

Some pages of this thesis may have been removed for copyright restrictions.

If you have discovered material in AURA which is unlawful e.g. breaches copyright, (either yours or that of a third party) or any other law, including but not limited to those relating to patent, trademark, confidentiality, data protection, obscenity, defamation, libel, then please read our [Takedown Policy](#) and [contact the service](#) immediately

**THE DEVELOPMENT OF CONSTRAINED SOURCE
LOCALISATION ALGORITHMS FOR HUMAN
BRAIN IMAGING**

ARJAN HILLEBRAND

Doctor of Philosophy

ASTON UNIVERSITY

April 2000

**THE DEVELOPMENT OF CONSTRAINED SOURCE
LOCALISATION ALGORITHMS FOR HUMAN
BRAIN IMAGING**

ARJAN HILLEBRAND

Doctor of Philosophy

ASTON UNIVERSITY

April 2000

This copy of the thesis has been supplied on condition that anyone who consults it is understood to recognise that its copyright rests with its author and that no quotation from the thesis and no information derived from it may be published without proper acknowledgement.

ASTON UNIVERSITY

THE DEVELOPMENT OF CONSTRAINED SOURCE
LOCALISATION ALGORITHMS FOR HUMAN
BRAIN IMAGING

ARJAN HILLEBRAND

Doctor of Philosophy

2000

This work sets out to evaluate the potential benefits and pit-falls in using *a priori* information to help solve the Magnetoencephalographic (MEG) inverse problem.

In chapter one the forward problem in MEG is introduced, together with a scheme that demonstrates how *a priori* information can be incorporated into the inverse problem.

Chapter two contains a literature review of techniques currently used to solve the inverse problem. Emphasis is put on the kind of *a priori* information that is used by each of these techniques and the ease with which additional constraints can be applied. The formalism of the FOCUSS algorithm is shown to allow for the incorporation of *a priori* information in an insightful and straightforward manner.

In chapter three it is described how anatomical constraints, in the form of a realistically shaped source space, can be extracted from a subject's Magnetic Resonance Image (MRI). The use of such constraints relies on accurate co-registration of the MEG and MRI co-ordinate systems. Variations of the two main co-registration approaches, based on fiducial markers or on surface matching, are described and the accuracy and robustness of a surface matching algorithm is evaluated.

Figures of merit introduced in chapter four are shown to give insight into the limitations of a typical measurement set-up and potential value of *a priori* information.

It is shown in chapter five that constrained dipole fitting and FOCUSS outperform unconstrained dipole fitting when data with low SNR is used. However, the effect of errors in the constraints can reduce this advantage.

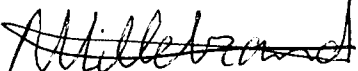
Finally, it is demonstrated in chapter six that the results of different localisation techniques give corroborative evidence about the location and activation sequence of the human visual cortical areas underlying the first 125ms of the visual magnetic evoked response recorded with a whole head neuromagnetometer.

Keywords: magnetoencephalography; inverse problem; MEG-MRI co-registration; realistically shaped source space; resolution matrix.

Acknowledgement

The work described in this thesis was carried out by the author, with the following exceptions:

Chapter 6. The stimulus was designed and programmed by Dr. G.R. Barnes and Dr. I.E. Holliday. MEG data collection was performed by Dr. G.R. Barnes, including signal averaging and filtering of the raw data. Also, figure 83 was kindly provided by Dr. G.R. Barnes. The author would like to thank the above persons for providing the MEG data and the graph.

Signed: 

This journey into science started by travelling to the UK; thanks go to my supervisor, Professor G.F.A. Harding, for providing the opportunity and research facilities.

In foreign places a good guide is a necessity. The guidance and support provided by my associate supervisor, Ian Holliday is greatly appreciated. Special thanks go to Gareth Barnes for keeping me on the right track and providing feedback on my achievements. Also many thanks for carefully reading this manuscript. Gratitude also goes to the other members of the MEG group, Paul Furlong and Fiona Fylan, for the support they have given.

A good expedition includes entertainment as well. Thanks to the many people that made working in the Clinical Neurophysiology Unit at Aston University so pleasurable, in particular the above mentioned persons and Ruth Webb, Catherine Suttle, Tracey Huaut, Ian Fawcett, Andrea Edson, Mitch Thomson and last but not least my fellow postgraduate students.

Finally, a journey is most enjoyed when travelled together, Fabiola, thanks for your love and support throughout the years.

TABLE OF CONTENTS

FIGURES	8
TABLES	17
1 INTRODUCTION	18
1.2 Objectives	21
1.3 Forward Problem	23
1.3.1 Introduction	23
1.3.2 Source Model	26
1.3.3 Volume Conductor Model	28
1.3.4 Theory of Volume Conducting	29
2 INVERSE PROBLEM	34
2.1 Introduction	34
2.2 Reconstruction of Discrete Sources	37
2.2.1 Equivalent Current Dipole Fitting	37
2.2.1.1 Assumptions	37
2.2.1.2 Finding Dipole Parameters	37
2.2.1.2.1 Cost-function and Validity of the Model	38
2.2.1.2.2 Optimisation Techniques	40
2.2.1.2.3 Noise Estimation	43
2.2.1.2.4 Dipole Models Using <i>a priori</i> Information	43
2.2.1.3 Limitations of Equivalent Current Dipole Fitting	44
2.2.2 Spatio-Temporal Dipole Models	46
2.2.2.1 Introduction	46
2.2.2.2 Assumptions	46
2.2.2.3 Selecting the Correct Model	48
2.2.2.4 Finding Dipole Positions	49
2.2.2.5 Limitations of Spatio-Temporal Dipole Models	54
2.2.3 Spatio-Temporal Decomposition	55
2.2.3.1 Singular Value Decomposition	56
2.2.3.2 Principal Component Analysis	57
2.2.3.2.1 Limitations	57
2.2.3.3 Independent Component Analysis	60
2.2.4 Recursive - Multiple Signal Classification (R-MUSIC)	62
2.2.4.1 Assumptions	63
2.2.4.2 Finding Dipole Positions	63
2.2.4.3 Limitations	68

