specific meshes required for EIT head imaging?*', Physiological Measurement, accepted for publication, January 2016

Chapter 4 M. Jehl, A. Dedner, T. Betcke, K. Aristovich, R. Klöfkorn, D. Holder. '*A fast parallel solver for the forward problem in Electrical Impedance Tomography*', IEEE Transactions on Biomedical Engineering, **62**(1), 126-137, January 2015, doi:10.1109/TBME.2014.2342280

Chapter 5 M. Jehl, J. Avery, E. Malone, D. Holder, T. Betcke. '*Correcting electrode modelling errors in EIT on realistic 3D head models*', Physiological Measurement, **36**(12), 2423-2442, October 2015, doi:10.1088/0967-3334/36/12/2423

Chapter 6 E. Malone, M. Jehl, S. Arridge, T. Betcke, D. Holder. '*Stroke type differentiation using EIT: evaluation of feasibility in a realistic head model*', *Physiological Measurement*, **35**(6), 1051-1066, May 2014, doi:10.1088/0967-3334/35/6/1051

Chapter 7 M. Jehl, D. Holder. '*Correction of electrode modelling in multi-frequency EIT imaging*', *Physiological Measurement*, accepted for publication, January 2016

Chapter 2

Literature Review

2.1 Electrical Impedance Tomography

2.1.1 Bioimpedance




Figure 2.1: The cell modelled as an electronic circuit - (a) A single cell can be modelled as a circuit of the extracellular resistance in parallel with the intracellular resistance and the capacitance of the cell membrane. **(b)** The frequency dependant complex impedance can be visualised in a Cole-Cole plot.

The *impedance* Z of a biological tissue characterises the electrical properties of this tissue when a current or a voltage is applied (Holder, 2004). The *resistance* R is a measure of the extent to which charge transport is opposed whereas the *capacitance* C describes the ability of the tissue to store electrical charges. Biological tissues conduct current through ion diffusion in the highly conductive extracellular and intracellular spaces. The lipid cell membrane can be modelled as a capacitance, a cell as a simple parallel circuit (figure 2.1a) and its impedance is thus

$$Z(\omega) = \frac{R_e R_i + \frac{R_e}{j\omega C_m}}{R_e + R_i + \frac{1}{j\omega C_m}}. \quad (2.1)$$

2. Literature Review

At low frequencies the cell membrane C_m is fully charged most of the time and thus most of the current flows through the extracellular space

$$\lim_{\omega \rightarrow 0} Z(\omega) = R_e, \quad (2.2)$$

while at high frequencies the capacitance is uncharged and the current can also flow through the cell

$$\lim_{\omega \rightarrow \infty} Z(\omega) = \frac{R_e R_i}{R_e + R_i}. \quad (2.3)$$

At intermediate frequencies the impedance has an imaginary component due to the influence of the capacitance. However, this phase shift is rarely used in EIT measurements because it is prone to noise from the instrumentation (stray capacitance, Saulnier, 2004). The real component can be calculated as

$$R(\omega) = \text{Re } Z = \frac{R_e R_i (R_e + R_i) + \frac{R_e}{\omega^2 C_m^2}}{(R_e + R_i)^2 + \frac{1}{\omega^2 C_m^2}}. \quad (2.4)$$

The characteristics of similar dielectrics have been studied extensively by Cole and Cole (1941) and this cell model is thus often mentioned alongside their research and the frequency response of the impedance shown in Cole-Cole plots (figure 2.1b). At a large scale, this simple cell model roughly explains the differences in observed tissue resistances. Dense tissues, such as fat, have a high resistance, while bodily fluids have a low resistance. The resistance increases with the length l and decreases with the cross-section A of the current path, and depends on the tissue specific *resistivity* ρ ,

$$R = \rho \frac{l}{A}. \quad (2.5)$$

The inverse of the resistivity is the *conductivity* σ given in S m^{-1} , which is the value generally reconstructed in EIT.

2.1.2 Measurements

Ohm's law $V = Z \cdot I$ already suggests which components constitute an EIT system (figure 2.2a)

- I - A precise current source is required to drive an alternating current with an amplitude which is limited by medical safety regulations.
- V - A voltmeter to measure the resulting in-phase surface potentials when current I is applied.
- Z - In order to image the impedance distribution within an object, voltages are measured for a set of independent current injections. Thus, an array of electrodes and a switching network is required.




Figure 2.2: EIT system and experimental setup - (a) The UCL Mk 2.5 system containing one current source, one voltmeter and a switching network (McEwan et al., 2006). (b) A classical EIT experiment, with a cylindrical saline tank and equispaced electrodes around the boundary.

The first EIT system was the Sheffield Mk 1 (Brown and Seagar, 1987), which applied a fixed 50 kHz current through adjacent pairs of 16 cylindrically arranged electrodes (similarly to the 32 electrodes in figure 2.2b). Voltages were also measured between adjacent electrodes. Since the Sheffield Mk 1, many other EIT systems have been developed to apply current and measure voltages at multiple frequencies or through multiple electrodes simultaneously, however there is always a trade-off between the versatility and speed, and the precision of the measurements. Theoretically, it is also possible to apply voltage and measure the resulting currents. This is rarely done, because the accurate measurement of currents is difficult (Saulnier, 2004).

Both main components of an EIT system introduce noise into measurements (Frangi et al., 2002). The current source has to be optimised for a certain load. Since the load in EIT is not accurately known, current divider effects introduce proportional (dependent on the signal size) noise into measurements (Boone and Holder, 1996a). The noise from the voltmeter is mainly due to thermal noise and is signal-independent (i.e. additive noise). Another source of uncertainty in EIT measurements is the interface between the electrodes and the skin. Because of electrochemical effects when electronic current in the system is transformed to ionic current in the body, a voltage drop is observed, which can differ significantly from one electrode to another (Boone and Holder, 1996b). If voltages were measured on the electrodes injecting current, the voltage drop caused by the contact impedance would introduce a large, unknown error. Therefore, EIT voltage measurements are usually restricted to electrodes not involved in current injection, so-called *four electrode measurements*.

Increasing the current amplitude leads to a better signal-to-noise ratio. However, the current which may be injected into a patient is limited by the IEC 60601 medical safety regulations (IEC, 2005). For recordings up to 1 kHz, the root mean square limit is $100\text{ }\mu\text{A}$ and above, the current is limited to $f \cdot 100\text{ }\mu\text{A}$ where f is the frequency in kHz.

2. Literature Review

2.1.3 Protocols

A very important part of EIT measurements is the protocol. A *current pattern* is a single current injection through either a pair of electrodes or multiple electrodes. The *current protocol* is a set of independent current patterns, thus it has $M - 1$ lines when M electrodes are used. *Single measurements* are always done on a pair of electrodes and the set of independent measurements performed for each applied current pattern is called the *measurement protocol*. The *full protocol* is the combination of the current protocol and the measurement protocol and defines the order in which single measurements are taken.

The choice of current patterns greatly influences the sensitivity of the measurements to perturbations in different locations. The adjacent protocol used in the Sheffield Mark 1 system has a very low sensitivity in the centre of the body since most of the current passes along the surface. Polar current injection through opposite electrodes improves the sensitivity in the centre, but reduces the number of independent current injections to $M/2$ when the electrodes are arranged geometrically around a cylindrical tank (which is still the standard test object in EIT). Thus, a ‘just off’ opposite current protocol has been suggested for these ‘quasi 2D’ applications with symmetrical electrode arrangement (Adler et al., 2011).

On the human head, which – unlike the thorax – is badly approximated by a cylinder, electrodes are commonly placed according to the so-called EEG 10-20 or EEG 10-10 positions (figure 2.3a), which are the standard in the electroencephalography community (Jurcak et al., 2007). Polar injections are even more important in head EIT, because the highly conductive scalp on top of the highly resistive skull result in adjacent injections being shunted through the scalp and not actually passing through the object of interest, the brain. A full protocol for 32 electrodes placed according to an extended EEG 10-20 system (figure 2.3b), was devised empirically by finding polar or near-polar electrodes for injections, and adjacent electrodes were used for measurements (Gibson, 2000). For systems with hard-wired measurement electrode combinations this *EEG31* protocol could not be used, and alternative protocols were investigated (Fabrizi et al., 2009).

The interest of using multiple electrode current injections was born from the idea of analytically deriving the optimal current patterns. Isaacson (1986) defined the *distinguishability* in EIT as the L^2 -norm of the boundary voltage changes caused by a change in conductivity $\sigma_2 - \sigma_1$. Using the distinguishability as the quality measure of an EIT measurement, the best current patterns are the eigenvectors corresponding to the largest eigenvalues of the difference of the Neumann-to-Dirichlet (NtD) or Current-to-Voltage maps for the two conductivity distributions σ_1 and σ_2 . On a finite element mesh with continuous Neumann boundary conditions the NtD map is simply the stiffness matrix weighted by the conductivity of the elements. For M discrete electrodes the NtD map can be computed with $M - 1$ forward solutions using the complete electrode model boundary conditions. For cylindrical or spherical geometry the optimal current patterns are *trigonometric patterns* with low spatial frequencies being most sensitive to conduc-

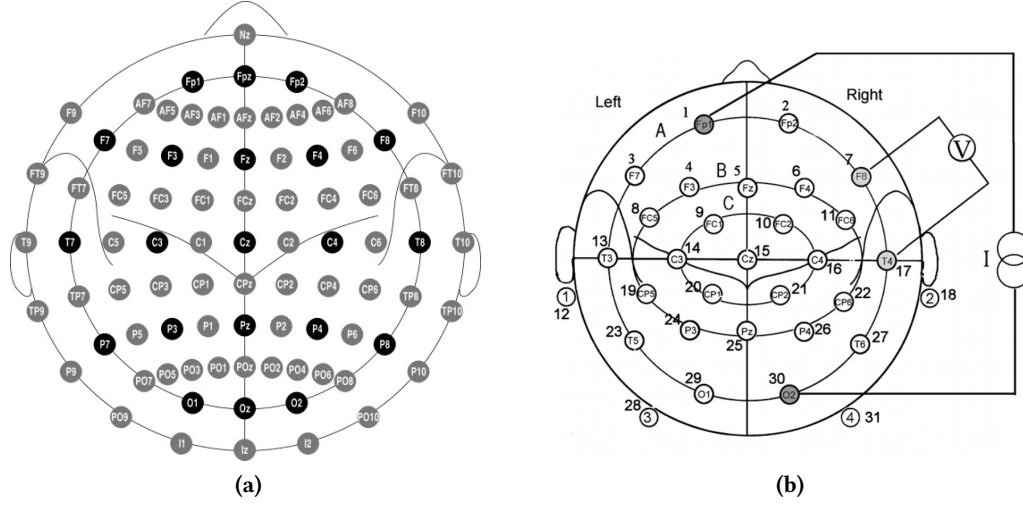


Figure 2.3: Electrode placement on the human head - (a) EEG 10-10 electrode positioning system with EEG 10-20 positions in black (Oostenveld and Praamstra, 2001). **(b)** Extended EEG 10-20 system, for even coverage of the head surface with 32 electrodes (Gibson, 2000).

tivity changes far from the boundary (figure 2.4) and high frequencies more sensitive to changes near the boundary (figure 2.5).




Figure 2.4: Optimal current patterns for central stroke - The optimal continuous current density patterns on the surface of the head were computed as the eigenvectors corresponding to the largest eigenvalues of the differences of the NtD maps with and without stroke.

However, there are two reasons why trigonometric patterns might still be a bad choice. The first reason is the formulation of the medical safety regulations IEC 60601 (Dybdahl, 2009), the interpretation of which is debated, but most likely limits the sum of all applied currents

2. Literature Review




Figure 2.5: Optimal current patterns for lateral stroke - The optimal current patterns for a lateral stroke have a high spatial frequency close to the stroke location.

(Lionheart et al., 2001). Trigonometric patterns are only optimal, if the power or the 2-norm of the applied current were limited. If, however, the sum of all applied currents is limited, then pairwise injection results in a better distinguishability. If the maximum injected current to a single electrode were limited, the so-called *Walsh patterns* would be best (Eyüboğlu and Pilkington, 1993). The second reason why parallel injection might result in worse images than pairwise injections, is that the electrodes used for voltage measurements are simultaneously used to inject current. Two electrode measurements are more sensitive to errors in electrode contact impedances than the commonly used four electrode measurements (Kolehmainen et al., 1997). To reduce these errors, contact impedances can be treated as unknowns and reconstructed together with the conductivities, in order to reduce the deterioration of image quality (Heikkinen et al., 2002). Methods for keeping the electrode contact impedance low include skin abrasion and applying conductive gels between electrodes and skin (McAdams et al., 1996).

2.1.4 Imaging Modalities

EIT image reconstruction methods can be categorised into three groups, absolute imaging, time- and frequency-difference. All methods rely on a computer model (the *mesh*) of the object under examination, which is used to simulate the expected electric potential distributions (the *forward solutions*) for the injected current patterns. In absolute imaging these simulated voltages are directly compared to the measured experimental voltages, and the model is iteratively adjusted

to reduce the differences between simulation and experiment. Time- and frequency difference data can either be reconstructed with linear algorithms (fast and stable), when the conductivity changes are small and localised, or iteratively, when this *linearity assumption* is invalid. For all reconstruction methods, the *Jacobian (or sensitivity) matrix* is required, which is computed based on the forward solutions. The Jacobian relates changes in conductivity to changes in measured voltages by linearising the computer model and informs the imaging algorithms of the search direction. The closer the mesh used for its computation is to the reality, the better are the resulting images.

Absolute Imaging

The most straightforward way to create an image from EIT measurements is to map the measured voltages directly to a conductivity distribution within the measured object (Vauhkonen et al., 1999). Absolute (or static) imaging is the most versatile imaging modality in EIT, since images can be produced at any point in time without the need for reference measurements. However, as was shown for instance by Kolehmainen et al. (1997), absolute imaging results in very noisy images if the model does not accurately match the reality. Because the inverse problem in EIT is non-linear and severely ill-posed, already small errors in the modelling lead to large artefacts in the reconstructed image. For these reasons, no satisfactory images have so far been produced with absolute imaging from clinical data or tank experiments with non-trivial geometry. To suppress image errors caused by modelling inaccuracies and systematic instrumentation noise, difference measurements are commonly used.

Time-Difference

Time-difference (TD) imaging compares the measured voltages at time t (the *data measurement*) to a *baseline (or reference) measurement* obtained at time t_0 and aims to explain the voltage difference by a conductivity change between these time points. Now, the mesh is only used to compute the relationship of conductivity changes to voltage changes (the Jacobian), but not the reference measurement. Therefore, large parts of the time-independent modelling and instrumentation errors are cancelled and the reconstructions are more stable. A small, localised change in conductivity $\delta\sigma$ between the two time points can be linearly reconstructed by inverting the Jacobian matrix \mathbf{J} in some way:

$$\delta\sigma = \mathbf{J}^{-1}(\mathbf{v}_t - \mathbf{v}_{t_0}). \quad (2.6)$$

Since TD requires a baseline measurement, its applications are limited to imaging physiological changes over a short period with a known starting time. However, EIT has the potential to

2. Literature Review

image conditions such as stroke or cancer which cannot be monitored during their build-up, but only after they occurred.

Frequency-Difference

Multi-frequency EIT (MFEIT), or EIT spectroscopy (EITS), relies on the observation that the impedances of different tissues have different frequency spectra and enable ‘one shot’ imaging without the drawbacks absolute imaging has. By measuring with different frequencies simultaneously or in rapid sequence, spectral data are obtained and can be reconstructed. *Simple* (Brown et al., 1995) and *weighted* (Seo et al., 2008; Jun et al., 2009) *frequency-difference* (FD) are the most basic reconstruction methods, where analogous to TD one measurement is taken as a baseline and a measurement at a different frequency is reconstructed by inverting the Jacobian matrix. However, linear FD can only reconstruct anomalies in a homogeneous background and its applications are thus very limited. More advanced, non-linear FD algorithms are required to enable MFEIT in realistic applications. Two such algorithms have recently been published. The *fraction reconstruction* algorithm (Malone et al., 2014) reconstructs tissue volume fractions instead of conductivities, thereby allowing the simultaneous use of measurements at multiple frequencies, provided that the conductivity spectra of the tissues are known. Uncertainty in the knowledge of the tissue spectra was incorporated in the *reconstruction-classification* algorithm (Malone et al., 2015), at the cost of having to optimise a second regularisation parameter.

2.1.5 Experiments

EIT applications and algorithms can be tested in three different environments. *Simulation studies* are useful to evaluate the feasibility of an application and the correct performance of a reconstruction algorithm, because they are fast, easy and the modelling errors and instrumentation noise can be controlled. In simulation studies, the experimental voltages are simulated on a finite element mesh, measurement noise is added to the voltages, and then they are reconstructed into images.

One step closer to real applications are *tank (or phantom) experiments*. A tank experiment is less controlled than a simulation study, but since the geometry is well known and the perturbation can be controlled accurately, reconstructions are generally reliable. A standard EIT tank is a cylindrical tank with electrodes in a circle around the boundary (figure 2.2b), which is usually filled with saline with controlled conductivity to represent the human thorax (Saulnier et al., 2001). Perturbations are placed into the cylinder, voltages are collected and images reconstructed with 2D or 3D reconstruction algorithms. For FD applications, the saline has been filled with carrot pieces, to achieve a background conductivity that changes with frequency (Packham et al., 2012). For head EIT, cylindrical tanks are not very useful, and several head shaped tanks

have been proposed. A tank modelled around the shape of a real skull (figure 2.6a, Tidswell et al., 2001a) was used successfully to reconstruct TD images, even though the simulated voltages were significantly different to the experimental ones (Avery, 2015). The difficulty of fitting a computer model to a tank was recently addressed by starting with the model, and 3D printing the tank according to this model (figure 2.6b, Avery, 2015). The conductivity of the skull was adapted to recent literature (Tang et al., 2008) by printing holes in the skull, that – once it was immersed in saline – gave it the desired conductivity values.




Figure 2.6: Head shaped tanks - (a) Head shaped tank with real skull (Tidswell et al., 2001a) and (b) 3D printed head shaped tank and skull (Avery, 2015).

The most challenging EIT experiments are *human studies*. Measurement errors are larger due to movement of the patient, high contact impedances between electrode and skin, unknown geometry and tissue conductivity values. Furthermore, the quality of the reconstructed images normally cannot easily be validated, since the exact conductivities or conductivity changes are not known. Consequently, the so far only established clinical application of EIT is the monitoring of lung ventilation, where the signal is very large (Frerichs, 2000; Gong et al., 2015).

2.1.6 Traumatic Brain Injury

Traumatic brain injury (TBI) is caused by sudden external forces to the head, and is most common in children and adolescents (NICE, 2014). It is a leading cause of death and disability worldwide, with 50 thousand deaths, 235 thousand hospitalisations and over a million emergency department visits in the United States alone (Langlois et al., 2006). Only around 5% of TBI patients present symptoms strong enough to be at risk of fatal acute intracranial complications (NICE, 2014). These patients are identified using the Glasgow Coma Scale and undergo a CT scan. Some patients with normal CT scans develop delayed onset haematoma in or around the brain (intracerebral haemorrhage and epidural haematoma, rarely subdural haematoma) up to a week after the initial trauma (Matsuda et al., 2008). Furthermore, patients who underwent surgery after the initial CT scan, can develop secondary bleedings (Xu et al., 2010).

2. Literature Review

Since it is not feasible to perform CT scans at regular intervals, EIT monitoring was proposed for patients at risk of developing delayed onset or secondary haematoma (Murphy et al., 1987; Xu et al., 2010). In recent feasibility studies on piglets with scalp electrodes, injected arterial blood was reproducibly detected (Dai et al., 2010; Manwaring et al., 2013). Human in vivo images were obtained during the standard twist-drill drainage treatment of subdural haematoma (Dai et al., 2013). However, Dai et al. (2013) did not image blood, but instead the influx and successive drainage of irrigating fluid (5% dextrose in water), which has a five times higher contrast when compared to healthy brain.

In EIT, conductivity changes caused by haemorrhage are usually modelled by replacing the conductivity of healthy brain with the conductivity of blood. This results in a conductivity increase of $\sim 180\%$ for intracerebral haemorrhages, as calculated from standard tissue conductivities (table 2.1). In the case of subdural haematoma, blood mixes with cerebrospinal fluid and the increased fluid volume compresses the brain. A localised compression of the brain is also observed for epidural haematoma. Consequently, intracranial haemorrhages can take very different shapes which have a different impact on the head conductivity and have to be modelled accordingly.

Tissue	Conductivity (S m^{-1})
White matter	0.15
Grey matter	0.3
Cerebrospinal fluid	1.79
Blood	0.7
Skull	0.018
Scalp	0.44

Table 2.1: Typical head tissue conductivities - Tissue conductivities at ~ 10 kHz compiled from literature by Horesh (2006).

2.1.7 Acute Stroke

Stroke is the third most common cause of death and the leading cause of adult disability in the UK (Power, 2004). Haemorrhagic stroke is caused by a bleeding in the brain and requires surgery for treatment. Ischaemic stroke is caused by an interruption of blood flow due to an embolism. Recombinant tissue plasminogen activator (rt-PA), a thrombolytic drug which dissolves the occlusion of the blocked blood vessel, has been approved in the UK in 2003 for treatment within 3 hours of the onset of an ischaemia. Moderate improvements for treatment with rt-PA have been observed up to 4.5 hours after the ischaemia (Stemer and Lyden, 2010). Before administration of the drug, the type of stroke has to be confirmed with a CT scan, since rt-PA can cause additional damage and even death in the case of a haemorrhage. Since 80% of all strokes are ischaemic, a quick identification can save many patients. However, due to the current necessity of patient transport to the nearest hospital and subsequent CT scan, only 2.5% to 6% of acute ischaemic stroke patients receive thrombolytic treatment within the required 3 hours (Power, 2004; Saver

et al., 2013), with expedited examination protocols around 13% (Sattin et al., 2006). Electrical Impedance Tomography could enable earlier stroke diagnosis and treatment in the ambulance or emergency unit, while the patient is being transported to a hospital or waiting for a CT scan (Holder and Tidswell, 2004). EIT will never compete with CT in terms of image quality, but its low cost and portability make it feasible to equip all ambulances with an EIT system.




Figure 2.7: Conductivity spectra of stroke and healthy brain - The different slopes of healthy and ischaemic brain tissue and blood might enable stroke type differentiation with MFEIT. This figure is adapted from Horesh (2006).

The lack of oxygen during an ischaemia causes the cells to swell due to energy failure, leading to an overall lower conductivity in the affected area. In the case of a haemorrhage, the abundance of blood increases the conductivity. The conductivity spectra for the different tissues in a healthy head and during stroke have been collected from literature by Horesh (2006) and show, that there is a significant difference between normal brain tissue, ischaemic brain tissue and blood (Fig. 2.7). The largest difference is at low frequencies, which has recently been confirmed by in vivo measurements (Dowrick et al., 2015). Dowrick et al. (2015) found, however, that the change of healthy brain was larger than that of ischaemic brain, and that the difference in change was $\sim 10\%$ below 250 Hz. Since the conductivity of the scalp at low frequencies is less than half compared to frequencies above 1 kHz (figure 6.1a), the amount of current shunted around the brain is lower at these frequencies. Feasibility studies have shown that MFEIT has the potential to differentiate between ischaemic and haemorrhagic stroke (Horesh et al., 2005; Packham et al., 2012).

Since stroke patients have generally not had EIT measurements of their head done before the stroke occurs, time-difference image reconstruction is not possible. Weighted frequency-difference is also not a viable option, since the background does not consist of a homogeneous tissue and the conductivity change over frequency can consequently not be described with

2. Literature Review

a linear model. Therefore, non-linear multi-frequency algorithms have to be used for stroke imaging (Malone et al., 2014).

The same application is envisaged for portable microwave imaging systems (Persson et al., 2014), which have an easy penetration of the skull. Microwave imaging relies on a similar ill-posed inverse problem than EIT and is consequently more stable for time-difference measurements than for absolute imaging (Scapaticci et al., 2014). Frequency-difference measurements are uncommon in microwave imaging, because the useful bandwidth in brain imaging applications is narrow (Scapaticci et al., 2012). The technique has been shown to work in phantom experiments with modelled strokes near the surface (Mohammed et al., 2014), but in vivo measurements are very noisy (Persson et al., 2014). Consequently, machine learning approaches are investigated as an alternative to imaging (Persson et al., 2014).

2.1.8 Other Head EIT Applications

Epilepsy

Epilepsy is a common neurological condition taking various forms, making characterisation difficult. Broadly, it can be categorised into *generalised* seizures, which engage networks spread over both hemispheres, and *focal* seizures, which are limited to one hemisphere and sometimes discretely localised (Berg et al., 2010). Focal seizures consistently start propagating from one location, while generalised seizure onset differs from one seizure to another. In focal epilepsy patients which are resistant to treatment with anti-epileptic drugs, resective surgery, in which the origin of the seizures is removed, is an established treatment (Schramm and Clusmann, 2008). However, surgery is only performed if the focus is accurately localised, surgically accessible and mostly redundant for brain function. The localisation of the origin of focal seizures is done with video-EEG monitoring, sometimes with intracranial EEG electrodes or depth electrodes (Rosenow and Lüders, 2001). EIT has been suggested for improving the depth localisation, which is relatively poor in EEG (Fabrizi et al., 2006). Since EIT can use the same electrodes than EEG, no additional preparatory or surgical procedures would be required. Conductivity changes during epilepsy are commonly attributed to a shrinkage of extracellular space due to cell swelling, and are in the order of percents (Holder and Tidswell, 2004). Recently, epileptic seizures have been successfully imaged in vivo in rat experiments (Vongerichten et al., 2016).

Fast Neural Activity

Imaging neural activity is of great interest to neuroscience, but methods are so far limited. fMRI can create images of the entire brain, but instead of detecting neuronal depolarisation and spiking directly, it shows the related changes in blood flow over seconds (Heeger and Ress, 2002).

Microelectrodes are invasive and currently cover a few cubic millimetres, while optical methods such as two-photon microscopy are limited to small volumes below a cubic millimetre (Aristovich et al., 2016). EIT has the potential to image the impedance decrease caused by the opening of ion channels during neuronal depolarisation (Holder, 1987). The magnitude of changes was modelled using cable theory (Liston et al., 2012) and the predicted changes in the order of 0.1% were validated in the rat cortex in vivo (Aristovich et al., 2016). Theoretically, with EIT it should be possible to image activity based conductivity changes throughout the brain with a resolution below half a millimetre (Aristovich et al., 2016).

2.2 Forward Problem

The *forward problem* in EIT is to compute the resulting electric potential on a domain, when a current is applied to the boundary of the domain. It is commonly solved by assuming that the magnetic field is negligible, therefore reducing Maxwell's equations to an elliptic boundary value problem with Neumann boundary conditions. This type of problem is well-known, and commonly solved using the finite element method (FEM). However, the electrode-skin interface cannot be modelled accurately with simple Neumann boundary conditions. Consequently, a Robin-type boundary condition called the *complete electrode model (CEM)* is used as a standard in EIT. The weak formulation of the elliptic problem with CEM is solved on a discretised model of the domain, by translating it into a system of linear equations with the Galerkin formulation. The electric potentials in the domain and on the electrodes are then found by inverting this system with a preconditioned conjugate gradient algorithm. These *forward solutions* are directly required in iterative image reconstruction algorithms, and indirectly for all inverse solutions in the form of the Jacobian matrix, which is a linearisation of the forward problem.

2.2.1 Maxwell's Equations

The physics behind EIT are governed by Maxwell's equations. Let E be the electric field, B the magnetic field, ρ the charge density, J the current density, μ the permeability and ϵ the permittivity. Then Maxwell's equations are:

$$\nabla \cdot E = \frac{\rho}{\epsilon} \quad (2.7a)$$

$$\nabla \cdot B = 0 \quad (2.7b)$$

$$\nabla \times E = -\frac{\partial B}{\partial t} \quad (2.7c)$$

$$\nabla \times B = \mu J + \mu\epsilon \frac{\partial E}{\partial t}. \quad (2.7d)$$

2. Literature Review

Taking the divergence of Ampère's Law (2.7d) and considering that $\nabla \cdot (\nabla \times B) = 0$ gives

$$\nabla \cdot J = -\frac{\partial \rho}{\partial t}. \quad (2.8)$$

This equation states that every change of charge within a domain Ω is accompanied by a flux across its closed surface $\partial\Omega$. Since, in EIT, there are no electric sources or drains, the right-hand side of (2.8) is zero. The current density can be written in terms of the electric field $J = \gamma E$, where the position dependent complex admittivity $\gamma(\mathbf{x}, \omega)$ is given for each position $\mathbf{x} \in \Omega$ as the sum of the conductivity σ and the permittivity ϵ multiplied by the frequency ω , $\gamma(\mathbf{x}, \omega) = \sigma(\mathbf{x}, \omega) + i\omega\epsilon(\mathbf{x}, \omega)$.

Because EIT applications usually operate at low frequencies, the change in the magnetic field $\frac{\partial B}{\partial t}$ is considered negligible (Lionheart et al., 2004), such that (2.7c) becomes $\nabla \times E = 0$ and $E = -\nabla u$, where u is the electric potential. Substituting E in (2.8) gives Laplace's equation

$$-\nabla \cdot (\gamma \nabla u) = 0, \quad (2.9)$$

the foundation of most mathematical EIT models. Soni et al. (2006) found that this quasi-static approximation of Maxwell's equations is valid up to approximately 1 MHz. Since, for practical reasons (Saulnier, 2004), most systems measure only the in-phase voltage, the admittivity is henceforth replaced with only the conductivity.

2.2.2 Mathematical Formulation of the Complete Electrode Model

Let the body under examination be a bounded domain $\Omega \subset \mathbb{R}^3$ with Lipschitz boundary $\Gamma = \partial\Omega$. Then the *continuum forward model* states that the electric potential distribution u in Ω induced by the application of the current density j on the boundary Γ solves

$$-\nabla \cdot (\sigma \nabla u) = 0 \quad \text{in } \Omega, \quad (2.10a)$$

$$\sigma \frac{\partial u}{\partial \nu} = j \quad \text{on } \Gamma, \quad (2.10b)$$

$$\int_{\Gamma} u \, d\Gamma = 0, \quad (2.10c)$$

where ν is the outward unit normal to Γ and with positive real conductivity $\sigma \in L^\infty(\Omega)$. This problem is known to have a unique solution $u \in H^1(\Omega)$ for all $j \in H^{-1/2}(\Gamma)$ satisfying $\int_{\Gamma} j \, d\Gamma = 0$, which can be indicated by a diamond in the space definition $H_\diamond^{-1/2}(\Gamma)$ (Kohn and Vogelius, 1985).

In practice, current is applied through a finite set $M \in \mathbb{N}$ electrodes $(\Gamma_m)_{m=1,\dots,M}$, rather than continuously on Γ . Γ_m is a non-zero connected open subset of Γ with Lipschitz boundary that does not overlap with any other electrode, meaning $\overline{\Gamma_m} \cap \overline{\Gamma_l} = \emptyset$ for $m \neq l$. The optimal

model for boundary conditions matching this experimental setup has been incrementally improved by Isaacson et al. around 1990 (for a review and experimental justification see Somersalo et al., 1992). The commonly used *complete electrode model* (CEM) accounts for two observed effects, electrodes shunting current due to their high conductivity and a voltage drop at the interface of the electrodes and the skin, which is due to an electrochemical effect. The applied current is constant on each electrode and vanishes on the surface between electrodes $\Gamma \setminus \bigcup_{m=1}^M \Gamma_m$. It can thus be represented as a vector $(I_m)_{m=1,\dots,M}$ in \mathbb{R}^M satisfying the condition $\sum_{m=1}^M I_m = 0$. According to the CEM, the voltage potential u then solves

$$-\nabla \cdot (\sigma \nabla u) = 0 \quad \text{in } \Omega \quad (2.11a)$$

$$\int_{\Gamma_m} \sigma \frac{\partial u}{\partial \nu} d\Gamma_m = I_m \quad m = 1, \dots, M \quad (2.11b)$$

$$u + z_m \sigma \frac{\partial u}{\partial \nu} = U_m \quad \text{on } \Gamma_m, \quad m = 1, \dots, M \quad (2.11c)$$

$$\sigma \frac{\partial u}{\partial \nu} = 0 \quad \text{on } \Gamma \setminus \bigcup_{m=1}^M \Gamma_m, \quad (2.11d)$$

where $(z_m)_{m=1,\dots,M} \in \mathbb{R}^M$ is the positive contact impedance, given in $\Omega \text{ m}^2$, and $(U_m)_{m=1,\dots,M} \in \mathbb{R}^M$ is the vector of the potentials on the electrodes abiding to the adapted grounding condition (2.10c), $\sum_{m=1}^M U_m = 0$. The CEM is proven to have a unique solution $u \in H^1(\Omega)$, which depends continuously on I_m (Somersalo et al., 1992).

2.2.3 Domain Discretisation

The forward problem can be solved analytically on very simple geometries (Isaacson, 1986). In order to solve it on non-symmetric and complicated realistic domains, the domains need to be discretised. Several numerical methods exist to solve the forward problem on discretised domains. The finite difference and finite volume methods usually require regular grids (Liszka and Orkisz, 1980; Eymard et al., 2000). While these methods are computationally efficient, the representation of curved structures is complicated by the regularity requirement (Lionheart et al., 2004). A more powerful, yet less intuitive, method is the finite element method (FEM, Brenner and Scott, 1994), which is the most commonly used in EIT. It does not require any regularity of the discretisation (the *finite element mesh*), and instead of evaluating local derivatives (finite difference) or fluxes (finite volume) it solves a *weak formulation*, which already includes the boundary conditions. As the number of elements in the mesh increases, the FEM solution approaches the real solution of the partial differential equation. On homogeneous domains, the boundary element method (BEM) is a viable alternative to FEM (Gençer and Tanzer, 1999). BEM only requires discretised surfaces and uses an analytical expression of the Green function between surfaces of different characteristics. For complicated geometries (such as the human head), the presence of many

2. Literature Review

different tissues diminishes the computational advantage of BEM over FEM (Lionheart et al., 2004).

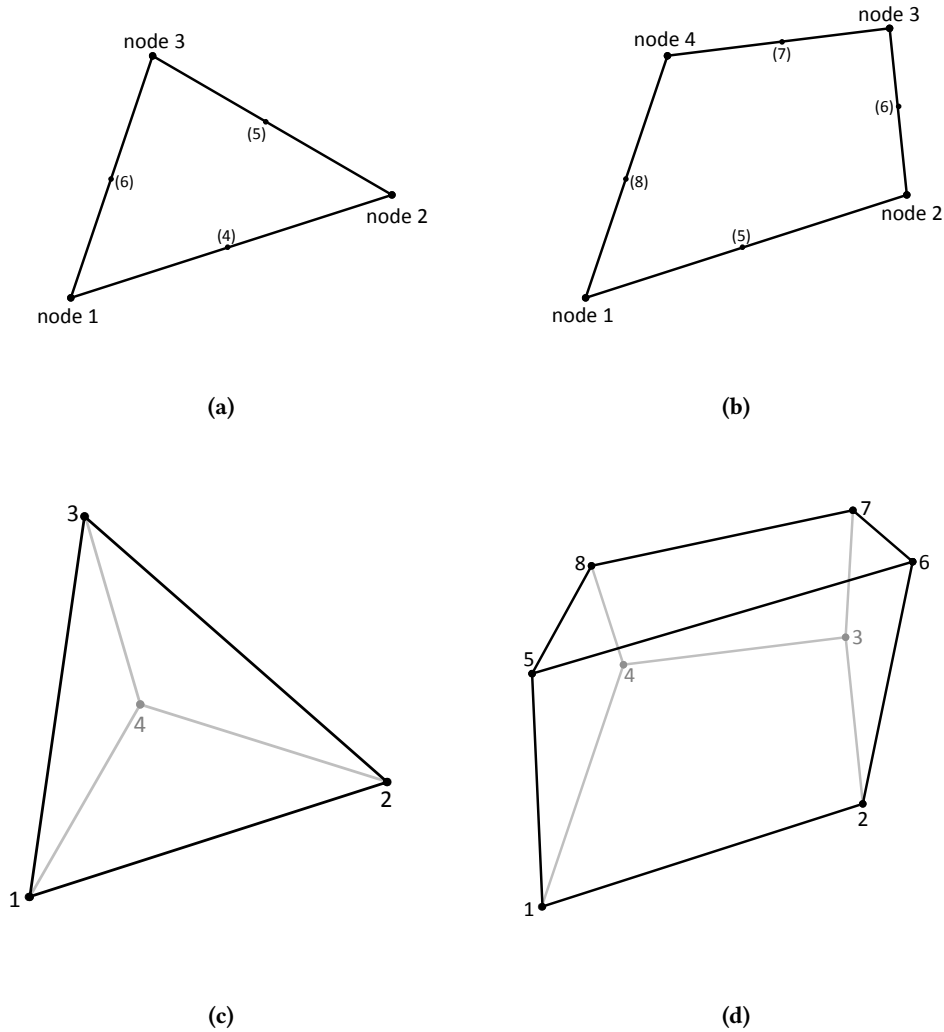


Figure 2.8: Most common finite element shapes - 2D finite elements are usually (a) triangles or (b) quadrilateral elements. The most common 3D finite elements are (c) tetrahedra and (d) hexahedra. For linear shape functions, only the vertices (corners) of the elements are considered as nodes. For quadratic shape functions, additionally the midpoints of the edges are used as nodes (numbering in brackets). More nodes on the edges and inside the elements are added, as the order of the shape functions is increased.

The most common finite elements in 2D are triangles and quadrangles, and in 3D tetrahedra and hexahedra (figure 2.8). While it is possible to mix different element shapes in one mesh, it is not commonly done. Methods for subdividing a domain into a mesh are numerous, but the aim of all methods is the same: to generate a mesh with non-overlapping elements filling the whole domain. Elements should have as few small angles as possible, since this aids accuracy and

convergence of the forward solutions (Shewchuk, 1997). Current meshing packages use different meshing strategies. While Netgen uses an advancing front algorithm (Schöberl, 1997), CGAL (The CGAL Project, 2015) first generates a surface Delaunay triangulation at the object surface and internal structural surfaces (Boissonnat and Oudot, 2005), followed by Delaunay refinement of the volumes between surfaces (Shewchuk, 1997). The fundamental principle of Delaunay triangulation and refinement is to generate triangles and tetrahedra, that do not contain vertices in the interior of their circumcircle or circumsphere (Shewchuk, 1997). In 2D, the Delaunay triangulation can be found by connecting points with neighbouring Voronoi polygons (Ho-Le, 1988).

After the initial meshing, the mesh usually has to be smoothed, in order to remove badly shaped elements. CGAL for instance includes four different mesh optimisation algorithms, which iteratively run through the mesh, moving nodes and splitting badly shaped elements. The *Lloyd optimisation* minimises a global energy defined as the L^1 -norm of the error when the square coordinate function x^2 is interpolated on the mesh domain. The energy gradient is set to zero by moving the vertices. Analogously, *optimal Delaunay triangulation* minimises global energy by updating the Delaunay triangulation instead of the vertices. Jointly with these two global optimisation methods, CGAL uses two local optimisers to remove the worst elements (*perturb* and *exude*).




Figure 2.9: Thorax model created from a 2D outline - Based on a 2D outline of the thorax and lungs segmented from an MRI scan, this 3D head model was generated with Netgen directly in EIDORS (Grychtol et al., 2012).



Figure 2.10: Head model based on manual MRI segmentation - This head model consisting of smooth curves (called NURBS) was manually created from MRI datasets (Tizzard et al., 2005).

For most EIT applications, meshes with relatively simple geometries (e.g. cylinders) are used, which can be created with Netgen directly in EIDORS (Grychtol and Adler, 2013). Three-dimensional thorax models based on two-dimensional outlines segmented from MRI scans have been created with the same method (Grychtol et al., 2012). The first meshes used for head EIT were either homogeneous spheres or concentric spheres with different conductivities

2. Literature Review

(Bagshaw et al., 2003). Since the observation, that more accurate head models (figure 2.10; Tizzard et al., 2005) result in better images (Bagshaw et al., 2003), there was an increased interest in generating more physiologically realistic head meshes. Vonach et al. (2012) presented a new, semi-automated workflow for creating head meshes based on either a CT or MRI scan. In a first step, brain, cerebrospinal fluid (CSF), skull and scalp were segmented from the available scan. For MRI scans, the open-source software BrainSuite was used for segmentation, while for CT scans an expectation maximisation algorithm was used to fit prior knowledge of brain, CSF and scalp shapes into the segmented skull with the registration tool NiftyReg (Modat et al., 2010). The created segmentations were then meshed as surfaces with the MeshLab software, and a tetrahedral mesh was generated from the surface meshes with the CUBIT meshing software. When comparing the thickness of the skull in the resulting meshes (figure 2.11), it becomes apparent that CT and MRI scans should be used conjointly to obtain more accurate tissue representations.