2.3	Reconstruction of Distributed Sources	70
2.3.1	Introduction	70
2.3.2	Minimum Norm Least Squares Approach.....	71
2.3.2.1	Assumptions	71
2.3.2.2	Formulation	76
2.3.2.3	Regularisation	82
2.3.2.4	Application	85
2.3.2.5	Limitations	85
2.3.3	Generalised Wiener-Helstrom Filter	90
2.3.4	Probabilistic Methods.....	91
2.3.4.1	Assumptions	92
2.3.4.2	Technique	94
2.3.4.3	Limitations	94
2.4	From Distributed Sources to Discrete Sources.....	95
2.4.1	Introduction	95
2.4.2	Iterative Methods.....	96
2.4.2.1	Minimum Dipole Solutions.....	98
2.4.2.2	Shrinking Ellipsoid Inverse.....	100
2.4.2.3	Focal Underdetermined System Solver (FOCUSS).....	101
2.4.2.3.1	Bias Correction	103
2.4.2.3.2	Relation with Statistical Methods	105
2.4.3	Other Methods.....	105
2.4.3.1	Low Resolution Electromagnetic Tomography (LORETA).....	107
2.4.3.2	Magnetic Field Tomography (MFT).....	109
2.4.3.3	A Bayesian Approach	109
2.5	Synthetic Aperture Magnetometry (SAM).....	113
2.6	Choosing the Inversion Technique.....	116
3	USING A REALISTICALLY SHAPED SOURCE SPACE	119
3.1	Introduction	119
3.2	Constructing a Realistically Shaped Source Space.....	121
3.2.1	Introduction.....	124
3.2.2	Segmentation.....	134
3.2.3	Surface Reconstruction	139
3.2.4	Creating and Displaying the Source Space	139
3.3	Methods of MEG/MRI Co-registration	142
3.3.1	Introduction	142
3.3.2	Using Landmarks	149
3.3.3	Surface Matching	149
3.4	Errors in Co-registration	153
3.4.1	Sources of Errors.....	156
3.4.2	Accuracy of Implemented Co-registration Procedure.....	162
3.4.2.1	Intrinsic Co-registration Accuracy.....	168
3.4.2.2	Influence of Number of Registration Points	169
3.4.2.3	Influence of Noise	171
3.4.2.4	Test-Retest Accuracy	171
3.5	Discussion	175

4	FIGURES OF MERIT FOR DISTRIBUTED SOURCE RECONSTRUCTION	179
4.1	Introduction	179
4.2	Theory.....	181
4.2.1	Decrease of Singular Values	181
4.2.2	Resolution Matrix.....	183
4.3	Usefulness of Figures of Merit in Conjunction With FOCUSS	195
4.3.1	Introduction.....	195
4.3.2	Evaluation With a Realistically Shaped Source Space.....	196
4.3.2.1	Groundwork	196
4.3.2.2	Figures of Merit.....	201
4.3.3	Relevance to FOCUSS.....	217
4.4	Discussion	224
5	EVALUATION OF SOURCE LOCALISATION ALGORITHMS WITH COMPUTER SIMULATIONS	227
5.1	Aim of This Study	227
5.2	Simulation Framework	228
5.3	Parameter Selection.....	237
5.3.1	Compound Versus Non-compound Version	237
5.3.2	Truncation Parameter	244
5.3.3	Gridspacing	252
5.4	Influence of Spatial Errors.....	268
5.4.1	Position Errors.....	268
5.4.2	Orientation Errors	273
5.5	Conclusions.....	278
6	SOURCE LOCALISATION FROM VISUAL EVOKED MEG DATA	280
6.1	Introduction	280
6.2	Methods	281
6.3	Results.....	286
6.3.1	Initial Inquiry of the Data.....	286
6.3.2	Source Reconstructions	293
6.4	Discussion	308
7	DISCUSSION AND RECOMMENDATIONS	311
	REFERENCES.....	314
	APPENDICES.....	326

Appendix A.....326

FIGURES

Chapter 1

figure 1-1: Comparison of spatial and temporal resolution of different brain imaging techniques. Note the superior temporal resolution of MEG. Source: CTF Systems, Inc.....	18
figure 1-2: Basic scheme to solve the inverse problem. Solving the inverse problem relies on the ability to solve the forward problem accurately. The forward problem can be solved when modelling assumptions are made about the sources in the brain and the head as a volume conductor. Furthermore, the sensor configuration must be encoded. Adding <i>a priori</i> knowledge at different stages can increase the accuracy of the solution of the inverse problem.....	20
figure 1-3: Schematic of a first order axial gradiometer consisting of two oppositely wound coils inductively coupled to the SQUID sensor.....	23
figure 1-4: Matrix representation of the recorded data (top panel). The signal vector contains the recorded values for a channel at all latencies. The sample vector contains the recorded values in all channels at a certain latency (after Hjorth, 1989). In the bottom panel a colour-coded topographic map of a sample vector is shown. In this case the sample vector is the measured field pattern at 110ms after onset of an auditory stimulus. Two dipolar field patterns are clearly visible in the data, suggesting that the magnetic field is induced by two areas of brain activity that can both be modelled by an equivalent current dipole in the left and right auditory cortex, respectively.....	25
figure 1-5: Schematic drawing of a nerve cell.	26
figure 1-6: An area of active cortex can be modelled as a layer of current dipoles. This layer of dipoles in turn can be modelled as a single equivalent current dipole when viewed from a large enough distance.	27
figure 1-7: Definition of the location of the current density and the measurement position with respect to the origin.....	31
figure 1-8: Definition of the radial and tangential components of a dipole (figure from Malmivuo <i>et al.</i> , 1997). A dipole with only a radial or tangential component is referred to as a radial or tangential dipole, respectively.	33

Chapter 2

figure 2-1: Hypothetical example of the minimisation process of the cost-function. Here, the cost-function is only optimised for one of the dipole parameters (r), the other parameters are kept constant. The numbers between brackets denote the iteration number in the optimisation process. The algorithm starts with the initial guess (r_0), traverses through the local (step 2) and global minimum (step 9) and finally (step 21) arrives back at the global minimum (r_g).	38
figure 2-2: Reduced chi-square for solutions obtained with 300 randomly chosen initial guesses, sorted in descending order. Figure modified from Huang <i>et al</i> (1998).....	50
figure 2-3: An example of a monophasic wavelet. It is defined by the four parameters: peak amplitude (PA), peak latency (PL), latency of the onset (ONL) and of the offset (OFFL).	52
figure 2-4: Plot of hypothetical principal correlations. The top panel shows the maximum principal correlation for each grid point ($= c_1$). It is clearly seen that with MUSIC and R-MUSIC it would be easy to localise peak A. Finding the other peaks in this figure (as is done with MUSIC) would be more cumbersome, because they are broader and lower. This could be due to noise or modelling errors. The bottom panel shows the second largest principal correlations ($= c_2$) that would be obtained with R-MUSIC after the	

first source is found from the top panel, using the first localised source as part of the source model. Finding the maximum in this plot is easy. Finding additional sources is therefore much easier with R-MUSIC than with MUSIC.	67
figure 2-5: Graphical representation of a slice through a part of the cortical surface, showing a sulcus. The distributed source, represented by the arrows, extends across both walls of the sulcus. The centre of gravity of the activity is indicated by the black dot. The result of localisation with a single dipole model, located in this centre of gravity, would give a meaningless result as the location is not even within the cortex.	70
figure 2-6: Projection of the total current vector \mathbf{J} into the subspace spanned by the leads L_1 and L_2 . The minimum norm estimate of the total current is given by $\mathbf{J}_{//}$. It can be seen that \mathbf{J}_{\perp} does not contribute to the minimum norm estimate, as $\langle \mathbf{J}_{\perp}, \mathbf{\Omega} \rangle = 0$. Generalisation to more than 2 sensors (leads) is straightforward. After Hämäläinen <i>et al</i> (1993).	72
figure 2-7: Example of an L-curve, with on the x-axis the normalised residual norm, $\ \mathbf{B} - \tilde{\mathbf{B}}\ $, and on the y-axis the normalised solution norm, $\ \hat{\mathbf{Q}}\ $. The number of singular values that were used decreases from left to right. One can see that the residual norm decreases when more singular values are used in the computations and that the solution norm increases at the same time. This is expected, as more sources are used to explain the data when the number of singular values that are used increases. The optimum truncation parameter is the singular value that lies near the corner of the L-curve. This example was obtained for a simulation set-up with the 19-channel CES system at Aston University.	79
figure 2-8: Detector plane, volume conductor and source space used in the simulations by Sekihara <i>et al</i> (1994). The detector plane consisted of 37 sensors.	89

Chapter 3

figure 3-1: Part of an MRI dataset displayed as a stack of images. Each image shows a coronal view of a cross section of a subject's head. Only every 10 th slice is displayed for the ease of interpretation of the figure (starting with the 70 th and stopping with the 110 th slice). In each slice only some contours are displayed, also for the ease of interpretation. Note that these contours are not part of the original MRI dataset, but are obtained after further image processing (see section 3.2.2).	122
figure 3-2: Diagram that shows how the relative contribution of different tissues to the recorded MR signal is affected by the RF pulse sequence (TE and TR). For example, long TR and TEs highlight differences in T2 between tissues (figure modified from Aine, 1995).	123
figure 3-3: Histogram of slice 100 from the MRI of the head of subject IEH. The peak on the left side is due to the background. The other two peaks represent mainly the grey and white matter.	126
figure 3-4: Example of the right side of a coronal slice (left hemisphere) with parts of the CSF marked (yellow). The sagittal and axial slices corresponding to the crosshairs are also shown. The computed contours (red) need modification in order to mark the outer boundary of the grey matter faithfully. The contours marking the inner and outer skull boundaries can easily be discarded.	133
figure 3-5: A manually edited contour where wrongly placed contours were removed and the different parts of other contours were correctly connected. This contour faithfully represents the outer boundary of the grey matter.	134
figure 3-6: Three contours in a coronal slice through the pole of the occipital cortex. Connecting these three contours to a single contour in an adjacent slice poses a problem for some surface reconstruction algorithms.	135
figure 3-7: Example of a global connection problem. A set of 6 contours in two parallel slices (a). There are several ways to connect the regions in adjacent slices ((b) and (c)) and form a reconstructed object. Intersecting both these reconstructed objects would return the original regions in the cross-section (from Geiger, 1993).	137
figure 3-8: Example of a local connection problem. The two contours can be connected in different ways. No discrimination can be made between the choice in the top panel and the one in the bottom panel, unless additional constraints are used (from Geiger, 1993).	138
figure 3-9: The triangulated cortical surface with the normals in the vertices displayed as arrows pointing in the direction of the normal in each vertex. The left occipital pole is here shown as viewed from the (left) back of the head.	139

figure 3-10: Two active brain areas, displayed in two different ways. The first way is to present the active area as a current dipole, represented by an arrow. The location and orientation of the arrow show the location and orientation of the current dipole. The strength of the dipole is represented by the size of the arrow. The second way to represent the activity is by colour coding each triangle face by the strength of the activity in the vertices of each triangle. In this case strong activity is represented by yellow, and weak activity by dark red. The elements with negligible activity are displayed in pink. 140