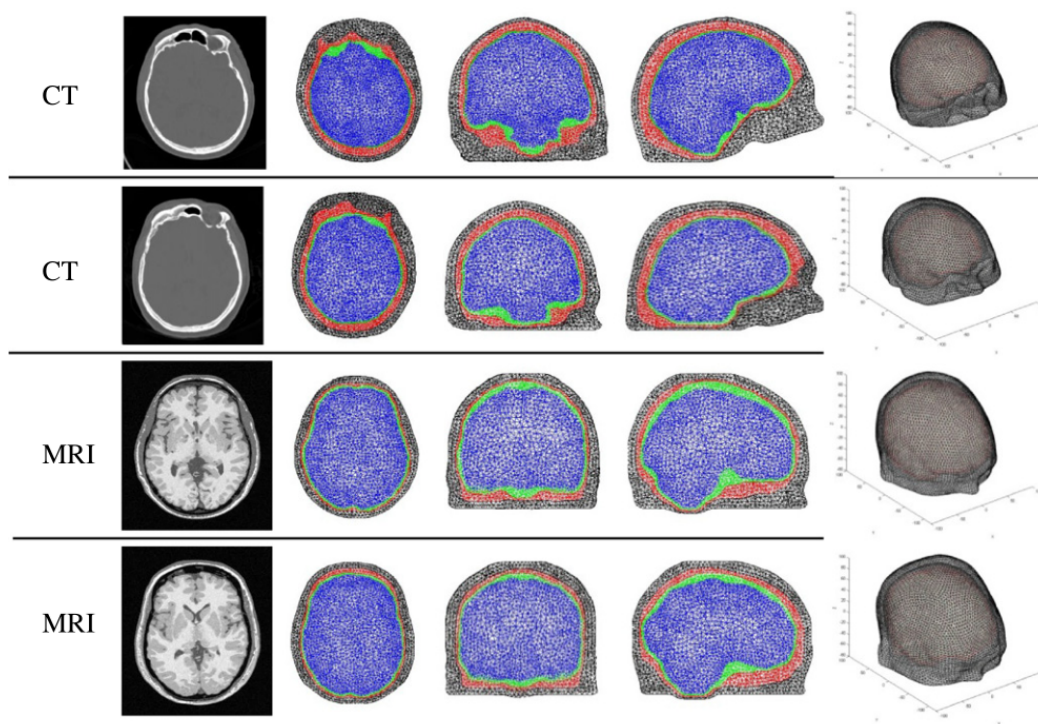


Figure 2.11: Meshes generated from either CT or MRI scan - Meshes created by Vonach et al. (2012) by either morphing a brain template model into a skull segmented from a CT scan or by estimating the skull shape from an MRI scan.

2.2.4 Weak Formulation

For the finite element method, a *weak formulation* of the problem is required. Multiplying (2.11a) by any test function $v \in H^1(\Omega)$ and integrating over the whole domain Ω gives

$$-\int_{\Omega} v \nabla \cdot (\sigma \nabla u) \, d\Omega = 0. \quad (2.12)$$

Integration by parts leads to

$$-\int_{\Omega} v \nabla \cdot (\sigma \nabla u) \, d\Omega = \int_{\Omega} \nabla v \cdot (\sigma \nabla u) \, d\Omega - \int_{\Gamma} v \left(\sigma \frac{\partial u}{\partial \nu} \right) \, d\Gamma, \quad (2.13)$$

where the outward normal derivative can be replaced with the potential difference between the electrode and the underlying nodes

$$\int_{\Gamma} v \left(\sigma \frac{\partial u}{\partial \nu} \right) \, d\Gamma = \sum_{m=1}^M \frac{1}{z_m} \int_{\Gamma_m} v (U_m - u) \, d\Gamma_m. \quad (2.14)$$

Applying (2.13) and (2.14) to (2.12) gives

$$\int_{\Omega} \nabla v \cdot (\sigma \nabla u) \, d\Omega = \sum_{m=1}^M \frac{1}{z_m} \int_{\Gamma_m} v (U_m - u) \, d\Gamma_m. \quad (2.15)$$

This is the weak formulation used in EIDORS (Polydorides and Lionheart, 2002). In DUNE however, one has to first calculate the nodal potentials within the mesh, and only after to compute the electric potential on the electrodes (chapter 4). Replacing U_m/z_m in (2.15) with the boundary condition of the CEM

$$I_m = U_m \frac{|\Gamma_m|}{z_m} - \frac{1}{z_m} \int_{\Gamma_m} u \, d\Gamma_m \quad (2.16)$$

gives the weak formulation used in PEITS

$$\begin{aligned} \int_{\Omega} \sigma \nabla v \cdot \nabla u \, d\Omega + \sum_{m=1}^M \frac{1}{z_m} \int_{\Gamma_m} v u \, d\Gamma_m - \sum_{m=1}^M \frac{1}{z_m |\Gamma_m|} \int_{\Gamma_m} v \, d\Gamma_m \int_{\Gamma_m} u \, d\Gamma_m \\ = \sum_{m=1}^M \frac{1}{|\Gamma_m|} \int_{\Gamma_m} v I_m \, d\Gamma_m. \end{aligned} \quad (2.17)$$

Proposition 2.2.1. *The system (2.17) is coercive if a ground condition is applied.*

Proof. To prove that this system is coercive, and thus uniquely solvable, it has to be shown that left-hand side (LHS) is positive. Replacing v by u gives

$$\text{LHS} = \int_{\Omega} \sigma \nabla u \cdot \nabla u \, d\Omega + \sum_{m=1}^M \frac{1}{z_m} \int_{\Gamma_m} u^2 \, d\Gamma_m - \sum_{m=1}^M \frac{1}{z_m |\Gamma_m|} \int_{\Gamma_m} u \, d\Gamma_m \int_{\Gamma_m} u \, d\Gamma_m. \quad (2.18)$$

2. Literature Review

The first term cannot be negative, thus it is sufficient to show that the second term is larger than the third term. By multiplying both with z_m , Cauchy-Schwarz gives

$$\int_{\Gamma_m} u^2 d\Gamma_m \geq \frac{1}{|\Gamma_m|} \left(\int_{\Gamma_m} u d\Gamma_m \right)^2 \quad (2.19)$$

with $u = \text{const.}$ leading to equality. Because a constant u sets the first term of the LHS to zero as well, the system is only positive semi-definite. Thus, an additional constraint is required to make the system uniquely solvable. Setting a ground condition does this. \square

2.2.5 Galerkin Formulation

The approximation to the correct solution u_{corr} on a finite element mesh with V nodes can be written as $u = \sum_{j=1}^V u_j \phi_j$, where the set of V *basis (or shape) functions* ϕ is defined as follows

$$\phi_j = \begin{cases} 1 & \text{on the node } j \\ 0 & \text{on all other nodes.} \end{cases} \quad (2.20)$$

Therefore, on each element linear shape functions approximate u_{corr} linearly from one vertex to another, while for quadratic shape functions an additional node is inserted in the midpoint of each edge (figure 2.12). Taking advantage of this new formulation for the approximation u (also called *trial function*) and replacing the test function v with the basis functions ϕ , the weak formulation as used in EIDORS, (2.15), can be written in Galerkin form as a set of $k = 1, \dots, V$ equations




Figure 2.12: Linear and quadratic shape functions - For quadratic shape functions, an additional node is inserted at the midpoint of each edge.

$$\sum_{j=1}^V u_j \left(\underbrace{\sigma \int_{\Omega} \nabla \phi_k \cdot \nabla \phi_j}_{\mathbf{A}_S} + \underbrace{\sum_{m=1}^M \frac{1}{z_m} \int_{\Gamma_m} \phi_k \cdot \phi_j}_{\mathbf{A}_{ii}} \right) + \sum_{m=1}^M U_m \left(\underbrace{-\frac{1}{z_m} \int_{\Gamma_m} \phi_k}_{\mathbf{A}_{ei}} \right) = 0. \quad (2.21)$$

From (2.16) an additional set of $m = 1, \dots, M$ equations is obtained, defining the electrode potential in dependence of the underlying nodal potentials, the contact impedance and the input current:

$$\sum_{j=1}^V u_j \left(\underbrace{-\frac{1}{z_m} \int_{\Gamma_m} \phi_j}_{\mathbf{A}_{ie}} \right) + U_m \underbrace{\frac{|\Gamma_m|}{z_m}}_{\mathbf{A}_{ee}} = I_m. \quad (2.22)$$

Combining (2.21) with (2.22) the linear system to be solved becomes

$$\begin{bmatrix} \mathbf{A}_S + \mathbf{A}_{ii} & \mathbf{A}_{ei} \\ \mathbf{A}_{ie} & \mathbf{A}_{ee} \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{U} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{I} \end{bmatrix}, \quad (2.23)$$

where \mathbf{u} is the vector of the nodal potentials u_j , \mathbf{U} the vector of the electrode potentials U_m and \mathbf{I} the current vector. The Galerkin formulation for the PEITS implementation (chapter 4) can be obtained, by proceeding accordingly with the weak formulation (2.17) as the starting point. In the following section, the resulting system of equations is simply written as $\mathbf{Ax} = \mathbf{b}$.

2.2.6 Numerical Algorithms and Preconditioning

The overwhelming majority of linear systems of equations stemming from the finite element method are solved with the conjugate gradient (CG) algorithm (very well explained in Shewchuk, 1994), because they are positive-definite, square, symmetric, large and sparse. To improve the convergence of CG on large problems, *preconditioning* should be used (Shewchuk, 1994). Preconditioning improves the condition number of the system matrix $\kappa(\mathbf{A})$ of the system

$$\mathbf{Ax} = \mathbf{b}, \quad (2.24)$$

by finding a symmetric, positive-definite matrix \mathbf{M} that approximates \mathbf{A} . If \mathbf{M} is easier to invert than \mathbf{A} and $\kappa(\mathbf{M}^{-1}\mathbf{A}) \ll \kappa(\mathbf{A})$, then the system

$$\mathbf{M}^{-1}\mathbf{Ax} = \mathbf{M}^{-1}\mathbf{b} \quad (2.25)$$

requires less CG iterations to converge than the original problem. The preconditioned CG method is then given by the algorithm

Algorithm 1 Preconditioned Conjugate Gradient

```

 $k = 0, \mathbf{x}_0 = \text{initial guess}$ 
 $\mathbf{r}_0 = \mathbf{b} - \mathbf{Ax}_0$ 
 $\mathbf{d}_0 = \mathbf{M}^{-1}\mathbf{r}_0$ 
while stopping criterion not fulfilled do
   $\tau_k = \frac{\mathbf{r}_k^\top \mathbf{M}^{-1} \mathbf{r}_k}{\mathbf{d}_k^\top \mathbf{A} \mathbf{d}_k}$ 
   $\mathbf{x}_{k+1} = \mathbf{x}_k + \tau_k \mathbf{d}_k$ 
   $\mathbf{r}_{k+1} = \mathbf{r}_k - \tau_k \mathbf{A} \mathbf{d}_k$ 
   $\beta_{k+1} = \frac{\mathbf{r}_{k+1}^\top \mathbf{M}^{-1} \mathbf{r}_{k+1}}{\mathbf{r}_k^\top \mathbf{M}^{-1} \mathbf{r}_k}$ 
   $\mathbf{d}_{k+1} = \mathbf{M}^{-1} \mathbf{r}_{k+1} + \beta_{k+1} \mathbf{d}_k$ 
   $k = k + 1$ 
end while

```

2. Literature Review

where the main difficulty is to find a preconditioner which improves the convergence well enough to make up for the cost of computing $\mathbf{M}^{-1}\mathbf{r}_k$ in each iteration. The aim of such a preconditioner is to move the eigenvalues as close together as possible. In that sense $\mathbf{M} = \mathbf{A}$ would be the perfect preconditioner, but since preconditioning is only considered because \mathbf{A} is difficult to invert, this is a very bad option.

A preconditioner which is very easy to invert is the *Jacobi preconditioner*, which is a diagonal matrix with the diagonal entries of \mathbf{A} (i.e. the matrix \mathbf{A} is scaled along the coordinate axes rather than along the eigenvector axes, if $\mathbf{M} = \mathbf{A}$ were used). Jacobi preconditioning already improves the speed of CG significantly. There are, however, more efficient preconditioners.

The *Cholesky decomposition* is a method to represent \mathbf{A} as the product of a lower triangular matrix \mathbf{L} and its transpose $\mathbf{A} = \mathbf{L}\mathbf{L}^\top$. $(\mathbf{L}\mathbf{L}^\top)^{-1}\mathbf{r}_k$ is then solved by forward and back substitution. For sparse systems such as the ones from FEM implementations the full Cholesky decomposition is rarely used. Instead, an *incomplete Cholesky decomposition* is used as preconditioner. This is an approximation \mathbf{K} similar to \mathbf{L} , which can for instance be obtained by the same algorithm, but instead of storing all entries of \mathbf{L} only the ones in the same position as the non-zero entries of \mathbf{A} are stored in \mathbf{K} . Cholesky decomposition is generally faster than the similar LU decomposition.

Multigrid algorithms are very effective for solving systems of discretised partial differential equations, and are based on the idea to support the convergence of the full problem by repeatedly adding a correction of the residual from a coarser version of the same problem. Intuitively this can be thought of as an iterative global communication of local errors. While high-frequency components of the residual are corrected on the fine level, the coarse level corrects for low-frequency components. This can be done in several hierarchical layers using recursion. *Geometric multigrid* (reviewed e.g. in Wesseling and Oosterlee, 2001) constructs several coarse grids based on the fine grid. Then the correction terms are computed using discrete systems constructed on these coarse grids. However, the construction of coarse grids can be difficult if the fine grid is unstructured and the underlying geometry is complicated. For these reasons *algebraic multigrid* methods (reviewed in Stüben, 2001) have been developed, which construct the coarser levels directly from the system matrix. This makes AMG very easy to use, as it can be applied to any discretised system without having to supply geometric informations. For $l = 1, \dots, L$ levels $\mathbf{A}_l = \mathbf{P}_{l-1}^\top \mathbf{A}_{l-1} \mathbf{P}_{l-1}$ with the *prolongation matrix* \mathbf{P}_{l-1} , a basic multigrid algorithm takes the following form:

Algorithm 2 Basic Multigrid Method

Perform some pre-smoothing iterations on $\mathbf{A}_l \mathbf{x}_l = \mathbf{b}_l$
 Restrict the residual to a coarser level: $\mathbf{b}_{l+1} = \mathbf{P}_l^T (\mathbf{b}_l - \mathbf{A}_l \mathbf{x}_l)$
if $l + 1 = L$ **then**
 Solve $\mathbf{A}_{l+1} \mathbf{x}_{l+1} = \mathbf{b}_{l+1}$ with a direct method
else
 Set the initial guess $\mathbf{x}_{l+1} = 0$
 Recursively call the Multigrid Algorithm: $\mathbf{x}_{l+1} = \text{MG}(\mathbf{x}_{l+1}, \mathbf{b}_{l+1})$
end if
 Prolong \mathbf{x}_{l+1} and add to \mathbf{x}_l : $\mathbf{x}_l = \mathbf{x}_l + \mathbf{P}_l \mathbf{x}_{l+1}$
 Perform some post-smoothing iterations on $\mathbf{A}_l \mathbf{x}_l = \mathbf{b}_l$

The pre- and post-smoothing iterations are also called *relaxation steps*. By reducing the tolerance of the MG algorithm, it can be used as a very efficient preconditioner for iterative Krylov subspace methods such as CG. Two fast and parallel implementations of AMG are BoomerAMG from Hypre (Henson and Yang, 2002) and ML from Trilinos (Tuminaro and Tong, 2000).

2.2.7 Derivation of the Jacobian Matrix

The *Jacobian (or sensitivity) matrix*, which is required in EIT image reconstructions, linearly relates a change in conductivity to a change in measured voltages. The entry in the r th row and the n th column is defined as

$$J_{r,n} = \frac{\partial V_r}{\partial \sigma_n}, \quad (2.26)$$

where V_r is the r th measured voltage and σ_n is the conductivity of finite element n . This derivation of a first order approximation of the Jacobian is adapted from Polydorides and Lionheart (2002). For a test function $v = u$, the weak form of the complete electrode model problem, (2.15), can be written as

$$\int_{\Omega} \sigma |\nabla u|^2 d\Omega = \int_{\Gamma} u \sigma \frac{\partial u}{\partial \nu} d\Gamma = \sum_{m=1}^M \int_{\Gamma_m} \left(U_m - z_m \sigma \frac{\partial u}{\partial \nu} \right) \sigma \frac{\partial u}{\partial \nu} d\Gamma_m, \quad (2.27)$$

which gives the power conservation formula, stating that the power input is either consumed in the domain Ω or on the contact impedance z_m :

$$\int_{\Omega} \sigma |\nabla u|^2 d\Omega + \sum_{m=1}^M \int_{\Gamma_m} z_m \left(\sigma \frac{\partial u}{\partial \nu} \right)^2 d\Gamma_m = \sum_{m=1}^M U_m I_m. \quad (2.28)$$

2. Literature Review

Looking at small perturbations $\sigma \rightarrow \sigma + \delta\sigma$, $u \rightarrow u + \delta u$ and $U_m \rightarrow U_m + \delta U_m$ and ignoring second-order terms, gives

$$\begin{aligned} \int_{\Omega} \delta\sigma |\nabla u|^2 d\Omega + 2 \int_{\Omega} \sigma \nabla u \cdot \nabla \delta u d\Omega \\ + 2 \sum_{m=1}^M \int_{\Gamma_m} z_m \left(\sigma \frac{\partial u}{\partial \nu} \right) \left(\delta\sigma \frac{\partial u}{\partial \nu} \right) d\Gamma_m = \sum_{m=1}^M I_m \delta U_m. \end{aligned} \quad (2.29)$$

Replacing

$$\left(\delta\sigma \frac{\partial u}{\partial \nu} \right) = \frac{1}{z_m} (\delta U_m - \delta u) \quad (2.30)$$

and using the weak formulation (2.27) with $v = \delta u$ instead of $v = u$

$$\int_{\Omega} \sigma \nabla u \cdot \nabla \delta u d\Omega = \int_{\Gamma} \delta u \sigma \frac{\partial u}{\partial \nu} d\Gamma \quad (2.31)$$

leads to

$$\begin{aligned} \int_{\Omega} \delta\sigma |\nabla u|^2 d\Omega + 2 \int_{\Gamma} \delta u \sigma \frac{\partial u}{\partial \nu} d\Gamma - 2 \sum_{m=1}^M \int_{\Gamma_m} \delta u \sigma \frac{\partial u}{\partial \nu} d\Gamma_m \\ + 2 \sum_{m=1}^M \delta U_m \int_{\Gamma_m} \sigma \frac{\partial u}{\partial \nu} d\Gamma_m = \sum_{m=1}^M I_m \delta U_m. \end{aligned} \quad (2.32)$$

The second and third term cancel and the fourth term is simply $2 \sum I_m \delta U_m$. Consequently, the total change in power is given as

$$\sum_{m=1}^M I_m \delta U_m = - \int_{\Omega} \delta\sigma |\nabla u|^2 d\Omega. \quad (2.33)$$

In order to find the change in voltage for a specific measurement r , the forward problem is solved once using a unit *measurement current* I^a between the two measurement electrodes (giving the *adjoint field* $u(I^a) = u^a$) and once using the actual driving current pattern I^d , giving u^d . Applying formula (2.33) to $u^d + u^a$ and $u^d - u^a$ and subtracting them from each other gives the formula for the adjoint field approach, which is commonly used to compute the Jacobian:

$$\delta V_r = \delta V_{da} = - \int_{\Omega} \delta\sigma \nabla u^d \cdot \nabla u^a d\Omega. \quad (2.34)$$

δV_r is then the linear approximation of the measured voltage change between the two measurement electrodes for a conductivity change $\sigma \rightarrow \sigma + \delta\sigma$, when current pattern I^d is injected (Polydorides, 2009).

2.2.8 Electrode Movement Jacobian Matrix

EIT reconstructions are not limited to conductivities, but also characteristics such as electrode positions, boundary shape and contact impedances can be recovered from EIT measurements (Lionheart, 1998; Nissinen et al., 2009; Dardé et al., 2013). Equivalently to conductivity reconstructions, recovery of another parameter relies on a Jacobian matrix describing the impact a change of the parameter has on the recorded voltages. The first approach to computing the Jacobian with respect to electrode movements was the perturbation technique, where the boundary nodes of the electrodes were moved in the forward model and the resulting change in voltages was simulated (Soleimani et al., 2006). Since the perturbation technique is very inefficient and movement of nodes is restricted on fine meshes, a direct formulation of the electrode movement Jacobian was proposed (Gómez-Laberge and Adler, 2008). The effect of movement of a vertex on the recorded voltages was studied analytically in the formulation of the discrete EIT forward problem, which resulted in a faster computation of the movement Jacobian. The approximation error approach is somewhat related to the perturbation technique, in that a Bayesian reconstruction framework is trained to correct for electrode position errors by simulating many different electrode positions (Nissinen et al., 2011).

The most fundamental and straightforward way of computing a Jacobian matrix describing voltage changes due to electrode boundary changes was proposed by Dardé et al. (2012). They derived an explicit formula for the Fréchet derivative with respect to the electrode boundary, which can be used to compute the Jacobian matrix with respect to electrode size, position and shape. The detailed derivation and proof of this electrode boundary Jacobian (EBJ) is given in their publication and only the parts relevant for the implementation of the EBJ are summarised here.

Electrode boundary changes can be characterised by C^1 vector fields on the electrode boundaries ∂E

$$\mathcal{B}_b = \{ \mathbf{v} \in C^1(\partial E, \mathbb{R}^p) \mid \|\mathbf{v}\|_{C^1(\partial E, \mathbb{R}^p)} < b \}, \quad (2.35)$$

where b is a radius larger than zero and p is the dimensionality of the problem, i.e. $p = 3$ for a realistic EIT application. For any $\mathbf{x} \in \partial E$, $\mathbf{P}_{\mathbf{x}}$ is defined as the orthogonal mapping of the ball $B_b(\mathbf{x}) = \{ \mathbf{z} \in \mathbb{R}^p \mid \|\mathbf{z} - \mathbf{x}\| < b \}$ onto the boundary Γ of the domain. Using these definitions, a modified boundary of electrode m can be formulated as

$$\partial E_m^{\mathbf{v}} = \{ \mathbf{z} \in \Gamma \mid \mathbf{z} = \mathbf{P}_{\mathbf{x}}(\mathbf{x} + \mathbf{v}(\mathbf{x})) \text{ for some } \mathbf{x} \in \partial E_m \}, \quad (2.36)$$

where $\mathbf{v} \in \mathcal{B}_b$ is the vector field defining the change in electrode boundary (e.g. change in electrode size, position or shape).

The measurement map including perturbed electrodes can then be considered as the operator

$$\mathbf{R} : (\mathbf{v}, I^d) \longrightarrow U^d(\mathbf{v}), \quad \mathcal{B}_b \times \mathbb{R}_{\diamond}^M \longrightarrow \mathbb{R}^M, \quad (2.37)$$

2. Literature Review

and its Fréchet derivative with respect to the vector field \mathbf{v} fulfils

$$\sum_{m=1}^M (\mathbf{R}'(\mathbf{v}, I^d))_m I_m^a = - \sum_{m=1}^M \frac{1}{z_m} \int_{\partial E_m} (\mathbf{v}_\tau \cdot \mathbf{n}_{\partial E_m}) (U_m^d - u^d) (U_m^a - u^a) \, ds, \quad (2.38)$$

where \mathbf{v}_τ is the component of \mathbf{v} which is tangential to Γ , $\mathbf{n}_{\partial E_m}$ the outward normal of ∂E_m which is tangential to Γ and (u^d, U^d) the solution to the unperturbed CEM forward problem corresponding to the drive current $I^d \in \mathbb{R}_\diamond^M$. The adjoint field (u^a, U^a) is the forward solution to the unit measurement current $I^a \in \mathbb{R}_\diamond^M$.

It is interesting to note, that the computation of the EBJ resembles the computation of the conductivity Jacobian (2.34) using the adjoint field method, in that the results of ‘drive current’ and ‘measurement current’ injections are used. Therefore, the calculation of the EBJ does not require any additional forward simulations to those performed for the calculation of the traditional Jacobian matrix.

2.2.9 Forward Solvers

Most research groups in EIT currently use the Electrical Impedance Tomography and Diffuse Optical Tomography Reconstruction Software EIDORS (Adler and Lionheart, 2006) which is programmed in Matlab (The MathWorks, US). EIDORS provides a set of useful features, such as 2D and 3D forward simulations and an extensive list of reconstruction algorithms, visualisation functions and more. Horesh et al. (2006) adapted EIDORS with different preconditioners and more efficient routines, resulting in a version called SuperSolver which is still used in our group at UCL. For large meshes, however, Matlab suffers from a lack of efficient parallel programming possibilities, which makes the computation of forward solutions a lengthy task. Additionally, the Matlab implementations are relatively excessive in terms of memory usage, thereby physically limiting the size of the meshes that can be used.

Borsic et al. (2010) moved the forward and the Jacobian calculations (but not the assembly of the system matrix) to the sparse parallel direct solver library PARDISO (Schenk, 2015) to surpass these limitations. They were able to improve the speed for forward simulations about 5.3 fold compared to Horesh et al. (2006) and used it on meshes with around half a million elements. Direct solvers require large amounts of memory, which usually limits the possible mesh size. Furthermore it is shown in chapter 4, that the assembly of a direct solver is much slower than that of a good preconditioner, resulting in faster execution times for iterative methods, depending on the number of unique current injection patterns. In particular, algebraic multigrid preconditioning has been shown to improve the solution time significantly (Soleimani et al., 2005). GPU based computations have already successfully been applied to the calculation of the Jacobian matrix (Borsic et al., 2012), where fast access to the memory is paramount.

A different approach to the forward modelling in EIT is to use boundary elements (Gençer

and Tanzer, 1999), a technique which requires the head to be modelled as enclosed surfaces of the different tissues with fixed conductivity. This works well for piecewise homogeneous media, but not for complicated heterogeneous geometries like the human head.

2.3 Inverse Problem

The *inverse problem* in EIT is to find an approximation to the correct conductivity distribution based on the experimentally acquired current-to-voltage map. It is ill-posed (reviewed e.g. in Borcea, 2002), which means that small changes in the measurements can lead to large changes in the reconstructed image. This makes the inverse problem very unstable in the presence of noise. Furthermore, it has only been proven to have a unique solution for isotropic conductivities and a full knowledge of the boundary (Sylvester and Uhlmann, 1987; Kohn and Vogelius, 1985). In practice, the application of current and measurement of voltages are restricted to a finite set of electrodes and human tissues are sometimes strongly anisotropic (i.e. their conductivity depends on the direction of the current). Therefore, it is not possible to precisely reconstruct the conductivities and instead the aim is to find an approximate solution using a stabilised reconstruction method. The reconstruction is commonly turned into a well-posed problem by introducing prior information through regularisation. Linear reconstruction methods can be used for time-difference and frequency-difference measurements of small, localised changes. For larger changes and absolute imaging, non-linear algorithms have to be used. This review of inversion methods draws from course notes of Arridge (2015).

2.3.1 Linear Methods Based on Singular Value Decomposition

Moore-Penrose Inverse

The Jacobian matrix \mathbf{J} linearises the forward problem $\mathbf{F}(\boldsymbol{\sigma}) = \mathbf{v}$ locally around the estimated baseline conductivity distribution $\boldsymbol{\sigma}_0$. A change in conductivity $\delta\boldsymbol{\sigma}$ can therefore be linearly approximated from a measured change in voltages $\delta\mathbf{v}$ by inverting the Jacobian. Since there are generally more conductivity values to reconstruct than there are measurements, the Moore-Penrose generalised inverse

$$\delta\boldsymbol{\sigma} = \mathbf{J}^\dagger \delta\mathbf{v} = (\mathbf{J}^T \mathbf{J})^{-1} \mathbf{J}^T \delta\mathbf{v} \quad (2.39)$$

is used, which gives the least squares solution

$$\delta\boldsymbol{\sigma} = \arg \min_{\delta\boldsymbol{\sigma}} \|\mathbf{J}\delta\boldsymbol{\sigma} - \delta\mathbf{v}\|. \quad (2.40)$$

2. Literature Review

As discussed in (Lionheart et al., 2004) this inversion is very unstable, because noise is amplified proportionally to the condition number κ of the Jacobian matrix, which is defined as

$$\kappa(\mathbf{J}) = \|\mathbf{J}\| \cdot \|\mathbf{J}^{-1}\|. \quad (2.41)$$

For the 2-norm this is equivalent to $\varsigma_{max}/\varsigma_{min}$, where ς are the singular values of \mathbf{J} . This condition number is very high for ill-posed problems like the EIT inverse and the Moore-Penrose inverse can therefore not be used. This characteristic is illustrated by writing the Moore-Penrose inverse in terms of the singular value decomposition.

Singular Value Decomposition

For any $\mathbf{J} \in \mathbb{R}^{r \times n}$, the non-negative definite Hermitian $\mathbf{J}^\top \mathbf{J}$ has a complete set of orthogonal eigenvectors $\mathbf{v}_i, i = 1, \dots, n$ with corresponding real eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$. Defining the singular values as $\varsigma_i = \sqrt{\lambda_i}$, the eigenvectors of the Hermitian matrix $\mathbf{J}\mathbf{J}^\top$ are given by $\mathbf{u}_i = \varsigma_i^{-1} \mathbf{J} \mathbf{v}_i$. If $\text{rank}(\mathbf{J}) = k < n$ then the eigenvectors $\mathbf{v}_{k+1}, \dots, \mathbf{v}_n$ form an orthonormal basis of $\text{null}(\mathbf{J})$ and the eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_k$ form a basis for $\text{range}(\mathbf{J})$. With $\mathbf{\Sigma}$ a diagonal matrix of the singular values with additional zero columns or rows to make it a $r \times n$ matrix, $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_r) \in \mathbb{R}^{r \times r}$ and $\mathbf{V} = (\mathbf{v}_1, \dots, \mathbf{v}_n) \in \mathbb{R}^{n \times n}$ the singular value decomposition of \mathbf{J} is given as

$$\mathbf{J} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top \quad \text{and} \quad \mathbf{U}^\top \mathbf{J} \mathbf{V} = \mathbf{\Sigma}. \quad (2.42)$$

The singular value decomposition (SVD) is expensive to calculate numerically, but once it is known the Moore-Penrose generalised inverse is easily computed:

$$\mathbf{J}^\dagger = \mathbf{V} \mathbf{\Sigma}^\dagger \mathbf{U}^\top, \quad (2.43)$$

where $\mathbf{\Sigma}^\dagger$ is simply $\mathbf{\Sigma}^\top$ with the inverted singular values on its diagonal. Therefore, noise containing image components with small singular values is strongly amplified with the Moore-Penrose generalised inverse.

The rate of the decay of the singular values corresponds to the ill-posedness of the problem. The faster the decay, the worse ill-posed the problem. The *truncated SVD* method utilises the singular value decomposition to determine the image components that can still be reliably reconstructed in the presence of noise in the data. The singular values are truncated at an appropriate level in order to stabilise the recon-




Figure 2.13: Singular values and Tikhonov filter - The blue curve shows the logarithmic decay of the singular values. The corresponding weightings for the Tikhonov filter were computed for two regularisation parameters, $\lambda^2 = 10^{-5}$ (red) and $\lambda^2 = 10^{-10}$ (yellow).

struction, i.e. the dimensions of \mathbf{V} , $\mathbf{\Sigma}^\dagger$ and \mathbf{U}^\top are reduced. Truncating the singular values is analogous to a low-pass filter in signal processing.

Alternatively, in the so-called *Tikhonov filtering*, the singular values are gradually dampened according to a *regularisation parameter* λ (figure 2.13). The Tikhonov filter is defined as

$$0 \leq q_\lambda(\varsigma_i^2) = \frac{\varsigma_i^2}{\varsigma_i^2 + \lambda^2} \leq 1 \quad (2.44)$$

and applied to the pseudo-inverse in a process also known as spectral filtering (Chung et al., 2011) this gives

$$\mathbf{J}_\lambda^\dagger = \sum_{i=1}^{\min(n,r)} \frac{q_\lambda \mathbf{v}_i \mathbf{u}_i^\top}{\varsigma_i} = (\mathbf{J}^\top \mathbf{J} + \lambda^2 \mathbf{I})^{-1} \mathbf{J}^\top. \quad (2.45)$$

2.3.2 Linear Methods With Variational Regularisation

In the previous section, two regularisation methods based on the SVD of the Jacobian matrix were introduced. For a large system it is usually impractical and computationally expensive to construct the SVD explicitly. Instead, one can formulate the problem as a minimisation of a variational form. Instead of explicitly solving

$$\delta \sigma_\lambda = \mathbf{J}_\lambda^\dagger \delta \mathbf{v} = \sum_{i=1}^{\min(n,r)} \frac{q_\lambda \mathbf{v}_i \mathbf{u}_i^\top}{\varsigma_i} \delta \mathbf{v} = (\mathbf{J}^\top \mathbf{J} + \lambda^2 \mathbf{I})^{-1} \mathbf{J}^\top \delta \mathbf{v}, \quad (2.46)$$

the following optimisation problem can be defined

$$\delta \sigma_\lambda = \arg \min_{\delta \sigma \in \mathbb{R}^n} \left[\Phi = \|\mathbf{J} \delta \sigma - \delta \mathbf{v}\|^2 + \lambda^2 \underbrace{\|\delta \sigma\|^2}_{\Psi(\delta \sigma)} \right]. \quad (2.47)$$

This problem can then be solved by iterative optimisation algorithms, such as steepest descent or conjugate gradients (the non-linear versions of which are explained in the next section).

Regularisation Terms

The regularisation term

$$\Psi(\delta \sigma) = \|\delta \sigma\|^2 \quad (2.48)$$

is called *zero-order Tikhonov*. Tikhonov regularisation can be generalised to any quadratic functional

$$\Psi(\delta \sigma) = \|\mathbf{D}^p \delta \sigma\|^2, \quad (2.49)$$

2. Literature Review

where D^p denotes the p th derivative and can be implemented in discrete domains as a finite difference matrix. $p = 1$ gives the first-order Tikhonov regularisation, which smooths the solution without dictating it to a mean embedded in the regularisation term. On small reconstruction meshes, problems with higher order Tikhonov regularisation can also be solved with direct inversion. For this, the SVD is replaced with a generalised singular value decomposition (gSVD) of the Jacobian and the regularisation matrix, giving $\mathbf{J} = \mathbf{U}\mathbf{\Lambda}\mathbf{X}^\top$ and $D^p = \mathbf{V}\mathbf{M}\mathbf{X}^\top$. Inserting the gSVD into the objective function and solving for the conductivity change gives

$$\delta\sigma_\lambda = \left(\mathbf{J}^\top \mathbf{J} + \lambda^2 D^{p\top} D^p \right)^{-1} \mathbf{J}^\top \delta\mathbf{v} = \mathbf{X} \left(\mathbf{\Lambda}^\top \mathbf{\Lambda} + \lambda^2 \mathbf{M}^\top \mathbf{M} \right)^{-1} \mathbf{\Lambda}^\top \mathbf{U}^\top \delta\mathbf{v}. \quad (2.50)$$

Since $\mathbf{\Lambda}$ and \mathbf{M} are both diagonal matrices with a limited number of generalised singular values (Hansen, 1994), they can easily be inverted.

Another common regularisation term is *total variation* (TV), which allows step changes in the solution by applying the 1-norm

$$\Psi(\delta\sigma) = |\mathbf{D}\delta\sigma|. \quad (2.51)$$

TV regularised problems cannot be solved by matrix inversion. Instead, a dual problem is defined in the *primal dual interior point method* (Lionheart et al., 2004; Borsic and Adler, 2012).

Regularisation Parameter Choice

The choice of the regularisation parameter λ is a trade-off between noise reduction and the accuracy of fitting the data. This trade-off is best illustrated with the *L-curve*, which constitutes one of the most commonly used methods of choosing the regularisation parameter (Hansen, 1994). The L-curve is a plot of the norm $\Psi(\delta\sigma_\lambda)$ of the regularised solution versus the residual norm $\|\mathbf{J}\delta\sigma_\lambda - \delta\mathbf{v}\|$. When plotting this curve in log-log scale it almost always has a sharp corner (an 'L') for ill-posed problems like EIT (figure 2.14). The regularisation parameter in this corner is an appropriate choice, and can be found in an automated way if the corner is pronounced enough. From the L-curve, it is easily seen how a large regularisation parameter reduces noise through stronger weighting of the regularisation. At the same time, a stronger regularisation biases the reconstruction and consequently increases the residual norm.




Figure 2.14: Characteristic L-curve shape - When plotting the norm of the regularisation term against the residual norm in log-log scale, ill-posed problems often result in a L-shaped curve.

Other methods exist to automate the choice of the regularisation parameter, such as the *discrepancy principle*, *Miller criterion*, *generalised cross-validation* or *unbiased predictive risk estimation* (Arridge, 2015).

2.3.3 Non-Linear Methods

Most non-linear methods involve the iterative search for the global minimum of the objective function

$$\sigma_\lambda = \arg \min_{\sigma \in \mathbb{R}^n} [\Phi = \|\mathbf{F}(\sigma) - \mathbf{v}\|^2 + \lambda^2 \Psi(\sigma)], \quad (2.52)$$

where the absolute values σ and \mathbf{v} are replaced by changes $\delta\sigma$ and $\delta\mathbf{v}$, if differential measurements are reconstructed. In each iteration, a *search direction* is computed and the algorithm moves in this direction by a certain distance (*step size*). With the new value for σ , the forward problem is re-evaluated and if the functional value is still above a defined stopping criterion, another iteration is performed.

Steepest Descent

The *steepest descent* method minimises the function $\Phi(\sigma)$ by finding the steepest descent direction in each iteration step k , which is the negative gradient $-\nabla\Phi$ at point σ_k . The gradient is linearly approximated with the Jacobian matrix \mathbf{J}_{σ_k} at this point

$$\nabla\Phi(\sigma_k) = \mathbf{J}_{\sigma_k}^\top (\mathbf{F}(\sigma_k) - \mathbf{v}) + \lambda^2 \nabla\Psi(\sigma_k). \quad (2.53)$$

and the solution is then updated according to step size τ_k

$$\sigma_{k+1} = \sigma_k - \tau_k \nabla\Phi(\Delta\sigma_k). \quad (2.54)$$

The convergence of the steepest descent method can be very slow for problems with an anisotropic optimisation surface (Shewchuk, 1994). Therefore the conjugate gradient method is almost always a better choice.

Conjugate Gradient

The steepest descent method builds a solution by iteratively moving in orthogonal directions, thereby often moving in the same direction several times. The *conjugate gradient* method aims at moving in each direction just once, which is achieved by Gram-Schmidt conjugation of the residuals, rather than the search directions (Shewchuk, 1997). The Polak-Ribière implementation of the non-linear CG algorithm is given as follows

2. Literature Review

Algorithm 3 Non-linear Conjugate Gradients (Polak-Ribière)

```

 $k = 0, \sigma_0 = \text{initial guess}$ 
 $d_0 = r_0 = -\nabla \Phi(\sigma_0)$ 
while stopping criterion not fulfilled do
     $\tau_k = \arg \min_{\tau} \Phi(\sigma_k + \tau d_k)$ 
     $\sigma_{k+1} = \sigma_k + \tau_k d_k$ 
     $r_{k+1} = -\nabla \Phi(\sigma_{k+1})$ 
     $\beta_{k+1} = \max \left\{ \frac{r_{k+1}^\top (r_{k+1} - r_k)}{r_k^\top r_k}, 0 \right\}$ 
     $d_{k+1} = r_{k+1} + \beta_{k+1} d_k$ 
     $k = k + 1$ 
end while

```

The advantage of the CG algorithm is a much faster convergence than for steepest descent and also that the Hessian matrix does not need to be calculated, when compared to Newton methods. This makes CG computationally and memory efficient and thus interesting for large systems.

Newton Methods

Since CG only uses the first-order derivative, its convergence can be slow. *Newton methods* try to find the unique minimiser of the quadratic approximation of the objective function

$$\Phi(\sigma_k + d_k) \approx \Phi(\sigma_k) + \nabla \Phi(\sigma_k)^\top d_k + \frac{1}{2} d_k^\top \nabla^2 \Phi(\sigma_k) d_k, \quad (2.55)$$

by setting the derivative to zero and therefore

$$\nabla^2 \Phi(\sigma_k) d_k = -\nabla \Phi(\sigma_k). \quad (2.56)$$

The new search direction is thus given by

$$d_k = -(\nabla^2 \Phi(\sigma_k))^{-1} \nabla \Phi(\sigma_k), \quad (2.57)$$

where the difficulty lies in the computation of the second derivative of the objective function, the so-called *Hessian* of Φ . The first and second derivatives are given by:

$$\nabla \Phi(\sigma_k) \longrightarrow \nabla \mathbf{F}^\top(\sigma_k) (\mathbf{F}(\sigma_k) - \mathbf{v}) + \lambda^2 \nabla \Psi(\sigma_k) \quad (2.58)$$

$$\nabla^2 \Phi(\sigma_k) \longrightarrow \nabla \mathbf{F}^\top(\sigma_k) \nabla \mathbf{F}(\sigma_k) + \nabla^2 \mathbf{F}^\top(\sigma_k) (\mathbf{F}(\sigma_k) - \mathbf{v}) + \lambda^2 \nabla^2 \Psi(\sigma_k). \quad (2.59)$$

To simplify the computation of the Hessian matrix, the *Gauss-Newton* method approximates it by omitting the second derivative of the forward map

$$\nabla^2 \Phi(\sigma_k) \approx \mathbf{J}_{\sigma_k}^\top \mathbf{J}_{\sigma_k} + \lambda^2 \nabla^2 \Psi(\sigma_k). \quad (2.60)$$

On large meshes the Hessian matrix cannot be held in memory, let alone be inverted. Consequently, Krylov subspace methods such as generalised minimal residuals (GMRes) or CG are used to solve the system

$$\nabla \Phi(\sigma_k) = - \left(\mathbf{J}_{\sigma_k}^\top \mathbf{J}_{\sigma_k} + \lambda^2 \nabla^2 \Psi(\sigma_k) \right) \mathbf{d}_k \quad (2.61)$$

$$= \mathbf{J}_{\sigma_k}^\top (\mathbf{J}_{\sigma_k} \mathbf{d}_k) - \lambda^2 \nabla^2 \Psi(\sigma_k) \mathbf{d}_k, \quad (2.62)$$

where the brackets highlight the order of the computations. This way, the Hessian is never computed explicitly, and only the result of its multiplication with a vector is stored. The conductivity is then iteratively updated along the computed search direction

$$\sigma_{k+1} = \sigma_k + \tau_k \mathbf{d}_k, \quad (2.63)$$

where the step size τ_k is either set to one, or chosen according to a line search in the *damped* Gauss-Newton method.

Line Search

All non-linear iterative methods profit from an accurately chosen step size τ_k , which can be found with a line search. In line search methods, there is a trade-off between the accuracy of the choice of τ_k which reduces the functional value substantially, and the time spent to compute this step size. Many methods exist to perform a line search, but they all contain the same two stages: in the first stage *brackets* are defined to limit the desirable step sizes and in the second stage a good step size is computed through bisections or interpolation (Nocedal and Wright, 1999).

Brent's line search method combines the stable bisection methods and the fast interpolation methods (Brent, 1973). It uses golden section bracketing within an initially defined interval $a < \tau_k < b$ known to contain a local minimum, and fits a parabola onto the three points. If the minimum of the parabola is between a and b , it is accepted as a new bound, otherwise the interval halves between a and b are again divided by the golden ratio. This process is continued until a predefined minimal distance between the brackets is reached.

2. Literature Review

Direct Methods

The alternative to non-linear iterative algorithms are direct methods. These have the advantage that they are faster than iterative methods. *Inclusion detection* methods search for discontinuities in the conductivity distribution. Instead of reconstructing the conductivity distribution itself, they are testing data on their plausibility when applied to a known background conductivity (which can consist of several tissues). Wherever the data are not plausible, it is assumed that there is a perturbation. The best known of these inclusion detection methods is the *factorisation method*, which has already been used in 3D EIT (Chaulet et al., 2014). Inclusion detection methods have the disadvantage that no knowledge of the conductivity of the perturbation is gained. They could, however, be used as preconditioners for other reconstruction algorithm, by indicating the region of interest.

Another direct method which has been applied to EIT is the *D-bar method* (Siltanen et al., 2000). So far it is mainly of interest for mathematicians, as it is practically limited to two-dimensional applications with simple geometries. Advantages of the D-bar method are that it is fast, gives an analytic solution rather than a finite element based approximation and that it finds the global minimum.

2.3.4 Multi-Frequency Reconstruction Methods

Since absolute imaging is very unstable with respect to modelling errors, difference imaging is currently seen as the only viable option to obtain reliable EIT images. For some diagnostic applications it is not possible to obtain reference measurements from before the change occurred (e.g. stroke, breast cancer). This has prompted researchers to investigate application of currents with different modulation frequencies. Different tissues have different conductivity spectra and could theoretically be distinguished based on the knowledge of these spectral informations. This is still a new field and there is a lack of mature reconstruction algorithms.

Simple Frequency-Difference

The simple frequency-difference (FD) algorithm is based on the assumptions that the background conductivity does not change over frequency and that the frequency dependence of the contact impedance and the EIT system can be neglected (e.g. Brown et al., 1995). Then the reconstruction is done analogous to time-difference EIT by inverting the Jacobian matrix $\mathbf{J}_{\sigma_{\omega_1}}$ computed for the conductivities at the reference frequency ω_1

$$\mathbf{v}_{\omega_2} - \mathbf{v}_{\omega_1} = \mathbf{J}_{\sigma_{\omega_1}} (\sigma_{\omega_2} - \sigma_{\omega_1}). \quad (2.64)$$

This basic approach results in very bad images when the background conductivity changes over frequency (Packham et al., 2012), which is why the weighted frequency-difference algorithm has been proposed.

Weighted Frequency-Difference

The weighted frequency-difference (WFD, Seo et al., 2008; Jun et al., 2009) is able to recover small perturbations with strongly frequency-dependent conductivity in a homogeneous, frequency-dependent background. Instead of taking the plain voltage differences between two frequencies, the voltages are weighted according to

$$\alpha = \frac{\mathbf{v}_{\omega_2}^\top \mathbf{v}_{\omega_1}}{\mathbf{v}_{\omega_1}^\top \mathbf{v}_{\omega_1}} \longrightarrow \delta \mathbf{v}_{WFD} = \mathbf{v}_{\omega_2} - \alpha \mathbf{v}_{\omega_1}. \quad (2.65)$$

In the absence of a perturbation, the conductivity change of the homogeneous background will be corrected for by α and consequently $\mathbf{v}_{WFD} = 0$. Therefore, the WFD fraction of the voltage changes consists of changes orthogonal to the homogeneous \mathbf{v}_{ω_1} , and can be reconstructed with a linear method. This algorithm has been applied to hemispherical tank data and was shown to be stable in the presence of modelling errors (Ahn et al., 2011). In cases where the linearity assumption is invalid, WFD has been shown to produce bad images (Malone et al., 2014; Jang and Seo, 2015).

Fraction Reconstruction

The fraction reconstruction algorithm was developed for diffuse optical tomography and was recently adapted to multi-frequency reconstruction in EIT (Malone et al., 2014). To enable the simultaneous use of conductivity measurements at multiple frequencies $\omega_i, i = 1, \dots, W$, the conductivity of each finite element $n = 1, \dots, N$ is split into two sets of variables: frequency-independent volume fractions f_{nj} of all expected tissues $t_j, j = 1, \dots, T$ and the conductivity spectra of these tissues ϵ_{ij} . The conductivity of the n th element at frequency ω_i is therefore described by

$$\sigma_n(\omega_i) = \sum_{j=1}^T f_{nj} \cdot \epsilon_{ij}, \quad (2.66)$$

where $0 \leq f_{nj} \leq 1$ and $\sum_{j=1}^T f_{nj} = 1$. The modified Jacobian matrix at each frequency is obtained using the chain rule

$$\frac{\partial \mathbf{F}(\boldsymbol{\sigma}_i)}{\partial \mathbf{f}_j} = \frac{\partial \mathbf{F}}{\partial \boldsymbol{\sigma}_i} \frac{\partial \boldsymbol{\sigma}_i}{\partial \mathbf{f}_j} = \frac{\partial \mathbf{F}}{\partial \boldsymbol{\sigma}_i} \epsilon_{ij} = \mathbf{J}_{\boldsymbol{\sigma}_i} \cdot \epsilon_{ij} = \mathbf{J}_{ij}. \quad (2.67)$$

The objective function to be minimised can then be expressed in terms of the fractions by

2. Literature Review

substituting the model for the conductivity (2.66):

$$\Phi(\mathbf{f}) = \sum_{i=1}^W \left\| \mathbf{F}\left(\sum_{j=1}^T \mathbf{f}_j \epsilon_{ij}\right) - \mathbf{v}_{\omega_i} \right\|^2 + \lambda^2 \Psi(\mathbf{f}), \quad (2.68)$$

where \mathbf{F} is the non-linear function translating conductivities to boundary voltages (the forward map), \mathbf{v} the measured voltages and λ the regularisation parameter. Further, $\mathbf{f}_j = \{f_{nj}; n = 1, \dots, N\} \in \mathbb{R}^{N \times 1}$ and $\mathbf{f} = \{\mathbf{f}_j; j = 1, \dots, T\} \in \mathbb{R}^{N \cdot T \times 1}$. In order to suppress frequency-independent modelling errors, changes in boundary voltages across frequencies can be taken instead of absolute data. The resulting objective function is

$$\Phi(\mathbf{f}) = \sum_{i=2}^W \left\| \mathbf{F}\left(\sum_{j=1}^T \mathbf{f}_j \epsilon_{ij}\right) - \mathbf{F}\left(\sum_{j=1}^T \mathbf{f}_j \epsilon_{1j}\right) - (\mathbf{v}_{\omega_i} - \mathbf{v}_{\omega_1}) \right\|^2 + \lambda^2 \Psi(\mathbf{f}) \quad (2.69)$$

with a first order Tikhonov regularization term of the form

$$\Psi(\mathbf{f}) = \sum_{j=1}^T \sum_{n=1}^N \sum_{l(n)} |f_{nj} - f_{l(n)j}|^2, \quad (2.70)$$

where $l(n)$ runs over all neighbours of the n th element. The constraint $\sum_{j=1}^T f_{nj} = 1$ is enforced by substituting $\mathbf{f}_1 = \mathbf{1} - \sum_{j=2}^T \mathbf{f}_j$ in the objective function.

The minimisation of the objective function (2.69) is performed by alternating steps of gradient projection and damped Gauss-Newton algorithms. The gradient projection (Nocedal and Wright, 1999) step is used to quickly move to the neighbourhood of the minimum, while considering the constraints on the fractions. This is done by computing the step sizes along the gradient, at which one of the fraction values reaches a constraint. The change of the objective function value along each of the resulting intervals is approximated quadratically using Taylor series. When the Taylor approximation finds a minimum on an interval, this so-called Cauchy point \mathbf{f}_c is chosen.

The subsequent Gauss-Newton step is only applied to the components that did not reach a constraint during the gradient projection. The search direction \mathbf{d} is calculated by solving

$$\nabla^2 \Phi(\mathbf{f}_c) \cdot \mathbf{d} = -\nabla \Phi(\mathbf{f}_c) \quad (2.71)$$

using a generalised minimal residual algorithm in order to avoid the explicit calculation of the Hessian matrix $\nabla^2 \Phi$ (section 2.3.3). The step width along direction \mathbf{d} is determined using the Brent line-search method (section 2.3.3; Brent, 1973) and the resulting minimum \mathbf{f}_g is projected back to the fraction constraints. The point, \mathbf{f}_c or \mathbf{f}_g , that gives a smaller function value is chosen for the next iteration.

Reconstruction-Classification

The reconstruction-classification algorithm was proposed as an alternative to the fraction reconstruction method, to address inaccurate knowledge of the conductivity spectra of the tissues (Malone et al., 2015). It relies on a Bayesian reconstruction framework, assigning probabilities for each element to be occupied by a specific tissue (the reconstruction step). The conductivity of each element at each frequency is then again drawn from a probability distribution. The reconstruction step is alternated with a classification step, in which the spectra of all elements are clustered in a W -dimensional scatter plot, where W is the number of frequencies used. Compared to the fraction reconstruction algorithm, the reconstruction-classification has two major disadvantages: the initialisation of the parameters requires prior knowledge of the perturbation location and a second regularisation parameter has to be set.

Chapter 3

Creation of Head Models from CT and MRI Scans

3.1 Overview

3.1.1 Introduction

Imaging in Electrical Impedance Tomography (EIT) relies on an initial model (the forward model) of the conductivity distribution in the object of interest (here the head). Because of the ill-posed image reconstruction problem, small measurement noise and forward modelling errors introduce large artefacts into the images and can make detection of physiological changes impossible. In static imaging, where no baseline measurement is available, it has been shown that inaccurate head models have a detrimental effect on image quality (chapter 5; Kolehmainen et al., 1997). Therefore, most EIT applications use time-difference (TD) data, based on the observation that most geometric and system related errors cancel out when a reference measurement is subtracted from the data measurement (Brown, 2003). TD applications of head EIT include the analysis of epileptic seizures (Bagshaw et al., 2003) and imaging of cortical activity (Tidswell et al., 2001b), but the focus of this chapter is the monitoring of patients with traumatic brain injury (TBI) for secondary bleeding (section 2.1.6; Xu et al., 2010).

The forward problem in EIT (section 2.2) is usually solved with the finite element method (FEM), where the continuous domain is approximated with piecewise linear (or quadratic or even higher-order) elements, which are assigned a fixed conductivity (Lionheart et al., 2004). Using the applied current as the Neumann boundary condition, the forward problem in EIT can be solved uniquely on the finite element mesh, and the solution approaches the real solution, as the size of the elements is reduced (and consequently the number of elements increases).

3. Head Models from CT and MRI Scans

The required resolution of the finite element mesh, to achieve a predefined accuracy of the forward solutions, can be found with local error estimates on individual elements (Johnson and MacLeod, 1994; Molinari et al., 2001; Sawicki and Okoniewski, 2010) or by analysing the convergence of the simulated voltages in dependence to regional element sizes (Aristovich et al., 2014). The influence of mesh discretisation on the inverse problem is more difficult to assess. While any influence can be excluded when the forward solutions converged, reconstructed images are generally more affected by modelling errors and noise, rather than by the accuracy of the forward solutions.

3.1.2 Background

It has previously been found that even in TDEIT, a head-shaped finite element model produces better images than spherical models (Bagshaw et al., 2003). Based on this finding, ways have been investigated to create patient specific head models based on either a CT or an MRI scan of the head (Tizzard et al., 2005; Vonach et al., 2012). However, it has not yet been ascertained if the differences between the geometries of the heads of different subjects introduce enough errors into reconstructed images to justify the use of patient specific head meshes. If not, then imaging studies of EIT of the adult head could all be undertaken with one anatomically accurate model of a generic head, which would represent a considerable saving in time and effort.

3.1.3 Purpose

The purpose of the work in this chapter was to assess how accurately a human head has to be modelled to enable EIT monitoring and imaging. The questions to be answered were: 1) How large are voltage and image errors introduced by using a generic head mesh, as opposed to a subject specific one? 2) Is a subject specific head mesh required for the detection of localised conductivity changes?

These were addressed in computer simulation based on accurate head meshes from MRI and CT scans of four patients, created using a novel method. Two conductivity perturbations with different magnitude, haemorrhage and ischaemia, were chosen to test the reconstruction quality. Boundary voltages and Jacobian matrices were computed on all four meshes, for ischaemia or haemorrhage at five different locations. Time-difference images were then reconstructed from the simulated noisy boundary voltages, using the four different meshes and a coarse, homogeneous head model. Analysis of the resulting TD image quality showed that even though the image quality was generally better with the correct mesh, the number of correct perturbation detections was not higher than with the other head meshes and even the homogeneous mesh.

3.2 Mesh Creation

The mesh generation pipeline used required a CT scan and an MRI scan of the same head, and the open source software suites Seg3D (CIBC, 2015), MIPAV (McAuliffe et al., 2001), CGAL (The CGAL Project, 2015) and MedInria (Toussaint et al., 2007).

3.2.1 Segmentation

Previous head segmentations for Electrical Impedance Tomography were based on either a CT scan and then morphing a brain into the shape of the skull (Vonach et al., 2012), or on MRI scans with poor sensitivity to the skull shape (Tizzard et al., 2005; Sadleir et al., 2010; Vonach et al., 2012). The method proposed by Vonach et al. (2012) did not allow for meshing of white matter or other tissues apart from brain, cerebrospinal fluid (CSF), skull and scalp. Also, it required that one tissue completely enclosed the inner one, which is not anatomically correct. The greatest difficulty in head segmentation for EIT is, however, the accurate modelling of the skull with respect to the soft tissues. The combination of the conductive scalp and highly resistive skull underneath means that a significant proportion of the injected current is shunted around the brain (Abascal et al., 2008). For this reason, CT scans were used in this study for accurate skull segmentations and MRI scans of the same heads for soft tissue segmentation. The segmentation was mostly manual and took around two working days per head. In the following, the approximate procedure is described, with segmented layers printed in *italics* and Seg3D tools in CAPITALS (figure 3.1).

- CT and MRI head scans of the same patient were loaded into Seg3D. The rotational mismatch was estimated and the MRI scan was then rotated to align with the CT scan using MIPAV. Once the two scans were aligned, all segmentation work was done manually in Seg3D. The resolution of the CT and MRI scans was increased to 512x512x512, using Gaussian interpolation (RESAMPLE tool). Increasing the number of pixels while smoothing the images was required to generate smooth surfaces in the resulting finite element meshes.
- The *skull* was segmented first by THRESHOLDING the CT scan at an appropriate intensity level (around one third of the maximum). The *diploë* was then found by FILLING HOLES in the *skull* and then REMOVING the *skull* from this newly created layer. The modelling of the *diploë* included anisotropy into the overall skull conductivity.
- The soft tissues were then extracted from the MRI scan. The first soft tissue to be segmented was the *white matter*, because it had a clearly visible contrast to the surrounding grey matter and could therefore be accurately found with the THRESHOLD tool (between around 40% and 50% intensity). Areas with similar intensity to the white matter were mostly outside the skull cavity, and could therefore be removed from the *white matter* by running REMOVE

3. Head Models from CT and MRI Scans

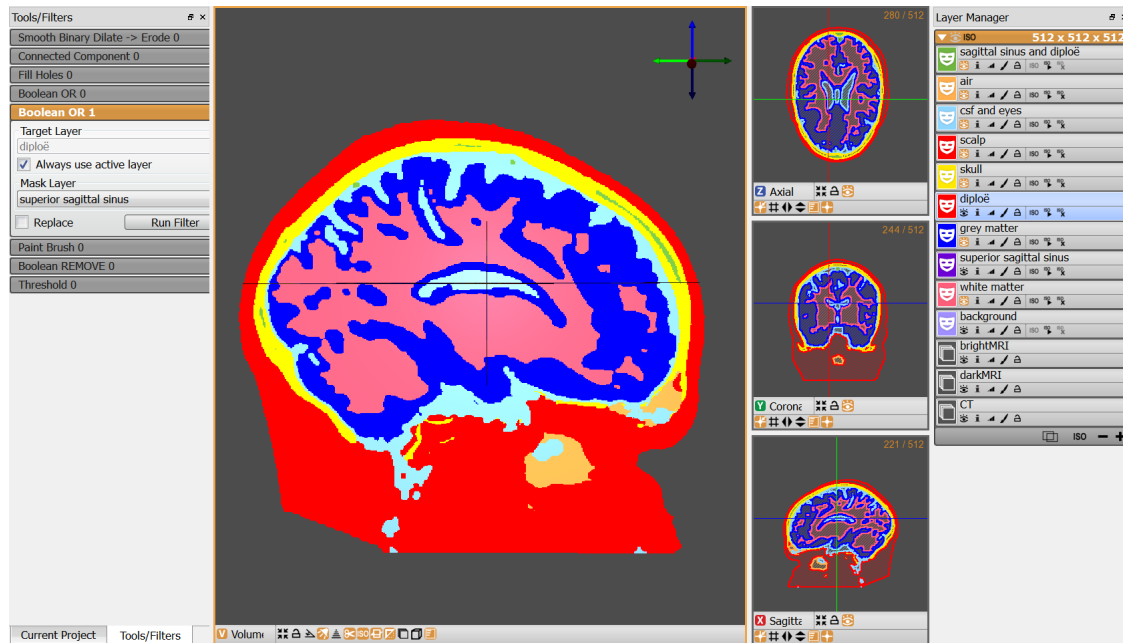


Figure 3.1: The Seg3D workspace after completed CT&MRI segmentation - Screenshot of the Seg3D workspace after completed segmentation of the CT and MRI scan. To the right are the created tissue layers and to the left the tools used.

with the *skull* segmentation as the mask, and then finding CONNECTED COMPONENTS on the *white matter*.

- Next, the *grey matter* was found in a similar manner to the *white matter*. Because the grey matter MRI intensity (around 30-40%) was similar to tissues in the lower part of the head, it was generally not possible to disconnect the actual *grey matter* from the tissues with similar intensity by simply removing the *skull*. Therefore, this was done manually using the PAINT BRUSH tool to remove the connections between the *grey matter* and other tissues, based on physiological atlases. This was the most time-consuming step in the head segmentation.
- The *superior sagittal sinus* had a high intensity on the MRI scans (50-70%) and was easily segmented by THRESHOLDING and finding CONNECTED COMPONENTS. It was added to the *diploë* layer.
- By THRESHOLDING the MRI scan at very low intensity (below 10%), the *background* was segmented. At the nose and ears, the air filled cavities were manually removed (PAINT BRUSH) from the *background* and assigned to a new segmentation layer *air*. This was done, in order not to have highly irregular surfaces in the resulting finite element meshes.
- The *cerebrospinal fluid (CSF)* and *eyes* were combined to one layer, and found by thresholding the MRI scan at an intensity level of 10-20% and subsequent subtraction of *grey matter*,

white matter, skull and diploë and superior sagittal sinus. Any unassigned gaps between *grey matter* and *skull* were filled with *CSF* by iteratively *DILATING* it and *REMOVING* the other layers to prevent overlap.

- All unassigned voxels were finally assigned to a *scalp* layer, including most of the lower head, nose and throat. This simplification was deemed acceptable, because these areas were far from the electrodes and the region of interest (i.e. the brain). The *scalp* was found by *THRESHOLDING* the MRI to include all relevant tissues (~15-70%) and then subtracting all other segmented layers. *DILATION* and *REMOVAL* of the *background* ensured that the *scalp* covered the surface of the segmented head.
- When all tissue layers were segmented, they were smoothed by running *SMOOTH BINARY DILATE* -> *ERODE* and overlaps were prevented with *BOOLEAN REMOVE*. To ensure that all pixels were assigned to a tissue, *CSF* and *scalp* were *DILATED* until they filled the *skull* and head, and then had the other layers *REMOVED* from them.
- In a final step, the segmentation was saved in .nrrd format and, using MedInria, translated into the .inr format required by the mesher.

3.2.2 Meshing

The high quality 3D mesh generator of the computational geometry algorithms library CGAL (The CGAL Project, 2015) was used to create tetrahedral meshes directly from the joint MRI&CT segmentation. The default 3D mesh generator of CGAL (written in C++) was adapted, by defining a problem specific element sizing field. This sizing field defined the desired finite element size throughout the head and was used to refine elements near the electrodes and towards the surface of the head. Specifically, all elements within a radius of 10 mm of the centre of an electrode were required to be 0.5 mm small, while the size of the other elements was assigned linearly from the surface to the centre of an ellipsoid fitted into the head (2 mm on the surface to 4 mm in the centre of the head). The effect of the mesh refinement around electrodes can be seen on the left of the slice taken of one head mesh (figure 3.2). The resulting meshes contained 3.3-3.5 million tetrahedral elements and their quality was controlled with the Joe-Liu quality measure (Liu and Joe, 1994)

$$q = 12 \frac{(3V)^{\frac{2}{3}}}{\sum_{0 \leq i < j \leq 3} l_{ij}^2}, \quad (3.1)$$

where q is the quality, V the element volume and l_{ij} the length of the edge between the i th and j th element vertex. All meshes had an average quality above 0.89 and not more than 7 elements with quality below 0.1 (with the worst element having a quality measure of 0.073).

3. Head Models from CT and MRI Scans




Figure 3.2: 3D mesh created from the segmentation shown in figure 3.1 - In this slice of the finite element mesh it can be seen, that elements towards the centre of the head are larger than the ones at the surface. Additionally, very fine elements at one spot at the back of the head indicate that this slice cuts through an electrode at this location.

After the manual location of nasion,inion and left and right ear tragus, the 32 electrode positions were computed automatically for each mesh using Matlab (MATLAB 2014b, The MathWorks, US). The electrode positions were defined according to the EEG 10-10 system (Nuwer et al., 1998) and matched the electrode assignment of the EasyCap EC40 (Brain

Products GmbH, Germany). The head dimensions were analysed based on the distances between EEG 10-10 electrode positions Fpz and Oz (length), T7 and T8 (width) and the height was found as the difference between maximum and minimum z dimensions of nodes in the meshes (table 3.1).

	mesh 1	mesh 2	mesh 3	mesh 4
length	18.25	18.48	18.20	19.47
width	14.89	15.52	15.19	15.07
height	19.22	20.10	15.88	18.67

Table 3.1: Dimensions of the four different heads used for meshing - All values are in centimetres.

3.3 Methods

3.3.1 Voltage Simulation

The conductivities used for the different tissues were 0.15 S m^{-1} for grey matter, 0.3 S m^{-1} for white matter, 1.79 S m^{-1} for CSF, 0.018 S m^{-1} for skull, 0.7 S m^{-1} for diploë and sagittal

sinus, 0.0001 S m^{-1} for air and 0.44 S m^{-1} for scalp, as compiled from literature for frequencies around 10 kHz (Horesh et al., 2006). The current level was $250 \mu\text{A}$ and contact impedances of all 32 electrodes were $z_m = 1 \text{ k}\Omega \cdot |E|$, with $|E|$ being the contact area of the electrodes with diameter 10 mm. The current injection pairs of electrodes were chosen to maximise the distance between electrodes by finding their maximum spanning tree. Measurements were made for each injection on all adjacent electrode pairs not involved in delivering the current, giving a total of 869 measured voltages from 31 independent current injections. All voltages and the Jacobian matrix were computed with PEITS (chapter 4), resulting in computation times of less than 11.5 minutes for voltages and Jacobian matrix on each healthy head model with ~ 3.4 million elements, and less than 3.5 minutes for the computation of only the voltages for each simulated stroke (on 12 processors with 20 MB cache each).

The five positions of the simulated perturbations were based on the EEG 10-10 nomenclature electrode positions. Anterior, central and posterior positions were 25%, 50% and 75% along a line between Fpz and Oz, respectively. The lateral and superior positions were halfway between the central position and T8 and Cz, respectively. The simulated strokes had a radius of 1.5 cm and conductivity 0.7 S m^{-1} for haemorrhage and 90% of the healthy conductivity for ischaemia (Horesh, 2006; Dowrick et al., 2015).

Three levels of noise were added to simulated voltages and were chosen to match, in ascending order, tank experiments (chapter 5), measurements on humans with relatively low noise and with higher noise (Goren et al., 2015). This noise is referred to as *system noise*, and accounts for noise from the instrumentation. The noise levels were $\varsigma_p = 0.006\%$ proportional noise and $\varsigma_a = 1 \mu\text{V}$ additive noise, $\varsigma_p = 0.01\%$ & $\varsigma_a = 2 \mu\text{V}$ and $\varsigma_p = 0.02\%$ & $\varsigma_a = 5 \mu\text{V}$, such that

$$V_{\text{with noise}} = V_{\text{no noise}}(1 + \text{rand}(\varsigma_p)) + \text{rand}(\varsigma_a), \quad (3.2)$$

where $\text{rand}(\varsigma)$ indicates random numbers drawn from a Gaussian distribution with zero mean and standard deviation ς .

3.3.2 Image Reconstruction

First order Tikhonov regularisation was used to bias the algorithm towards finding small connected perturbations. All images were created with a standard least-squares minimisation using generalised singular value decomposition (gSVD). The advantage of using gSVD was, that it only had to be computed once for each reconstruction mesh and could then be used for all image reconstruction by simple matrix multiplication (chapter 2.3).

To reduce the computational cost of calculating the gSVD and to prevent the ‘inverse crime’ (Lionheart et al., 2004), much smaller hexahedral meshes (3-4 thousand elements of $1 \text{ cm} \times 1 \text{ cm} \times 1 \text{ cm}$) were used for the image reconstructions. The Jacobian matrices were computed on the fine meshes and then projected onto the geometrically regular cubes of the

3. Head Models from CT and MRI Scans

hexahedral mesh. The Laplacian matrices for the first order Tikhonov regularisation were computed on the hexahedral meshes.

For all reconstructed images, the regularisation factor was kept constant for each noise level, and was $\lambda = 5 \cdot 10^{-4}$ for low noise, $\lambda = 1 \cdot 10^{-3}$ for medium noise and $\lambda = 2 \cdot 10^{-3}$ for high noise. The value for the lowest noise was chosen manually based on the shape of the L-curve (Hansen, 1994). As the noise was roughly doubled, the regularisation factor was also doubled to account for the higher noise level. The colour bars of the images were scaled according to the largest reconstructed change in the whole mesh. Therefore images of slices sometimes do not contain the maximum value indicated in the colour bar.

3.3.3 Image Quality Measures

Three metrics were used to objectively evaluate the image errors. They were all 0% in ideal images. The overall error was expressed as the sum of these three metrics in percent, and image quality as its reciprocal. Acceptable quality images were empirically found to have image error $< 100\%$ or image quality > 0.01 , respectively.

The volume P corresponding to the reconstructed perturbation was identified as the largest connected cluster of voxels with at least 75% of the maximum absolute change in the image. The region of interest (ROI) was defined as the largest connected cluster of voxels with 50% of the maximum of the simulated conductivity change. The used error metrics were then

- Location error: ratio between the distance $\|(x_P, y_P, z_P)\|$ of the centre of mass of the reconstructed perturbation P from the actual perturbation location, and the average dimension of the mesh $\text{mean}(d_x, d_y, d_z)$

$$\frac{\|(x_P, y_P, z_P)\|}{\text{mean}(d_x, d_y, d_z)}. \quad (3.3)$$

- ROI contrast: ratio between the average reconstructed change ($d\sigma_r$) outside the region of interest and the average reconstructed change in the region of interest

$$\frac{|\text{mean}_{\Omega \setminus \text{ROI}}(d\sigma_r)|}{|\text{mean}_{\text{ROI}}(d\sigma_r)|}. \quad (3.4)$$

- ROI noise: ratio between the absolute changes outside the region of interest and the absolute changes within the ROI

$$\frac{\text{mean}_{\Omega \setminus \text{ROI}}(|d\sigma_r|)}{\text{mean}_{\text{ROI}}(|d\sigma_r|)}. \quad (3.5)$$

The overall image quality was then defined as the inverse of the sum of the three error measures

3. Head Models from CT and MRI Scans

in percent (e.g. the image quality for 10% location error, 5% ROI contrast error and 25% ROI noise would be $1/(10 + 5 + 25) = 0.025$). Exemplary image reconstructions with differently assessed quality illustrate that up to 100% image errors it was possible to detect the stroke visually (figure 3.3). Between 100-120% some strokes could still be seen, while others could not. If the image errors were larger than 120%, it was no longer possible to detect strokes. 100% image errors were therefore defined as the limit of stroke detectability. The occurrence of a reconstruction with more than 120% errors was highlighted in tables 3.2 and 3.3 by colouring the entry red.




Figure 3.3: Illustration of image quality measures - (a) haemorrhage with 20% image errors, (b) ischaemia with 61% image errors, (c) 105% image errors and (d) 108% image errors. Black squares indicate the outline of the simulated stroke, thresholded at 50% of the maximum conductivity change in the hexahedral mesh.

Voltage errors between two measurements $\mathbf{v}_1 - \mathbf{v}_2$ were assessed with two metrics. The 2-norm was computed as $(\sum_r (\mathbf{v}_1 - \mathbf{v}_2)_r^2)^{1/2}$ and the average proportional errors as $\text{mean}_r(|\mathbf{v}_1 - \mathbf{v}_2|/|\mathbf{v}_1|)$, where $r = 1, \dots, 869$ indexes the individual measured voltages.

3.4 Results

3.4.1 Analysis of the Image Quality

Was the image quality better with the correct mesh?

For reconstruction on the correct mesh, the image quality metric was on average 25% better, than when different meshes were used for simulation and reconstruction. Sometimes, image reconstructions on the wrong mesh were better than reconstructions on the correct mesh. The image quality of reconstructions on the correct mesh was plotted against the average image quality of reconstructions on the other three head meshes (figure 3.4). When filtering out reconstructions that were unsuccessful in both modalities (i.e. quality below $1/100\% = 0.01$ for both), then the following ratios of better reconstruction with the correct mesh versus better reconstruction with the other meshes were obtained: 32:8 at low noise, 22:16 at medium noise and 20:10 at high noise. An example of a reconstructed image on the correct mesh and the same reconstruction on a different mesh is shown in figure 3.5.

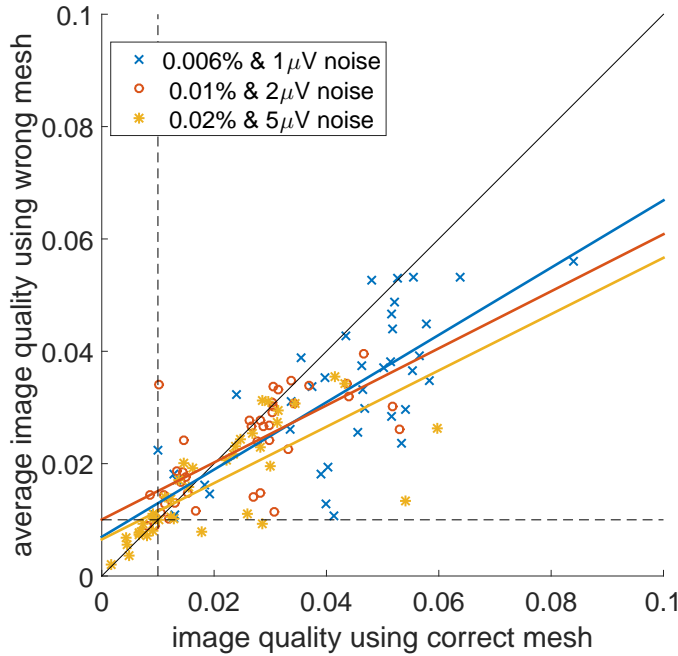


Figure 3.4: Image quality for correct reconstruction mesh vs. different head meshes - Image quality comparison for reconstructions on correct mesh and on different head meshes. The solid lines are the linear fit of the data points with the same colour. For low noise, the improved quality using the correct mesh is clearly visible (blue crosses). The difference in image quality between correct mesh and other mesh reconstructions reduces with increasing noise. For each noise level there are 40 data points, corresponding to four meshes used for simulation, five stroke positions and two stroke types.

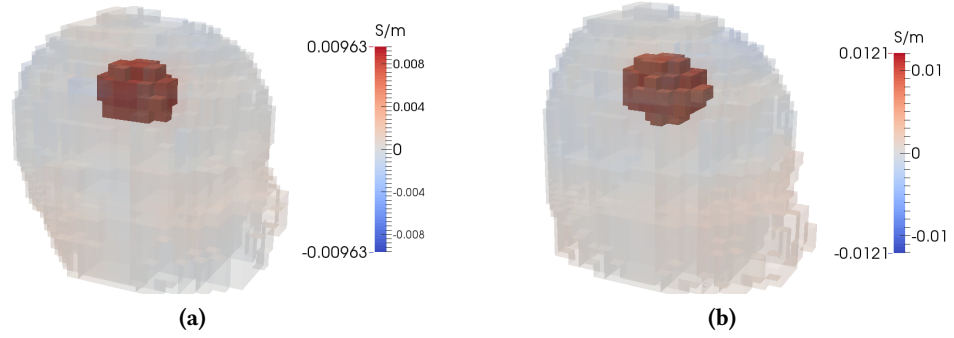


Figure 3.5: Exemplary reconstructions with the correct and a wrong mesh - Example of reconstructions of a haemorrhage in the superior position simulated on mesh 1 and reconstructed on (a) mesh 1 and (b) mesh 2. The images were thresholded at 66% of the maximum reconstructed change.

Were stroke detection rates higher with the correct mesh?

While using the correct mesh for reconstruction increased image quality on average, the total amount of successful stroke detections was not improved (figure 3.4). The number of cases where reconstructions with the correct mesh were successful, while reconstructions with the wrong meshes were not, was 3 (all at the highest noise level). On the other hand, there were 5 cases where only the reconstructions on the wrong meshes gave an acceptable image quality (2 at the highest noise level, 1 at medium noise and 2 at low noise).

Which factors did stroke detection rates depend on?

In this section, we compare the influence of mesh differences, stroke position, stroke type and system noise level on stroke detection rates. The system noise was found to have the strongest impact on image quality, and mostly affected reconstructions of ischaemia.

		simulation			
		mesh 1	mesh 2	mesh 3	mesh 4
reconstruction	mesh 1	9/9/9	10/8/6	10/8/5	10/8/5
	mesh 2	8/9/10	9/10/5	10/10/5	9/8/4
	mesh 3	10/9/10	10/10/7	10/9/7	10/9/5
	mesh 4	10/9/10	10/10/7	10/9/8	10/9/6
	h. 100k	10/8/8	10/8/5	10/8/8	9/8/9

Table 3.2: Stroke detection statistics for reconstruction and simulation meshes - Number of successful reconstructions using different voltage simulation and image reconstruction mesh pairs. The first value of each entry corresponds to small noise on the voltages, the second to medium noise and the last to high noise. The maximum score for each noise level is 10, which means that all five positions and both stroke types were visible in the reconstructed image. 0 – 100% image error were considered successful and everything above unsuccessful. If all images had an error of less than 120% the number is black, otherwise red. The last row corresponds to reconstructions on a coarse, homogeneous mesh.

3. Head Models from CT and MRI Scans

No clear differences were found between the meshes used for voltage simulation and image reconstruction (table 3.2). The only surprise was, that even at the highest noise level, almost all reconstructions from voltages simulated on mesh 1 were successful. No explanation was found for this observation. The least successful reconstructions of haemorrhages were observed for the posterior position, while the lateral position was the least successful for ischaemia recovery (table 3.3). The added system noise had a very strong influence on the successful reconstructions of ischaemic strokes. At the medium noise level, 66 out of 80 ischaemias were detected, while with the highest noise this number dropped to 33 (table 3.3).

	0.006% & 1 μ V	0.01% & 2 μ V	0.02% & 5 μ V
posterior	16/16	14/15	13/8
central	16/15	16/15	16/7
anterior	16/14	16/15	16/7
superior	16/14	16/5	16/4
lateral	16/16	16/16	16/7

Table 3.3: Stroke detection statistics for noise levels and positions - Number of successful reconstructions for different noise levels and stroke positions. The first value corresponds to haemorrhage, the second to ischaemia. The maximum score for each stroke type is 16, which means that all combinations of voltage simulation and image reconstruction meshes resulted in a good image. 0 – 100% image error were considered successful and everything above unsuccessful. If all images had an error of less than 120% the number is black, otherwise red.

Out of curiosity, time-difference reconstructions were also made on a coarse 100 thousand element head mesh with homogeneous conductivity of 0.3 S m^{-1} (last row of table 3.2). Surprisingly, most of these reconstructions were still good enough to identify the stroke, suggesting that TD reconstructions are more stable to geometrical errors than previously thought (Bagshaw et al., 2003).

3.4.2 Analysis of the Voltage Errors

How large were voltage changes caused by the perturbations?

The average 2-norm of the voltage change caused by haemorrhages was 301 μ V (1.3% proportional) and for ischaemias 27 μ V (0.1% proportional). These values were obtained by averaging the changes of the respective stroke types for all five positions. The position of the stroke had only a small influence on the signal. The largest changes were caused by posterior strokes, and were approximately 1.5 times larger than the smallest changes, which were caused by central strokes.

How large were voltage errors introduced by wrong meshes?

Baseline voltage differences were compared between all four head meshes, and the average for these six comparisons was computed. The average voltage differences were large with 2-norms between 16 mV–35 mV (59-100% proportional).

Time-difference voltage changes were compared between all four head meshes for perturbations simulated in posterior position. The average TD voltage differences were 146 μ V (0.63% proportional) for haemorrhage and 13 μ V (0.07% proportional) for ischaemia.

How did this compare to the introduced system noise?

The average 2-norm of the added proportional and additive noise was 28 μ V (5.6% proportional) for low noise, 59 μ V (11.3% proportional) for medium noise and 148 μ V (28.3% proportional) for the highest used noise level. These values were computed based on the noise added to the baseline. Consequently, time-difference data contained a combination of the noise on the baseline and equivalent noise on the perturbation voltages.

Comparison of the voltage differences

In time-difference imaging, the geometrical errors were in the same range as the signal. On the other hand, the proportional and additive noise added to the simulated voltages was 2-10 times smaller than the haemorrhagic signal. This means that reconstructions of haemorrhagic strokes had a ten times better signal to system noise ratio than reconstructions of ischaemic strokes, while the signal to mesh difference ratios were comparable (figure 3.6).

Successful reconstructions of ischaemic strokes reduced strongly between medium and high system noise. The 2-norm of the signal was two times smaller than the 2-norm of the medium noise, suggesting this ratio as a good rule of thumb for estimating feasibility of head EIT applications. The observed strong dependence of the image quality on system noise masked the positive effect of using the correct head mesh.

3.5 Discussion

3.5.1 Mesh Creation

The created head meshes were anatomically more accurate than previous EIT head models, due to the use of both CT and MRI scans, instead of just one of the two (Tizzard et al., 2005; Vonach et al., 2012). The quality of the meshes was further improved by using the established open source CGAL mesh generator directly on the segmentation, thereby skipping previously required

3. Head Models from CT and MRI Scans

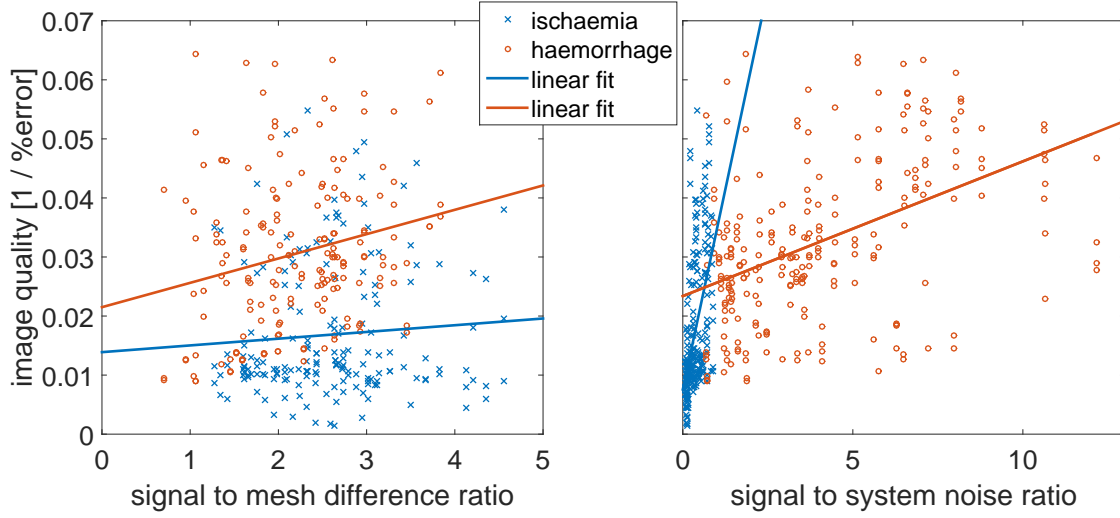


Figure 3.6: Image quality dependence on mesh differences and system noise - The plot on the left shows the effects of using a generic head mesh for reconstructions. The signal to mesh difference ratio was computed as the 2-norm of the voltage changes caused by the simulated stroke, divided by the 2-norm of the difference in time-difference voltage changes on the mesh used for voltage simulation and the mesh used for reconstruction. The plot on the right shows the image quality with respect to signal to system noise ratio for both types of stroke. For each stroke type there are 240 data points, corresponding to four meshes used for simulation, four meshes used for reconstruction, five stroke positions and three noise levels. However, the signal to mesh difference ratio was infinite if the same mesh was used for simulation and reconstruction, which is why only 180 data points for each stroke type are visible on the left plot.

separate steps of enclosed 3D surface generation and meshing (Vonach et al., 2012). Problem specific element sizing fields were defined, in order to reduce element size near the electrodes, where electric potential gradients were large. Once the manual CT and MRI segmentations were performed, meshes with variable element sizes could be created by simply varying the CGAL settings in a text file.