figure 3-11: An active brain area, represented by the red current dipole, is displayed on the cortical surface. The source is visible because a part of the outer surface of the head is 'cut-away'. A presentation like this shows the location of an active brain areas with respect to (sulcal) anatomy and its location with respect to the outer surface of the head. 141

figure 3-12: Plot of the ratio of $\sqrt{1/(n+1)}$ and $\sqrt{1/n}$, with n the number of fiducial points. The minimum number of fiducial points needed for target registration is 3. The graph shows that there is a large decrease of the registration error when the number of fiducial points is increased from 3 to 4 or 5. A further increase of the number of points reduces the registration error with a relatively small amount, as is illustrated by the flattening of the curve. 146

figure 3-13: Schematic of the bite-bar system used in Aston. It contains four small holes that are digitised in the MEG co-ordinate system. The markers can also be identified in the MRI when filled with oil, enabling the co-registration of the MEG and MRI co-ordinate systems. 147

figure 3-14: Schematic of the co-registration procedure. The reference coils determine the location of the MEG sensors with respect to the head. The digitised location of the reference coils (or any other chosen fixed reference system) defines the MEG co-ordinate system. Digitisation of the scalp surface within this reference system gives a set of 'hat' points in MEG co-ordinates. A set of 'head' points in MRI co-ordinates is obtained from the subject's MRI through segmentation and boundary extraction. Surface matching then maps the anatomical information from the subjects MRI to the MEG co-ordinate system. Combining all these steps gives the location of the sensors and the anatomical information in the same co-ordinate system. 156

figure 3-15: Example of co-registration of a set of hat points obtained with the 3D digitiser (dots) and the head surface obtained from the subjects MRI. 159

figure 3-16: Plot of the cost-function after several rotations of the set of hat-points from the perfect location. The left panel shows the cost-function for rotation in the y-direction and the right panel the cost-function for rotation in the z-direction. 163

figure 3-17: The mean distance between the head and hat points and the value of the total squared distance at the start and after convergence of the surface matching algorithm. The simulations were repeated 10 times. For each run an initial guess for the three translation and rotation parameters was obtained from a Gaussian distribution with zero mean and a standard deviation of 2mm and 2° respectively. Note the strong correlation between the mean distance and the cost-function. 165

figure 3-18: The mean distance between the head and hat points and the value of the total squared distance at the start and after convergence of the surface matching algorithm. The simulations were repeated 10 times. For each run an initial guess for the three translation and rotation parameters was obtained from a Gaussian distribution with zero mean and a standard deviation of 2mm and 2° respectively. The initial guess for the scaling parameters were taken from a Gaussian distribution with a mean of 1.9 and 1.1mm for the slice thickness and pixel dimension, respectively, and a standard deviation of 0.01mm..... 166

figure 3-19: The mean distance between the head and hat points due to noise alone and at the start and end of the surface matching algorithm. The simulations were repeated 10 times. For each run an initial guess for the three translation and rotation parameters was obtained from a Gaussian distribution with zero mean and a standard deviation of 2mm and 2° respectively. The noise that was added to the location of each hat point was obtained from a Gaussian distribution with zero mean and a standard deviation of 2mm. 170

figure 3-20: Variability due to repeated co-registration for different regions of the brain. The total RMS error is given in the top panel and the standard deviation in the z-direction is given in the bottom panel. Note that the largest errors were found towards the back of the head and that these errors were mainly caused by the variability in the z-direction. 173

Chapter 4

- figure 4-1: Decay pattern of the singular values for several set-ups with different numbers of sensors. With low SNR data only a few singular values are larger than the singular values due to the noise, for set-ups with a few as well as many sensors. The cut-off level is indicated with the line annotated with *Low SNR*. On the other hand, with high SNR data many more singular values are larger than the noise singular values and the difference in the number of singular values that is usable between a set-up with a few or many sensors is large. The cut-off level for high SNR data is indicated with the line annotated with *High SNR* (modified from Gençer and Williamson (1998)). 183
- figure 4-2: Resolution kernels for three different linear estimators for the target point A. The reconstructed activity in A will be influenced by activity from other sources if their value in the resolution kernel is non-zero. For example, activity in B will contribute to the reconstructed activity in A or due to slightly stronger activity in B, so the spatial resolution of the linear estimator is related to the width of the main peak in the resolution kernel. When the Minimum Norm estimator is used, activity in C will have a negative contribution to the activity in A. This could cancel the activity of source A when a strong source is present in C. Activity in C has a small positive contribution to the activity in A when the B&G estimator is used and no contribution if the WROP estimator is used. Hence, the optimal linear estimator should have a narrow peak around the target source (influence of source B is minimal and therefore high spatial resolution) and small sidelobes (influence of distant sources like C is small). From Grave de Peralta Menendez *et al* (1997), see original paper for details about the linear estimators. 187
- figure 4-3: Configuration of the 19 channel CES system (ideal). The space between two sensors is 29mm. 196
- figure 4-4: An MRI slice in which the sources from the realistically shape source space are overlaid. Each source was manually classified as either being gyral (red), sulcal (yellow) or 'undefined' (cyan). 197
- figure 4-5: Histogram of the angles between the sphere radius and the orientations of gyral (left panel) and sulcal (right panel) sources. 198
- figure 4-6: Surface map of the l_2 -norm of the lead field of each element in the source space, viewed from the left (left panel) and right (right panel). The scale is in fTesla. 199
- figure 4-7: Surface map of the l_2 -norm of the lead field of each element in the source space, viewed from left (left panel) and right (right panel). The surface maps are displayed on a logarithmic scale. 199
- figure 4-8: Surface maps of the resolution kernels for a superficial (top row), shallow (middle row) and deep (bottom row) source. The left column contains views from the left and the right column views from the right. The location of the source for which the resolution kernel was computed is displayed as a red arrow. The surface maps for the columns in the resolution matrix were exactly equal to the resolution kernels displayed here, due to the symmetry in the resolution matrix for the minimum norm estimator. Note that the resolution kernel was closest to the ideal one for the superficial source and was worst for the deep source. 203
- figure 4-9: Resolution kernels for a superficial (top row, left), shallow (top row, right) and deep (bottom row, left) source. Each resolution kernel was sorted for the distance between each source and the target source. Hence, the first value on the left in each plot is the value of the resolution kernel for the target source. The plots of the columns in the resolution matrix were exactly equal to the plots for the resolution kernels displayed here, due to the symmetry in the resolution matrix for the minimum norm estimator. The same resolution kernels were displayed as surface maps in figure 4-8. 204
- figure 4-10: Surface maps of the resolution kernels for a superficial (top row), shallow (middle row) and deep (bottom row) source. The left column contains views from the left and the right column views from the right. The location of the source for which the resolution kernel was computed is displayed as a red arrow. 207
- figure 4-11: Resolution kernels for a superficial (top row, left), shallow (middle row, right) and deep (bottom row, left) source. Each resolution kernel was sorted for the distance between each source and the source for which the resolution kernel was given. The same resolution kernels were displayed as surface maps in figure 4-10. 208
- figure 4-12: The top row contains the surface map (logarithmic scale) of the *Source Identifiability* viewed from the left (left panel) and the right (right panel). In the middle row a histogram of the logarithm of *Source Identifiability* is given for both gyral (left in left panel) and sulcal sources (right in left panel) and a plot of the *Source Identifiability* versus the angle that the source orientation formed with the sphere radius (right panel). In the bottom row, the left and right panel show the *Source Identifiability* versus the l_2 -norm of the lead field (Gain) and the average distance to the sensors from each element in the source space, respectively. 209
- figure 4-13: The top row contains the surface map of the *Bias in Dipole Location* viewed from the left (left panel) and the right (right panel). In the middle row a histogram of the *Bias in Dipole Location* is given

for both gyral (left in left panel) and sulcal sources (right in left panel) and a plot of the <i>Bias in Dipole Location</i> versus the angle that the source orientation formed with the sphere radius (right panel). In the bottom row, the left and right panel show the <i>Bias in Dipole Location</i> versus the l_2 -norm of the lead field (Gain) and the average distance to the sensors from each element in the source space, respectively.....	210
figure 4-14: The top row contains the surface map (logarithmic scale) of the <i>Expected squared error due to noise</i> viewed from the left (left panel) and the right (right panel). In the middle row a histogram of the logarithm of the <i>Expected squared error due to noise</i> is given for both gyral (left in left panel) and sulcal sources (right in left panel) and a plot of the <i>Expected squared error due to noise</i> versus the angle that the source orientation formed with the sphere radius (right panel). In the bottom row, the left and right panel show the <i>Expected squared error due to noise</i> versus the l_2 -norm of the lead field (Gain) and the average distance to the sensors from each element in the source space, respectively.....	211
figure 4-15: The surface maps of the <i>Source Identifiability</i> (logarithmic scale), <i>Bias in Dipole Location</i> and <i>Expected squared error due to noise</i> (logarithmic scale) are given in the top, middle and bottom row respectively. The left column contains views from the left and the right column contains views from the right.....	215
figure 4-16: Plot of the <i>Bias in Dipole Location</i> versus the angle that the source orientation formed with the sphere radius (top, left panel), versus the l_2 -norm of the lead field (top, right panel) and versus the average sensor distance (bottom, left panel). The bottom right panel shows the <i>Expected squared error due to noise</i> versus the l_2 -norm of the lead field (Gain).....	216
figure 4-17: Examples of impulse responses typically found with FOCUSS when the identity matrix is used as initial weighting matrix. The source spaces are viewed from the right-back. The left panel shows an ideal impulse response. The one element with a value of 1 is at the location of the original source (indicated with the red arrow). Note that no colourbar is given for this figure; dark blue represents a value of zero. The right panel shows an impulse response for which the maximum value is not in the same location as the original source. Additionally, there are non-zero elements surrounding the original source. Note however that the vast majority of elements in the impulse response have a value of zero.	217
figure 4-18: Weighting values selected to emphasise a target area. The value of the weights fell off with the distance from the target source, following a 3D-Gaussian distribution.	218
figure 4-19: Minimum norm weighting. Top panel: histogram of the values of the <i>Bias in Dipole Location</i> as obtained from the resolution matrix. Bottom panel: histogram of the values of the <i>Bias in Dipole Location</i> as obtained from concatenating the impulse responses of FOCUSS. Note the huge reduction of the number of sources with a large <i>Bias in Dipole Location</i> for FOCUSS.....	220
figure 4-20: Smoothed norm weighting. Top panel: histogram of the values of the <i>Bias in Dipole Location</i> as obtained from the resolution matrix. Bottom panel: histogram of the values of the <i>Bias in Dipole Location</i> as obtained from concatenating the impulse responses of FOCUSS.	221
figure 4-21: Target area weighting. The <i>Bias in Dipole Location</i> as obtained from the resolution matrix. Note the small <i>Bias in Dipole Location</i> for sources in the target area (the target area was defined in figure 4-18).....	222
figure 4-22: Comparison of target area weighting with smoothed norm weighting. The <i>Bias in Dipole Location</i> was obtained from the resolution matrix. Top panel: histogram of the values of the <i>Bias in Dipole Location</i> for target area weighting. Bottom panel: histogram of the values of the <i>Bias in Dipole Location</i> for smoothed norm weighting. The same scale on the x-axis was used for the ease of comparison.....	223
figure 4-23: Comparison of target area weighting with smoothed norm weighting. The <i>Bias in Dipole Location</i> was obtained concatenating the impulse responses of FOCUSS. Top panel: histogram of the values of the <i>Bias in Dipole Location</i> for target area weighting. Bottom panel: histogram of the values of the <i>Bias in Dipole Location</i> for smoothed norm weighting. The same scale on the x-axis was used for the ease of comparison. Note however the difference in scale with figure 4-22.....	223