3.5.2 Impact of Mesh Differences

Voltage differences between meshes, voltage changes caused by the simulated perturbations and the degree of added system noise were quantified by the 2-norm. The image quality of the reconstructions was assessed objectively by predefined metrics and a limit of detectability was set at 100% image errors. From these objective measures, the following answers were found to the questions posed at the beginning:

1. **How large are voltage and image errors introduced by using a generic head mesh, as opposed to a patient specific one?** Time-difference voltage errors between head meshes were of the same order than voltage changes caused by the perturbations. The image quality was on average 25% better, when the reconstructions were made on the correct mesh. When the added system noise was more than twice as high as the signal, then the stroke detection rate decreased significantly.

2. **Is a patient specific head mesh required for the detection of localised conductivity changes?** While using the correct mesh for reconstructions improved the overall image quality, stroke detection rates were not improved. In three cases, only the correct mesh lead to a correct stroke identification, while there were five cases, where the reconstruction on the correct mesh failed and the average image errors on the other meshes were below 100%.

Finally, it was observed that even a coarse, homogeneous reconstruction mesh resulted in an acceptable image quality in most cases.

3.5.3 Conclusion

For monitoring, it was found not to be essential to use patient specific head meshes for image reconstructions. The results of this study, however, suggest that the noise on the measured voltages has to be kept very low to enable reliable detection of small conductivity changes in the head, such as ischaemia, epileptic seizures or even fast neural activity. Ways in achieving this may include long averaging, correction for small electrode movements (chapter 5) and reduction of system errors (Fitzgerald et al., 2002).

An EIT head imaging application in which it is very unlikely to have access to patient specific head meshes, is the differentiation of acute ischaemic and haemorrhagic strokes (Holder and Tidswell, 2004). Since a baseline measurement is not available in this case, time-difference (TD) reconstructions are not available and multi-frequency (MF) methods have to be used (Malone et al., 2014). Previous results and preliminary work suggest, that MF reconstructions are much more sensitive to geometrical errors than TD (chapter 6), but more analysis needs to be done to give a definitive answer, whether patient-specific meshes are required in MF imaging.

To find the optimal finite element size for head meshes, a convergence analysis (Aristovich et al., 2014) or local error estimates (Sawicki and Okoniewski, 2010) should be performed. Indications from unpublished results suggest that the forward solutions converge at human head mesh sizes above 10 million elements, but a detailed study was not possible at time of writing due to memory requirements of the CGAL mesher exceeding 24 GB at around 15 million elements. Preliminary results suggest, that for mesh sizes between 3 and 5 million elements, the discretisation errors are approximately half as big as the smallest noise level added to the simulated voltages in this chapter. To solve the forward problem on meshes with a couple of million elements, new efficient software is needed to reduce memory requirements and computation time.

Chapter 4

A Fast Parallel Solver for the Forward Problem in EIT (PEITS)

4.1 Overview

4.1.1 Introduction

In the previous chapter, the requirement for patient specific head meshes for time-difference EIT was studied. To obtain accurate forward solutions, particularly on the human head with its complicated structure with strongly varying conductivity, fine meshes had to be used. It was shown, that time-difference (TD) image reconstructions do not necessarily require an accurate head model, because most modelling errors are cancelled in difference measurements. In TD applications with high precision requirements, it is however desirable to achieve the best possible image, and therefore use the best mesh available for the Jacobian matrix computation. One of these applications is the imaging of fast neural activity in the rat cortex (Aristovich et al., 2016).

Where TD imaging is not possible due to the lack of a baseline measurement, multi-frequency (MF) or absolute image reconstructions are required. MF reconstruction methods have stronger requirements on the mesh accuracy than TD, as observed in chapter 6. Since absolute imaging directly images the difference between simulated and measured voltages, the mesh precision requirements are the highest in this modality (chapter 5). MF and absolute reconstruction algorithms are usually iterative, and therefore require multiple forward simulations and Jacobian matrix computations.

The main limitations in using finer meshes are computation time and memory requirements. Consequently, simulation studies and accurate image reconstructions in EIT depend on fast and memory efficient forward solvers.

4. A Fast Parallel Forward Solver (PEITS)

4.1.2 Background

A detailed summary of the forward solvers for EIT is given in section 2.2.9. Most research groups in EIT currently use the Matlab (The MathWorks, US) based Electrical Impedance Tomography and Diffuse Optical Tomography Reconstruction Software EIDORS (Adler and Lionheart, 2006). Horesh et al. (2006) adapted EIDORS with different preconditioners and more efficient routines, resulting in a version called SuperSolver which is still used in our group at UCL. For large meshes, however, Matlab suffers from a lack of efficient parallel programming possibilities, which makes the computation of forward solutions a lengthy task. Additionally, the Matlab implementations are relatively excessive in terms of memory usage, thereby physically limiting the size of the meshes that can be used.

Previous efforts to speed up the forward simulations include using the parallel direct solver library PARDISO (Borsic et al., 2010), calculating the Jacobian matrix on a GPU (Borsic et al., 2012) and using AMG preconditioning (Soleimani et al., 2005). These all addressed specific steps of the forward modelling, and no complete parallel software package for system assembly, forward simulations and Jacobian computation has so far been developed.

4.1.3 Purpose

The aim of the work presented in this chapter, was to build a comprehensive software package for the fast, parallel and memory-efficient computation of forward solutions on parallel machines and clusters, in order to reduce the computational time of the image reconstruction process in EIT. The solver, abbreviated hereinafter PEITS (Parallel EIT Solver), is flexible in terms of preconditioning options and inversion algorithms, and it includes the Jacobian matrix computation.

After a short review of the used numerical libraries and fundamental mathematical problem formulation, the implementation of PEITS is described in chronological steps of computation. The performance of each step is analysed separately and different preconditioners and inversion methods are compared. The overall performance for practical applications is assessed and compared to EIDORS. Finally, first and second order shape functions are compared. A user guide to the solver is found in Appendix A.

4.2 Technical Background

4.2.1 DUNE and PETSc

The solver was implemented in the Distributed and Unified Numerics Environment (DUNE), which is a grid-based C++ toolbox for solving partial differential equations. DUNE includes the discretisation module DUNE-FEM, which allows implementations of finite element solvers

for parallel computers. It provides functions to implement local grid adaptivity, dynamic load balancing and higher order discretisation schemes (Dedner et al., 2010). Apart from native implementations of conjugate gradients solvers it also provides interfaces to the solvers and preconditioners of the DUNE-ISTL module, UFMPACK (Davis, 2004) for unsymmetric problems and PETSc (Balay et al., 2015), which has an extensive collection of solvers and preconditioners. DUNE-FEM supports two types of parallelism, MPI and pthread. DUNE is licensed under the GNU General Public Licence version 2.0 and thus free to use for everyone.

The decision to use DUNE-FEM was made, because it is a fast, template-based and thus versatile C++ library that allowed the implementation of the complete electrode model (CEM), which has an uncommon weak formulation and is thus not easily implementable in most finite element libraries. Furthermore, DUNE-FEM provides an interface to many different preconditioner and solver libraries and supports tetrahedral elements. The module is still in development and was flexibly adjusted to the requirements of this project.

PETSc is a C++ library providing data structures and routines for the solution of systems obtained from discretisation of partial differential equations. Its focus is on providing scalable parallel routines supporting MPI, pthread and GPU parallelism. DUNE interfaces to PETSc in order to use the linear system solvers included in PETSc:

- *Krylov Methods* - PETSc has more than a dozen Krylov solvers available, such as CG, Bi-CG, Bi-CG-stab and GMRES.
- *Preconditioners* - Many preconditioners are provided, some from external packages. For non-parallel applications there are implementations for incomplete LU decomposition and successive over-relaxation. Supporting parallelism, the most interesting preconditioners are the 'classic' algebraic multigrid from Hypre (Falgout, 2015) and the smoothed aggregation AMG from Trilinos (Gee et al., 2006).
- *Direct Solvers* - Some direct solvers are available, e.g. MUMPS (Amestoy et al., 2006).

4.2.2 Complete Electrode Model and Jacobian Matrix

The derivation of the weak formulation of the CEM, which was used in PEITS, is given in section 2.2.4. It does not directly solve for $(u, U) \in (H^1(\Omega) \oplus \mathbb{R}^M)/\mathbb{R}$, but instead first solves for the internal potential u and then computes the electrode potentials U_m in a second step

$$\int_{\Omega} \sigma \nabla v \nabla u + \sum_{m=1}^M \frac{1}{z_m} \int_{\Gamma_m} v u - \sum_{m=1}^M \frac{1}{z_m |\Gamma_m|} \int_{\Gamma_m} v \int_{\Gamma_m} u = \sum_{m=1}^M \frac{1}{|\Gamma_m|} \int_{\Gamma_m} v I_m. \quad (4.1)$$

One important observation in this weak formulation is the uncommon third term of type $\int_{\Gamma_m} u \int_{\Gamma_m} v$. To facilitate the parallel assembly of this term and to reduce communication

4. A Fast Parallel Forward Solver (PEITS)

between processes, it has to be ensured that each electrode is not split onto different partitions (see Section 4.3.1 for the implementation).

Different grounding conditions can be applied. One option would be to set the average surface potential to zero (2.10c), but for ease of implementation it was decided to set one surface node to 0 V by applying a Dirichlet boundary condition. It is important to note, that the traditional CEM grounding condition $\sum_{m=1}^M U_m = 0$ is not easily implemented here, since the electrode potentials are only computed in a later step. Once the forward solutions are computed, most EIT inversion algorithms require the *Jacobian matrix* which translates a change in conductivity to a change in measured voltages by linearisation at the simulated conductivity distribution. The lead (or adjoint) fields method (section 2.2.7) was used to compute the Jacobian

$$\delta V_{da} = - \int_{\Omega} \delta \sigma \nabla u(I^d) \cdot \nabla u(I^a) d\Omega, \quad (4.2)$$

where $u(I^d) \in H^1(\Omega)$ is the electric potential emerging when the drive current I^d is applied to the electrodes and $u(I^a) \in H^1(\Omega)$ the electric potential when a unit current is applied to the two measurement electrodes. $\delta V_{da} \in \mathbb{R}$ is then the linearly approximated voltage change between the two measurement electrodes when the conductivity changes by $\delta \sigma \in L^\infty(\Omega)$.

4.2.3 Methods

Unless otherwise noted, all run times were computed on a head mesh with different conductivities for the scalp, skull, cerebro-spinal fluid (CSF), white matter, grey matter and superior sagittal sinus (figure 4.1). The meshes were created from a CT and MRI scan of the same patients head as described in chapter 3. The assembly of the system matrix will be slower, the more elements are part of an electrode. To ensure that the results presented here can be compared to each other, the ratio of electrode elements to other elements was kept fixed by having a constant element size throughout the mesh. For real applications it is advantageous to refine the mesh around the electrodes and use larger elements towards the center of the head.

To measure the parallel scalability of the code, the efficiency was calculated as follows for p parallel processes

$$\text{strong scaling efficiency}(p) = \frac{\text{runtime}(1)}{\text{runtime}(p) \cdot p}. \quad (4.3)$$

This efficiency is a value for the strong scaling. For the weak scaling it should be shown how much more elements can be computed in the same time by using more processors. The following definition for the efficiency of the weak scaling was used, where x is a fixed number of elements and $\text{runtime}(px)$ means the time it takes to compute $p \cdot x$ elements on p processors

$$\text{weak scaling efficiency}(p) = \frac{\text{runtime}(x)}{\text{runtime}(p \cdot x)}. \quad (4.4)$$




Figure 4.1: Layered cut through a 5 million element head mesh - The mesh was created with CGAL from a segmentation of a CT and an MRI scan of the same person and contains the following tissues: scalp, skull, cerebro-spinal fluid, grey and white matter, and also parts of the superior sagittal sinus and air cavities, which are not visible in this image.

All run times were taken on a cluster with 5 nodes. Each node had two 6 core 2.40 GHz Intel Xeon processors with 12 MB cache and a total of 192 GB of memory. The nodes were connected by a dedicated 1 GB Ethernet switch. PETSc version 3.4.2 and Zoltan version 3.6 were used.

4.2.4 Overall Structure of PEITS

The code is structured in different files that contain classes, structs and functions for specific tasks. The main file is `dune_peits.cc`, which performs the following important steps in this order:

- Loading the mesh and partitioning it. If the mesh was already partitioned before, those partitions are loaded by the parallel processes directly.
- The electrode positions are loaded into a struct. This struct has query functions that evaluate if a given element belongs to an electrode.
- The current protocol is read from the specified file. Upon reading the protocol, it is

4. A Fast Parallel Forward Solver (PEITS)

disassembled into unique injections. The solution for each unique current injection is computed just once. This reduces the number of required forward solutions for a standard EIT protocol with around 1000 lines to around 60.

- The system matrix is assembled. The function which computes the system matrix entries is located in the file `elliptic.hh`.
- In a for-loop the following steps are performed for each unique current injection:
 - The right-hand side of the weak formulation is assembled using a function in `rhs.hh`.
 - The CG solver computes the resulting electric potentials and the result is stored in a vector.
 - If specified in the parameter file, the first solution is written to a VTK file for visual inspection.
- In a second for-loop the following steps are performed for each line in the current protocol:
 - Trace back which forward solutions correspond to the drive current and measurement current of this protocol line.
 - Compute the measured voltage and save to a binary file.
 - Compute the row of the Jacobian matrix using the forward solutions for the drive and measurement current and save it to a binary file.

4.3 Implementation of the Functional Steps

4.3.1 Parallel Substructuring

When using a finite element mesh for the first time, it needs to be partitioned evenly onto the used number of parallel processes. DUNE-FEM has a partitioning tool available, but this tool does not enable the user to guide the load-balancing by fixing regions to a specific process. PEITS requires the electrodes to be on only one process each, since this facilitates the correct system matrix assembly and minimises communication between processes. A well documented and established library supporting user defined load-balancing is Zoltan (Devine et al., 2006). Zoltan is a parallel C, C++ and Fortran 90 library with a simple object-based interface that is easily adapted to many applications. The Zoltan tool employed here is called hypergraph partitioning. A hypergraph interpretation of a finite element mesh has the following appearance: each element is a vertex of the hypergraph and the element together with all its neighbours forms one hyperedge. Thus, for a mesh with N elements, the hypergraph has N vertices and N hyperedges, which correspond to the communication requirements of the parallel program. The

parallel hypergraph partitioning (PHG) tool in Zoltan allows the user to assign different weights to hyperedges and to fix selected vertices to one section.

When a mesh is loaded in PEITS for the first time it is initially partitioned by the load-balancing function of DUNE-FEM. This first step is very time consuming, because it is done serially and can not be sped up by employing more processors. The resulting parts are then made accessible to Zoltan by translating them into the hypergraph format Zoltan requires. Defining the electrode areas with the query functions `ZOLTAN_NUM_FIXED_OBJ_FN` and `ZOLTAN_FIXED_OBJ_LIST_FN`, Zoltan's

	partitioning		loading	
	2m	15m	2m	15m
12 processes	214	2137	5.7	52
24 processes	250	2405	2.89	26.8
48 processes	241	2400	1.45	12.8

Table 4.1: Total time required for partitioning and loading the meshes - All times are in seconds. The times for the partitioning include the loading of the mesh, the initial load-balancing, the Zoltan repartitioning and the second load-balancing in DUNE-FEM. Since the parallel substructuring has to be done only once for each mesh, the performance of this operation is not critical.

PHG partitioner is applied to the mesh in order to optimise the load-balancing while ensuring that each electrode is assigned to one process only. Zoltan will return a list of elements that need to be moved from one part to another part. This list is subsequently applied using the load-balancing function of DUNE-FEM (figure 4.2). The number of elements which need to be transferred between processes varies, but tends to increase as the number of processes increases. Because these three steps are time consuming (table 4.1) and need to be done only once for each finite element mesh and number of parallel processes, the resulting mesh parts are written into separate DUNE grid files (DGF). These DGF can then be loaded in parallel for each subsequent use of the same mesh. The loading of these partitions took less than a minute for a 15 million element mesh and scales almost linearly, since each processor loads its mesh partition separately (table 4.1).

4.3.2 Assembly of the System Matrix

The assembly of the system matrix is done in two mesh iterations. The first iteration stores all elements that belong to an electrode in a two-dimensional vector `electrodeElements` of length $M \in \mathbb{N}$ equal to the number of electrodes with `electrodeElements[m]` containing all elements that constitute electrode m . Furthermore, the overall electrode areas are computed and stored in a vector. Storing the electrode elements is essential to reduce the time for the matrix assembly and later on for the computation of the electrode potentials.

During the second iteration, the Laplacian term $\int_{\Omega_n} \sigma_n \nabla v \cdot \nabla v \, d\Omega_n$ is added to the system matrix for each element Ω_n for $n = 1, \dots, N$ where $N \in \mathbb{N}$ is the number of elements in the mesh. If element Ω_n constitutes part of an electrode Γ_m , a nested loop iterates over all elements

4. A Fast Parallel Forward Solver (PEITS)




Figure 4.2: Partitions before and after Zoltan load-balancing - This is a relatively small head mesh with two million elements which was partitioned into four sections. Some electrodes were split onto different processors by the load-balancing of DUNE-FEM. The Zoltan partitioner minimises the number of elements that have to be moved from one process to another, while optimising the partitions and fixing the electrode areas to one process.

of that electrode `electrodeElements[m]`, writing each elements contribution to the third term of weak formulation (4.1) $\frac{1}{z_m |\Gamma_m|} \int_{\Gamma_m} v d\Gamma_m \int_{\Gamma_m} v d\Gamma_m$ into the system matrix. Additionally, for each element under an electrode the second term of the weak formulation $\frac{1}{z_m} \int_{\Gamma_m} v^2 d\Gamma_m$ is assembled.

In a last step, one surface node of the mesh is assigned a Dirichlet boundary condition to make the problem uniquely solvable. The coordinates of the Dirichlet nodes have to be given to the solver in a mesh specific parameter file. The solver then evaluates for each node if the coordinates match the given ones and applies the Dirichlet boundary conditions if they do.

Since the numerical solvers available in the PETSc library (Balay et al., 2015) are used, the system matrix is directly assembled in the native PETSc sparse row matrix format `MATMPIAIJ`. The difficulty using this format lies in the correct memory pre-allocation. If PETSc has to re-allocate more memory during the matrix assembly, the performance can decrease by more than a factor of 50. Due to the DUNE-FEM implementation of the boundaries between parallel partitions, only one side ever knows that there are neighbouring elements. Thus, a perfect pre-allocation can only be achieved by a process in charge of all its interfaces. All other processes underestimate the number of entries in the rows corresponding to communication with elements belonging to other processes. Thus, it is not easily possible to perfectly pre-allocate the memory for the PETSc sparse row matrix structure in our application. Instead, a first estimate of the maximum number of non-zeros per row is pre-allocated for each mesh and the solver is subsequently run with the option `-info` which outputs the precise maximal number of non-zeros per row per process and the required number of `mallocs` required. Based on this output, the pre-allocation can be further optimised such that no additional `malloc` is required during subsequent solves.

4. A Fast Parallel Forward Solver (PEITS)

For realistic head meshes with sizes of up to 15 million elements, 100 diagonal and 40 off-diagonal entries per row were allocated with the function `MatMPIAIJSetPreallocation(mat, 100, PETSC_NULL, 40, PETSC_NULL)` in the corresponding file in the DUNE-FEM library `\dune\fem\misc\petsc\petsccommon.hh`. This approach allocates far more entries than are actually used, but the performance of the matrix assembly is not significantly decreased.

The parallel efficiency of the matrix assembly drops down to around 0.4 for small meshes and 0.5 for large meshes (figure 4.3a). Even though the efficiency drops, the absolute times to assemble the system matrix still improve on 60 parallel processes (table 4.2). The weak scaling (figure 4.3b) indicates that a load of around 0.5 million elements per processor is optimal.

	2m	15m
1 process	27.89	231.41
2 processes	15.71	126.07
5 processes	6.56	54.04
10 processes	3.74	30.19
20 processes	2.23	16.93
40 processes	1.34	10.27
60 processes	1.11	7.3

Table 4.2: Total time required to assemble the system matrix - All times are in seconds.




Figure 4.3: Strong and weak scaling of the system matrix assembly - (a) As expected, the efficiency of the matrix assembly drops when more processors are used, since the ratio of computation to communication is getting smaller. On the larger mesh this effect is smaller than on the coarse mesh. (b) For half a million elements per process the ratio of communication over computation scaled best.

4.3.3 Preconditioning

Multigrid methods are known to be very efficient solvers for elliptic boundary value problems, such as the Laplace problem solved in EIT. The underlying principle of multigrid methods is to use several layers of coarseness to guide information exchange rather than having elements exchange information only locally to their direct neighbours. Two general approaches to multigrid methods are geometric multigrid (GMG) and algebraic multigrid (AMG). GMG relies

4. A Fast Parallel Forward Solver (PEITS)

on coarser finite element meshes with the same geometry, which are not easily created in head EIT because the complicated geometry cannot simply be coarsened. AMG on the other hand does not require any geometric information and constructs the coarser levels directly from the system matrix, which makes it very adaptable to different problems. With a reduced tolerance, multigrid methods can efficiently be used as preconditioners for iterative solvers.

Since the problem (4.1) is a Laplace problem with a compact perturbation, AMG is a very efficient preconditioner in EIT. Through PETSc PEITS has two very good AMG implementations available, BoomerAMG from Hypre (Falgout, 2015) and ML from Trilinos (Gee et al., 2006). The performance of the two AMG implementations was compared on two different mesh sizes using the default settings of the respective preconditioner. The results in table 4.3 show, that ML was faster for our application.

	2m elements				15m elements			
	assembly		solving		assembly		solving	
	ML	BAMG	ML	BAMG	ML	BAMG	ML	BAMG
5 processes	1.1	4.4	3.5	4.4	13.8	55.0	79.0	64.7
10 processes	0.7	3.3	2.3	3.1	5.8	34.1	45.8	43.5
20 processes	0.62	5.54	0.60	2.34	3.3	36.7	22.3	24.7
40 processes	1.02	8.3	0.58	2.1	3.7	41.5	13.0	14.4
60 processes	4.8	10.5	1.3	2.53	4.5	39.4	9.6	10.5

Table 4.3: Times for the assembly of the preconditioner and subsequent CG solutions - All times are in seconds. The assembly of the ML preconditioner was significantly faster than for BoomerAMG (BAMG). Also, the performance of both preconditioners dropped with more than 40 parallel processes on the smaller mesh. This was most likely due to the increasing communication requirements between processors.

The efficiency of the setup of the ML preconditioner was tested on two different sized finite element meshes. As expected, a better parallel performance was observed on the large mesh, due to the larger ratio of computation over communication (figure 4.4a). The weak scaling (figure 4.4b) indicates that a load of approximately 0.5 – 1 million elements per processor is optimal.

4.3.4 Solver

The optimal multigrid preconditioning can be used with all Krylov subspace solvers available in PETSc, which are Conjugate Gradients (CG), Generalised Minimal Residual (GMRES) and two CG algorithms for non-symmetric problems BiCG and BiCGstab. Since the system matrix is positive symmetric, CG is generally used in PEITS. The stopping criterion for the CG solver was set to a relative residuum of 10^{-12} . On the two million element mesh,

4. A Fast Parallel Forward Solver (PEITS)




Figure 4.4: Strong and weak scaling of the assembly of the AMG preconditioner ML - (a) Since the computational cost was larger on the fine mesh, the increasing communication volume between parallel processes had a smaller influence on the overall efficiency of the preconditioner assembly when compared to the coarse mesh. (b) With an average load of half a million elements, the ML preconditioner assembly scaled best.

large and reproducible fluctuations in the efficiency were observed (figure 4.5a), which most likely emerged from different communication requirements and cache usage for the different partitions. This effect was less visible on the large mesh since the computation load was proportionally larger than the access to the memory and communication between processors. The absolute run times on 60 processors were 1.3 seconds on two million elements and 9.6 seconds on 15 million (table 4.4). The weak efficiency (figure 4.5b) indicates that the optimal load per processor was around 0.5 million elements. The weak scaling was not as good as the strong scaling, because the CG solver generally had a slower convergence rate on the larger meshes (table 4.5). The required iterations on larger meshes correlate very well with the decrease in efficiency of the weak scaling (figure 4.5b).

	2m	15m
1 process	18.5	399
2 processes	8.3	169
5 processes	3.6	77
10 processes	2.25	45.8
20 processes	0.6	22.3
40 processes	0.58	13.0
60 processes	1.3	9.6

Table 4.4: Time for one forward solution for CG with ML preconditioning - All times are in seconds.

	335k	656k	1m	2m	5m	8m	15m
Iterations	44	41	34	35	46	43	63

Table 4.5: Iterations of the CG solver on different mesh sizes - The number of iterations was not dependent on how many parallel processes were used.

A possible alternative to using an AMG preconditioned CG algorithm is to set up a direct solver. A direct solver takes very long to assemble, but reduces subsequent solutions to mere

4. A Fast Parallel Forward Solver (PEITS)




Figure 4.5: Strong and weak scaling of the CG solver with ML preconditioning - (a) On the small mesh the run times show reproducible large fluctuations, which were most likely caused by different cache efficiency for different partitions. On the larger mesh, where the computational load is larger these cache effects are not visible. (b) An average load of half a million elements per processor lead to the optimal ratio of computation over communication. However, there was a significant drop in efficiency on many parallel processes. This can be explained by different convergence rates of the solver on the different meshes, as can be seen in table 4.5.

forward and back substitutions. Thus, if many forward solutions are required (i.e. many electrodes are used) then a direct solver might be faster. PETSc interfaces to the MUMPS direct solver (Amestoy et al., 2006), using it as a preconditioner for the CG solver. This reduces each solution to one or two iterations. The assembly took very long on large meshes and the CG solver with MUMPS preconditioning did not scale well on many processes (table 4.6). The weak scaling of MUMPS was much worse than that of the MG preconditioners, meaning that for large problems the number of required forward solutions for the direct solver to be faster increased (table 4.7). Furthermore, the strong scaling was worse than that of ML as well, such that for many parallel processes AMG was always the better choice (as indicated with the minus symbol in the last row in table 4.7).

	2m		15m	
	assembly	solve	assembly	solve
1 process	1388.5	3.5	91375	48.9
10 processes	258.0	0.94	28969	15.9
50 processes	128.9	1.0	9469	16.3

Table 4.6: Performance of the MUMPS direct solver - All times are in seconds. Especially the assembly times on the larger mesh illustrate well why the direct solver is not a valuable alternative for our application. Also, the solving times did not scale very well on many parallel processes.

	2m	15m
1 process	93	261
10 processes	197	969
50 processes	-	-

Table 4.7: Number of forward solutions for MUMPS to be faster than ML - The number of forward solutions to be computed, in order for MUMPS to be faster. On 50 processes ML was always faster than MUMPS.

It was concluded, that for this application a direct solver was only worth considering on relatively small problems with many electrodes, leading to more than one hundred independent current injections. Most applications will be solved faster by using ML as a preconditioner.

4.3.5 Jacobian Calculation

The Jacobian matrix is computed based on the adjoint field method (4.2). In PEITS the computations are handled by a struct `JacobianRowCalculator`, which computes the local stiffness matrices of all elements in the constructor and stores them. When PEITS then iterates over all lines of the measurement protocol, the member function `getJacobianRow` is called with the voltage distributions for both drive and measurement current as arguments. `getJacobianRow` then iterates over all elements and computes two matrix-vector products in each step to obtain the local entry of the row of the Jacobian matrix. Since this process requires no communication between processors at all, a very good parallel efficiency can be expected. We observed that the efficiency was reliably larger than 1, going up to more than 2 in one case (figure 4.6a). This was most likely due to a more efficient use of cache memory, since for the larger mesh the efficiency kept improving for more processors while it remains around 1.4 for the smaller mesh, indicating that already all local stiffness matrices were stored in cache. The weak scaling of the Jacobian matrix computation (figure 4.6b) indicates that the optimal load per processor was around half a million elements.

	2m	15m
1 process	1.44	20.3
2 processes	0.57	8.6
5 processes	0.22	3.8
10 processes	0.11	1.85
20 processes	0.057	0.84
40 processes	0.026	0.34
60 processes	0.017	0.16

Table 4.8: Time taken for the computation of one row of the Jacobian matrix - All times are in seconds.

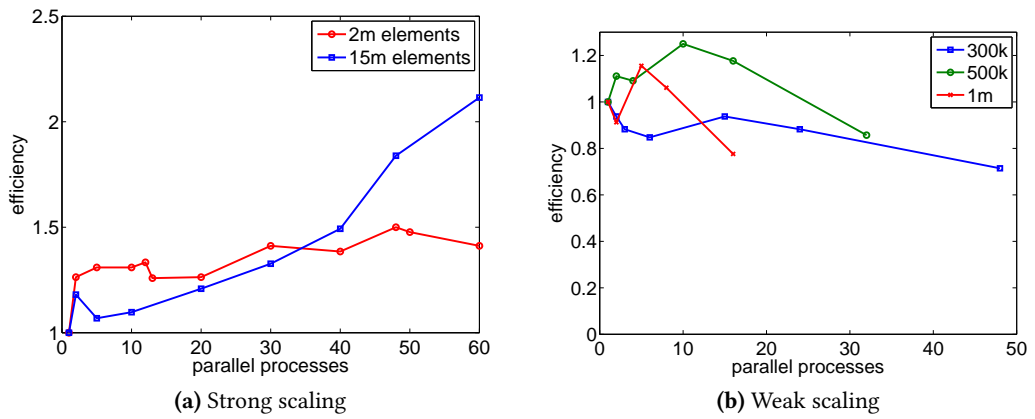


Figure 4.6: Strong and weak scaling of the computation of the Jacobian matrix - (a) The efficiency of the computation of a Jacobian row increased, as the number of processors increased. This was most likely due to a more efficient usage of the cache. On the larger mesh, where the computational load is larger, these cache effects were more visible on many processors. (b) The weak efficiency decreased on large meshes, most likely due to the memory access. When the processes were not evenly distributed over all cluster nodes the weak efficiency dropped earlier, supporting this claim.

4. A Fast Parallel Forward Solver (PEITS)

4.3.6 Verification of Correct Performance

The correctness of the results of the forward problem was verified in two ways. First, simulations were done on meshes of cubes of varying size, number of elements, conductivities and contact impedances of the two electrodes, which were placed on opposite sides of the cube. The results were then compared to the analytical solution and were precise up to computer precision. To verify that PEITS gave correct results on more complicated shapes, the results of simulations on a head shaped mesh were compared both to the version of EIDORS currently used in



Figure 4.7: Simulated electric potential in a 2 million element mesh of the head - A current is applied from the front of the head to the back of the head and the computed voltage distribution on a slice through the head mesh is shown. As expected, the electric potential drops sharply at the skull due to its low conductivity.

our group and to real measurements in a saline filled tank. They matched the computed results by computer precision and the experimental results closely. Figure 4.7 shows the resulting simulated electric potential distribution when a current was applied from the front of the head to the back of the head. It is visible how the potential drops at the highly resistive skull.

4.4 Performance

4.4.1 Total Run Times with 1st Order Elements

A common forward problem in EIT with pairwise current injection requires around 60 forward solutions for the unique drive and measurement current injections and around 1000 current protocol steps, i.e. 1000 lines in the Jacobian matrix. Therefore, the total runtime for PEITS can be estimated by adding up the times for the components which have to be done once per execution (loading partitions, finding electrode elements and areas, matrix assembly, preconditioning, computing local stiffness matrices used for the Jacobian calculations) with 60 times the time it took for one forward solution and 1000 times the time for one Jacobian row computation. Based on the run times shown in the previous sections, the total runtime for a common application of the solver was estimated (table 4.9). It was found that the overall efficiency scaled very well, in parts due to the high efficiency of the Jacobian matrix calculation (figure 4.8).

4. A Fast Parallel Forward Solver (PEITS)

	2m	15m
1 processes	3089.5	48644
2 processes	1335.6	22457
5 processes	498.2	6714.6
10 processes	254.3	3357.8
20 processes	117.1	1948.6
40 processes	78.3	1071.1
60 processes	86.7	842.4

Table 4.9: Total estimated run times for a protocol with 1000 lines - All times are in seconds and include the loading of the mesh partitions, system matrix assembly, setup of the preconditioner, 60 forward solutions and computing 1000 rows of the Jacobian matrix.

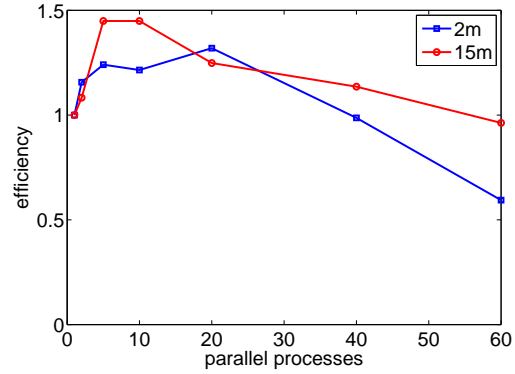


Figure 4.8: Estimated efficiency of the total run time of a realistic EIT protocol - Strong scaling based on the estimated run times shown in table 4.9.

4.4.2 Comparison to EIDORS

Since EIDORS is the most prevalent software for the computation of the forward solutions in EIT, the performance of PEITS was compared to EIDORS on Matlab. The new EIDORS version 3.7.1 has recently been released and was used to do the timings. For this section the performance of EIDORS on a 2 million element head mesh was compared with the performance of the PEITS on the same mesh in serial (table 4.10). To make the comparison valid, Matlab's multi-threading routines were disabled by setting `maxNumCompThreads(1)`. Matlab version R2013a was used.

The standard solver in EIDORS is Matlab's backslash operator, which took approximately 1936 seconds for the direct solver to be assembled and around 12.5 seconds for each unique current pattern solved with it. The mumps direct solver was faster for each solve (3.5 seconds) as well as for its assembly (1389 s). To compare iterative solvers, an incomplete LU decomposition was used as a preconditioner for a conjugate gradient solver in Matlab and the ML preconditioned CG solver in PEITS, since these combinations were the best for the respective solvers. While the Matlab routine `ilu` was very quick (0.8 s), each successive solve with `pcg` took 39 seconds. In DUNE, the assembly of the AMG preconditioner took 5.4 seconds and each solve 18.5 s.

	Matlab	PEITS
matrix assembly	128	27.8
preconditioner assembly	0.8	5.4
CG solve step	39	18.5
Jacobian row calculation	0.3	1.4

Table 4.10: Comparison of EIDORS/SuperSolver and PEITS - All times are in seconds for an execution on one processor.

The assembly of the system matrix took 128 seconds in EIDORS. This is difficult to compare to the assembly in PEITS, since EIDORS creates more data structures for later use (plotting, inverse,

4. A Fast Parallel Forward Solver (PEITS)

...) and also assigns the electrode areas and ground indices differently. PEITS is less flexible and focuses only on the forward problem, which is one of the reason why it was more than twice as fast for the matrix assembly (27.8 seconds).

The mesh was loaded much faster in Matlab (0.5 seconds as opposed to 61 seconds), because Matlab uses a compressed binary data format (.mat file) whereas PEITS is currently still using ASCII data files. It is planned to switch to binary data files in future.

The calculation of a single row of the Jacobian was very difficult to compare since EIDORS and PEITS use completely different approaches. While PEITS uses the adjoint field method, EIDORS applies the derivative form (Yorkey, 1990). Thus, the total time it takes to compute a Jacobian matrix with either 1, 7 or 259 lines was compared. In EIDORS it took 2433, 2442 and 2817 seconds, respectively, and in PEITS 1.44, 10.1 and 373 seconds. The matrix based approach of EIDORS is less dependent on the number of protocol lines. However, the memory usage is much higher and becomes inhibitive for bigger meshes and longer protocols. For 2 million elements and 258 protocol steps the memory usage of EIDORS during the Jacobian calculation was 150 GB. This is why the time to compute the Jacobian was also measured with a Matlab based adjoint field method which has been implemented in our group (used for instance in Horesh et al., 2006). This implementation turned out to be extremely quick and much more memory efficient than the EIDORS implementation. The computation times for 1, 7 and 259 lines were 1.47, 4.3 and 77.3 seconds, which compares to PEITS on approximately 4 parallel processes. It was significantly faster than PEITS, because a generic derivative matrix was constructed. For successive protocol lines it was then multiplied by the nodal potentials of the different forward solutions to get the gradients. PEITS on the other hand iterates over all elements and multiplies the local potentials with the gradients element wise for each protocol line.

To summarise, for a typical EIT application with 60 forward solutions and 1000 current protocol lines on a two million element mesh with an iterative solver EIDORS took around 6469 seconds and PEITS on one processor around 3090 seconds. In parallel, the performance of PEITS improves linearly with the number of processors, while the parallel efficiency of Matlab is lower. Furthermore, due to more efficient memory usage, PEITS allows for the use of much larger finite element meshes.

4.4.3 Comparison with 2nd Order Elements

In PEITS, it is very straight forward to switch to quadratic (or even cubic) shape functions. Only one environment variable needs to be changed, when compiling the solver. For smooth functions the use of higher order shape functions achieves the same precision of the solution obtained with 1st order elements, but with a much smaller mesh, thereby possibly reducing computation time. In order to get a rough estimate of the ratio of the required element sizes to get the same precision, cube shaped meshes were created with regular tetrahedral elements of

size $h = 1/10, 1/20, 1/30, \dots$ for a cube with dimensions $(1, 1, 1)$. Two electrodes were assigned to the central square with area $1/25$ on opposite sides of the cube and a current of $133 \mu\text{A}$ was applied. The cube had a uniform conductivity of 0.3 S m^{-1} . The convergence of the simulated voltage between these two electrodes was plotted against the element size for first and second order polynomials. For a similar accuracy, the number of elements could be reduced by a factor of 67 when second order elements were used (figure 4.9). The slope of both curves is the same, because the solution is in $H^1(\Omega)$ and the convergence rate of second order elements is therefore the same than that of linear elements.




Figure 4.9: Comparison of 1st and 2nd order shape functions - By using quadratic shape functions the same convergence could be achieved with around 67 times less elements.

The run times of the different parts of the solver were compared on the $24'576'000$ linear element mesh and on the $384'000$ element mesh with quadratic shape functions. The result on the small mesh with second order shape functions was 0.2679887 V and on the large linear mesh 0.2679873 V , meaning the second order mesh was a bit closer to the value the simulations converged to. All parts of the solver were much faster on the small mesh of second order elements (table 4.11). For this test example the PETSc pre-allocation was set to 2500 diagonal and 2500 off-diagonal entries per row to account for the electrodes on the 25 million element mesh.

4.4.4 Two Applications of PEITS

To illustrate the range of applications envisaged for the presented software, two works using PEITS are highlighted. The first was a simulation study evaluating the feasibility of detecting two different types of stroke in the human head using EIT measurements at different frequencies for the injection current (chapter 6). In the second application, PEITS was used to compute the

4. A Fast Parallel Forward Solver (PEITS)

	serial		20 processes	
	1st ; 24'576k	2nd ; 384k	1st ; 24'576k	2nd ; 384k
loading partitions	357.7	4.78	21.33	0.42
matrix assembly	372.4	39.8	112.02	6.76
AMG assembly	38.8	6.52	14.53	1.70
solve	64.2	14.6	10.61	3.27
Jacobian row	9.94	0.32	0.59	0.02

Table 4.11: Run time comparison 1st vs. 2nd order elements - All times are given in seconds. To achieve the same precision using 2nd order elements, the mesh could have around 64 times less elements. For this ratio, the performance of the solver was compared: the loading of the partitioned mesh, the assembly of the system matrix, the assembly of the ML preconditioner, one solve and one Jacobian line computation. It is apparent that second order shape functions reduced computation time in this test example.

forward solutions and the Jacobian matrix on a 7 million element mesh of the rat brain, in order to reconstruct neural activity from EIT measurements on the brain of a living rat (Aristovich et al., 2014).

The main difficulty in stroke type detection with multi-frequency EIT is that the finite element model never accurately matches the measurement setup. These modelling errors can introduce large artefacts into the reconstructed images. To distinguish the main sources of artefacts, boundary voltages were simulated in the presence of three different types of modelling errors on a fine mesh and images were reconstructed on a coarse, modelling error free mesh. Forward solutions were computed at 12 different frequencies for three different modelling errors with two different standard deviations each, and this for ischaemic and haemorrhagic stroke at two different locations in the head. This means that the study involved the computation of $288 \cdot N$ forward solutions, where N was the number of independent current injections (in our case 31). To compute that many forward solutions on a 5 million element mesh in Matlab would have taken $288 \cdot 30$ minutes. This estimate shows that simulation studies of this scale were previously not feasible. Using PEITS, the time for the forward simulations was reduced to $288 \cdot 1.7$ minutes on a workstation with two eight-core 2.4GHz Intel Xeon CPUs with 20MB cache each.

For EIT applications with high precision requirements, a very fine mesh is required for the forward computations. In the second application the aim was to image fast neural activity in the rat cortex using a planar electrode array, which was surgically applied directly to the brain surface. A convergence study on the required finite element size was performed as follows. Iteratively, the element size was reduced and ten meshes were created using the same settings. Then the differences in simulated voltages on these meshes were compared to the differences of the next coarser meshes. Once the variability between meshes of the same resolution was the same than the variability between different mesh sizes, the optimal mesh size was reached (figure 4.10). It was found that the required mesh size for this application was 8 million elements.

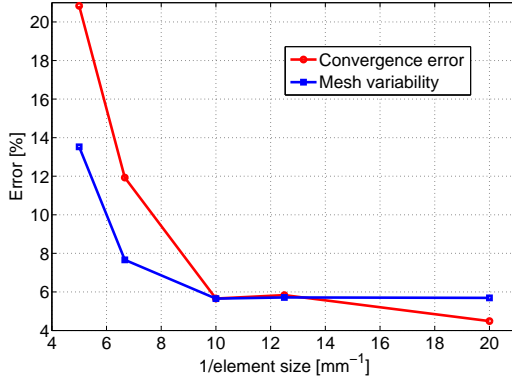


Figure 4.10: Convergence of simulated voltages by reducing element size - The maximal relative error in simulated boundary voltages on meshes with the same element size (variability) was compared to the same error between meshes with different element size (convergence). 0.1 mm was found to be the optimal element size for this application.




Figure 4.11: Three dimensional image of conductivity increase due to neural activity - This image shows the conductivity increase caused by opening ion channels during neural activity, which was induced by whisker stimulation of a rat. The activity patterns match the literature and correlate with intrinsic optics and local field potentials (Aristovich et al., 2016).

Due to the large number of electrodes, the measurement protocol (and thus the Jacobian matrix) was very long and to compute all forward solutions and the Jacobian matrix in Matlab would have taken around 16.5 hours. By using PEITS, this was reduced to just about an hour on the same workstation used for the first application. Therefore, it was possible to make reasonably quick informed decisions about the quality of acquired data and experimental paradigm changes. The iterative process of improving experimental procedures was sped up by more than a factor of 16 and for the first time high-quality images of fast neural activity in the rat cortex using EIT could be reconstructed with the methods outlined in (Aristovich et al., 2014) (figure 4.11).

4.5 Discussion

It was shown, that PEITS significantly reduces the computation time of forward solutions in EIT. To facilitate the use of the solver, Matlab functions are provided to write a mesh in DGF format and call the solver with different settings. This makes it possible to run PEITS from Matlab by calling just one function `run_forward_solver()`, which returns the Jacobian matrix and the measured voltages.

PEITS is actively used for forward simulations on 3-15 million element meshes. The computed results are used for the image reconstruction from experimental data. Furthermore, simulation studies are performed on models with changing settings such as electrode position, tissue conductivity, contact impedances and more, with the aim to characterise different sources of image artefacts caused by modelling errors.

Particularly in the use of adaptive mesh optimisation there is room for improvements, such

4. A Fast Parallel Forward Solver (PEITS)

as using second order elements in regions where the solution is in H^2 or refining the mesh around the electrodes based on local error estimates (Sawicki and Okoniewski, 2010; Johnson and MacLeod, 1994; Molinari et al., 2001). DUNE-FEM natively supports such local grid refinements.

In comparison to Matlab, PEITS is already faster in serial in every step except for the computation of the Jacobian matrix. This is where we see the largest potential speed improvements, especially since this is the step that is repeated around 1000 times for a typical EIT application. Several approaches to speed up the Jacobian computation are known. One would be to use a coarser mesh to compute the Jacobian on (as shown in Adler et al., 2008; Borsic et al., 2010), since most reconstruction algorithms do not (and do not need to) rely on a fine mesh. Another one would be to run the Jacobian matrix calculation on a GPU, where memory access is much faster. This has already been shown to improve the speed significantly (Borsic et al., 2012).

Chapter 5

Correction for Electrode Movement in Time-Difference EIT

5.1 Overview

5.1.1 Introduction

In chapter 3 a novel method for making accurate head models was presented. The created head models were used to evaluate the influence of geometrical errors on time-difference (TD) image reconstructions. With the fast forward solver described in chapter 4, a tool is now available to efficiently simulate other sources of image artefacts on accurate head models. While time-difference reconstructions are relatively stable in the presence of static modelling errors and instrumentation noise, changes between reference and data measurement strongly affect the image quality, due to the ill-posed inverse problem in EIT. Consequently, there is a lot of interest in finding ways to account for different sources of errors, such as varying electrode contact impedances, object boundary shape and electrode positions.

The consideration of changes in the boundary shape of the imaged object is of particular interest for thoracic imaging studies (e.g. lung ventilation, gastric emptying or heart cycles), because breathing changes the shape of the thorax significantly. Therefore, a simultaneous reconstruction of the conductivities and the boundary shape would improve the image quality (Boyle et al., 2012; Dardé et al., 2013). In TD head imaging, such temporal changes of the geometry do not occur. The focus of this chapter is on the long term monitoring of haemorrhagic transformations (Xu et al., 2010), an application where electrodes are kept on the head of a patient for several hours. In measurements over such an extended period of time, changes in electrode contact impedance and small electrode movements can be expected. Contact impedances were

5. Electrode Model Correction in TDEIT

found to have particularly severe effects on image quality for current injection schemes where voltages are measured on injecting electrodes (Kolehmainen et al., 1997), but less so when injecting current through electrode pairs while measuring on all other electrodes (four electrode measurements, i.e. McEwan et al. (2007)).

Electrode movements between reference and data measurement, however, are long known to have a strong impact on image quality (Barber and Brown, 1988).

5.1.2 Background

Initial approaches to correcting electrode positions include differential approximations of the electrode displacement Jacobian (Soleimani et al., 2006), direct methods based on the mesh geometry (Gómez-Laberge and Adler, 2008) and the approximation error approach (Nissinen et al., 2008). The most fundamental and straightforward way of computing a Jacobian matrix predicting voltage changes due to electrode boundary changes was shown by Dardé et al. (2012). They derived an explicit formula for the Fréchet derivative with respect to the electrode boundary, which can be used to compute the Jacobian matrix with respect to electrode size, position and shape. Both, Soleimani et al. (2006) and Dardé et al. (2012), applied electrode movement corrections to cylindrical 3D saline phantoms with positive results, but the performance in anatomically realistic 3D problems has not been studied. A more in-depth review of previous work on electrode movement correction is given in section 2.2.8.

5.1.3 Purpose

The purpose of the work in this chapter was to develop a method for correction of errors in electrode modelling in EIT of the human head. The questions to be answered were: 1) Which aspect of the electrodes - position, size, shape - most affects the boundary voltages and how accurately does the electrode boundary Jacobian matrix describe the influence of this aspect? 2) Does the simultaneous reconstruction of conductivities and electrode parameters remove image artefacts originating from imprecise electrode modelling in simulations and experiments? 3) Does the correction work in absolute imaging, where electrodes are modified iteratively?

These were addressed using a fast calculation of the Jacobian matrix with respect to electrode boundary perturbations, which was computed on a 4 million element mesh in less than a minute (described in Section 5.2). Characterising the different aspects of the electrodes, it was found that the electrode position was the dominating variable and that it was accurately modelled by the electrode boundary Jacobian (Section 5.3). The correction of electrode positions from time-difference data was therefore applied to an anatomically realistic 3D head model in simulations (Section 5.4) and experiments on a 3D printed head shaped saline tank with skull (Section 5.5). Finally, the correction for electrode movements was applied to absolute imaging (Section 5.6).

5.2 Electrode Boundary Jacobian Implementation

5.2.1 Mathematical Formulation

The detailed derivation and proof of the electrode boundary Jacobian (EBJ), which describes voltage changes caused by electrode boundary movement, is found in (Dardé et al., 2012), and a summary of the parts relevant for the implementation is given in section 2.2.8. Electrode boundary changes can be characterised by C^1 vector fields \mathbf{v} on the electrode boundaries ∂E_m and the measurement map including perturbed electrodes can then be considered as the operator

$$\mathbf{R} : (\mathbf{v}, I^d) \longrightarrow U^d(\mathbf{v}), \quad \mathcal{B}_b \times \mathbb{R}_\diamond^M \longrightarrow \mathbb{R}^M. \quad (5.1)$$

The Fréchet derivative with respect to the vector field \mathbf{v} fulfils

$$\sum_{m=1}^M (\mathbf{R}'(\mathbf{v}, I^d))_m I_m^a = - \sum_{m=1}^M \frac{1}{z_m} \int_{\partial E_m} (\mathbf{v}_\tau \cdot \mathbf{n}_{\partial E_m}) (U_m^d - u^d) (U_m^a - u^a) \, ds, \quad (5.2)$$

where \mathbf{v}_τ is the component of \mathbf{v} which is tangential to Γ , $\mathbf{n}_{\partial E}$ the outward normal of ∂E_m which is tangential to Γ and (u^d, U^d) the solution to the unperturbed CEM forward problem corresponding to the drive current $I^d \in \mathbb{R}_\diamond^M$. The adjoint field (u^a, U^a) is the forward solution to the unit measurement current $I^a \in \mathbb{R}_\diamond^M$.

It is interesting to note, that the computation of the EBJ resembles the computation of the conductivity Jacobian using the adjoint field method (section 2.2.7), in that the results of ‘drive current’ and ‘measurement current’ injections are used. Therefore, the calculation of the EBJ does not require any additional forward simulations to those performed for the calculation of the traditional Jacobian matrix.

5.2.2 Implementation

To describe the movement of an electrode along the surface of a mesh, a surface coordinate system is required. On the human head mesh, a surface coordinate system $[x_s \times y_s]$ was created as follows. Starting from the surface node with the largest y-coordinate (top of the head - (0, 0) in figure 5.2), the algorithm moved along the surface in x- and z-direction and found surface nodes that were a defined distance $d \pm \varepsilon$ away from the previous point, where ε was chosen large enough that surface nodes were found. Distance d was chosen such that the surface coordinate points were far enough from each other to describe the surface curvature without significant influence of the mesh discretisation. If no more surface nodes could be found that either fulfilled these requirements or were above a certain y-coordinate limit, the x_s - and y_s -axis were defined by the found points. The four quadrants were then filled with points the same way, by averaging

5. Electrode Model Correction in TDEIT




Figure 5.1: Surface coordinate system on a 800k mesh with one refined electrode - (a) The discrete points of the surface coordinate system in red and (b) a 10 mm displacement of electrode number 30 along the x_s -coordinate (arrow) seen as indicated by the box in (a).

the coordinates of all surface vertices within the area spanned by $d \pm \varepsilon$ (figure 5.2). To describe directions of movement for each electrode, the four closest surface coordinate points to that electrode were found (figure 5.1). A mapping back from surface coordinates to mesh nodes was not required. This approach had the advantage that the surface coordinate system could be stored together with the mesh and used for any electrode positions. The Jacobian matrix with respect to electrode radius was computed with a vector field \mathbf{v} pointing from the electrode center to the center of each electrode boundary edge.

For small meshes, the electrode boundary Jacobian (EBJ) was calculated in Matlab using a modified version of EIDORS (Adler and Lionheart, 2006). On the 4 million element mesh used for simulations, the EBJ was computed with PERTS (chapter 4) in order to reduce computing time and memory requirements. The implementation was similar in Matlab and PERTS: for each electrode, a sparse template matrix was constructed to store the contributions of drive and measurement fields. For each injection and measurement electrode pair this template matrix could then be multiplied by the electric potential fields of the drive and measurement currents to obtain the **EBJ** entry for this electrode for this line of the protocol. The integral along the electrode boundary (5.2) can be decomposed into a sum of the integrals over each element edge e , from node i_1 to node i_2 . To compute this integral, the drive and measurement potential field values at the two nodes and the corresponding electrode m have to be known. In




Figure 5.2: Generation of the surface coordinate system - Starting from the top of the head (0, 0), equispaced surface nodes within a tolerance ε were found in x and z direction. The coordinates of all surface nodes within this tolerance were averaged to define the surface coordinate system.

this implementation, Gauss-Lobatto quadrature with three points was used to approximate the integral along one edge. The Gauss-Lobatto weights on a unit interval were given as $w_1 = 1/6$, $w_2 = 4/6$ and $w_3 = 1/6$, and therefore the integral along edge e could be approximated as

$$\begin{aligned} \frac{1}{|e|} \int_e (U_m^d - u^d)(U_m^a - u^a) \, ds &\approx \frac{1}{6} (U_m^d - u_{i_1}^d)(U_m^a - u_{i_1}^a) \\ &+ \frac{4}{6} (U_m^d - u_{i_m}^d)(U_m^a - u_{i_m}^a) + \frac{1}{6} (U_m^d - u_{i_2}^d)(U_m^a - u_{i_2}^a). \end{aligned} \quad (5.3)$$

Since linear shape functions were used, the voltage at the midpoint i_m was given by $u_{i_m} = \frac{1}{2}(u_{i_1} + u_{i_2})$. Calculating out and writing as a matrix vector multiplication, the right hand side (5.3) can be written as

$$f(\mathbf{d}_e, \mathbf{a}_e) = \mathbf{d}_e^\top \begin{bmatrix} 1/3 & 1/6 & -1/2 \\ 1/6 & 1/3 & -1/2 \\ -1/2 & -1/2 & 1 \end{bmatrix} \mathbf{a}_e,$$

with drive voltages $\mathbf{d}_e^\top = [u_{i_1}^d, u_{i_2}^d, U_m^d]^\top$ and adjacent voltages $\mathbf{a}_e^\top = [u_{i_1}^a, u_{i_2}^a, U_m^a]^\top$. For the integral along one edge, the dot product of the outward normal and the permutation vector field, $(\mathbf{v} \cdot \mathbf{n}_e)$, was kept fixed. Therefore, the electrode Jacobian entry for one electrode for one current injection and voltage measurement electrode pair can be written as the sum of the integral approximations along all electrode boundary edges

$$\begin{aligned} \text{EBJ}(m, \mathbf{d}, \mathbf{a}) &= \sum_e \frac{-|e|}{z_m |\Gamma_m|} (\mathbf{n}_e \cdot \mathbf{v}) \mathbf{d}_e^\top \begin{bmatrix} 1/3 & 1/6 & -1/2 \\ 1/6 & 1/3 & -1/2 \\ -1/2 & -1/2 & 1 \end{bmatrix} \mathbf{a}_e \\ &= \mathbf{d}^\top \mathbf{M}_{temp}(m) \mathbf{a}. \end{aligned} \quad (5.4)$$

Consequently, one sparse $[N + M \times N + M]$ template matrix \mathbf{M}_{temp} could be created for each electrode by summing the contributions of all electrode boundary edges into the corresponding indices i_1 , i_2 and i_{elec} . N are the number of nodes in the mesh, M the number of electrodes and R the number of protocol lines (i.e. number of injections times number of measurements at each injection). These M template matrices could then be multiplied with all used combinations of $\mathbf{d}^\top = [\mathbf{u}^d, \mathbf{U}^d]^\top$ and $\mathbf{a}^\top = [\mathbf{u}^a, \mathbf{U}^a]^\top$ to obtain the $[R \times M]$ electrode boundary Jacobian with respect to vector field \mathbf{v} .

Once the template matrices were set up for all electrodes, the drive and measurement (adjacent) fields to multiply them with were the same forward solutions that were required to compute the ‘traditional’ Jacobian matrix and therefore no additional forward solutions had to be computed. The computation of one line of the **EBJ** took on average 0.03 seconds for a 4m element mesh on ten processors using PEITS, whereas one line of \mathbf{J}_σ took around 0.1 seconds

5. Electrode Model Correction in TDEIT

to compute. The additional computation of the EBJ therefore had a minimal effect on the computation time.

5.3 Electrode Boundary Jacobian Characteristics

5.3.1 Simulation Parameters and Analysis Methods

In order to study the characteristics of the voltage changes due to electrode boundary changes, a simulation study was performed on an 800 thousand element mesh (figure 5.1a). The mesh had refined elements (element size 3 mm instead of 6 mm) in the region where electrodes were located and very fine elements in the vicinity of one electrode (element size 0.4 mm). Current injection was then either simulated to be through a polar electrode pair involving the refined electrode and a coarse electrode on the opposite side of the head or through adjacent electrodes, again a coarse and the fine one. A representative set of ten different measurement electrode pairs (that were not injecting current) were chosen. The ten measured voltages were simulated for the following permutations: x_s -coordinate, y_s -coordinate or electrode diameter change, polar or adjacent injection, fine or coarse electrode. The diameter of 7 mm was altered between ± 1.5 mm. Electrode movements were simulated in the range of 0.1 mm–10 mm in both directions and both surface coordinate dimensions. The mesh was not altered to model the electrode movement. Instead, the assignment of surface facets to the electrodes was used to move the electrodes. This approach was simpler and could account for larger movements of the electrodes.

Subsequently, the simulated voltage changes due to the changes in electrode boundaries were scaled by the changes predicted by the EBJ. This scaling was done in order to be able to show all voltage changes together in one figure, even though they were of different magnitude. The plots of the EBJ accuracy have the simulated electrode change in millimetres on the x-axis and the electrode movement the EBJ would recover — given the simulated voltage difference — on the y-axis (in millimetres as well). Therefore an optimal prediction of voltage change would result in all 10 curves being perfectly diagonal (black dashed bar in figures 5.3 and 5.4). The absolute values of the EBJ with respect to position (both x_s and y_s) and the EBJ with respect to radius were taken in order to determine which factor was more important. The accuracy of the three different EBJs for the three studied electrodes for adjacent and polar injections was evaluated by linearly interpolating the voltages from simulated electrode boundary changes and comparing the slope to the value of the EBJ.

The simulations were performed using a current level of 133 μA and contact impedance of the electrodes (with area $|E|$) of $z_m = 1 \text{ k}\Omega \cdot |E|$ on a realistic head mesh with the following layers and corresponding conductivities: scalp (0.44 S m^{-1}), skull (0.018 S m^{-1}), dura mater (0.44 S m^{-1}), cerebrospinal fluid (1.79 S m^{-1}), grey matter (0.3 S m^{-1}) and white matter (0.15 S m^{-1}). The

mesh was created as described in chapter 3. The ten simulated measurements were the first ten lines of the ‘EEG31’ protocol (Tidswell et al., 2001b) and for the adjacent injection one of the injecting electrodes was substituted with an adjacent one.

5.3.2 Voltage Dependence on Electrode Characteristics

Comparison of x_s , y_s and d Prediction of EBJ

The EBJ accurately predicted voltage changes caused by electrode movement (figure 5.3). The changes caused by electrode diameter variations were overestimated by the EBJ, and were significantly smaller than the changes caused by electrode movement (colour bars in figure 5.3).




Figure 5.3: Ratio of voltage change over EBJ prediction for adjacent injection - (a) For movement of the fine electrode along x_s , **(b)** movement along y_s and **(c)** change in electrode diameter (different scaling of x- and y-axis). The dashed lines represent a perfect prediction of the voltage changes by the EBJ and the colours represent the amplitude of the changes (i.e. the value of the EBJ: large changes are blue, small changes yellow).

Effect of Mesh Refinement

The mesh refinement around an electrode did not limit the precision of the EBJ. However, if the electrode positions were corrected on a coarse mesh then the discretisation error was bigger (the big steps in figure 5.4b as opposed to the small steps in figure 5.4a). This was to be expected, since the boundary of the electrode could only change in intervals equal to the element size.

Comparison of Polar and Adjacent Current Injection

The voltage changes caused by electrode movement were only minimally different between adjacent and polar current injection (figures 5.3b and 5.4a). This matched expectations, considering that the formula for the computation of the EBJ only depends on electric potential differences at the electrode boundary.

5. Electrode Model Correction in TDEIT




Figure 5.4: Same than figure 5.3 for polar current injection - EBJ accuracy for movement along y_s (a) on the fine electrode and (b) on the coarse electrode.

5.3.3 Precision of Electrode Boundary Jacobian

The slopes of simulated voltage changes due to electrode boundary changes were linearly fitted and compared to the EBJ values, to get a measure of the EBJ precision (first three rows of table 5.1). The voltage changes due to electrode movement were well approximated by the EBJ, with most values having around 15% mismatch. The predicted changes due to electrode size changes were on average 690% off. The average absolute values of the EBJ were compared for the different types of electrode perturbation (movement along x_s and y_s and change in radius r) to illustrate that the changes due to electrode size were significantly smaller than the ones caused by movement (last three rows of table 5.1).

	polar injection		adjacent injection	
	coarse	fine	coarse	fine
x_s error	10.4%	23.3%	12.4%	23.3%
y_s error	10.2%	7.0%	46.0%	7.0%
r error	698%	1372%	30.2%	663%
$\text{abs}(\text{EBJ}_{x_s})$	9.8 mV/m	19.0 mV/m	8.3 mV/m	18.8 mV/m
$\text{abs}(\text{EBJ}_{y_s})$	16.1 mV/m	25.6 mV/m	18.2 mV/m	25.3 mV/m
$\text{abs}(\text{EBJ}_r)$	0.3 mV/m	1.1 mV/m	1.3 mV/m	1.1 mV/m

Table 5.1: Accuracy and amplitudes of the EBJ - The accuracy and the amplitude of the EBJ are shown for movements along x_s and y_s and electrode radius changes for polar and adjacent current injection. All values are the average of 10 different measurements.

5.4 Simulation Study

5.4.1 Time-Difference Image Reconstruction Algorithm

In both the simulation study and its experimental validation, spherical perturbations in the brain were used to represent a stroke. Therefore, prior information could be introduced into the reconstruction by using first order Tikhonov regularisation to bias the algorithm towards finding small connected perturbations. All time-difference images were created with a standard least-squares minimisation using generalised singular value decomposition (gSVD). Many other algorithms have been successfully applied to EIT data (Lionheart et al., 2004), and could equally well be used here for simultaneous reconstruction of conductivity changes and electrode movements.

The image reconstruction problem can be described by the minimisation of the cost functional

$$\Phi(\mathbf{x}) = \frac{1}{2}(\mathbf{F}(\mathbf{x}) - \delta\mathbf{v})^\top (\mathbf{F}(\mathbf{x}) - \delta\mathbf{v}) + \frac{1}{2}\lambda^2 \mathbf{x}^\top \mathbf{D}^\top \mathbf{\Gamma}_x^{-1} \mathbf{D} \mathbf{x} \quad (5.5)$$

with $\mathbf{x} = [\delta\boldsymbol{\sigma}, \delta\mathbf{p}]^\top$ the change in conductivity and electrode positions to be reconstructed, $\delta\mathbf{v}$ the voltage difference between the two measurements, $\mathbf{F}(\mathbf{x})$ a non-linear function relating conductivity and electrode position changes to voltage changes, $\mathbf{\Gamma}_x$ the expected variance of the change of the reconstructed variables (conductivities $\mathbf{\Gamma}_\sigma$ and positions $\mathbf{\Gamma}_p$)

$$\mathbf{\Gamma}_x = \begin{bmatrix} \mathbf{\Gamma}_\sigma & 0 \\ 0 & \mathbf{\Gamma}_p \end{bmatrix} = \begin{bmatrix} \text{std}_\sigma \cdot \mathbf{I} & 0 \\ 0 & \text{std}_p \cdot \mathbf{I} \end{bmatrix}^2 = \boldsymbol{\Sigma}_x \boldsymbol{\Sigma}_x \quad (5.6)$$

and the first order Tikhonov regularisation term for the conductivities combined with zero order Tikhonov for the electrode positions

$$\mathbf{D}^\top \mathbf{D} = \begin{bmatrix} \mathbf{L}^\top \mathbf{L} & 0 \\ 0 & \mathbf{I} \end{bmatrix}. \quad (5.7)$$

The cost function 5.5 was minimised with the gSVD of \mathbf{J} and $\boldsymbol{\Sigma}_x^{-1} \mathbf{D}$ (section 2.3.2).

To reduce the computational cost of calculating the gSVD and to prevent the ‘inverse crime’ (Lionheart et al., 2004), a much smaller hexahedral mesh (2526 elements of $1 \text{ cm} \times 1 \text{ cm} \times 1 \text{ cm}$) was used for the image reconstructions (figure 5.5a). The Jacobian matrix \mathbf{J}_σ , which was computed on the fine mesh, was summed into geometrically regular cubes \mathbf{J}_{hex} and the Laplacian matrix \mathbf{L} for the first order Tikhonov regularisation was computed on this hexahedral mesh. In order to simultaneously reconstruct conductivity changes and electrode movements, the ‘traditional’ Jacobian matrix \mathbf{J}_{hex} was combined with the \mathbf{EBJ} as $\mathbf{J} = \begin{bmatrix} \mathbf{J}_{hex} & \mathbf{EBJ} \end{bmatrix}$. Equivalently, if only the conductivity changes or only the electrode movements were reconstructed, only the

5. Electrode Model Correction in TDEIT

relevant Jacobian matrix and the corresponding part of $\Sigma_x^{-1}\mathbf{D}$ was used.

The expected standard deviation of the conductivity was set to $\text{std}_\sigma = 0.1 \text{ S m}^{-1}$ and for the electrode positions to $\text{std}_p = 1 \text{ mm}$. For all reconstructed images and electrode positions, the regularisation factor $\lambda^2 = 2.1 \cdot 10^{-7}$ was kept constant. The value was chosen based on the shape of the L-curve (Hansen, 1994), however the L-curve was not pronounced enough to choose λ in an automated way. The solution norm $\|\mathbf{x}\|_2$ was generally around 0.1 and the residual norm $\|\mathbf{J}\mathbf{x} - \delta\mathbf{v}\|_2$ around $1.0 \cdot 10^{-4}$. The colour bar of all images was scaled according to the largest reconstructed change in the whole mesh. Therefore images of slices sometimes do not contain the maximum value of the colour bar.

5.4.2 Image Error Quantification

In order to evaluate the image quality objectively, three image error quantification measures were applied. The volume P corresponding to the reconstructed perturbation was identified as the largest connected cluster of voxels with 50% of the maximum of the image. The region of interest (ROI) was defined as the largest connected cluster of voxels with 50% of the maximum of the simulated conductivity change.

- Localisation error: ratio between the distance $\|(x_P, y_P, z_P)\|$ of the centre of mass of the reconstructed perturbation P from the actual perturbation location, and the average dimension of the mesh $\text{mean}(d_x, d_y, d_z)$

$$\frac{\|(x_P, y_P, z_P)\|}{\text{mean}(d_x, d_y, d_z)}. \quad (5.8)$$

- ROI change: difference of the average value of the reconstructed image ($d\sigma_r$) in the ROI and the average value of the actual perturbation ($d\sigma_a$), divided by the average value of the perturbation

$$\frac{|f_{tikh} \cdot \text{mean}_{\text{ROI}}(d\sigma_r) - \text{mean}_{\text{ROI}}(d\sigma_a)|}{|\text{mean}_{\text{ROI}}(d\sigma_a)|}, \quad (5.9)$$

where the Tikhonov smoothing correction factor f_{tikh} corrects for the effect, that the reconstructed perturbation is larger in volume and smaller in amplitude by scaling the reconstructed conductivity change within the ROI ($f_{tikh} = 10$ was used throughout).

- ROI noise: noise-to-signal ratio of the reconstructed image ($d\sigma_r$), computed as the ratio between the standard deviation (std) outside the ROI and the average value in the ROI

$$\frac{\text{std}_{\Omega \setminus \text{ROI}}(d\sigma_r)}{\text{mean}_{\text{ROI}}(d\sigma_r)}. \quad (5.10)$$

5.4.3 Simulation Parameters

\mathbf{J}_σ and **EBJ** for the image reconstructions in the simulation study and tank experiments were computed on the same 4 million element human head mesh with a homogeneous saline background of 0.4 S m^{-1} conductivity and a realistic human skull with variable conductivities between 0.0094 S m^{-1} and 0.025 S m^{-1} (Tang et al., 2008). The current level was $250 \mu\text{A}$, contact impedances of all electrodes were set to $z_c = 220 \Omega \cdot |E|$ and the diameter of the 32 electrodes was 10 mm. The perturbation had a radius of 1.5 cm and a conductivity of 0.0001 S m^{-1} for plastic and 0.36 S m^{-1} for sponge (figure 5.5b). The injecting pairs of electrodes were chosen to maximise the distance between electrodes by finding the maximum spanning tree, weighted by the inter-electrode distances. Measurements were made for each injection on all adjacent electrode pairs not involved in delivering the current, giving a total of 869 measured voltages from 31 independent current injections. All voltages, the conductivity Jacobian matrix and the EBJ were computed with PEITS (chapter 4).

The level of noise added to simulated voltages was chosen to match the experiments and consisted of $\varsigma_p = 0.006\%$ proportional noise and $\varsigma_a = 1 \mu\text{V}$ additive noise, such that

$$V_{\text{with noise}} = V_{\text{no noise}}(1 + \text{rand}(\varsigma_p)) + \text{rand}(\varsigma_a), \quad (5.11)$$

where $\text{rand}(\varsigma)$ indicates random numbers drawn from a Gaussian distribution with zero mean and standard deviation ς .




Figure 5.5: Slices through the coarse hexahedral mesh used for reconstruction - (a) Baseline conductivities and (b) conductivity change when a plastic perturbation of 0.0001 S m^{-1} was inserted in the back of the head.

5.4.4 Electrode Position Recovery

To reconstruct conductivities and electrode positions simultaneously, the algorithm outlined in section 5.4.1 was used. Analogously, if only electrode positions were recovered, then the full Jacobian was replaced by only the EBJ. The performance of the EBJ for electrode position

5. Electrode Model Correction in TDEIT

recovery was validated with three different recovery modalities: 1) only the EBJ was used to recover electrode positions when the electrodes moved and the conductivities did not change; 2) only EBJ was used when electrode positions and conductivity (plastic ball in back of the head) changed; 3) the full Jacobian matrix was used to reconstruct conductivities and electrode movement at the same time when both electrode positions and conductivity have changed. All these differently recovered electrode positions were plotted together with the actual electrode movement for two cases, one where only one electrode on the back of the head moved by 5 mm (figure 5.6a) and the other one where all electrodes were moved along both x_s and y_s by random values with standard deviation 1 mm (figure 5.6b).

The 2-norm of the difference of recovered movement versus actual movement shows that the simultaneous reconstruction of a conductivity perturbation and electrode movements (last row in table 5.2) was close to that of the sole recovery of the electrode positions when no perturbation was introduced (first row in table 5.2). Contrastingly, if the conductivity changed simultaneously with the electrode positions, the recovery of only the electrode positions resulted in a significant over-correction (second row in table 5.2), in particular close to where the perturbation was introduced (numbers 41 to 64 in figure 5.6 corresponding to x_s and y_s coordinates of electrodes 21 to 32 towards the back of the head).

	Elec. 24 x_s : 5 mm	STD: 1 mm
EBJ no perturbation	1.1 mm	3.8 mm
EBJ with perturbation	11.3 mm	10.3 mm
J with perturbation	1.9 mm	4.7 mm

Table 5.2: 2-norm of the electrode movement recovery error - 2-norm of the difference of recovered and simulated electrode movement. Simultaneous recovery of conductivity changes and electrode movement (last row) was almost as accurate as sole recovery of electrode movements.

5.4.5 Images

The best possible images (i.e. in the presence of no electrode movement) that were achievable with time-difference reconstructions without electrode correction (figure 5.7a) and with electrode correction (figure 5.7b) were qualitatively similar. The simultaneous reconstruction of conductivity and electrode positions slightly reduced the contrast and the precision of the imaged perturbation.

If one electrode in the top back of the head was moved by 5 mm, the simple time-difference reconstruction resulted in an unsurprisingly noisy image (figure 5.9a) with large artefacts around the moved electrode. Simultaneous reconstruction of the electrode positions (figure 5.9b) restored the quality of the image almost to the ideal case (figure 5.7a).

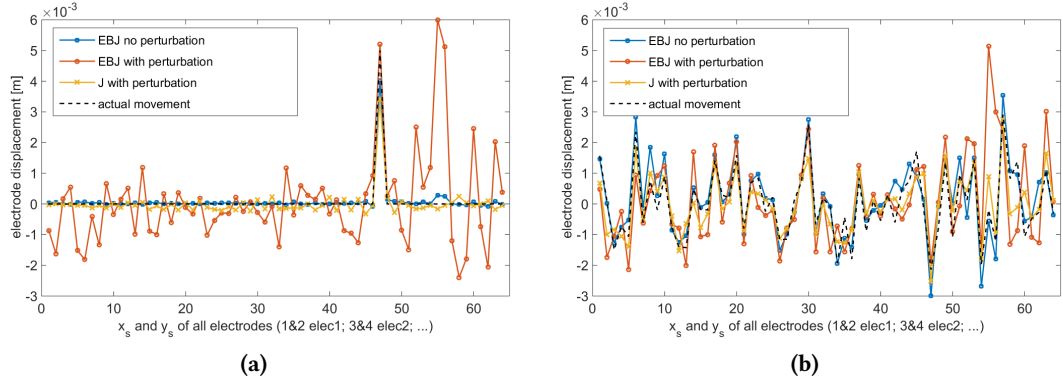


Figure 5.6: Recovered electrode movement - For three different cases: only EBJ was used and only the electrode positions changed, only EBJ was used and additionally to the electrode movement a perturbation was inserted in the back of the head, and the complete Jacobian \mathbf{J} was used when both electrodes and conductivity changed. In (a) an electrode on the back of the head was moved along x_s by 5 mm and in (b) all electrodes were moved along x_s and y_s by a random value with standard deviation 1 mm. On the x axis of these two plots the entries 1 and 2 correspond to the x_s and y_s coordinate of electrode 1, then the next two entries correspond to electrode 2 and so on.




Figure 5.7: Reconstructions when electrodes have not moved - A perturbation in the back of the head (black outline) was simulated with the correct electrode positions. (a) The best achievable reconstruction without using the EBJ and (b) the best achievable reconstruction with EBJ, demonstrating the inevitable decrease in image quality compared to (a), since electrodes have not moved.

5. Electrode Model Correction in TDEIT

As with a large movement of one electrode, smaller movements of all electrodes had a detrimental effect on the image quality for simple time-difference imaging (figure 5.10a). The image reconstructed with the full Jacobian was again largely unaffected by the electrode movement (figure 5.10b). All these findings were summarised for three different positions of the perturbation by assessing the image quality according to the previously defined error metrics (figure 5.11). The larger localisation error (and inherently larger ROI errors) for a perturbation on the side can be explained by the reduced sensitivity of the current protocol used for the simulations (figure 5.8).




Figure 5.8: Electrodes connected by the 31 unique current injections - All electrodes and injections were projected onto the x-z plane and the electrodes were coloured according to their maximal EBJ entry. Because the used protocol was optimised according to electrode distances, most injections connected the front (top of this figure) to the back of the head. The used perturbation locations are indicated by the green stars.




Figure 5.9: Reconstructions when one electrode moved by 5 mm - Electrode 24 (top back of the head) was moved 5 mm along x_s . (a) Reconstruction without using the EBJ and (b) reconstruction with EBJ, demonstrating a decrease in artefact around electrode 24 and improved localisation of the perturbation.

All results shown so far were obtained with a plastic perturbation, which provided a large contrast. When imaging a sponge ball with 90% of the saline conductivity instead of a plastic ball with effectively zero conductivity, the reduced signal-to-noise ratio introduced artefacts into the image of larger magnitude than the perturbation (figure 5.12). However, even when reconstructing a 10% conductivity change inside the highly resistive skull, the use of the EBJ significantly improved the image quality (figure 5.13).




Figure 5.10: Reconstructions when all electrodes moved around 1 mm - Electrodes were moved along both surface coordinates x_s and y_s by normally distributed random distances with standard deviation 1 mm. (a) Reconstruction without using the EBJ and (b) reconstruction with EBJ, demonstrating decreased electrode artefacts and improved perturbation localisation.

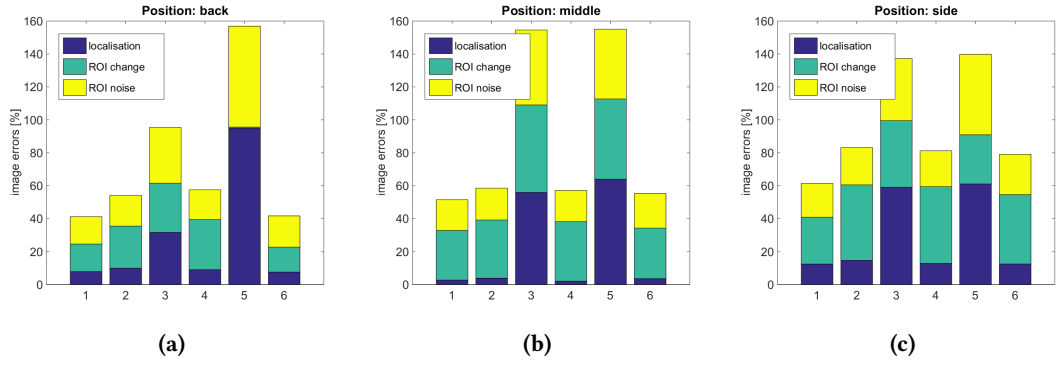


Figure 5.11: Image errors for three different perturbation locations - (a) Perturbation in the back, (b) in the middle and (c) on the side. 1 and 2 are the error metrics for reconstructions without and with electrode correction when the electrodes have not moved. 3 and 4 are the metrics for reconstructions without and with electrode correction when electrode 24 in the top back of the head was moved along x_s by 5 mm. And 5 and 6 are the same metrics in the case of random movements of all electrodes with a standard deviation of 1 mm.




Figure 5.12: Reconstructions of a simulated sponge in the back of the head - Reconstructions are made with electrode movement correction (a) when electrodes have not moved, (b) when electrode 24 in the back top of the head was moved along x_s by 5 mm and (c) when all electrodes were moved along x_s and y_s by random distances with standard deviation of 1 mm.

5. Electrode Model Correction in TDEIT

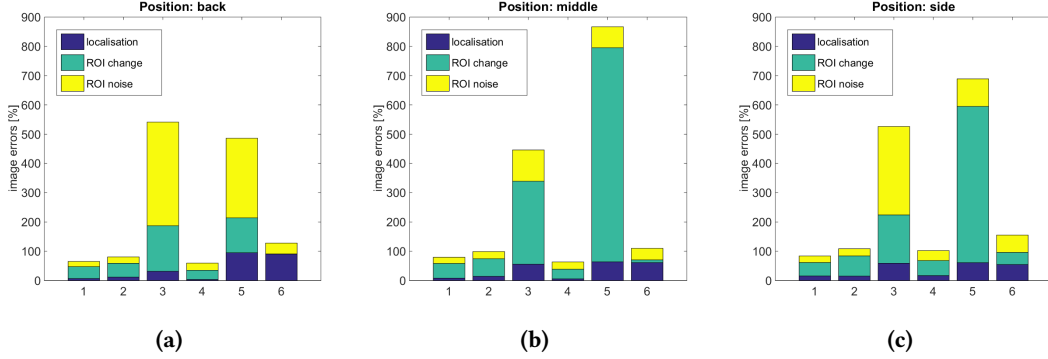


Figure 5.13: Image errors for reconstructions of a sponge - (a) Sponge in the back, (b) in the middle and (c) on the side. 1 and 2 are the error metrics for reconstructions without and with electrode correction when the electrodes have not moved. 3 and 4 are the metrics for reconstructions without and with electrode correction when electrode 24 in the top back of the head was moved along x_s by 5 mm. And 5 and 6 are the same metrics in the case of random movements of all electrodes with a standard deviation of 1 mm.

5.5 Experimental Validation

5.5.1 Experimental Setup

Two saline tanks were printed using the 3D printer Makerbot Replicator 2 from Makerbot Ind. The model for the tank was created from the same MRI segmentation used for the mesh creation, whereas the skull was segmented from a corresponding CT scan. The skull model was further edited by introducing small holes, such that the conductivity matched that of a real skull (Avery, 2015). One tank was printed with a modified EEG 10-20 electrode placement (Avery, 2015; Tidswell et al., 2001a) and the other tank had perturbed electrode positions with random displacements along x_s and y_s with standard deviation of 1 mm, matching the electrode movement simulated in section 5.4.

For the recordings, the tanks were filled with 0.4 S m^{-1} saline (figure 5.14). The electrode contact impedances were measured to be $220 \Omega \cdot |E|$. Current was injected at 1.76 kHz with $250 \mu\text{A}$ amplitude, to approximately match the allowed current level in human measurements (Dybdahl, 2009). A slightly modified 32-channel version of the KHU Mark 2.5 system (Wi et al., 2014) was used for the recordings and each measurement was repeated 20 times over the course of approximately one minute. Plastic perturbations with 3 cm diameter were placed in approximately the same locations used in the simulations.

5.5.2 Images

With baseline and perturbation measurements in the same saline tank, both using \mathbf{J}_{hex} (figure 5.15a) and using \mathbf{J} (figure 5.15b) resulted in a reconstruction of the perturbation with only minimally worse quality than for simulated noisy voltages (first two bars in the quality measures




Figure 5.14: The KHU Mark 2.5 system and the 3D printed head shaped tank with skull - The skull (blue) was printed with small holes in it, to give it a realistic conductivity when it is immersed in saline.

in figure 5.16). However, when reconstructing a perturbation measured in the tank with the moved electrodes, then large artefacts in the outer layers of the head near the skull were observed (figures 5.15c and 5.15d). These artefacts strongly influenced the localisation errors (figure 5.16), where an artefact was interpreted as the reconstructed perturbation. The drop in image quality was most likely caused by small differences in skull placement in the two different saline tanks.




Figure 5.15: Experimental images in the head tank with realistic skull - The perturbation was a plastic ball in the back of the 3D printed head tank. (a) The reconstruction without EBJ when baseline and perturbation measurements were done in the tank with the correct electrode positions and (b) the same reconstruction including EBJ. (c) The reconstruction without EBJ when the baseline was measured in the tank with the correct electrode positions and the perturbation measurement was done in the tank with the electrodes shifted by random values with standard deviation 1 mm in both surface directions and (d) the same reconstruction with EBJ.

Simultaneous reconstruction of conductivities and electrode movements improved the image

5. Electrode Model Correction in TDEIT

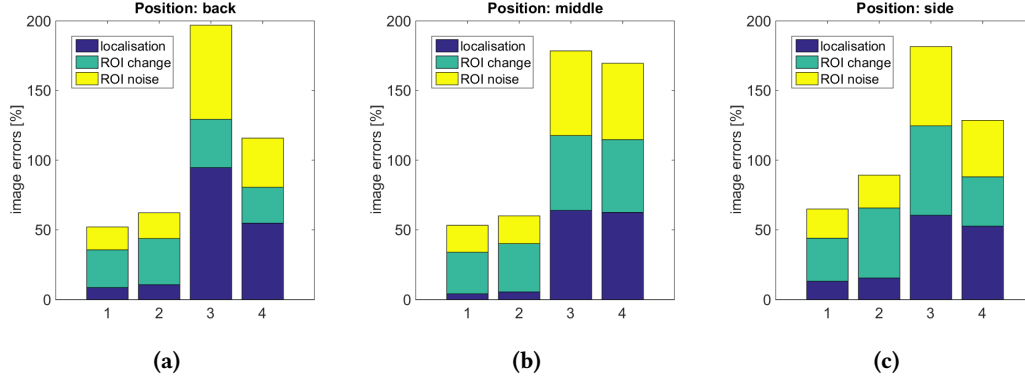


Figure 5.16: Image error quantification results for experimental images - (a) Plastic perturbation in the back, (b) in the middle and (c) on the side. 1 and 2 are the error metrics for reconstructions without and with electrode correction when the electrodes have not moved. And 3 and 4 are the same metrics in the case of random movements of all electrodes with a standard deviation of 1 mm.

visibly (figure 5.15d), but did not remove all artefacts. Since electrode movement cannot account for all voltage changes, artefacts caused by other effects (such as skull position in the two tanks) remained in the conductivity image. Parts of the voltage differences caused by these geometrical differences were, however, pulled into the electrode movement recovery. Therefore, the precision of the recovered electrode movement between the two 3D printed tanks was lower than in simulations (error 2-norm 8.0 mm instead of 4.7 mm in simulations).