Chapter 5

figure 5-1: Source space modeled as a cortical fold viewed from the top (left panel) and right (right panel). The channel array is depicted by the filled circles and the centers of the target sources on the gyrus and the sulcal wall are indicated by arrows.	229
figure 5-2: The mean location error [mm] and the standard deviation of the location error [mm] for different Monte Carlo realisation. The results obtained with FOCUSS, unconstrained dipole fitting and	

constrained dipole fitting (A) are given in the top, middle and bottom row, respectively (see text for details about the algorithms). These figures were obtained by adding the solution found for a Monte Carlo realisation to the solutions obtained from previous realisations and computing the error measures. The results displayed were obtained for simulated data with an SNR of 1.5 for the target source on the gyrus. Note the stabilisation of the results as the number of realisations increased..... 231

figure 5-3: Fieldmap produced by the target source on the gyrus (left panel) and sulcal wall (right panel), respectively. The dataset has an SNR of 3 and the source activity is near its maximum. Note that the spatial coverage of the sensor array is big enough to incorporate the extrema in the dipolar magnetic field pattern. 235

figure 5-4: Comparison of compound and non-compound version of FOCUSS, unconstrained dipole fit and constrained dipole fit algorithms for data with different SNRs. Target source is the source on the gyrus. The top panel shows the mean of the distance to the target source for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note the increase in accuracy and the decrease of the spread in the solutions with increasing SNR and the relatively poor performance of FOCUSS..... 239

figure 5-5: Comparison of compound and non-compound version of FOCUSS, unconstrained dipole fit and constrained dipole fit algorithms for data with different SNRs. Target source is the source on the sulcal wall. The top panel shows the mean distance to the target, gyrus, source for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance..... 241

figure 5-6: Solutions obtained from the Monte Carlo simulations with data with an SNR of 1.5, using the compound and non-compound version of FOCUSS. Target source is the source on the sulcal wall. Top figure: view from the right-back. Bottom figure: view from the right. Note the smaller number of unique solutions for the compound version. Also note that the non-compound solutions contain many sources on the wrong side of the gyrus. 242

figure 5-7: L-curves for different iterations of the FOCUSS algorithm. Note the improvement in the shape of the L-curve in the first 4 iterations and the loss of the L-shape in the graphs for iterations 5 to 8. 246

figure 5-8: Timecourses typically found when different values of α were used to set the truncation level for noise contaminated data. The location of the constructed sources is given in the figures. The location of the target source on the gyrus was (-0.137 -0.002 0.043) [m]..... 248

figure 5-9: Normalised error measures for different values of α . The legend denotes the error measure that was used to obtain the corresponding curve. Only the solutions that were different from solutions found for other values of α are displayed. Hence, although only six solutions are displayed here, more values of α were used by FOCUSS. Note that the solution with minimum χ_r^2 was at a different α to the solution that gave the minimum localisation error and the solution selected on the basis of the timecourses..... 248

figure 5-10: The mean χ_r^2 for solutions with 1, 2, or 3 sources, respectively. For a dataset the solutions obtained with different truncation levels were grouped on the basis of the number of sources in the solution. Then, the mean value of χ_r^2 was computed for each set of solutions. Note that the correct number of sources was 1, but that the lowest mean χ_r^2 was found for solutions with three sources. ... 249

figure 5-11: Source space viewed from the back (slightly shifted to the right). The channel array is depicted by the filled circles and the centers of the target sources on the gyrus and the sulcal wall are indicated by arrows. Note that the average distance to the sensors is approximately equal for both target sources. ... 253

figure 5-12: Influence of grid spacing on source reconstructions with FOCUSS for data with different SNRs. The results for the unconstrained dipole fits are shown for comparison. The top panel shows the mean distance to the target, gyrus, source for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note the increase in the mean distance error and the standard deviation of the distance error due to an increase in gridspacing..... 255

figure 5-13: Influence of grid spacing on source reconstruction with FOCUSS for data with different SNRs. The results for the unconstrained dipole fits are shown for comparison. The top panel shows the mean distance to the target, sulcal, source for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note that the results were almost independent of the gridspacing..... 256

figure 5-14: Influence of grid spacing on constrained dipole fits for data with different SNRs. The results for the unconstrained dipole fits are shown for comparison. The top panel shows the mean distance to the target, gyrus, source for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note that the influence of the gridspacing was largest for data with low SNR..... 258

figure 5-15: Influence of grid spacing on constrained dipole fits (with the final solution constrained to the mesh) for data with different SNRs. The results for the unconstrained dipole fits are shown for