5.6 Absolute Reconstructions

5.6.1 Absolute Reconstruction Algorithm

In contrast to time-difference imaging, absolute reconstructions do not require a baseline measurement. In many realistic applications, absolute imaging fails – among other reasons – because of small differences between the model and the imaged object, resulting in large image artefacts (Kolehmainen et al., 1997; Heikkinen et al., 2002; Malone et al., 2014). Since the 3D printed saline tank and skull were very close to the corresponding mesh, a trial how the correction of electrode positions performed in iterative absolute reconstructions was considered worthwhile. A generalised minimal residual algorithm with tolerance of $5 \cdot 10^{-11}$ was used to solve the standard Gauss-Newton problem

$$\nabla \Phi(\mathbf{x}_k) = -\mathbf{J}^\top(\mathbf{x}_k)(\mathbf{J}(\mathbf{x}_k)\mathbf{d}_k) - \lambda^2 \mathbf{D}^\top \mathbf{\Gamma}_x^{-1} \mathbf{D} \mathbf{d}_k \quad (5.12)$$

for the search direction \mathbf{d}_k . The reconstructed change in conductivity and electrode positions was then updated in each iteration

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \tau_k \mathbf{d}_k, \quad (5.13)$$

where τ_k is the step size and was optimised using the Brent line search method (section 2.3.3; Brent, 1973) with a gold-section bracketing loop to find the Brent abscissae. Analogous to the time-difference approach, the reconstructions were weighted by the expected variance. In iterative reconstructions it is important to ensure the positivity of the elements' conductivity, which was done here by using the substitution

$$\mathbf{y} = \begin{bmatrix} \log(\text{std}_\sigma^{-1} \boldsymbol{\sigma}) \\ \text{std}_p^{-1} \mathbf{p} \end{bmatrix}, \quad (5.14)$$

where $\mathbf{x} = [\boldsymbol{\sigma}, \mathbf{p}]^\top$. Applying this substitution to (5.12) moved the scaling by the expected standard deviation of the variables from the regularisation term to the Jacobian matrix, making the Jacobian dimensionless. According to the chain rule, the logarithmic part of the substitution results in a multiplication of the Jacobian entries corresponding to conductivities by e^σ . After \mathbf{y}_{k+1} was found using Gauss-Newton, the electrode positions and conductivities were updated according to the corresponding \mathbf{x}_{k+1} .

The regularisation parameter for all absolute reconstructions was set to $\lambda^2 = 10^{-11}$ and the stopping criterion was set to $< 10^{-7}$ change between iterations with a maximum of 6 iterations. As a trade-off between computational efficiency and precision of the forward solutions, a tetrahedral mesh with 1 million elements was used for the absolute reconstructions.

5.6.2 Images

All absolute reconstructions from simulated data had a large positive artefact in the front of the head (figure 5.17a), where the sensitivity was high. When the full Jacobian matrix was used, the algorithm unsuccessfully tried to reduce this artefact by moving the electrodes in the front of the head (left side of figure 5.17b). When evaluating the recovery of electrode movements, this 'baseline' movement was therefore subtracted from the recovered movement of the electrodes when they have actually moved. Doing this, a good performance of the iterative electrode

	Elec. 24 x_s : 5 mm	STD: 1 mm
Back	2.5 mm	4.7 mm
Middle	2.0 mm	4.5 mm
Side	3.3 mm	4.5 mm

Table 5.3: 2-norm of the electrode position recovery error
 - The error was defined as the difference of recovered electrode movement and actual electrode movement, minus the 'baseline' movement that corresponds to the algorithm trying to reduce the artefact in the front of the head.

5. Electrode Model Correction in TDEIT

movement correction was found (table 5.3, as compared to the last row in table 5.2). The large artefact resulted in a poor image quantification score for all reconstructed images, with or without inclusion of the EBJ. Consequently, the difference the electrode correction made is illustrated by showing one reconstruction without and one with EBJ, both thresholded to one third of the maximum reconstructed change (figure 5.18).

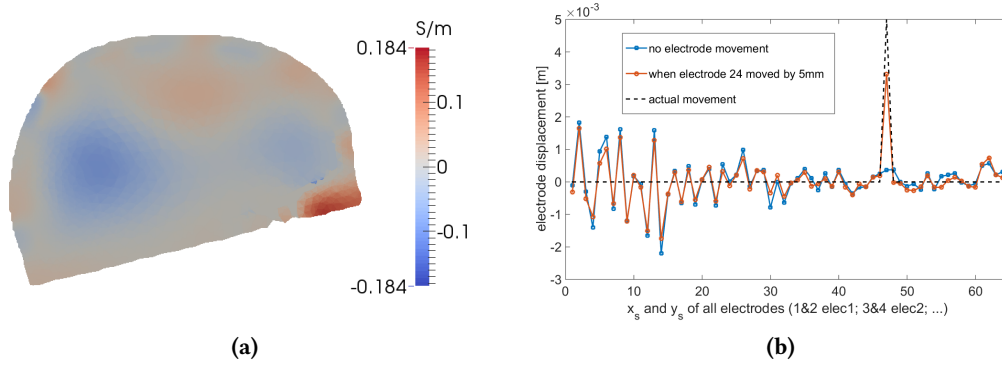


Figure 5.17: Absolute reconstruction and absolute electrode position recovery - (a) Conventional absolute reconstruction (no electrode correction) of a simulated plastic perturbation in the back of the head. Voltages were simulated with original electrode positions. (b) Comparison of the recovered electrode movement (using the full Jacobian) in absolute imaging when using the original electrode positions, and when electrode 24 was moved by 5 mm along x_s (dashed black line).




Figure 5.18: Absolute reconstructions with 1 mm electrode movement - Electrode positions were changed by normally distributed random distances with standard deviation 1 mm along both surface dimensions x_s and y_s . (a) Reconstruction without using the EBJ and (b) reconstruction with EBJ. Both reconstructions were thresholded at $\pm 33\%$ of the maximum change.

We were not able to reconstruct meaningful absolute images from the experimental data. Even on the tank with no electrode movement, the reconstructions did not show the perturbation, suggesting that the differences between the mesh and the tank were too large to enable absolute imaging.

5.7 Discussion

5.7.1 Electrode Boundary Jacobian Characteristics

Measured voltages changed linearly with electrode movement over 1 cm and were accurately predicted by the EBJ (the linearity was already observed multiple times, as reviewed in Kolehmainen et al. (1997)). There were outliers where some voltages did behave highly non-linearly (figure 5.3a), but as seen from the colour coding these outliers were of small amplitude. Therefore they were not expected to adversely influence the recovery of electrode movements. Changes in electrode diameter had a less predictable influence on the measured voltages (figure 5.3c). Firstly, the EBJ predicted changes were between one and two orders of magnitude smaller than for electrode movements and, secondly, the observed voltage changes were even smaller than the EBJ predictions. Changes in electrode shape can be viewed as a combination of a change in diameter and the discretisation errors seen in the electrode movement on the coarse electrode (figure 5.4b). Therefore these changes can be expected to be small, even though this was not analysed specifically. It was concluded, that changes in electrode size and shape could be ignored in the presence of electrode movement, and that they were not well approximated by the EBJ. The focus of the simulation study and the experiments was therefore laid on the correction of electrode movements, which had the strongest impact on the image quality.

The mesh refinement around an electrode had no influence on the precision of the EBJ (figure 5.4), which indicated that the meshes did not have to be altered for the use of electrode movement correction. Electrodes could be moved by simply assigning different surface facets to the electrode, without introducing errors larger than the discretisation errors caused by the element size. This facilitated the electrode movement correction significantly, since the mesh remained unchanged and the electrodes could be moved by more than the size of the finite elements.

5.7.2 Simulation Study

Time-difference reconstructions were very stable and gave good images in all cases. Using the EBJ, the reconstruction algorithm could account for large electrode movements and thereby reduce the image noise notably. In the presence of image artefacts which were not caused by electrode movement, using the EBJ resulted in a less precise recovery of the electrode positions. This was because the reconstruction algorithm reduced the cost functional by moving parts of the image artefact into the electrode positions. However, since 64 parameters constituting the positions of the 32 electrodes could not account for these artefacts, they were still visible in the resulting conductivity image.

5. Electrode Model Correction in TDEIT

5.7.3 Experimental Validation

Time-difference reconstructions in the printed head tank with the correct electrode positions gave good results. Using the baseline measurement with the correct tank and the perturbation measurement on the tank with changed electrode positions resulted in noisier images than in the corresponding simulations. The additional noise was most likely caused by small geometrical differences in the skull positioning in the two different tanks, slightly different saline levels between the measurements and ambient and system temperature differences. While such errors were not present in conventional time-difference measurements, they manifested themselves when two different tanks were used with a half an hour break in between setting up each tank. Still, the simultaneous reconstruction of conductivity and electrode movement significantly improved the reconstructions and allowed for the detection of the perturbation.

5.7.4 Absolute Reconstructions

The iterative absolute recovery of electrode positions was accurate and strongly reduced the electrode modelling related artefacts. Still, the absolute conductivity reconstructions with or without electrode correction showed a systematic artefact in the front of the head. This was in a region of strong sensitivity and might have been exacerbated by the comparative imprecision of the one million element mesh used for the absolute reconstructions. Moving to a finer, more precise mesh would result in very long computation times even for a single image reconstruction.

We were not able to get meaningful absolute reconstructions from the experimental data, suggesting that the simulations and experiments were still too different even though the geometry of the tank and inserted skull were printed with a precision of 0.2 mm. However, the placing of the skull in the tank was less precise than this. More advanced absolute reconstruction algorithms might be able to retrieve more information from the collected experimental data by including geometrical uncertainties around the skull into the imaging method.

5.7.5 Conclusion

The Fréchet derivative of the EIT forward problem with respect to electrode boundary changes was applied to a realistic three-dimensional model of the human head, using a fast implementation. This allowed for the simultaneous recovery of conductivities and electrode positions, and therefore made image reconstructions more stable in the presence of electrode movement. It was found that:

1. The electrode position has a much stronger effect on the boundary voltages than the electrode diameter and shape. The voltages changed linearly for electrode movements up to 1 cm and were accurately predicted by the EBJ.

2. The simultaneous reconstruction of conductivity and electrode positions worked very well. Image artefacts caused by electrode movements were reliably removed and the electrode positions were accurately corrected. The positive results of the simulation study were confirmed in experiments on a 3D-printed saline tank containing a realistic 3D-printed skull.
3. While the correction of electrode movements was accurate in iterative absolute imaging as well, the absolute conductivity reconstructions from simulated data had a large artefact in the front of the head. No meaningful absolute images could be reconstructed on the 3D-printed saline tank.

The proposed application of the EBJ performs well in time-difference reconstructions and can be used for long term monitoring of physiological changes. For stroke type differentiation, time-difference measurements are not possible. Therefore, multi-frequency (MF) reconstruction algorithms have to be used (Malone et al., 2014; Jun et al., 2009). An evaluation of the influence of modelling errors in MF imaging could show, whether electrode modelling errors are equally detrimental for image quality and should be accounted for in the reconstruction.

Chapter 6

Feasibility of Stroke Imaging with Multi-Frequency Tissue Fraction Reconstructions

6.1 Overview

6.1.1 Introduction

In chapter 3, the influence of geometrical modelling errors on time-difference (TD) reconstructions was investigated. In the previous chapter the analysis of modelling errors in TD imaging was extended to the electrode modelling and it was found that electrode movement caused large artefacts. A computationally efficient correction method for electrode modelling errors was implemented and tested on a realistic head. The electrode correction was stable in the presence of large electrode movements and even worked in an iterative absolute reconstruction algorithm.

Multi-frequency Electrical Impedance Tomography (MFEIT) is a method for imaging biological tissues with frequency-dependent conductivity, which does not require reference measurements. Therefore MF imaging can be used as a diagnostic tool for multiple applications, where reference measurements are not usually available (Malich et al., 2003; Hampshire et al., 1995; Brown et al., 1995; Shi et al., 2008), including stroke type differentiation (Horesh et al., 2005; Romsauerova et al., 2006; Packham et al., 2012). While in ischaemic strokes, cell swelling caused by energy failure results in a conductivity decrease (Hansen and Olsen, 1980; Holder, 1992), in haemorrhagic strokes, the increased blood volume results in higher conductivity. Knowledge of the conductivity spectra of ischaemic and haemorrhagic tissue compared to healthy brain tissue allows for their differentiation in MFEIT.

Like TD reconstructions, MF imaging has the advantage over absolute imaging, that some

6. Feasibility of Stroke Imaging with MFEIT

modelling and instrumentation errors are cancelled by using differential measurements. A novel, non-linear method for performing MF reconstructions simultaneously from measurements at multiple (as opposed to two) frequencies was recently proposed (Malone et al., 2014). The inverse problem was modified by substituting conductivity with an auxiliary variable, which depends on the conductivity spectra of the tissues. The new variable, the volume fraction of each tissue in each voxel, describes the physical distribution of tissues in the domain, and is independent of frequency.

So far, it is unknown how stable non-linear multi-frequency reconstructions are with respect to modelling errors, such as contact impedances, electrode positions, geometry and tissue conductivity spectra.

6.1.2 Background

A first MF reconstruction algorithm was proposed for homogeneous, frequency independent backgrounds (Brown et al., 1995), and was found to be relatively stable in the presence of geometrical errors in 2D. Later, this linear method was extended to be compatible with frequency dependent, but still homogeneous, background media (Seo et al., 2008; Jun et al., 2009) and named weighted frequency difference algorithm (WFD). WFD was then applied to a simple hemispherical 3D model and tested with regards to different modelling errors (Ahn et al., 2011). It was shown to cope well with wrongly modelled electrode positions and mesh geometries.

For non-homogeneous backgrounds, the linearity assumption of the WFD algorithm is no longer valid and non-linear reconstruction methods have to be used (Malone et al., 2014). Two non-linear MF algorithms have been proposed, one was more stable with respect to tissue spectrum errors (Malone et al., 2015) than the fraction reconstruction method (Malone et al., 2014), but at the cost of having to optimize an additional parameter. Both non-linear methods were only tested on simple cylindrical shapes and their stability with respect to modelling errors was not assessed.

6.1.3 Purpose

In this chapter, the fraction reconstruction algorithm was for the first time applied to a realistic three-dimensional model of the head. The purpose of this simulation study was to evaluate the feasibility of stroke type differentiation with MFEIT. The influence of imprecise modelling was evaluated for four common sources of image artefacts in MFEIT (McEwan et al., 2007): mesh discretisation, electrode positions, electrode contact impedance and tissue conductivity spectra.

Voltages were simulated on a head mesh with realistic geometries for brain, skull and scalp, and modelling errors that can be expected with modern experimental setup were included. Images were then reconstructed from the simulated noisy voltages with the fraction reconstruction

method and objective metrics for image errors were computed.

6.2 Methods

6.2.1 Model and Tissue Conductivity Spectra

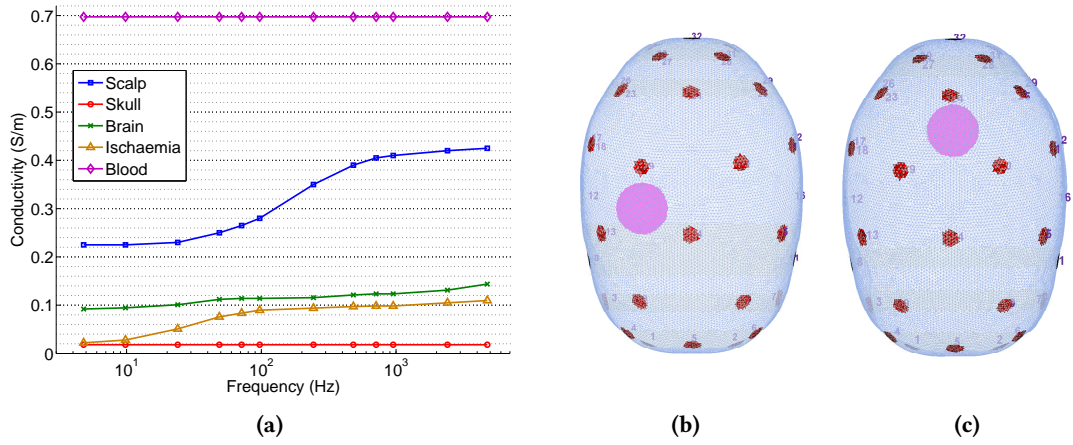


Figure 6.1: Conductivity spectra and simulated stroke positions - (a) Conductivity spectra of tissues in the head and view from the top on the (b) lateral and (c) posterior stroke position.

A three dimensional model of a human head was used to simulate EIT data. The model comprised three layers, corresponding to the scalp, skull, and brain. For simplicity, in this first feasibility study the model did not include the cerebro-spinal fluid, a simplification which is commonly made in EIT research. A fine mesh with ~ 5 million elements was used to simulate the boundary voltages, and a coarse mesh with ~ 180 thousand elements was used to reconstruct the images, in order to avoid the inverse crime (Lionheart et al., 2004). A spherical perturbation of diameter 3 cm was placed in two different positions inside the brain: lateral and posterior (figures 6.1b and 6.1c). The conductivity spectra of the tissues (scalp, skull, brain, ischaemic brain, blood – figure 6.1a) were obtained from the literature (Romsauerova et al., 2006; Horesh et al., 2005). In order to simulate an ischaemic stroke, the conductivity of the perturbation was set to the conductivity of ischaemic brain approximately one hour after onset, and in order to simulate a haemorrhagic stroke, the conductivity of the perturbation was set to the conductivity of blood. Twelve frequencies were chosen in the range 5 Hz–5 kHz based on the observation that the slopes of different tissues are most separate in this region (figure 2.7), and the boundary voltages were simulated for each frequency.

32 electrodes of diameter 10 mm were placed on the scalp in the same positions used to acquire EEG measurements and injecting electrode pairs were chosen to maximize the distance

6. Feasibility of Stroke Imaging with MFEIT

between the electrodes. This was done by finding the maximum spanning tree of the electrodes, weighted by the distance between the electrodes. The electrodes were modelled using the complete electrode model (CEM), and the contact impedance was set to $1 \text{ k}\Omega \cdot |E|$ for the electrode contact surfaces $|E|$. The amplitude of the current was set to $140 \mu\text{A}$ and voltage measurements were made for each injection on all adjacent electrode pairs not involved in delivering the current. The total number of measurements acquired for each frequency was 869. The boundary voltages were computed with PEITS (chapter 4) to reduce computation time for forward simulations. On all 16 cores of a workstation with two 2.4GHz Intel Xeon CPUs with eight cores and 20MB cache each, the computation time for the 31 forward solutions on the fine 5 million element mesh, which were required for the simulation of each frequency with each modelling error, was less than 2 minutes.

6.2.2 Image Reconstruction

Images were reconstructed using the fraction reconstruction method (section 2.3.4; Malone et al., 2014). For all frequencies $\{\omega_i; i = 1, \dots, W\}$ and tissues $\{t_j; j = 1, \dots, T\}$, the conductivity of a tissue sample $\epsilon_{ij} = \sigma^{t_j}(\omega_i)$ was assumed to be known. The conductivity of finite element n at a certain frequency $\sigma_n(\omega_i)$ was then modelled as the linear combination of the conductivities of individual tissues $\{\epsilon_{ij}; j = 1, \dots, T\}$:

$$\sigma_n(\omega_i) = \sum_{j=1}^T f_{nj} \cdot \epsilon_{ij}, \quad (6.1)$$

with $0 \leq f_{nj} \leq 1$ and $\sum_{j=1}^T f_{nj} = 1$. The weightings f_{nj} of the linear combination are called *fractions*, and were defined for each tissue and element. The fractions represent the physical distribution of the tissues, and are independent of the frequency at which the data is acquired.

For $\mathbf{f}_j = \{f_{nj}; n = 1, \dots, N\} \in \mathbb{R}^{N \times 1}$ and $\mathbf{f} = \{\mathbf{f}_j; j = 1, \dots, T\} \in \mathbb{R}^{N \times T \times 1}$, the objective function was

$$\Phi(\mathbf{f}) = \sum_{i=2}^W \left\| \sum_{j=1}^T \mathbf{J}_{ij} \mathbf{f}_j \epsilon_{ij} - \sum_{j=1}^T \mathbf{J}_{1j} \mathbf{f}_j \epsilon_{1j} - (\mathbf{v}_{\omega_i} - \mathbf{v}_{\omega_1}) \right\|^2 + \lambda^2 \Psi(\mathbf{f}) \quad (6.2)$$

with a first order Tikhonov regularization term of the form

$$\Psi(\mathbf{f}) = \sum_{j=1}^T \sum_{n=1}^N \sum_{l(n)} |f_{nj} - f_{l(n)j}|^2, \quad (6.3)$$

where $l(n)$ runs over all neighbours of the n th element.

The fractions were recovered simultaneously for all tissues and elements by minimizing

the objective function $\Phi(\mathbf{f})$, using alternating steps of gradient projection and damped Gauss-Newton algorithms (section 2.3.4). The skull and scalp were fixed in place, and it was assumed that the area inside the skull was occupied by either the brain or the stroke. The initial guess was the healthy brain. The optimal regularization parameter was approximated by computing the L-curve for one step of Gauss-Newton descent. The corner of the L-curve was selected for the first step of the reconstruction, and the value was halved after each iteration for ischaemic strokes and, divided by 3 for the haemorrhages. The different reduction of the regularisation parameter was originally thought necessary, but was since found not to affect reconstructions significantly. The maximum number of iterations for the reconstruction algorithm was set to ten.

6.2.3 Image Error Quantification

An image error quantification method was devised to evaluate the images objectively. The image quality was assessed in terms of the ability to distinguish an anomaly (the stroke) from a background (the brain). For this reason the fraction \mathbf{f}_s corresponding to the tissue that makes up the anomaly was assessed. The volume P corresponding to the reconstructed perturbation was identified as the largest connected cluster of voxels with values larger than 50% of the maximum of the image. Three measures of image errors were chosen:

1. Location error: ratio between the distance $\|(x_P, y_P, z_P)\|$ of the centre of mass of the reconstructed perturbation P from the actual position, and the average dimension of the head $\text{mean}(d_x, d_y, d_z)$

$$\frac{\|(x_P, y_P, z_P)\|}{\text{mean}(d_x, d_y, d_z)}. \quad (6.4)$$

2. Shape error: ratio of the difference between the dimensions of the simulated (s_x, s_y, s_z) and reconstructed perturbation (r_x, r_y, r_z) , and the dimensions of the simulated perturbation

$$\frac{\|(r_x - s_x, r_y - s_y, r_z - s_z)\|}{\|(s_x, s_y, s_z)\|}. \quad (6.5)$$

3. Image noise: inverse of the contrast-to-noise ratio between the perturbation P and the background B

$$\frac{\text{std}(f_s^B)}{|\bar{f}_s^P - \bar{f}_s^B|}, \quad (6.6)$$

where \bar{f}_s^P and \bar{f}_s^B are the mean intensities of the perturbation and background and std the standard deviation.

6. Feasibility of Stroke Imaging with MFEIT

6.2.4 Error Simulation

Instrumentation noise

Instrumentation noise was added to all simulated voltages and was chosen to match that of measurements acquired using the KHU Mark 2.5 EIT system (Oh et al., 2011) on a human subject and averaged over 64 frames. This noise level is achievable with most EIT measurement systems and can be reduced by use of better instrumentation. The standard deviation of the proportional noise was $\varsigma_p = 0.02\%$ and the standard deviation of the additive noise was $\varsigma_a = 5 \mu\text{V}$, such that

$$V_{\text{with noise}} = V_{\text{no noise}} (1 + \text{rand}(\varsigma_p)) + \text{rand}(\varsigma_a), \quad (6.7)$$

where $\text{rand}(\varsigma)$ indicates a random number drawn from a Gaussian distribution with zero mean and standard deviation ς . The additive noise was dominating these simulated EIT measurements.

Mesh Discretisation

In order to validate the reconstruction method on the realistic head model, images were reconstructed from simulated data without the addition of modelling errors except those due to mesh discretisation and noise. The data were simulated using the fine 5 million element mesh and the images were reconstructed using the coarse 180 thousand element mesh. Mesh discretisation errors could be reduced by using a finer mesh for reconstructions, at the cost of longer computation times and memory requirements.

Electrode Positions

Electrode positions can currently be measured to around 1 mm precision using photogrammetry (Qian and Sheng, 2011). Other technologies, such as the commercial MicroScribe, laser 3D scanners, or electrode helmets, can achieve an even higher precision in electrode localisation. To demonstrate the importance of using such localisation technologies, electrode position errors of around 1 mm and 2 mm were simulated. Given that the electrodes were represented on a discrete mesh, the shape and size of the electrodes did also change when an error was added to the position of the centre. This could have been corrected for by refining the mesh. However, it was decided not to refine the mesh around the electrodes, since a coarse representation of the electrodes constituted an unpredictable error source on electrode shape and contact area, like the application of conductive gel in an experimental setup (Kolehmainen et al., 1997; Boyle and Adler, 2011). The positional errors were added by perturbing the (x, y, z) coordinates of all electrodes with random numbers drawn from zero-mean Gaussian distributions with standard deviations of 0.25 mm and 0.5 mm. Deviations of up to 3 times the standard deviation

of the error were expected in the majority of cases. Therefore the overall displacement of each electrode was normally less than or equal to

- $\sqrt{3(3 \cdot 0.25)^2} \approx 1.3$ mm for a standard deviation of 0.25 mm;
- $\sqrt{3(3 \cdot 0.5)^2} \approx 2.6$ mm for a standard deviation of 0.5 mm.

Tissue Spectra

Prior spectral information is imprecise because tissue conductivity is influenced by anisotropy, inhomogeneity and temperature. The resulting effect of these factors is difficult to predict. Therefore, the frequency-dependent contribution of those factors were simulated for a reasonable and a "worst-case" level of spectral errors, 1% and 5% (Edd et al., 2005). The errors were added to the tissue conductivities ϵ_{ij} independently for each frequency i and each tissue type j . It is important to note here that the multi-frequency reconstruction algorithm used in this study is insensitive to conductivity changes with a flat frequency-spectrum. Therefore, only frequency-dependent errors had to be considered, which constitute a small fraction of the above mentioned error sources.

Contact Impedances

Contact impedance errors were chosen to match experiments on human skin. It was assumed that all electrodes had sufficiently low contact impedance. In an experimental setup, this is equivalent to discarding any electrodes with near-infinite impedance, which may have detached from the head. The two levels of contact impedance errors had standard deviation 20%, corresponding to suboptimal measurement setups, and 50%, in which case the electrodes would normally be re-applied in experiments. These two high levels of error were chosen to illustrate that the contact impedance mostly influences the instrumentation, rather than the modelling and image reconstruction. This is because a high transfer impedance introduces higher instrumentation noise (McEwan et al., 2007).

6.3 Results

6.3.1 Mesh Discretisation

The voltages were simulated on the fine mesh, noise was added to the data, and then the images were reconstructed on the coarse mesh. For comparison, the process was repeated using data simulated on the coarse mesh (figure 6.2). The discretisation introduced average electrode area

6. Feasibility of Stroke Imaging with MFEIT

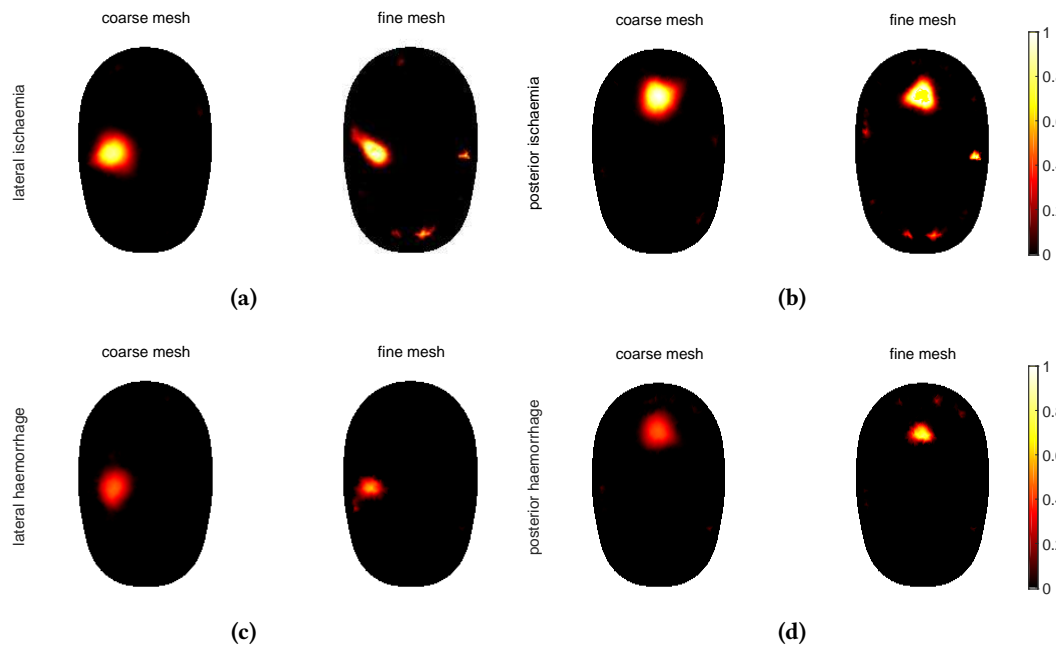


Figure 6.2: Effect of mesh discretisation on image reconstructions - Images reconstructed on the coarse mesh from data simulated on the coarse and the fine mesh for (a) lateral ischaemic stroke, (b) posterior ischaemic stroke, (c) lateral haemorrhagic stroke and (d) posterior haemorrhagic stroke.

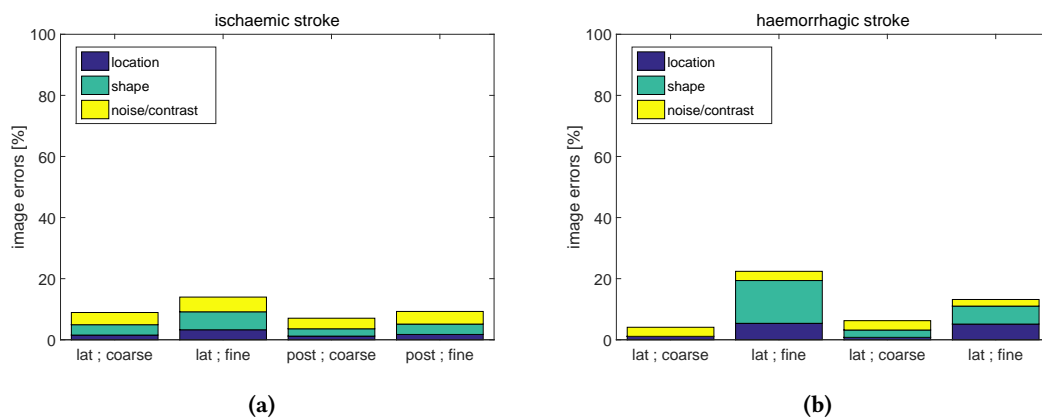


Figure 6.3: Quantification of image errors caused by mesh discretisation - Image errors in reconstructions on the coarse mesh, for (a) ischaemic strokes and (b) haemorrhagic strokes simulated on the coarse and the fine mesh.




Figure 6.4: Electrode areas on the fine and the coarse mesh - The voltages were simulated on the fine 5 million element mesh and the images were reconstructed on the coarse 180 thousand element mesh.

errors of 0.56 mm^2 , corresponding to $\sim 1.4\%$ of the average electrode area (figure 6.4). Image error measures for the reconstructions indicated, that images from data simulated and reconstructed on the same mesh were superior to those obtained from data simulated on the fine mesh (figure 6.3). The reconstructions of ischaemic strokes had a larger contrast than for haemorrhagic stroke and images obtained for the posterior position were usually better than for the lateral position.

6.3.2 Erroneous Electrode Positions

The images were reconstructed assuming that the electrodes were at the original positions (figure 6.5) and image error measures were computed for each reconstruction (figure 6.6). Imaging was not even possible with small modelling errors of 0.25 mm electrode position mismatch.

6.3.3 Erroneous Tissue Conductivity Spectra

Errors were added to the conductivity values of the model to simulate a deviation from the literature values. The images were reconstructed using the original values for the conductivities of the brain and stroke (figure 6.7), and error measures were computed for each image (figure 6.8). The perturbation was recovered successfully in all cases with 1% spectral errors, but for 5% spectral errors imaging was unsuccessful. Plotting the frequency-difference spectra for brain, ischaemic brain and blood together with the associated error bars (figure 6.9) illustrates the degree of the used spectral errors. The error bars indicate the minimum and maximum limit within which the majority of the errors were drawn, given by ± 3 times the standard deviation.

6. Feasibility of Stroke Imaging with MFEIT

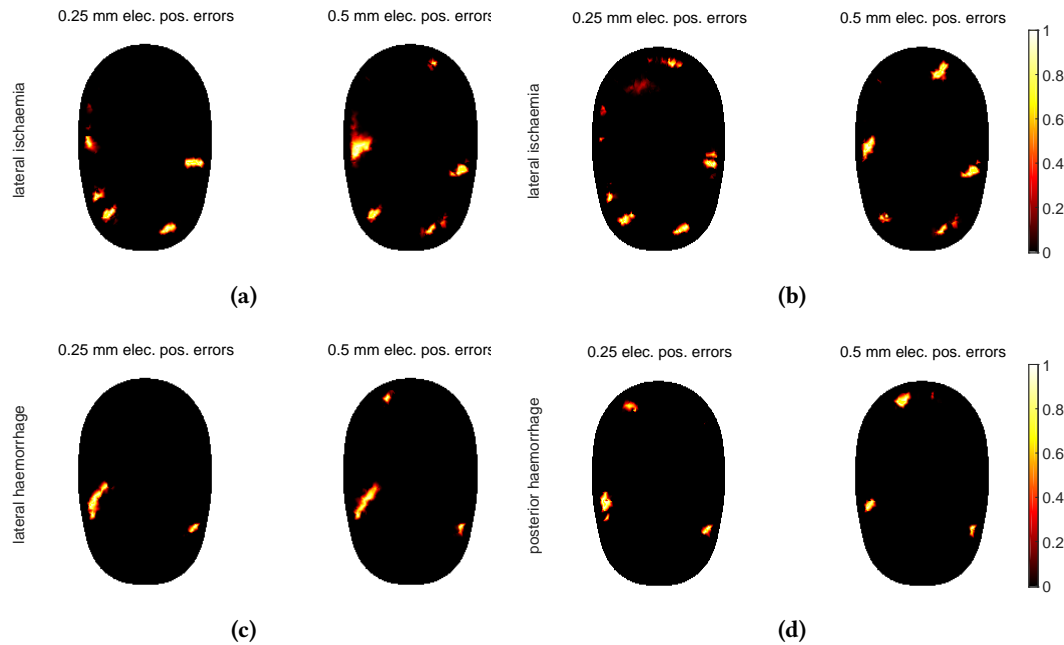


Figure 6.5: Image reconstructions with inaccurately modelled electrode positions - Image reconstructions from voltages simulated with errors of 0.25 mm and 0.5 mm standard deviation added to all three coordinates of the electrode position. (a) Lateral ischaemic stroke, (b) posterior ischaemic stroke, (c) lateral haemorrhagic stroke and (d) posterior haemorrhagic stroke.

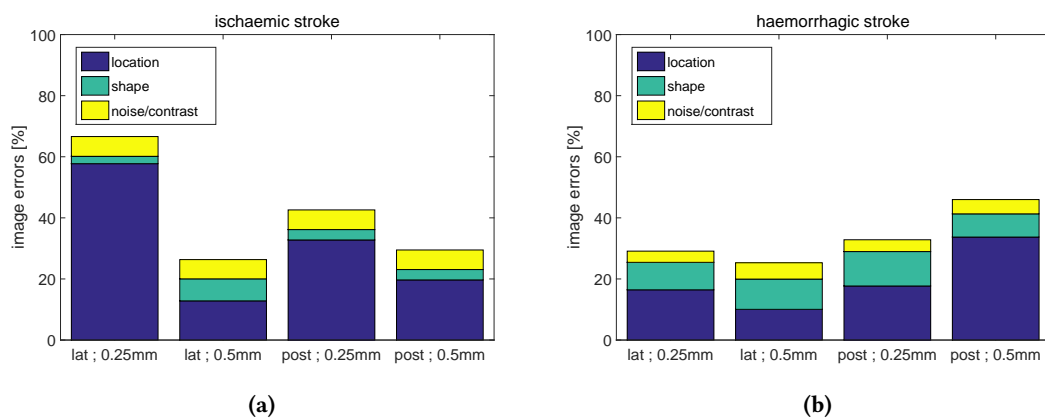


Figure 6.6: Quantification of image errors caused by inaccurate electrode modelling - Positional errors of 0.25 mm and 0.5 mm standard deviation were added to all three coordinates of the electrode position and voltages were simulated with two positions for (a) ischaemic strokes and (b) haemorrhagic strokes.

6. Feasibility of Stroke Imaging with MFEIT

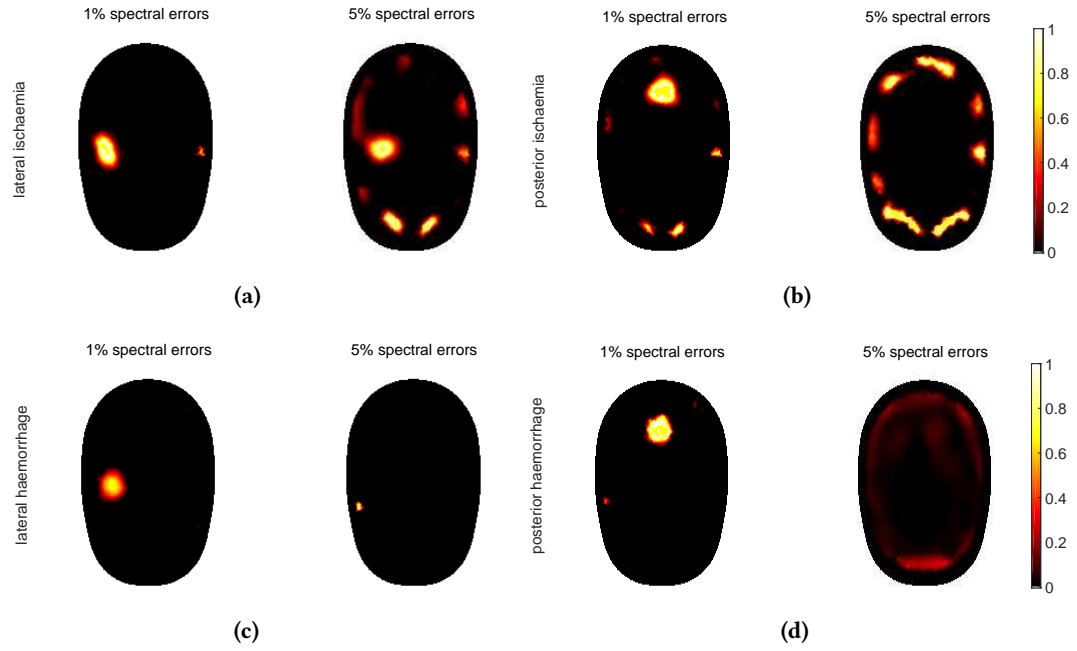


Figure 6.7: Image reconstructions with erroneous tissue spectra - Images were reconstructed from voltages simulated with 1% and 5% errors on the tissue conductivities for (a) lateral ischaemic stroke, (b) posterior ischaemic stroke, (c) lateral haemorrhagic stroke and (d) posterior haemorrhagic stroke.

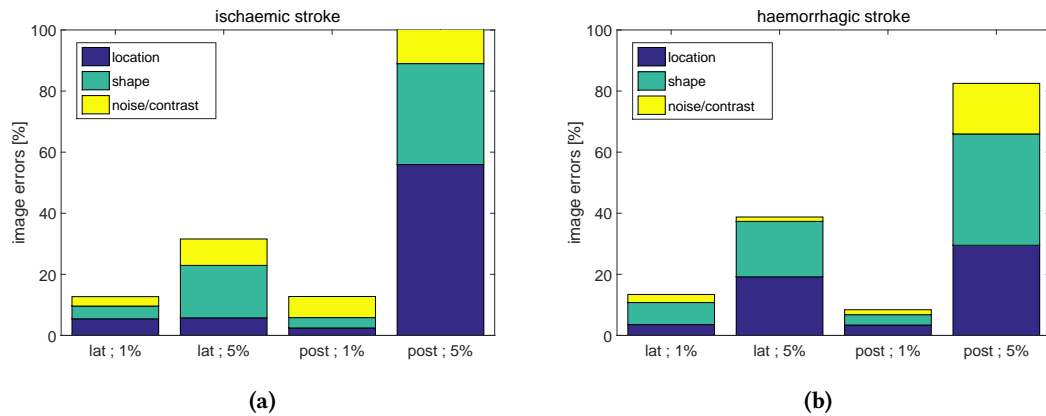


Figure 6.8: Quantification of image errors caused by spectral errors - Errors on the conductivity spectra of the tissues of standard deviation 1% and 5% were simulated for (a) ischaemic strokes and (b) haemorrhagic strokes.

6. Feasibility of Stroke Imaging with MFEIT




Figure 6.9: Conductivity differences and noise with respect to the lowest frequency - For each tissue the conductivity differences with respect to the lowest frequency are shown with the associated error bars for (a) 1% and (b) 5% errors added to the absolute spectra. The error bars represent the minimum and maximum limits within which the errors on the relative spectra were drawn. The errors were added to the absolute values of the conductivity, therefore the variance was proportional to the conductivity value.

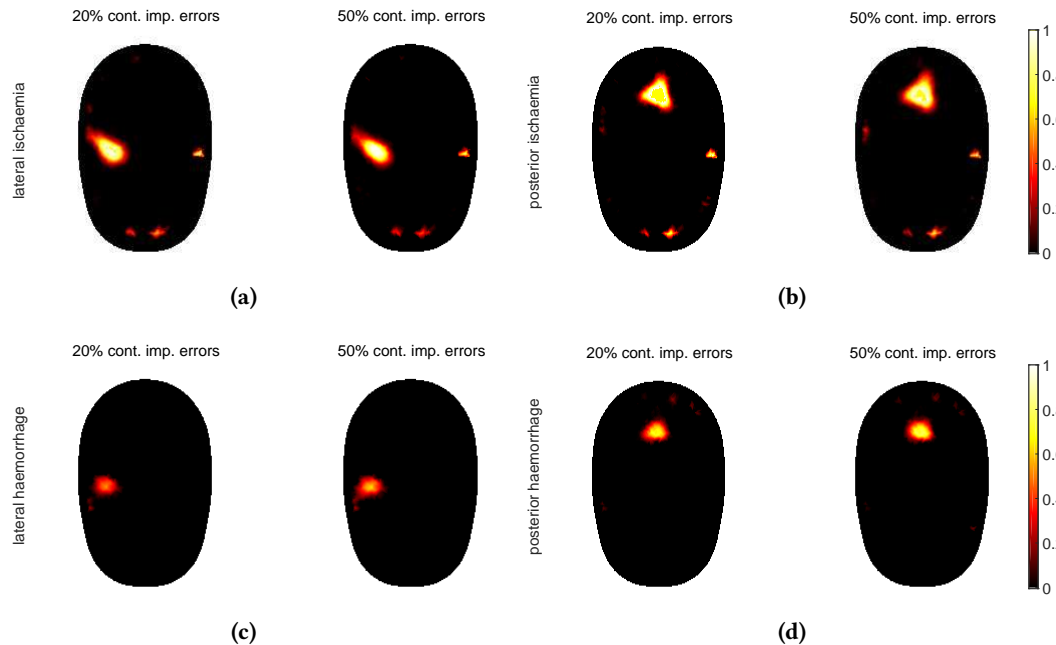


Figure 6.10: Image reconstructions with inaccurately modelled contact impedances - Voltages were simulated with errors of 20% and 50% standard deviation added to the electrode contact impedances and reconstructions were made with the original values. (a) Lateral ischaemic stroke, (b) posterior ischaemic stroke, (c) lateral haemorrhagic stroke and (d) posterior haemorrhagic stroke.

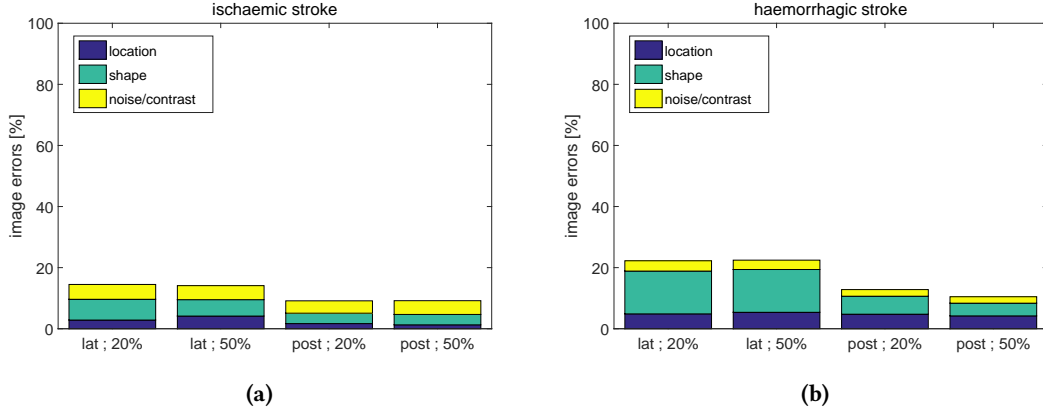


Figure 6.11: Quantification of image errors caused by contact impedance errors - Errors of 20% and 50% standard deviation were added to the electrode contact impedances for (a) ischaemic strokes and (b) haemorrhagic strokes.

6.3.4 Erroneous Electrode Contact Impedances

Images were reconstructed assuming a value of $1 \text{ k}\Omega \cdot |E|$ for the contact impedances of the electrodes (figure 6.10), and error measures were computed for each image (figure 6.11). The images were nearly unchanged by the introduction of up to 50% errors on the contact impedances.

6.4 Discussion

6.4.1 Mesh Discretisation

The reduction in image quality between the case of data simulated on the coarse and fine mesh was primarily caused by the modelling of the electrodes and skull. Given the different resolutions, the shape and size of the electrodes and the skull differed between the two meshes. The purpose of using a fine mesh to simulate the data, and a coarse mesh to reconstruct the image, was to make the simulation more realistic. In the case of imaging a real human head, the size and thickness of the skull is not usually known accurately. Furthermore, the modelling of the electrodes on the mesh is rarely precise in reality. Therefore it was necessary to consider these discrepancies for realistic simulations.

6.4.2 Erroneous Electrode Positions

Errors added to the electrode positions severely affected the image quality. This highlights the importance of registering the positions of the electrodes accurately. Simultaneous recovery of fraction values and electrode positions (as shown in chapter 5 for time-difference imaging), could further improve stability of MFEIT in the presence of electrode modelling errors.

6. Feasibility of Stroke Imaging with MFEIT

6.4.3 Erroneous Tissue Conductivity Spectra

The fraction reconstruction method requires knowledge of the conductivity spectra of all tissues, and these were assumed to be fixed and exact. The performance of the algorithm was therefore diminished, if the used spectral information was incorrect (figure 6.8). The tissues were distinguished on the basis of the respective change in conductivity between the lowest and the other frequencies. If a random error was added to the absolute spectrum, then the error on the difference in the spectrum with respect to the lowest frequency was given by the sum of the absolute errors. For 1% error, all the spectra were distinct, but for 5% error, the spectra overlapped for some or all frequencies (figure 6.9). For this reason it was not possible to locate the haemorrhagic stroke in the case of 5% error added to the conductivities (figure 6.7), and the ischaemic stroke was only identified in the lateral position (figure 6.7a). For haemorrhagic strokes, the addition of a proportional 5% error caused a large degree of uncertainty because the absolute value of the conductivity of blood was large.

6.4.4 Erroneous Electrode Contact Impedances

The effect of the errors added to the contact impedance values was very small, because four electrode measurements were performed (McEwan et al., 2007). If voltages were measured on electrodes injecting current, much larger errors would be expected (Kolehmainen et al., 1997).

6.4.5 Technical Remarks

Ideally, several images would have been created for each noise level in order to characterise the effect of modelling errors over a large number of samples. The computational expense of multiple repetitions was prohibitive, in that reconstruction of a single image took 5-6 hours. However, given that the electrode specific errors (contact impedance and position) were sampled on the 32 electrodes individually, this provides a sufficiently large number of samples to give a reasonable characterization of the influence of the noise. Likewise, errors on the tissue spectra were added independently to each tissue and at each frequency, thereby describing the effect of spectral errors reasonably well. The conclusions derived from this relatively small number of images appear to be valid in principle. In future, examination of more permutations in simulation and tank studies may allow the definition of quantitative limits for the acceptable variation of each parameter.

6.4.6 Conclusion

The fraction reconstruction method using spectral constraints was applied to a numerical head phantom with realistic conductivities for skull, scalp and brain tissue. Instrumentation noise and

modelling errors were added to investigate the robustness of this MFEIT method. The results showed a varying degree of sensitivity to different sources of modelling errors.

- Mesh discretisation lead to image artefacts near the skull and electrodes.
- The method was highly sensitive to errors in the position of the electrodes.
- If the respective spectra were sufficiently distinct, tissues could be distinguished even in the presence of spectral inaccuracies.
- The method was highly robust to errors in the modelled contact impedances.

Further work is required to improve the image quality in the presence of modelling errors. The artefacts caused by the discretisation of the skull and electrodes in the coarse reconstruction mesh, may be reduced by simultaneously reconstructing the brain and the skull. Using a finer mesh for reconstructions would further allow for a more accurate modelling of the skull and electrodes, at the cost of longer computation times. Inaccurate modelling of the electrodes had the strongest impact on image quality. To reduce the sensitivity to electrode position errors, a simultaneous correction of the modelled electrode positions should be included into the fraction reconstruction method.

Chapter 7

Correction of Electrode Modelling Errors in Multi-Frequency Tissue Fraction Reconstruction

7.1 Overview

7.1.1 Introduction

In the previous chapter, the feasibility of imaging stroke types with multi-frequency EIT (MFEIT) imaging was evaluated. The influence of common modelling errors — mesh discretisation, electrode positions, tissue conductivity spectra and contact impedances — on the image quality was assessed, and the accurate modelling of electrode positions was found to be very important. Electrode positions can currently be measured to around 1 mm precision using photogrammetry (Qian and Sheng, 2011). Other technologies, such as the commercial MicroScribe, laser 3D scanners, or electrode helmets (Avery, 2015), can achieve an even higher precision in electrode localisation.

While electrode position modelling errors in this range were found not to affect weighted frequency difference imaging (Ahn et al., 2011), this linear method cannot be used for stroke type differentiation because of the inhomogeneous head geometry (Malone et al., 2014). The non-linear tissue fraction reconstruction algorithm used in chapter 6, was already strongly affected by electrode position errors below 1 mm.

7. Electrode Model Correction in MFEIT

7.1.2 Background

MFEIT is a relatively new field, and not much work has been done on making multi-frequency reconstructions more resilient to modelling errors. The linear weighted-frequency difference algorithm (Jun et al., 2009) is inherently stable to modelling errors (Ahn et al., 2011). A non-linear MF reconstruction algorithm was recently proposed (Malone et al., 2015), which is more stable with respect to spectral errors when compared to the fraction reconstruction method. This, however, comes at the cost of having a second parameter — apart from the regularisation factor —, which has to be optimized. Furthermore, it has been shown in chapter 6, that tissues are well differentiated by the fraction reconstruction algorithm (Malone et al., 2014) as long as their conductivity spectra do not overlap.

The correction for electrode modelling errors has been used in time-difference (Soleimani et al., 2006) and absolute imaging (Dardé et al., 2012) on cylindrical domains and on realistic head shapes (chapter 5). It was shown to be very stable, both for linear and iterative reconstruction algorithms, and could therefore be beneficial to MF reconstructions in the presence of electrode modelling inaccuracies.

7.1.3 Purpose

The purpose of this study was to evaluate the performance of simultaneous recovery of electrode positions and conductivity spectrum changes. Specifically, the following two questions were addressed: 1) Does the simultaneous recovery of conductivity changes and electrode modelling errors remove image artefacts caused by inaccurately modelled electrode positions? 2) At which magnitude of electrode position errors does the proposed algorithm begin to fail?

To answer these questions, multi-frequency boundary voltages were simulated on a fine 5 million element head mesh with different levels of electrode position errors. Reconstructions were made with and without the proposed addition of electrode modelling corrections and the resulting images were compared. It was found, that the proposed algorithm could stably recover simulated strokes in the presence of electrode modelling errors of up to 1.5 mm standard deviation.

7.2 Methods

7.2.1 Tissue Fraction Reconstruction With Electrode Position Correction

The general structure of the used multi-frequency reconstruction algorithm was identical to the one described in section 2.3.4 (originally published by Malone et al., 2014). To be able to use conductivity measurements at different frequencies $\omega_i, i = 1, \dots, W$ simultaneously in a single

image reconstruction, the conductivity spectrum of one element in the mesh was described as a linear combination of the known spectra of present tissues $t_j, j = 1, \dots, T$. Instead of reconstructing conductivities directly, the prior knowledge of the conductivity spectrum ϵ_{ij} of all tissues was therefore used to assign fractions f_{nj} of these tissues to all finite elements $n = 1, \dots, N$, such that

$$\sigma_n(\omega_i) = \sum_{j=1}^T f_{nj} \cdot \epsilon_{ij}, \quad (7.1)$$

where $0 \leq f_{nj} \leq 1$ and $\sum_{j=1}^T f_{nj} = 1$. The modified Jacobian matrix at each frequency was obtained using the chain rule

$$\frac{\partial \mathbf{F}(\boldsymbol{\sigma}_i)}{\partial \mathbf{f}_j} = \frac{\partial \mathbf{F}}{\partial \boldsymbol{\sigma}_i} \frac{\partial \boldsymbol{\sigma}_i}{\partial \mathbf{f}_j} = \frac{\partial \mathbf{F}}{\partial \boldsymbol{\sigma}_i} \epsilon_{ij} = \mathbf{J}_{\boldsymbol{\sigma}_i} \cdot \epsilon_{ij} = \mathbf{J}_{ij} \in \mathbb{R}^{R \times N}, \quad (7.2)$$

where $r = 1, \dots, R$ are the lines in the protocol, i.e. the combinations of different current injection and voltage measurement electrode pairs d and m . The Jacobian matrix $\mathbf{J}_{\boldsymbol{\sigma}_i}$ relating voltage changes to changes in conductivity $\boldsymbol{\sigma}_i$ at each frequency ω_i , was computed using the adjoint fields method (section 2.2.7), giving one entry for element n as

$$\mathbf{J}_{rn} = - \int_{\Omega} \nabla u^d \cdot \nabla u^a dV, \quad (7.3)$$

where $u^d \in H^1(\Omega)$ is the electric potential emerging when the drive current I^d is applied to the electrodes and $u^a \in H^1(\Omega)$ the electric potential when a unit current is applied to the two measurement electrodes.

In order to correct for wrongly modelled electrode positions, this ‘traditional’ Jacobian matrix was augmented by an electrode boundary Jacobian $\mathbf{EBJ} \in \mathbb{R}^{R \times M}$, relating electrode boundary changes (in our case electrode movement) to voltage changes. Given a continuous vector field \mathbf{v} on the boundary of electrode $m = 1, \dots, M$, one entry of the \mathbf{EBJ} along this vector field was computed similarly to the conductivity Jacobian (chapter 5; Dardé et al., 2012)

$$\mathbf{EBJ}_{rm} = - \frac{1}{z_m} \int_{\partial E_m} (\mathbf{v} \cdot \mathbf{n}_{\partial E}) (U_m^d - u^d) (U_m^a - u^a) ds, \quad (7.4)$$

where z_m is the contact impedance, U_m^d and U_m^a the drive and measurement electrode potentials and $\mathbf{n}_{\partial E}$ the outward normal of the electrode boundary, tangential to the head surface. The vector field \mathbf{v} describes the studied change in the electrode boundary, e.g. to compute the \mathbf{EBJ} with respect to movement along one direction, the vector field was chosen to point homogeneously in that direction. The Jacobian matrices were combined $\boldsymbol{\Upsilon}_i = [\mathbf{J}_{i1}, \dots, \mathbf{J}_{iT}, \mathbf{EBJ}_i]$ and the unknown electrode position errors \mathbf{p} were appended to the vector of the tissue fractions to be recovered $\mathbf{x} = [\mathbf{f}, \mathbf{p}]^\top$, where $\mathbf{f} = [\mathbf{f}_1, \dots, \mathbf{f}_T]$ and \mathbf{p} consisted of two variables per electrode that described electrode movements along both surface directions x_s and y_s , $\mathbf{p} =$

7. Electrode Model Correction in MFEIT

$[p_{x_s}^1, p_{y_s}^1, p_{x_s}^2, p_{y_s}^2, \dots]$. The regularised objective function to be minimised was then

$$\Phi(\mathbf{x}) = \sum_{i=2}^W \left\| (\Upsilon_i - \Upsilon_1) \mathbf{x} - (\mathbf{v}_{\omega_i} - \mathbf{v}_{\omega_1}) \right\|^2 + \lambda^2 \Psi(\mathbf{x}), \quad (7.5)$$

with \mathbf{v}_{ω_i} being the boundary voltages measured at frequency ω_i and the regularisation term

$$\Psi(\mathbf{x}) = \mathbf{x}^\top \Sigma_x^\top \mathbf{D}^\top \mathbf{D} \Sigma_x \mathbf{x}. \quad (7.6)$$

The regularisation matrix \mathbf{D} was composed of one Laplacian matrix per recovered tissue and one identity matrix for the electrode movement variables. All components were scaled according to the expected standard deviation of the corresponding variable changes $\text{std}_f = 0.01$ and $\text{std}_p = 1 \text{ mm}$

$$\mathbf{D} = \begin{bmatrix} \mathbf{L} & 0 & 0 & 0 \\ 0 & \ddots & 0 & 0 \\ 0 & 0 & \mathbf{L} & 0 \\ 0 & 0 & 0 & \mathbf{I} \end{bmatrix}; \quad \Sigma_x = \begin{bmatrix} \text{std}_f^{-1} \cdot \mathbf{I} & 0 & 0 & 0 \\ 0 & \ddots & 0 & 0 \\ 0 & 0 & \text{std}_f^{-1} \cdot \mathbf{I} & 0 \\ 0 & 0 & 0 & \text{std}_p^{-1} \cdot \mathbf{I} \end{bmatrix}. \quad (7.7)$$

The minimisation of the objective function (7.5) was performed by alternating steps of gradient projection and damped Gauss-Newton algorithms (section 2.3.4). After two iterations of gradient projection and Gauss-Newton, the electrode positions had normally converged and were subsequently kept fixed for the remaining iterations. The number of iterations of this reconstruction algorithm was fixed to ten for all image reconstructions. To avoid the inverse crime (Lionheart et al., 2004) and speed up image reconstruction, all reconstruction were made on a coarse 180 thousand element mesh on which the skull and scalp were kept fixed and it was assumed the inside of the skull was occupied by either the brain or the stroke with the initial guess being the healthy brain. The regularisation parameter of $\lambda^2 = 8 \cdot 10^{-10}$ was chosen empirically and was the same for all reconstructed images. After each iteration of gradient projection and Gauss-Newton minimisation, the regularisation parameter was halved in the case of ischaemias, and divided by three for haemorrhages.

7.2.2 Data Simulation

The simulation parameters in this study were identical to the ones used in the preliminary feasibility study (chapter 6). Boundary voltages were computed on a fine 5 million element mesh, which was created from a CT scan of a human head and included three homogeneous tissues: brain, skull and scalp. Computations of the boundary voltages were done with PERTS (chapter 4) on all 16 processors of a workstation with two 2.4 GHz Intel Xeon CPUs with eight cores and

20 MB cache each. It took less than 2 minutes to compute the required 31 forward solutions for each frequency. 32 electrodes of diameter 10 mm were placed on the surface of the model in the same positions used to acquire EEG measurements (Tidswell et al., 2001b). The electrodes were modelled using the complete electrode model (CEM), and the contact impedance was set to $1 \text{ k}\Omega \cdot |E|$ for all electrodes, where $|E|$ was the electrode area. The amplitude of the current was set to $140 \mu\text{A}$ and twelve frequencies between 5 Hz and 5 kHz were used (figure 6.1a). 31 linearly independent current injection pairs were created by finding the maximum spanning tree of the electrode positions. Voltage measurements were made for each injection on all adjacent pairs not involved in delivering current. The total number of measurements acquired for each frequency was 869.

Strokes were simulated by changing the conductivities of all elements within a 1.5 cm radius of the stroke location. Simulated locations were set in a posterior (figure 6.1c) and lateral position in the head (figure 6.1b), and stroke conductivities were set to the spectral values of ischaemic tissue (figure 6.1a) or to the conductivity of blood, 0.697 S m^{-1} , for haemorrhage (Horesh et al., 2006). Both, proportional and additive noise, was added to all simulated voltages:

$$V_{\text{with noise}} = V_{\text{no noise}} (1 + \text{rand}(\varsigma_p)) + \text{rand}(\varsigma_a), \quad (7.8)$$

where $\text{rand}(\varsigma)$ indicates a random number drawn from a Gaussian distribution with zero mean and standard deviation ς . The standard deviation of the proportional noise was $\varsigma_p = 0.02\%$ and the standard deviation of the additive noise was $\varsigma_a = 5 \mu\text{V}$, which correspond to human experiments (Goren et al., 2015).

Electrode position errors were created by drawing two random numbers for each electrode from Gaussian distributions with zero mean and standard deviations 0.5 mm, 1 mm, 1.5 mm and 2 mm. According to these drawn random numbers, the electrodes were then moved along a two dimensional surface coordinate system (chapter 5). Therefore, the overall electrode position errors were the combined errors of the displacement along the two surface dimensions. The electrode movement was implemented by assigning different surface facets to the electrodes, such that the mesh did not have to be changed.

7.2.3 Image Error Quantification

The quality of an image was objectively quantified using the same metrics than in the previous chapter. First, the images reconstructed on the 180 thousand element mesh were averaged onto cubic voxels with 0.5 cm sides. The volume P corresponding to the reconstructed perturbation was identified as the largest connected cluster of voxels with values larger than 50% of the maximum of the image. Three measures of image errors were defined:

1. Location error: ratio between the distance $\|(x_P, y_P, z_P)\|$ of the centre of mass of the

7. Electrode Model Correction in MFEIT

reconstructed perturbation P from the actual position, and the average dimension of the head $\text{mean}(d_x, d_y, d_z)$

$$\frac{\|(x_P, y_P, z_P)\|}{\text{mean}(d_x, d_y, d_z)}. \quad (7.9)$$

2. Shape error: ratio of the difference between the dimensions of the simulated (s_x, s_y, s_z) and reconstructed perturbation (r_x, r_y, r_z) , and the dimensions of the simulated perturbation

$$\frac{\|(r_x - s_x, r_y - s_y, r_z - s_z)\|}{\|(s_x, s_y, s_z)\|}. \quad (7.10)$$

3. Image noise: inverse of the contrast-to-noise ratio between the perturbation P and the background B

$$\frac{\text{std}(f_s^B)}{|\bar{f}_s^P - \bar{f}_s^B|}, \quad (7.11)$$

where \bar{f}_s^P and \bar{f}_s^B are the mean intensities of the perturbation and background and std the standard deviation.

7.3 Results

7.3.1 Multi-Frequency Tissue Fraction Reconstructions

When electrode positions were modelled accurately in the reconstruction algorithm, image reconstructions with electrode modelling correction were slightly worse than without correction (figure 7.1a). This was expected, since the reconstruction algorithm simply aims to minimise the value of the objective value, with no knowledge whether electrodes or conductivities caused the voltage differences. Consequently, it tried to minimise some voltage differences caused by conductivity changes by moving the electrodes away from the correct position. The average image error of reconstructions without electrode correction was 10% and with correction 17%; particularly the reconstruction of the lateral haemorrhage was not good (24%).

With the traditional fraction reconstruction MFEIT algorithm, already 0.5 mm of electrode position modelling errors made stroke detection impossible (figure 7.2a). The average image error of the reconstructions without electrode modelling correction was 55% (figure 7.2b).

Simultaneous recovery of stroke tissue fractions and electrode positions significantly improved image quality in the presence of electrode modelling errors (figure 7.3a). The average image error of the reconstructed images with correction was 23% (figure 7.3b), excluding the three outliers (numbers 5, 10 and 14) only 16%. Remarkably, all three bad reconstructions occurred for ischaemic strokes, suggesting that local conductivity spectrum changes caused by ischaemia were more difficult to separate from electrode position errors than for haemorrhage.

The shape error of the reconstruction of a lateral ischaemia with electrode movement of 2 mm (number 13) is misleadingly small, because the recovered perturbation was a diagonal disc with very similar x-y-z dimensions than the simulated stroke.

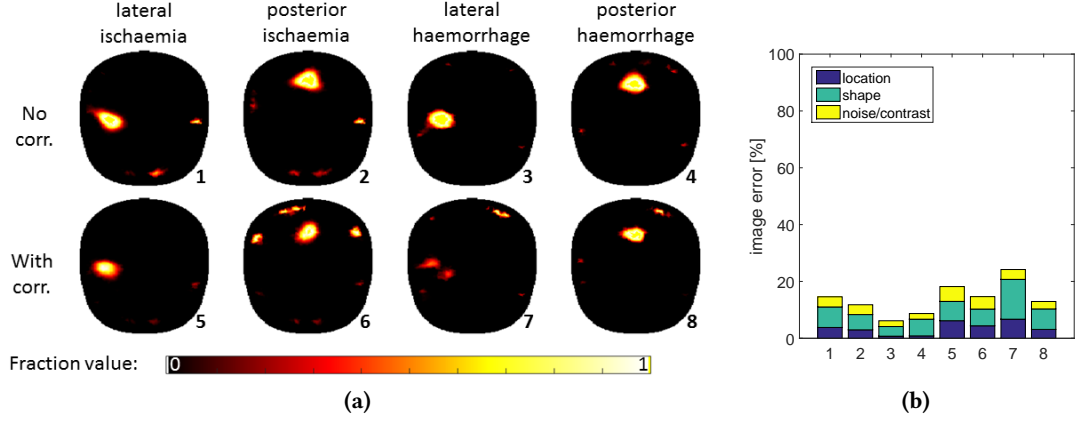


Figure 7.1: Image reconstructions without electrode modelling errors - (a) Multi-frequency fraction reconstructions of strokes without and with electrode modelling correction when electrodes were modelled accurately and (b) the corresponding image error measures.

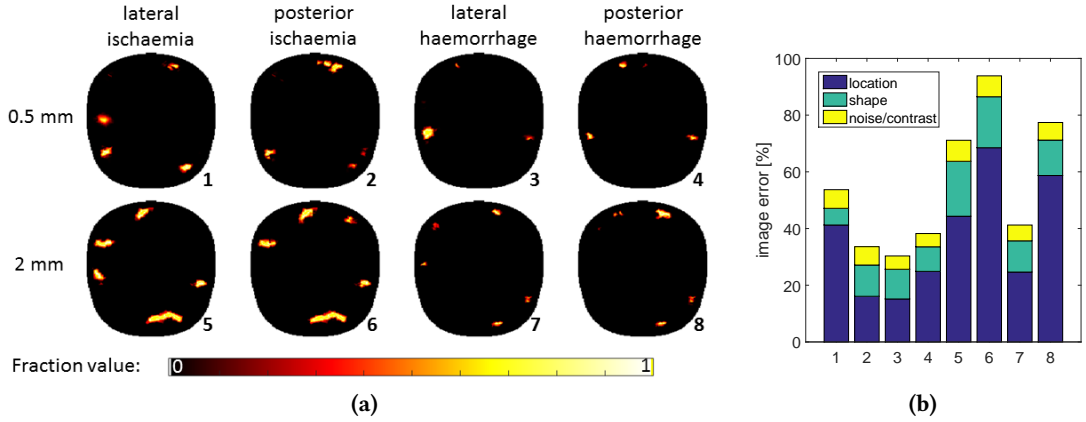


Figure 7.2: Traditional MF fraction reconstructions with electrode modelling errors - (a) Multi-frequency fraction reconstructions of strokes without electrode modelling correction for two different levels of electrode position errors and (b) the corresponding image error measures.

7.3.2 Electrode Placement Correction

The 2-norm of the difference between recovered and simulated electrode position mismatch was computed. While the electrode position correction was more accurate for ischaemic strokes when electrodes were correctly modelled, for position mismatch of 1 mm standard deviation the correction was better in the presence of haemorrhages (table 7.1). The accuracy of the electrode position recovery tended to correlate with the quality of the reconstructed image (figures 7.1a

7. Electrode Model Correction in MFEIT

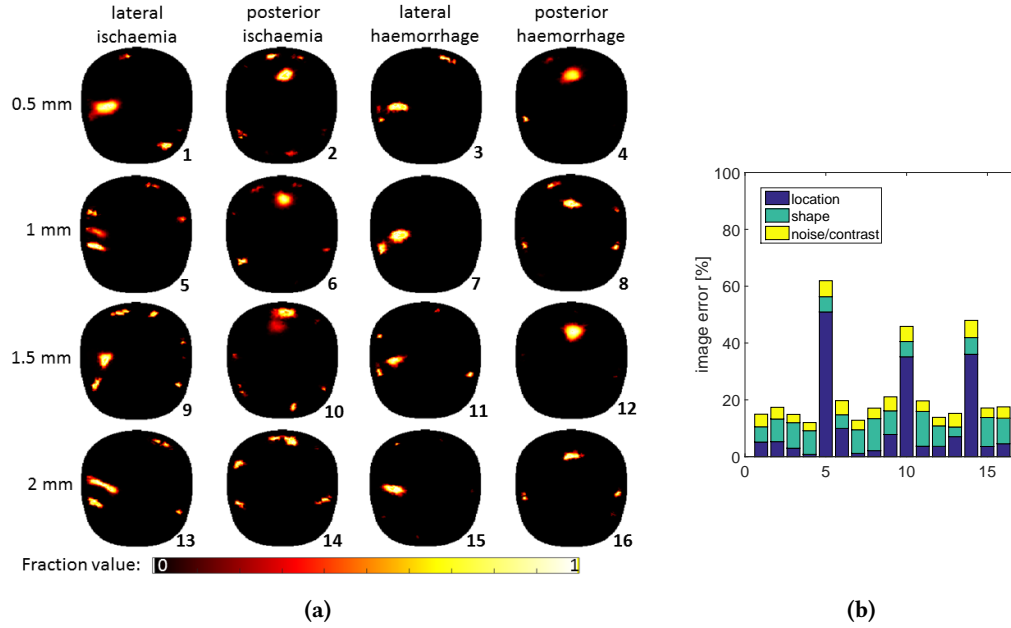


Figure 7.3: Novel MF fraction reconstructions correcting for electrode modelling errors - (a) Multi-frequency fraction reconstructions of strokes using electrode correction for four different levels of electrode position modelling errors and (b) the corresponding image error measures.

and 7.3a). This is intuitive and was already observed for time-difference electrode movement corrections (chapter 5).

	ischaemia		haemorrhage	
	lateral	posterior	lateral	posterior
0 mm	$0.8 \cdot 10^{-3}$	$1.6 \cdot 10^{-3}$	$1.7 \cdot 10^{-3}$	$3.7 \cdot 10^{-3}$
1 mm	$13.0 \cdot 10^{-3}$	$11.2 \cdot 10^{-3}$	$7.1 \cdot 10^{-3}$	$5.8 \cdot 10^{-3}$

Table 7.1: Accuracy of electrode position recovery - 2-norm of the difference in simulated and recovered electrode position errors, when electrodes were accurately modelled in the reconstruction algorithm (first row) and when there was a position mismatch of 1 mm standard deviation along both surface dimensions (second row).

The reason for the worse 2-norm of ischaemic electrode position correction of 1 mm errors, was electrode 17 (entries 33 and 34 on x-axis of figure 7.4). Interestingly, after one iteration of the proposed algorithm, the electrode position recovery of this electrode was still accurate. Only in the second iteration, the electrode was moved several millimetres in both surface dimensions. Since electrode 17 was located on the side, 5.4 cm from the centre of the lateral ischaemia, this affected the reconstruction of the lateral ischaemia more than the reconstruction of the posterior ischaemia (figure 7.3a).




Figure 7.4: Recovered electrode position errors for different stroke types and positions - Recovery of electrode position modelling errors for different stroke types and positions, when electrodes were simulated with standard deviation of 1 mm positional errors (black dashed line). Along the x-axis are the movement components along both surface dimensions for each electrode, i.e. 1&2 are x_s and y_s of electrode 1, 3&4 for electrode 2 and so on.

7.4 Discussion

7.4.1 Electrode Modelling Correction in Multi-Frequency Reconstruction

The simultaneous recovery of stroke tissue fractions and electrode positions significantly improved image quality in the presence of electrode modelling errors. For movements between 0.5 mm and 2 mm standard deviation along both surface dimensions, the average image error was 23% compared to 55% without electrode correction. While reconstructions of haemorrhagic strokes were even successful in the presence of 2 mm of electrode errors, ischaemia detection was less reliable from 1 mm onwards.

7.4.2 Technical Remarks

Ideally, several images would have been created for each electrode movement level, in order to characterise the effect of electrode modelling errors over a number of samples. However, the computational expense of multiple repetitions was prohibitive, in that reconstruction of a single image took around 6 hours. For the same reason, only two stroke positions were studied. Nonetheless, the produced images clearly illustrate the advantage of simultaneous tissue fraction and electrode position recovery.

7.4.3 Conclusion

Simultaneous iterative electrode position correction with the fraction reconstruction method using spectral constraints was applied to a numerical head phantom with realistic conductivities. Realistic noise was added to the simulated voltages to investigate the robustness of the proposed method. The results show that

1. the simultaneous recovery of tissue volume fractions and electrode position errors removed most image artefacts caused by inaccurately modelled electrodes.
2. while haemorrhagic strokes could be reconstructed with electrode position errors up to 2 mm standard deviation in both surface dimensions, the reconstruction of ischaemic strokes was less reliable from electrode movements of 1 mm onwards.

Further work is required to understand why ischaemic stroke reconstructions were less reliable with the proposed method, and to correct for it. Additionally, it has so far not been studied how stable non-linear multi-frequency reconstruction methods are in the presence of geometric modelling errors, such as skull shape (chapter 3). The mesh used for this simulation study had the same geometry than the 3D printed head shaped tank used in chapter 5, such that the presented results can be verified in tank experiments.

Chapter 8

Conclusion

8.1 Summary

In chapter 3, a new method for creating head meshes was proposed. The new meshing procedure enabled the creation of head models with more physiological accuracy than previously achieved. Once the manual segmentation was done, meshes with different element density — globally, as well as locally — could easily be created. Four meshes of different heads were then used to study, whether subject specific head models are required for time-difference imaging of stroke. While it was concluded that the image quality profited from accurate models, the detection of localised conductivity changes in the range of 100% and 10% was not improved.

To address the computational requirements concomitant with the use of accurate head models, a fast and memory efficient parallel solver for the forward problem (PEITS) was implemented in chapter 4. With PEITS, forward solutions, the conventional conductivity Jacobian matrix and the Jacobian matrix with respect to electrode movement could be computed on many processors in parallel, thereby increasing the speed of computations up to 60 times compared to the currently most prevalent EIDORS forward solver.

In chapter 5, a fast implementation for the calculation of the electrode boundary Jacobian matrix (EBJ) was used on a human head model. The implementation was based on an analytical formulation for the EBJ and was efficient due to the use of template matrices storing the contributions of drive and adjoint fields. The EBJ was shown to accurately model electrode movement, but overestimated the already small effect electrode size had on the measured voltages. Because electrode position errors were found to be the most important factor in electrode modelling, electrode movement correction was applied to time-difference and absolute imaging on a realistic head geometry. The simultaneous recovery of conductivity changes and electrode movement removed most image artefacts caused by inaccurate electrode modelling in

8. Conclusion

simulation studies and tank experiments. All absolute reconstructions from simulated voltages had a large artefact in the front of the head, which was probably caused by mesh discretisation errors. Absolute reconstructions from experimental data were unsuccessful.

The feasibility of acute stroke imaging with non-linear multi-frequency EIT (MFEIT) was evaluated in chapter 6. Measurements were simulated in the presence of different types of modelling errors: electrode positions, tissue conductivity spectra and contact impedances. Images were reconstructed from the noisy simulated voltages with the fraction reconstruction algorithm and the image quality was objectively assessed. While erroneously modelled contact impedances did not markedly affect images, spectral errors of 5% resulting in spectral overlap and electrode position errors of 0.25 mm made reliable stroke imaging impossible.

In chapter 7, the findings of the previous two chapters were combined, in order to iteratively update the electrode modelling during MFEIT reconstructions. This modified fraction reconstruction algorithm was shown to be considerably more stable in the presence of electrode modelling errors. Haemorrhages could be imaged with electrode position errors of 2 mm, while ischaemia detection deteriorated from errors of 1 mm upwards.

8.2 Limitations

The manual segmentation presented in chapter 3 is time-consuming and requires logical skills in applying the tools provided by Seg3D efficiently. Therefore, the presented meshing pipeline is far from the end goal of ‘pushing a button and getting a good quality mesh’. Whether subject specific head models are required, was only studied for the application of stroke detection in time-difference imaging. Other TD applications, where the number and shape of the perturbations are of interest, will likely have different mesh accuracy requirements. Consequently, similar studies should be performed for applications where the characteristics of the perturbation and image resolution requirements are different, such as imaging fast neural activity in the rat brain.

The fast forward solver described in chapter 4 (PEITS) is currently not straight-forward to install on Windows and OS X. Since it is mostly aimed at high parallelism on workstations or clusters, which are usually UNIX-based, this is not essential. However, the lower memory consumption compared to EIDORS would make PEITS attractive also for desktop use, and an easier installation procedure is essential in increasing the PEITS uptake. The requirements for mesh refinement can only inadequately be studied with PEITS, by creating meshes with different element sizes and studying convergence of voltages. Therefore, an implementation of adaptive mesh refinement based on local error estimates in PEITS would help obtaining more accurate mesh precision requirements.

The implementation of the electrode movement correction on realistic heads presented in chapter 5 relies on a surface coordinate system. This was created with a relatively basic

algorithm moving along the surface and pooling vertices within defined distances from each other. Parametrising a roughly spherical 3D surface with a 2D coordinate system was difficult, and the proposed system only worked reliably on the upper two thirds of the head. Below, the x_s and y_s axes started overlapping and the algorithm stopped adding new coordinate points. Consequently, a more reliable system of generating surface coordinates for various types of geometries would facilitate the use of electrode movement correction.

In chapter 6, differences between the 5 million element mesh used for voltage simulations and the 180 thousand element mesh for the image reconstruction introduced visible image artefacts. Consequently, differences between patients are likely to have a stronger impact on non-linear multi-frequency reconstructions, than on time-difference imaging (chapter 3). A thorough analysis of the influence of using a generic head model in MFEIT should be performed to validate this assumption.

The fraction reconstruction method with electrode position correction, which was used in chapter 7, was much more stable in the presence of modelling errors than the original implementation. However, particularly the imaging of ischaemia was still not fully reliable. More work is required to determine, why ischaemic strokes were more challenging to detect than haemorrhages and to find ways to account for it. Furthermore, the fraction reconstruction algorithm — with or without electrode position correction — is relatively slow, as it takes around 6 hours to create one image. This severely limits the extent of simulation studies and currently excludes diagnostic applications, where quick results are required. With a more efficient implementation including better parallel routines, the speed could reasonably be doubled. However, for further reductions in computation time, structural changes in the algorithm will be required.

8.3 Outlook

The meshing procedure presented in this thesis is time-intensive and requires mostly manual segmentation of CT and MRI scans. The resulting meshes, however, are of very good quality, unparalleled by current automatic segmentation algorithms and meshing pipelines. Even though the meshing would profit from automatic segmentation, this is currently not the most important challenge in EIT, since meshes only have to be created once and can then be used for the development of many different applications. There are still open questions as to the required mesh refinement for the forward solutions to converge, as well as the precision requirements on the forward solutions to exclude easily avoidable discretisation errors in EIT image reconstructions. Future analysis of these points will be facilitated by the fast parallel forward solver PEITS, which was presented in chapter 4.

As shown by results in this thesis, time-difference EIT imaging on patients should from

8. Conclusion

now on standardly include electrode movement correction, to exclude artefacts from inaccurate electrode modelling at the cost of minimally reduced image quality. Since the movement recovery combines naturally with all conventional EIT image reconstruction methods, its inclusion involves very little work, but adds significant stability in the presence of electrode movements. The ease of use of the movement correction can, however, be improved by finding more efficient methods for defining a surface coordinate system. TDEIT applications, where the appearance of one local conductivity change has to be detected, are further facilitated by the finding presented here, that subject specific head meshes are not necessarily required.

Multi-frequency imaging is still a relatively new EIT modality, and fast and reliable algorithms need to be developed. One step towards more stable MFEIT imaging was the inclusion of electrode model adaptivity presented in this thesis. In a next stage, the influence of using a generic head model on MFEIT image quality should be assessed. As long as subject specific head meshes are required, acute stroke diagnosis is not feasible and monitoring of TBI patients is complicated by the requirement of segmenting and meshing each head from CT and MRI scans. If it is found that generic head models are sufficient, or when MF algorithms have been developed that can cope with geometrical errors, then MFEIT can be validated in clinical studies and promises to be a valuable diagnostic tool. In the meantime, tank experiments with known geometry can be used to test systems, and algorithms with real data in a relatively controlled environment.

In applications that only require a *yes* or *no* answer, machine learning (ML) is a potentially valuable alternative to MFEIT imaging. In ML, an algorithm is trained to differentiate different groups of data by showing it many examples. Based on the presented data, it finds distinguishing features (i.e. *learns* them) that are not easily visible to a human. The two head EIT applications this thesis focused on, fall into this category of situations where a binary answer might be sufficient. In acute stroke type detection, the question would be: does the patient have an ischaemic stroke or not? Should thrombolysis be performed or not? And for the monitoring of TBI patients, ML could answer the question: did a bleeding develop?

Machine learning has two advantages over MF imaging: 1) By reducing the variables to be reconstructed from the number of elements in the mesh to one binary value, the problem becomes better (if not well) posed. 2) ML is a booming research field with many available algorithms, that can basically be used out of the box on the unprocessed multi-frequency data. The obvious downside of using ML is that an image is no longer available, and that physicians have to rely on the confidence value the algorithm computed. If an image is required to localise the stroke and correlate the location with the symptoms a patient presents, the result from the machine learning could still be used to evaluate, whether the data contain enough information to make imaging feasible.

ML has already been applied to stroke type detection with microwave tomography by Persson et al. (2014). At 99.9% sensitivity to detect haemorrhages, they diagnosed around 30% of the ischaemas. However, the microwave-based imaging technique is still at an early development

stage and the data collected on patients look very noisy. Given the achieved performance of the ML algorithm with these noisy microwave data, MFEIT machine learning has the potential to expedite diagnosis and treatment for many patients with ischaemic stroke.