comparison. The top panel shows the mean distance to the target, gyral, source for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance.	259
figure 5-16: Influence of grid spacing on constrained dipole fits for data with different SNRs. The results for the unconstrained dipole fits are shown for comparison. The top panel shows the mean distance to the target, sulcal, source for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note that the results were almost independent of the gridspacing.	260
figure 5-17: Influence of grid spacing on constrained dipole fits (with the final solution constrained to the mesh) for data with different SNRs. The results for the unconstrained dipole fits are shown for comparison. The top panel shows the mean distance to the target, sulcal, source for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance.	261
figure 5-18: Rotated mesh with FOCUSS. Influence of grid spacing on the sulcal and gyral target sources for data with an SNR of 1.5. The results for the unconstrained dipole fits are shown for comparison. The top panel shows the mean distance to the target sources for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note that the results for the gyral source were influenced in a different way as the results for the sulcal source.	263
figure 5-19: Rotated mesh with constrained dipole fits. Influence of grid spacing on the sulcal and gyral target sources for data with an SNR of 1.5. The results for the unconstrained dipole fits are shown for comparison. The top panel shows the mean distance to the target sources for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance.	264
figure 5-20: Influence of displacement errors on the source reconstruction with the three different algorithms. The target source is the gyral source. The top panel shows the mean distance to the target sources for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note that the mesh perturbations only had a large influence on CDIPB.	271
figure 5-21: Influence of displacement errors on the source reconstruction with the three different algorithms. The target source is the sulcal source. The top panel shows the mean distance to the target sources for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note that the mesh perturbations only had a large influence on CDIPB.	272
figure 5-22: Influence of errors in the orientation of the surface normals on the source reconstruction with the three different algorithms. The target source is the gyral source. The top panel shows the mean distance to the target sources for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note the increase in performance for FOCUSS and the decrease in performance for the constrained dipole fits, due to the perturbations of the surface normals.	274
figure 5-23: Influence of errors in the orientation of the surface normals on the source reconstruction with the three different algorithms. The target source is the sulcal source. The top panel shows the mean distance to the target sources for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance.	275

Chapter 6

figure 6-1: Gaussian windowed patch of sinusoidal grating with spatial frequency of 2cpd and fixation spot in top right corner. The standard deviation of the Gaussian window is 1.67 degrees. Note that this visual stimulus was viewed through a mirror during the MEG experiment, thus it was presented in the lower right visual field.	281
figure 6-2: Vertices of the realistically shaped source space shown in a coronal slice from the subject's MRI	283
figure 6-3: Figures of merit for the measurement set-up during the recordings. Minimum norm weighting was used in the computations of these figures of merit. Top row: <i>Source Identifiability</i> displayed on a logarithmic scale, viewed from the left-back (left panel) and the right-back (right panel). Bottom row, left panel: <i>Bias in Dipole Location</i> . Bottom, row right panel: <i>Expected squared error due to noise</i> , displayed on a logarithmic scale. The patterns for the figures of merit displayed in the bottom row were very similar when viewed from the right and were therefore omitted. Note the low <i>Source Identifiability</i> for deep sources located towards the midline. Also note the difference in <i>Source Identifiability</i> between sulcal and gyral sources on the left. Furthermore, it is striking that the <i>Bias in Dipole Location</i> was negligible for most elements in the source space. The <i>Expected squared error due to noise</i> displayed a random pattern.	287

figure 6-4: Signal to noise ratio over the latency range from 71 to 122ms. The channels were sorted for SNR. Note that almost half of the sensors had an SNR that was larger than 2.....	288
figure 6-5: Top panel: Averaged evoked responses recorded for each sensor. The signals are displayed on a schematic of the subject's head (nose pointing upwards and the left and right ear on the left and right, respectively). Note the stimulus related response that is visible in the occipital sensors and the responses of opposite polarity over the left and right side of the head at a later latency. Bottom panel: The same schematic of the head, displaying an interpolated field map of the recorded magnetic field at a latency of 87ms. The circles denote the sensor locations. A dipolar field pattern is visible over the left occipital pole, suggesting that there was an active source in the visual cortex at this latency.....	290
figure 6-6: Averaged evoked responses recorded in the sensors placed over the occipital cortex. An evoked response triggered by the presentation of the visual stimulus is clearly visible.....	291
figure 6-7: The timecourses for the occipital channels overlaid. Only the early visual evoked response is shown (60-130ms). The timecourses show a complicated activity pattern. The sensors with the strongest signals show a positive peak in activity at about 75 and 110ms and a negative peak at about 95ms. Note however that the maxima in each sensor were shifted in time with respect to one another.	292
figure 6-8: Summation over time of the $\ \text{minimum norm solutions} \ $, with different levels of singular value truncation. From left to right, top to bottom, the truncation level was set at the 3 rd , 10 th , 30 th and 43 rd singular value, respectively. Noise truncation at intermediate levels gave very similar results to those shown here.	293
figure 6-9: Left panel: Activation curves of the 100 strongest sources in the minimum norm solution. All these elements are located in two hotspots in the occipital pole (see figure 6-8), apart from a few spurious sources. Right panel: Activation curves of the elements in the 3 rd hotspot in the minimum norm solution. Note that these sources have an activity pattern similar to the sources in the occipital pole as well as similar peak strength.	294
figure 6-10: Locations of a single moving dipole fitted to the data over the latency range 72-122ms, forming a 'trace' of dipole locations. The trace is formed by the dipole moving from an anterior-inferior position to a more posterior-superior area in the left hemisphere and then to the right hemisphere and even more posterior.	295
figure 6-11: Sources reconstructed with FOCUSS, using a range of truncation levels and different initial weightings. Top row: minimum norm weighting. Middle row: smoothed norm weighting. Bottom row: average distance weighting. Note that with all three weighting matrices two separate clusters of sources were found, most clearly defined in the axial and sagittal views. The spread in the solutions was largest in the superior-inferior direction, especially for the anterior cluster. The posterior cluster also showed a lateral prolongation.	296
figure 6-12: Best solution found with FOCUSS, using minimum norm weighting for source reconstruction over the latency range from 71 to 122ms. The top panel shows the source locations and the bottom panel the activation curves and the reduced chi-square error function. Two sources were reconstructed in the occipital pole (dipole 2 and 3), with dipole 2 located towards the left. This source had the strongest activity, followed by the activity of dipole 3. The peak activity of these two sources was shifted in time, with the dipole located towards the midline reaching maximum activity first. A third source with smaller activity, dipole 1, was located in a more anterior-superior region. This dipole peaked at a later latency.	297
figure 6-13: Best solution found with FOCUSS, using smoothed norm weighting. for source reconstruction over the latency range from 71 to 122ms. The top panel shows the source locations and the bottom panel the activation curves and the reduced chi-square error function. Two sources were reconstructed in the occipital pole, with dipole 1 located towards the left. This source had the strongest activity. The peak activity of the two sources was shifted in time.	298
figure 6-14: Best solution found with FOCUSS, using average distance weighting for source reconstruction over the latency range from 71 to 122ms. The top panel shows the source locations and the bottom panel the activation curves and the reduced chi-square error function. Two sources were reconstructed in the occipital pole, with dipole 1 located towards the left. This source had the strongest activity. The peak activity of the two sources was shifted in time, with the dipole located towards the midline reaching maximum activity first.	299
figure 6-15: Unconstrained single dipole solutions. Top row: Location of all sources found with the multi-start approach. Middle row, left panel: Activation curve and error function for the best solution. Note the tight cluster formed by the solutions. Middle row, right panel: Best solution overlaid on the subject's MRI (axial slice). Bottom row: Best solution overlaid on a coronal (left panel) and sagittal slice (right panel).	301
figure 6-16: Constrained single dipole solutions. Above: Location of all sources found with the multi-start approach. Left: Activation curve and error function for the best solution.	302