8.4 Guidelines

Based on the findings in this thesis, practical guidelines for EIT head imaging are briefly summarised in this section.

For simulation studies, voltages should be simulated on a fine head mesh to exclude discretisation errors. On 4 million element head meshes, the discretisation errors were estimated to be approximately half the typical instrumentation noise in tank experiments. This would suggest, that a 3-5 million element head mesh should be sufficient for most simulation studies. The forward solutions on a 4 million element mesh can be efficiently computed on PEITS, and then realistic instrumentation noise should be added: around 0.006% proportional and 1 μV additive noise for tank experiments, 0.01% and 2 μV for human measurements with low noise or 0.02% and 5 μV for higher noise.

Time-difference reconstructions from simulated and tank data do not require electrode movement correction, unless the influence of electrode movement is studied. Reconstructions from human experiments should be done with electrode movement correction, since it reduces noise caused by patient movement and electrode drift. If the aim of the TD reconstructions is to detect one localised conductivity change, then a subject specific head model is not required. If, however, the image quality is of utmost importance, then an accurate head model is recommended.

Multi-frequency imaging with the fraction reconstruction algorithm requires more accurate head models than time-difference imaging. At the same time, the reconstruction mesh should not be too fine, because a single image reconstruction already takes several hours on a 200 thousand element mesh. A mesh of this size, ideally reflecting the geometry accurately, is therefore recommended for MFEIT. For human studies, it is again strongly recommended to include electrode position recovery methods into the reconstruction algorithm, since this stabilises reconstructions significantly in realistic experiments.

Appendix A

PEITS User Guide

This appendix guides the reader through the installation process of the PEITS forward solver and explains how it is used. Even though the solver was successfully installed on Windows and OS X, the instructions here only work reliably on UNIX.

A.1 Downloading and Installing the Required Modules

Create a new folder — for instance named PEITS (Parallel EIT Solver) —, where the modules will be installed into. First, the external libraries should be installed. PETSc can be downloaded with Git,

```
git clone -b maint https://bitbucket.org/petsc/petsc petsc
```

and then configured with many external libraries:

```
./configure --prefix=/home/username/PEITS/petscBUILD --with-x=0 --with-  
debugging=0 -CFLAGS="-O3 -DNDEBUG -ffast-math" --with-parmetis=1  
--download-parmetis=yes --with-hypre=1 --download-hypre=yes  
--with-superlu_dist=1 --download-superlu_dist=yes --with-mumps=1  
--download-mumps=yes --with-mt=1 --download-mt=yes --with-metis=1  
--download-metis=yes --download-scalapack=yes --download-blacs=yes
```

In a final step it has to be built:

```
make all test  
make install
```

Most external libraries are installed now, only Zoltan is left. Zoltan can be downloaded from <http://www.cs.sandia.gov/Zoltan/> and then installed:

```
./configure --prefix=/home/username/PEITS/Zoltan_v3.8/BUILD_DIR  
--with-parmetis --with-parmetis-incdir="/home/username/PEITS/  
petscBUILD/include" --with-parmetis-libdir="/home/username/PEITS/
```

A. PEITS User Guide

```
    petscBUILD/lib“  
    make everything  
    make install
```

Now that the external libraries are installed, the actual EIT solver with all the general DUNE modules can be downloaded:

```
git clone http://users.dune-project.org/repositories/projects/  
dune-peits.git
```

In the file `config.opts`, the paths to PETSc and Zoltan have to be adapted to the local machine. Then, the installation of the solver can be done with just one command:

```
sh INSTALL
```

In the generated output it should be controlled that all DUNE modules found METIS, MPI, PARMETIS, PETSc and Zoltan. If not then the configuration options in `config.opts` have to be corrected. Now it should be possible to run the code in the folder `PEITS/dune-peits/src/` with the command

```
mpirun -np 2 ./dune_peits
```

where `-np` specifies the number of processors the solver should be run on. If the solver needs unreasonably long for the assembly of the system matrix, then the pre-allocation of memory in PETSc needs to be adjusted. In file `dune-fem-1.4.0/dune/fem/misc/petsc/petsccommon.hh` the number of allocated non-zeros can be changed in the command `MatMPIAIJSetPreallocation(mat,100,PETSC_NULL,40,PETSC_NULL)`. A safe way of adjusting this is to use very high numbers (e.g. 1000 and 150) and then running the solver with the option `-info`, which outputs the precise number of non-zeros required on the used mesh. Also, on some meshes ML preconditioning fails on some numbers of parallel processes. If this happens, either the number of processes can be changed or hypre preconditioning can be used.

A.2 Running the Solver with Different Settings

All regularly used settings are defined in `dune-peits/data/parameter`. The most important settings are the following:

- `fem.io.loadPartitions` - if the used mesh was already partitioned before this should be set to `true`, if it is used for the first time to `false`
- `fem.io.do_elec_volts` - if the measured boundary voltages should be saved into a binary file or not
- `fem.io.do_jacobian` - if the Jacobian matrix is computed and saved into a binary file
- `fem.io.do_electrode_jacobian` - if the electrode movement Jacobian is computed and appended to the Jacobian matrix binary file

- `mesh` - the mesh that is used
- `fem.io.load_sigma_separately` - if this is set to `true`, then the conductivities of the mesh are loaded from a separate file specified in the next setting
- `fem.io.separate_sigma_file` - the binary file with the conductivity values for all elements
- `conductivities` - a file containing the conductivities for the different tissues in the mesh
- `petsc.preconditioning.method` - which preconditioner should be used. Interesting options are `ml`, `hypre` and as a direct solver `mumps`
- `petsc.ksp_solver.method` - which solver should be used (default is conjugate gradients)
- `poisson.solve_reps` - precision of solutions
- `poisson.solve_iter` - maximum iterations of the solver
- `current.protocol` - the injection and measurement protocol file to be used
- `mesh.perturbation` - `true` if a spherical perturbation should be simulated
- `mesh.perturbation.mult0Rabs` - if the conductivity of the perturbation should be a multiplication of the normal conductivity or an absolute value
- `mesh.perturbation.value` - the value that is either multiplied to the normal conductivity of the relevant elements or assigned as their conductivity
- `mesh.perturbation.radius` - the radius of the perturbation in millimetres
- `mesh.perturbation.pos_C` - the C-coordinate of the perturbation in meters

Rarely changed settings are defined in `dune-peits/data/standardparams` out of which the following are the most interesting:

- `write.vtk` - writes the solution of the first protocol line into a `vtk` file for illustration. This is also helpful to visually evaluate the result of the Zoltan load-balancing, since the additionally saved rank variable encodes the partitions.
- `output.print_electrode_voltages` - prints the electrode areas and voltages of each protocol step into the terminal. This is useful to spot mistakes in the electrode assignment.
- `fem.io.write_only_measured_voltage` - there are two modes how the electrode voltages can be written to the binary file. If this variable is set to `true`, then for each protocol step only the voltage measured between the two defined measurement electrodes is written. If it is `false`, then all electrode potentials are written.

- `fem.uniform_conductivity` - if `true`, then the conductivities of the mesh are overwritten by a constant value which is defined in `fem.uniform_conductivity_value`.
- `contact.impedance` - the contact impedance of the electrodes, given in Ohm. This will internally be multiplied with the electrode area.
- `input.current` - the current level in given in Ampère.
- `electrode.diameter` - the diameter of the electrodes in millimetres.

A.3 Making a Mesh Accessible to the Solver

The solver only accepts finite element meshes saved in the DUNE Grid Format (DGF). A Matlab function `dune_exporter` to write a mesh to this format is provided in the folder `dune-peits/matlab`. It takes seven arguments:

- `Nodes` - an $N \times 3$ matrix with the coordinates in meters, where N is the number of nodes.
- `Tetra` - a $T \times 4$ matrix defining the T tetrahedra by their four nodes.
- `sigma` - a $T \times 1$ matrix with the conductivity value in $S\,m^{-1}$ or tissue index for each element. If it is an integer it is interpreted as a tissue index and assigned a conductivity value that is set in the conductivities file specified in the parameter file.
- `filepath` - the path where the file should be saved to, i.e. `'/home/username/PEITS/dune-peits/data/'`.
- `filename` - the name of the new mesh, e.g. `'new_mesh.dgf'`.
- `electrodepositions` - an $E \times 3$ matrix with the approximate coordinates of the E electrodes in meters.
- `groundposition` - a 1×3 vector with the approximate coordinates of where a node should be set to ground.

The `dune_exporter` function writes three files, the DGF file with the mesh, a `param_new_mesh` file with mesh specific parameters and an `electrode_positions_new_mesh` file with the precise electrode positions. The precise electrode positions are used by the solver to find surface triangles that have their center within a surrounding sphere. If the electrode movement Jacobian (chapter 5) should be computed, then the function `generate_surface_coords` can be used to create the surface coordinate system (`eps` and `d` have to be adapted to the mesh resolution), which can then be written into a text file (e.g. `fprintf`). The name of the surface coordinate file then needs to be updated in the mesh specific parameter file `param_new_mesh`.

A.4 Reading the Results into Matlab

All binary files the solver writes have a timestamp in their name for identification. The `electrodevoltagesTIMESTAMP.bin` file can be loaded in Matlab with the function `dune-peits/matlab/load_electrode_voltages_binary.m`. The Jacobian matrix has to be read in Matlab using the following four steps:

```
[ elementID, unsorted_sigma ] = load_sigma_vector_binary('sigmavector
TIMESTAMP.bin');
unsorted_jacobian = load_jacobian_binary('jacobianTIMESTAMP.bin');
jacobian = zeros(size(unsorted_jacobian));
jacobian(:,elementIDs) = unsorted_jacobian;
```

The sorting is necessary because DUNE rearranges the elements when loading the mesh. Thus, both the columns of the Jacobian matrix and the entries in the vector with the conductivity values of the elements have to be sorted to match the original mesh.

A.5 Running the Solver from Matlab

To facilitate the use of the solver, two functions are provided that allow Matlab users to define the solver settings and run the solver directly from Matlab code. The function `dune-peits/matlab/set_forward_default_values.m` returns a struct with default settings for the solver, which are well documented in the code. Once the `forward_settings` struct is set up and adjusted, the solver can be called with the function `dune-peits/matlab/run_forward_solver.m`:

```
forward_settings = set_forward_default_values();
[v,J,output] = run_forward_solver(forward_settings, processes);
```

The electrode voltages are written to `v`, `J` is the Jacobian matrix, `output` contains the run time output of the solver and `processes` defines on how many parallel processes the solver should be run.

Bibliography

- Abascal J, Arridge S, Atkinson D, Horesh R, Fabrizi L, De Lucia M, Horesh L, Bayford R & Holder D (2008). Use of anisotropic modelling in electrical impedance tomography: description of method and preliminary assessment of utility in imaging brain function in the adult human head. *NeuroImage* **43**, 258–268. 47
- Adler A, Borsic A, Polydorides N & Lionheart W (2008). Simple FEMs aren't as good as we thought: experiences developing EIDORS v3.3. In *Proc. 9th Conf. EIT (Hannover, NH, USA)*. 80
- Adler A, Gaggero P & Maimaitijiang Y (2011). Adjacent stimulation and measurement patterns considered harmful. *Physiological Measurement* **32**, 731–744. 10
- Adler A & Lionheart W (2006). Uses and abuses of EIDORS: an extensible software base for EIT. *Physiological Measurement* **27**, 25–42. 32, 62, 84
- Ahn S, Oh T, Jun S, Seo J & Woo E (2011). Validation of weighted frequency-difference EIT using a three-dimensional hemisphere model and phantom. *Physiological Measurement* **32**, 1663–1680. 41, 106, 121, 122
- Amestoy P, Guermouche A, L'Excellent J & Pralet S (2006). Hybrid scheduling for the parallel solution of linear systems. *Parallel Computing* **32**, 136–156. 63, 72
- Aristovich K, Packham B, Koo H, Sato Dos Santos G, McEvoy A & Holder D (2016). Imaging fast electrical activity in the brain with electrical impedance tomography. *NeuroImage* **124**, 204–213. 19, 61, 79
- Aristovich K, Sato Dos Santos G, Packham B & Holder D (2014). A method for reconstructing tomographic images of evoked neural activity with electrical impedance tomography using intracranial planar arrays. *Physiological Measurement* **35**, 1095–1109. 46, 59, 78, 79
- Arridge S (2015). *Optimisation and Inverse Problems in Imaging*. Course Notes for GV08, University College London. 33, 37
- Avery J (2015). Improving Electrical Impedance Tomography of brain function with a novel servo-controlled electrode helmet. Ph.D. diss., University College London. 15, 96, 121
- Bagshaw A, Liston A, Bayford R, Tizzard A, Gibson A, Tidswell A, Sparkes M, Dehghani H, Binnie C & Holder D (2003). Electrical impedance tomography of human brain function using reconstruction algorithms based on the finite element method. *NeuroImage* **20**, 752–764. 24, 45, 46, 56
- Balay S, Abhyankar S, Adams M, Brown J, Brune P, Buschelman K, Dalcin L, Eijkhout V, Gropp W, Kaushik D, Knepley M, McInnes L, Rupp K, Smith B, Zampini S & Zhang H (2015). PETSc Website <http://www.mcs.anl.gov/petsc> Accessed: 16-12-2015. 63, 68
- Barber D & Brown B (1988). Errors in reconstruction of resistivity images using a linear reconstruction technique. *Clinical Physics and Physiological Measurement* **9**, 101–104. 82
- Berg A, Berkovic S, Brodie M, Buchhalter J, Cross J, Van Emde Boas W, Engel J, French J, Glauser T, Mathern G, Moshé S, Nordli D, Plouin P & Scheffer I (2010). Revised terminology and concepts for organization of seizures and epilepsies: Report of the ILAE Commission on Classification and Terminology, 2005-2009. *Epilepsia* **51**, 676–685. 18
- Boissonnat J & Oudot S (2005). Provably good sampling and meshing of surfaces. *Graphical Models* **67**, 405–451. 23
- Boone K & Holder D (1996a). Current approaches to analogue instrumentation design in electrical impedance tomography. *Physiological Measurement* **17**, 229–247. 9

Bibliography

- Boone K & Holder D (1996b). Effect of skin impedance on image quality and variability in electrical impedance tomography: a model study. *Medical and Biological Engineering and Computing* **34**, 351–354. 9
- Borcea L (2002). Electrical impedance tomography. *Inverse Problems* **18**, 99–136. 33
- Borsic A & Adler A (2012). A primal-dual interior-point framework for using the L1 or L2 norm on the data and regularization terms of inverse problems. *Inverse Problems* **28**, 095011. 36
- Borsic A, Attardo E & Halter R (2012). Multi-GPU Jacobian accelerated computing for soft-field tomography. *Physiological Measurement* **33**, 1703–1715. 32, 62, 80
- Borsic A, Halter R, Wan Y, Hartov A & Paulsen K (2010). Electrical impedance tomography reconstruction for three-dimensional imaging of the prostate. *Physiological Measurement* **31**, 1–16. 32, 62, 80
- Boyle A & Adler A (2011). The impact of electrode area, contact impedance and boundary shape on EIT images. *Physiological Measurement* **32**, 745–754. 110
- Boyle A, Adler A & Lionheart W (2012). Shape deformation in two-dimensional electrical impedance tomography. *IEEE Transactions on Medical Imaging* **31**, 2185–2193. 81
- Brenner S & Scott L (1994). *Mathematical theory of finite element methods*. Springer Verlag. 21
- Brent R (1973). *Algorithms for Minimization Without Derivatives*. Prentice-Hall, Englewood Cliffs, New Jersey. 39, 42, 99
- Brown B (2003). Electrical impedance tomography (EIT): a review. *Journal of Medical Engineering & Technology* **27**, 97–108. 1, 45
- Brown B, Leathard A, Lu L, Wang W & Hampshire A (1995). Measured and expected Cole parameters from electrical impedance tomographic spectroscopy images of the human thorax. *Physiological Measurement* **16**, 57–67. 3, 14, 40, 105, 106
- Brown B & Seagar A (1987). The Sheffield data collection system. *Clinical Physics and Physiological Measurement* **8**, 91–97. 9
- Calderón A (2006). On an inverse boundary value problem. *Computational & Applied Mathematics* **25**, 133–138. 2
- Chaulet N, Arridge S, Betcke T & Holder D (2014). The factorization method for three dimensional electrical impedance tomography. *Inverse Problems* **30**, 045005. 40
- Chung J, Chung M & O’Leary D (2011). Designing optimal spectral filters for inverse problems. *SIAM Journal on Scientific Computing* **33**, 3132–3152. 35
- CIBC (2015). Seg3D: Volumetric Image Segmentation and Visualization. Scientific Computing and Imaging Institute (SCI) <http://www.seg3d.org> Accessed: 16-12-2015. 47
- Cole K & Cole R (1941). Dispersion and Absorption in Dielectrics I. Alternating Current Characteristics. *The Journal of Chemical Physics* **9**, 341–351. 8
- Dai M, Li B, Hu S, Xu C, Yang B, Li J, Fu F, Fei Z & Dong X (2013). In vivo imaging of twist drill drainage for subdural hematoma: a clinical feasibility study on electrical impedance tomography for measuring intracranial bleeding in humans. *PloS one* **8**, e55020. 16
- Dai M, Wang L, Xu C, Li L, Gao G & Dong X (2010). Real-time imaging of subarachnoid hemorrhage in piglets with electrical impedance tomography. *Physiological Measurement* **31**, 1229–1239. 16
- Dardé J, Hakula H, Hyvönen N & Staboulis S (2012). Fine-tuning electrode information in electrical impedance tomography. *Inverse Problems and Imaging* **6**, 399–421. 31, 82, 83, 122, 123
- Dardé J, Hyvönen N, Seppänen A & Staboulis S (2013). Simultaneous recovery of admittivity and body shape in electrical impedance tomography: an experimental evaluation. *Inverse Problems* **29**, 085004. 31, 81
- Davis T (2004). A column pre-ordering strategy for the unsymmetric-pattern multifrontal method. *ACM Transactions on Mathematical Software* **30**, 165–195. 63
- Dedner A, Klöfkor R, Nolte M & Ohlberger M (2010). A generic interface for parallel and adaptive scientific

- computing: abstraction principles and the Dune-Fem module. *Computing* **90**, 165–196. 63
- Devine K, Boman E, Heaphy R, Bisseling R & Catalyurek U (2006). Parallel hypergraph partitioning for scientific computing. In *Proc. Parallel & Distributed Processing Symposium*, p. 10. IEEE. 66
- Dowrick T, Blochet C & Holder D (2015). *In vivo* bioimpedance measurement of healthy and ischaemic rat brain: implications for stroke imaging using electrical impedance tomography. *Physiological Measurement* **36**, 1273–1282. 17, 51
- Dybdahl K (2009). Patient protection at risk in IEC 60601-1 3rd edition. *Medical Device Technology* **20**, 11, 96
- Edd J, Horowitz L & Rubinsky B (2005). Temperature dependence of tissue impedivity in electrical impedance tomography of cryosurgery. *IEEE Transactions on Biomedical Engineering* **52**, 695–701. 111
- Eymard R, Gallouët T & Herbin R (2000). Finite Volume Methods In *Handbook of Numerical Analysis*, Vol. 7, pp. 713–1018. North-Holland, Amsterdam. 21
- Eyüboğlu B & Pilkington T (1993). Comments on distinguishability in electrical impedance imaging. *IEEE Transactions on Biomedical Engineering* **40**, 1328–1330. 12
- Fabrizi L, McEwan A, Oh T, Woo E & Holder D (2009). An electrode addressing protocol for imaging brain function with electrical impedance tomography using a 16-channel semi-parallel system. *Physiological Measurement* **30**, 85–101. 10
- Fabrizi L, Sparkes M, Horesh L, Abascal J, McEwan A, Bayford R, Elwes R, Binnie C & Holder D (2006). Factors limiting the application of electrical impedance tomography for identification of regional conductivity changes using scalp electrodes during epileptic seizures in humans. *Physiological Measurement* **27**, 163–174. 18
- Falgout R (2015). Hypre website <http://www.llnl.gov/CASC/hypre> Accessed: 16-12-2015. 63, 70
- Fitzgerald A, Holder D, Eadie L, Hare C & Bayford R (2002). A comparison of techniques to optimize measurement of voltage changes in electrical impedance tomography by minimizing phase shift errors. *IEEE Transactions on Medical Imaging* **21**, 668–675. 59
- Frangi A, Riu P, Rosell J & Viergever M (2002). Propagation of measurement noise through backprojection reconstruction in electrical impedance tomography. *IEEE Transactions on Medical Imaging* **21**, 566–578. 9
- Frerichs I (2000). Electrical impedance tomography (EIT) in applications related to lung and ventilation: a review of experimental and clinical activities. *Physiological Measurement* **21**, 1–21. 1, 15
- Gee M, Siefert C, Hu J, Tuminaro R & Sala M (2006). ML 5.0 smoothed aggregation user's guide. Technical report SAND2006-2649, Sandia National Laboratories. 63, 70
- Gençer N & Tanzer I (1999). Forward problem solution of electromagnetic source imaging using a new BEM formulation with high-order elements. *Physics in Medicine and Biology* **44**, 2275–2287. 21, 32
- Gibson A (2000). Electrical impedance tomography of human brain function Ph.D. diss., University College London. 10, 11
- Gómez-Laberge C & Adler A (2008). Direct EIT Jacobian calculations for conductivity change and electrode movement. *Physiological Measurement* **29**, 89–99. 31, 82
- Gong B, Krueger-Ziolek S, Moeller K, Schullcke B & Zhao Z (2015). Electrical impedance tomography: functional lung imaging on its way to clinical practice? *Expert Review of Respiratory Medicine* **6348**, 1–17. 15
- Goren N, Avery J & Holder D (2015). Feasibility study for monitoring stroke and TBI patients. In Solà J, Braun F & Adler A, editors, *Proc. 16th Int. Conf. on Biomed. Appl. of EIT*, p. 71. 51, 125
- Grychtol B & Adler A (2013). FEM electrode refinement for electrical impedance tomography In EBMS, editor, *Proc. Ann. Int. Conf. IEEE Eng. Med. & Biol. Soc.*, pp. 6429–6432. 23

Bibliography

- Grychtol B, Lionheart W, Wolf G, Bodenstein M & Adler A (2012). Impact of model shape mismatch on reconstruction quality in Electrical Impedance Tomography. *IEEE Transactions on Medical Imaging* **31**, 1754–1760. 23
- Hampshire A, Smallwood R, Brown B & Primhak R (1995). Multifrequency and parametric EIT images of neonatal lungs. *Physiological Measurement* **16**, 175–189. 105
- Hansen A & Olsen C (1980). Brain extracellular space during spreading depression and ischemia. *Acta Physiologica Scandinavica* **108**, 355–365. 105
- Hansen P (1994). Regularization Tools: A Matlab Package for Analysis and Solution of Discrete Ill-Posed Problems. *Numerical Algorithms* **6**, 1–35. 36, 52, 90
- Heeger D & Ress D (2002). What does fMRI tell us about neuronal activity? *Nature Reviews Neuroscience* **3**, 142–151. 18
- Heikkinen L, Vilhunen T, West R & Vauhkonen M (2002). Simultaneous reconstruction of electrode contact impedances and internal electrical properties: II. Laboratory experiments. *Measurement Science and Technology* **13**, 1855–1861. 12, 98
- Henson V & Yang U (2002). BoomerAMG: A parallel algebraic multigrid solver and preconditioner. *Applied Numerical Mathematics* **41**, 155–177. 29
- Ho-Le K (1988). Finite element mesh generation methods: a review and classification. *Computer-Aided Design* **20**, 27–38. 23
- Holder D (1987). Feasibility of developing a method of imaging neuronal activity in the human brain: a theoretical review. *Medical and Biological Engineering and Computing* **25**, 2–11. 19
- Holder D (1992). Detection of cerebral ischaemia in the anaesthetised rat by impedance measurement with scalp electrodes: implications for non-invasive imaging of stroke by electrical impedance tomography. *Clinical Physics and Physiological Measurement* **13**, 63–75. 105
- Holder D (2004). Appendix A: Brief Introduction to Bioimpedance. In Holder D, editor, *Electrical Impedance Tomography: Methods, History and Applications*, pp. 411–422. Taylor & Francis. 7
- Holder D & Tidswell T (2004). Electrical Impedance Tomography of Brain Function. In Holder D, editor, *Electrical Impedance Tomography: Methods, History and Applications*, chapter 4, pp. 127–166. Taylor & Francis. 1, 2, 17, 18, 59
- Horesh L (2006). Some Novel Approaches in Modelling and Image Reconstruction for Multi Frequency Electrical Impedance Tomography of the Human Brain. Ph.D. diss., University College London. 16, 17, 51
- Horesh L, Gilad O, Romsauerova A, McEwan A, Arridge S & Holder D (2005). Stroke type differentiation by multi-frequency electrical impedance tomography – a feasibility study. In *Proc. 3rd Europ. Med. & Biol. Eng. Conf., Prague*, pp. 1252–1256. 17, 105, 107
- Horesh L, Schweiger M, Bollhöfer M, Douiri A, Holder D & Arridge S (2006). Multilevel preconditioning for 3D large-scale soft field medical applications modelling. *Int. J. Inf. Syst. Sci.* **2**, 532–556. 32, 51, 62, 76, 125
- IEC (2005). *IEC 60601 Medical Electrical Equipment – Part 1: General requirement for basic safety and essential performance*. International Electrotechnical Commission, Geneva. 9
- Isaacson D (1986). Distinguishability of conductivities by electric current computed tomography. *IEEE Transactions on Medical Imaging* **5**, 91–95. 10, 21
- Jang J & Seo J (2015). Detection of admittivity anomaly on high-contrast heterogeneous backgrounds using frequency difference EIT. *Physiological Measurement* **36**, 1179–1192. 41
- Johnson C & MacLeod R (1994). Nonuniform spatial mesh adaptation using a posteriori error estimates: applications to forward and inverse problems. *Applied Numerical Mathematics* **14**, 311–326. 46, 80
- Jun S, Kuen J, Lee J, Woo E, Holder D & Seo J (2009). Frequency-difference EIT (fdEIT) using weighted difference and equivalent homogeneous admittivity: validation by simulation and tank experiment. *Physiological Measurement* **30**, 1087–1099. 14, 41, 103, 106, 122

- Jurcak V, Tsuzuki D & Dan I (2007). 10/20, 10/10, and 10/5 systems revisited: Their validity as relative head-surface-based positioning systems. *NeuroImage* **34**, 1600–1611. 10
- Kohn R & Vogelius M (1985). Determining conductivity by boundary measurements II. Interior results. *Communications on Pure and Applied Mathematics* **38**, 643–667. 20, 33
- Kolehmainen V, Vauhkonen M, Karjalainen P & Kaipio J (1997). Assessment of errors in static electrical impedance tomography with adjacent and trigonometric current patterns. *Physiological Measurement* **18**, 289–303. 12, 13, 45, 82, 98, 101, 110, 118
- Langlois J, Rutland-Brown W & Wald M (2006). The epidemiology and impact of traumatic brain injury: a brief overview. *The Journal of Head Trauma Rehabilitation* **21**, 375–378. 15
- Lionheart W, Polydorides N & Borsic A (2004). The reconstruction problem. In Holder D, editor, *Electrical Impedance Tomography: Methods, History and Applications*, chapter 1, pp. 3–64. Taylor & Francis. 20, 21, 22, 34, 36, 45, 51, 89, 107, 124
- Lionheart W, Kaipio J & McLeod C (2001). Generalized optimal current patterns and electrical safety in EIT. *Physiological Measurement* **22**, 85–90. 12
- Lionheart W (1998). Boundary shape and electrical impedance tomography. *Inverse Problems* **139**, 139–147. 31
- Liston A, Bayford R & Holder D (2012). A cable theory based biophysical model of resistance change in crab peripheral nerve and human cerebral cortex during neuronal depolarisation: implications for electrical impedance tomography of fast neural activity in the brain. *Medical and Biological Engineering and Computing* **50**, 425–37. 19
- Liszka T & Orkisz J (1980). The finite difference method at arbitrary irregular grids and its application in applied mechanics. *Computers & Structures* **11**, 83–95. 21
- Liu A & Joe B (1994). Relationship between tetrahedron shape measures. *BIT Numerical Mathematics* **34**, 268–287. 49
- Malich A, Böhm T, Facius M, Kleinteich I, Fleck M, Sauner D, Anderson R & Kaiser W (2003). Electrical impedance scanning as a new imaging modality in breast cancer detection — a short review of clinical value on breast application, limitations and perspectives. *Nuclear Instruments and Methods in Physics Research* **497**, 75–81. 105
- Malone E, Sato Dos Santos G, Holder D & Arridge S (2014). Multifrequency Electrical Impedance Tomography using spectral constraints. *IEEE Transactions on Medical Imaging* **33**, 340–350. 14, 18, 41, 59, 98, 103, 106, 108, 121, 122
- Malone E, Sato Dos Santos G, Holder D & Arridge S (2015). A reconstruction-classification method for multifrequency electrical impedance tomography. *IEEE Transactions on Medical Imaging* **34**, 1486–1497. 14, 43, 106, 122
- Manwaring P, Moodie K, Hartov A, Manwaring K & Halter R (2013). Intracranial electrical impedance tomography: A method of continuous monitoring in an animal model of head trauma. *Anesthesia and Analgesia* **117**, 866–875. 16
- Matsuda W, Sugimoto K, Sato N, Watanabe T, Fujimoto A & Matsumura A (2008). Delayed onset of post-traumatic acute subdural hematoma after mild head injury with normal computed tomography: a case report and brief review. *The Journal of Trauma* **65**, 461–463. 15
- McAdams E, Jossinet J, Lacknermeier A & Risacher F (1996). Factors affecting electrode-gel-skin interface impedance in electrical impedance tomography. *Medical and Biological Engineering and Computing* **34**, 397–408. 12
- McAuliffe M, Lalonde F, McGarry D, Gandler W, Csaky K & Trus B (2001). Medical Image Processing, Analysis and Visualization in Clinical Research. In *Proc. 14th IEEE Symp. on Computer-Based Medical Systems*, pp. 381–386. IEEE Comput. Soc. 47

Bibliography

- McEwan A, Cusick G & Holder D (2007). A review of errors in multi-frequency EIT instrumentation. *Physiological Measurement* **28**, 197–215. 82, 106, 111, 118
- McEwan A, Romsauerova A, Yerworth R, Horesh L, Bayford R & Holder D (2006). Design and calibration of a compact multi-frequency EIT system for acute stroke imaging. *Physiological Measurement* **27**, 199–210. 9
- Modat M, Ridgway G, Taylor Z, Lehmann M, Barnes J, Hawkes D, Fox N & Ourselin S (2010). Fast free-form deformation using graphics processing units. *Computer Methods and Programs in Biomedicine* **98**, 278–284. 24
- Mohammed B, Abbosh A, Mustafa S & Ireland D (2014). Microwave system for head imaging. *IEEE Transactions on Instrumentation and Measurement* **63**, 117–123. 18
- Molinari M, Cox S, Blott B & Daniell G (2001). Adaptive mesh refinement techniques for electrical impedance tomography. *Physiological Measurement* **22**, 91–96. 46, 80
- Murphy D, Burton P, Coombs R, Tarassenko L & Rolfe P (1987). Impedance imaging in the newborn. *Clinical Physics and Physiological Measurement* **8**, 131–140. 16
- NICE (2014). *Head Injury – Methods, evidence and recommendations*. National Institute for Health and Care Excellence, United Kingdom. 15
- Nissinen A, Heikkinen L & Kaipio J (2008). The Bayesian approximation error approach for electrical impedance tomography – experimental results. *Measurement Science and Technology* **19**, 015501. 82
- Nissinen A, Heikkinen L, Kolehmainen V & Kaipio J (2009). Compensation of errors due to discretization, domain truncation and unknown contact impedances in electrical impedance tomography. *Measurement Science and Technology* **20**. 31
- Nissinen A, Kolehmainen V & Kaipio J (2011). Compensation of modelling errors due to unknown domain boundary in electrical impedance tomography. *IEEE Transactions on Medical Imaging* **30**, 231–242. 31
- Nocedal J & Wright S (1999). *Numerical optimization*. Springer Series in Operations Research and Financial Engineering, New York. 39, 42
- Nuwer M, Comi G, Emerson R, Fuglsang-Frederiksen A, Guérit JM, Hinrichs H, Ikeda A, Luccas F & Rappelsburger P (1998). IFCN standards for digital recording of clinical EEG. *Electroencephalography and Clinical Neurophysiology* **106**, 259–261. 50
- Oh T, Wi H, Kim D, Yoo P & Woo E (2011). A fully parallel multi-frequency EIT system with flexible electrode configuration: KHU Mark2. *Physiological Measurement* **32**, 835–849. 110
- Oostenveld R & Praamstra P (2001). The five percent electrode system for high-resolution EEG and ERP measurements. *Clinical Neurophysiology* **112**, 713–719. 11
- Packham B, Koo H, Romsauerova A, Ahn S, McEwan A, Jun S & Holder D (2012). Comparison of frequency difference reconstruction algorithms for the detection of acute stroke using EIT in a realistic head-shaped tank. *Physiological Measurement* **33**, 767–786. 14, 17, 41, 105
- Persson M, Fhager A, Trefná H, Yu Y, McKelvey T, Pegenius G, Karlsson J & Elam M (2014). Microwave-based stroke diagnosis making global prehospital thrombolytic treatment possible. *IEEE Transactions on Biomedical Engineering* **61**, 2806–2817. 18, 134
- Polydorides N (2009). Linearization Error in Electrical Impedance Tomography. *Progress In Electromagnetics Research* **93**, 323–337. 30
- Polydorides N & Lionheart W (2002). A Matlab toolkit for three-dimensional electrical impedance tomography: a contribution to the Electrical Impedance and Diffuse Optical Reconstruction Software project. *Measurement Science and Technology* **13**, 1871–1883. 25, 29
- Power M (2004). An update on thrombolysis for acute ischaemic stroke. *Advances in Clinical Neuroscience and Rehabilitation* **4**, 36–37. 16
- Qian S & Sheng Y (2011). A single camera photogrammetry system for multi-angle fast localization of EEG electrodes. *Annals of Biomedical Engineering* **39**, 2844–2856. 110, 121

- Romsauerova A, McEwan A, Horesh L, Yerworth R, Bayford R & Holder D (2006). Multi-frequency electrical impedance tomography (EIT) of the adult human head: initial findings in brain tumours, arteriovenous malformations and chronic stroke, development of an analysis method and calibration. *Physiological Measurement* **27**, 147–161. 105, 107
- Rosenow F & Lüders H (2001). Presurgical evaluation of epilepsy patients. *Brain* **124**, 1683–1700. 18
- Sadleir R, Vannorsdall T, Schretlen D & Gordon B (2010). Transcranial direct current stimulation (tDCS) in a realistic head model. *NeuroImage* **51**, 1310–1318. 47
- Sattin J, Olson S, Liu L, Raman R & Lyden P (2006). An Expedited Code Stroke Protocol Is Feasible and Safe. *Stroke* **37**, 2935–2939. 17
- Saulnier G (2004). Eit instrumentation. In Holder D, editor, *Electrical Impedance Tomography: Methods, History and Applications*, chapter 2, pp. 65–104. Taylor & Francis. 8, 9, 20
- Saulnier G, Blue R, Newell J, Isaacson D & Edic P (2001). Electrical impedance tomography. *IEEE Signal Processing Magazine* **18**, 31–43. 14
- Saver J, Fonarow G, Smith E, Reeves M, Grau-Sepulveda M, Pan W, Olson D, Hernandez A, Peterson E & Schwamm L (2013). Time to Treatment With Intravenous Tissue Plasminogen Activator and Outcome From Acute Ischemic Stroke. *The Journal of the American Medical Association* **309**, 2480–2488. 16
- Sawicki B & Okoniewski M (2010). Adaptive mesh refinement techniques for 3-D skin electrode modeling. *IEEE Transactions on Biomedical Engineering* **57**, 528–533. 46, 59, 80
- Scapaticci R, Bucci O, Catapano I & Crocco L (2014). Differential Microwave Imaging for Brain Stroke Followup. *International Journal of Antennas and Propagation* **2014**, 1–11. 18
- Scapaticci R, Di Donato L, Catapano I & Crocco L (2012). A feasibility study on microwave imaging for brain stroke monitoring. *Progress In Electromagnetics Research B* **40**, 305–324. 18
- Schenk O (2015). PARDISO Website <http://www.pardiso-project.org> Accessed: 16-12-2015. 32
- Schöberl J (1997). NETGEN - An Advancing Front 2D/3D Mesh Generator Based on Abstract Rules. *Computing and Visualization in Science* **1**, 41–52. 23
- Schramm J & Clusmann H (2008). The surgery of epilepsy. *Neurosurgery* **62**, 463–481. 18
- Seo J, Lee J, Kim S, Zribi H & Woo E (2008). Frequency-difference electrical impedance tomography (fdEIT): algorithm development and feasibility study. *Physiological Measurement* **29**, 929–944. 14, 41, 106
- Shewchuk J (1994). An introduction to the conjugate gradient method without the agonizing pain. *Carnegie Mellon University, Pittsburgh, PA*. 27, 37
- Shewchuk J (1997). Delaunay refinement mesh generation. *DTIC Document*. 23, 37
- Shi X, You F, Fu F, Liu R, You Y, Dai M & Dong X (2008). Preliminary research on monitoring of cerebral ischemia using electrical impedance tomography technique. In *Proc. IEEE Eng. Med. Biol. Soc.*, pp. 1188–1191. 105
- Siltanen S, Mueller J & Isaacson D (2000). An implementation of the reconstruction algorithm of a nachman for the 2d inverse conductivity problem. *Inverse Problems* **16**, 681–699. 40
- Soleimani M, Gómez-Laberge C & Adler A (2006). Imaging of conductivity changes and electrode movement in EIT. *Physiological Measurement* **27**, 103–113. 31, 82, 122
- Soleimani M, Powell C & Polydorides N (2005). Improving the forward solver for the complete electrode model in EIT using algebraic multigrid. *IEEE Transactions on Medical Imaging* **24**, 577–583. 32, 62
- Somersalo E, Cheney M & Isaacson D (1992). Existence and Uniqueness for Electrode Models for Electric Current Computed Tomography. *SIAM Journal on Applied Mathematics* **52**, 1023–1040. 21
- Soni N, Paulsen K, Dehghani H & Hartov A (2006). Finite element implementation of Maxwell's equations

Bibliography

- for image reconstruction in electrical impedance tomography. *IEEE Transactions on Medical Imaging* **25**, 55–61. 20
- Stemer A & Lyden P (2010). Evolution of the thrombolytic treatment window for acute ischemic stroke. *Current neurology and neuroscience reports* **10**, 29–33. 16
- Stüben K (2001). A review of algebraic multigrid. *Journal of Computational and Applied Mathematics* **128**, 281–309. 28
- Sylvester J & Uhlmann G (1987). A global uniqueness theorem for an inverse boundary value problem. *Annals of Mathematics* **125**, 153–169. 33
- Tang C, You F, Cheng G, Gao D, Fu F, Yang G & Dong X (2008). Correlation between structure and resistivity variations of the live human skull. *IEEE Transactions on Biomedical Engineering* **55**, 2286–2292. 15, 91
- The CGAL Project (2015). *CGAL User and Reference Manual*. CGAL Editorial Board, 4.7 edition. 23, 47, 49
- Tidswell A, Gibson A, Bayford R & Holder D (2001a). Validation of a 3D reconstruction algorithm for EIT of human brain function in a realistic head-shaped tank. *Physiological Measurement* **22**, 177–185. 15, 96
- Tidswell T, Gibson A, Bayford R & Holder D (2001b). Three-dimensional electrical impedance tomography of human brain activity. *NeuroImage* **13**, 283–294. 45, 87, 125
- Tizzard A, Horesh L, Yerworth R, Holder D & Bayford R (2005). Generating accurate finite element meshes for the forward model of the human head in EIT. *Physiological Measurement* **26**, 251–261. 23, 24, 46, 47, 57
- Toussaint N, Souplet J & Fillard P (2007). MedINRIA: Medical image navigation and research tool by INRIA. *Proc. of MICCAI* **7**. 47
- Tuminaro R & Tong C (2000). Parallel smoothed aggregation multigrid: aggregation strategies on massively parallel machines. In *Proc. Conf. on Supercomp.*, p. 5. IEEE Computer Society. 29
- Vauhkonen P, Vauhkonen M, Savolainen T & Kaipio J (1999). Static Three-Dimensional Electrical Impedance Tomography. *Annals of the New York Academy of Sciences* **873**, 472–481. 13
- Vonach M, Marson B, Yun M, Cardoso J, Modat M, Ourselin S & Holder D (2012). A method for rapid production of subject specific finite element meshes for electrical impedance tomography of the human head. *Physiological Measurement* **33**, 801–816. 24, 46, 47, 57, 58
- Vongerichten A, Sato dos Santos G, Aristovich K, Avery J, McEvoy A, Walker M & Holder D (2016). Characterisation and imaging of cortical impedance changes during interictal and ictal activity in the anaesthetised rat. *NeuroImage* **124**, 813–823. 18
- Wesseling P & Oosterlee C (2001). Geometric multigrid with applications to computational fluid dynamics. *Journal of Computational and Applied Mathematics* **128**, 311–334. 28
- Wi H, Sohal H, McEwan A, Woo E & Oh T (2014). Multi-frequency electrical impedance tomography system with automatic self-calibration for long-term monitoring. *IEEE Transactions on Biomedical Circuits and Systems* **8**, 119–128. 96
- Xu C, Wang L, Shi X, You F, Fu F, Liu R, Dai M, Zhao Z, Gao G & Dong X (2010). Real-time imaging and detection of intracranial haemorrhage by electrical impedance tomography in a piglet model. *Journal of International Medical Research* **38**, 1596–1604. 15, 16, 45, 81
- Yorkey T (1990). Electrical impedance tomography with piecewise polynomial conductivities. *Journal of Computational Physics* **91**, 344–360. 76
- Zou Y & Guo Z (2003). A review of electrical impedance techniques for breast cancer detection. *Medical engineering & physics* **25**, 79–90. 2