figure 6-17: Recorded magnetic field, magnetic field induced by the (unconstrained) dipole model and the difference between the recorded and calculated magnetic fields. Left panel: field maps at 83ms. Right panel: field maps at 108ms. Note the dipolar patterns in the difference maps, indicating that an extra dipole should be added to the source model at these latencies.....	303
figure 6-18: Unconstrained two-dipole solutions. Above: Locations of all sources found with the multi-start approach. Left: Activation curves and error function for the best solution. Note the two separate clusters formed by the multi-start approach.....	304
figure 6-19: Recorded magnetic field, magnetic field induced by the (unconstrained) two-dipole model and the difference between the recorded and calculated magnetic fields. Left panel: field maps at 83ms. Right panel: field maps at 108ms. Note that the difference maps did not contain a distinct dipolar pattern (compare with figure 6-17).	305
figure 6-20: Constrained two-dipole solutions. Above: Location of all sources found with the multi-start approach. Left: Activation curve and error function for the best solution.	306
figure 6-21: Solution for the constrained dipole fit algorithm (constrained to lie within the mesh) overlaid on the subject's MRI. Top row, from left to right: dipole 1 in an axial, coronal and sagittal slice, respectively. Bottom row, from left to right: dipole 2 in an axial, coronal and sagittal slice, respectively.	307
figure 6-22: Sources reconstructed for the experiment described in Barnes <i>et al</i> (1998). Note that the locations of these sources correspond closely to the sources reconstructed in the study reported in this chapter (compare with figure 6-12), despite the use of a different (19 channel CES) MEG system. Figure taken from Barnes <i>et al</i> , 1998.....	309

TABLES

Chapter 3

table 3-1: Mean location and standard deviation of nasion (NA), left pre-auricular (PAL) and right pre-auricular (PAR) as obtained by 10 times digitising the location of a coil at these positions. Each coil was localised twice in a run (NA1 and NA2 etc.).	157
table 3-2: Difference in mean location of the different coils (definitions as in table 3-1) in the x, y and z direction and the distance ϵ between the mean locations. Values were obtained by repeatedly (N=10) digitising the location of a coil at these positions. Each coil was localised twice in a run (NA1 and NA2 etc.).	158
table 3-3: Mean distance from the digitised location of a coil to the mean location as obtained from 10 repeats. Experiment and definitions as in table 3-1.	158
table 3-4: Standard deviation of the repeated (N=10) digitisation of the 4 base points in the bite-bar.	158
table 3-5: Value of the cost-function (total squared distance, d^2) after optimisation of the translation parameter (z direction). The initial guess for this parameter was off by 1-10mm in the z-direction.	164
table 3-6: Value of the cost-function (total squared distance, d^2) after optimisation of the rotation parameter (z direction). The initial guess for this parameter was off by 1-9 degrees in the z-direction.	164
table 3-7: Influence of the number of head points on co-registration accuracy. The mean distance per hat point after co-registration is given in the table. Note that there was a decrease in accuracy when the number of points was increased from 100 to 200 and from 300 to 400. However, the trend was that the accuracy increased when the number of points was increased.	168

Chapter 4

table 4-1: Standard deviation of the figures of merit for the gyral and sulcal sources.	212
---	-----

Chapter 5

table 5-1: Mean distance error for solutions found with FOCUSS for Monte Carlo simulations (SNR of 1.5). For each dataset a set of solutions is obtained for different truncation levels. The best solution is then selected with each method. This is repeated for all datasets in the Monte Carlo simulations and the mean location error is computed for each method. Note that the method based on the timecourses outperforms the method based on the value of χ_r^2 .	249
--	-----

1 Introduction

There is an ongoing interest in the functioning of the normal and pathological human brain. Electroencephalography (EEG) and Magnetoencephalography (MEG) can be used to localise areas in the brain that are electrically active, and can therefore be used to obtain information about the functioning of the brain. The advantage of both EEG and MEG above other brain imaging techniques like functional Magnetic Resonance Imaging (fMRI) and Positron Emission Tomography (PET) is the higher temporal resolution (figure 1-1). The MEG and EEG are directly related to the electrical activity of the neurons in the brain whereas fMRI and PET are based on measures of an increase in metabolism due to neuronal activity. These processes are inherently slower and are only an indirect measure of the electrical activity of the neurons. With EEG the electric potential distribution on the scalp and with MEG the magnetic field outside the head is measured. The advantage of

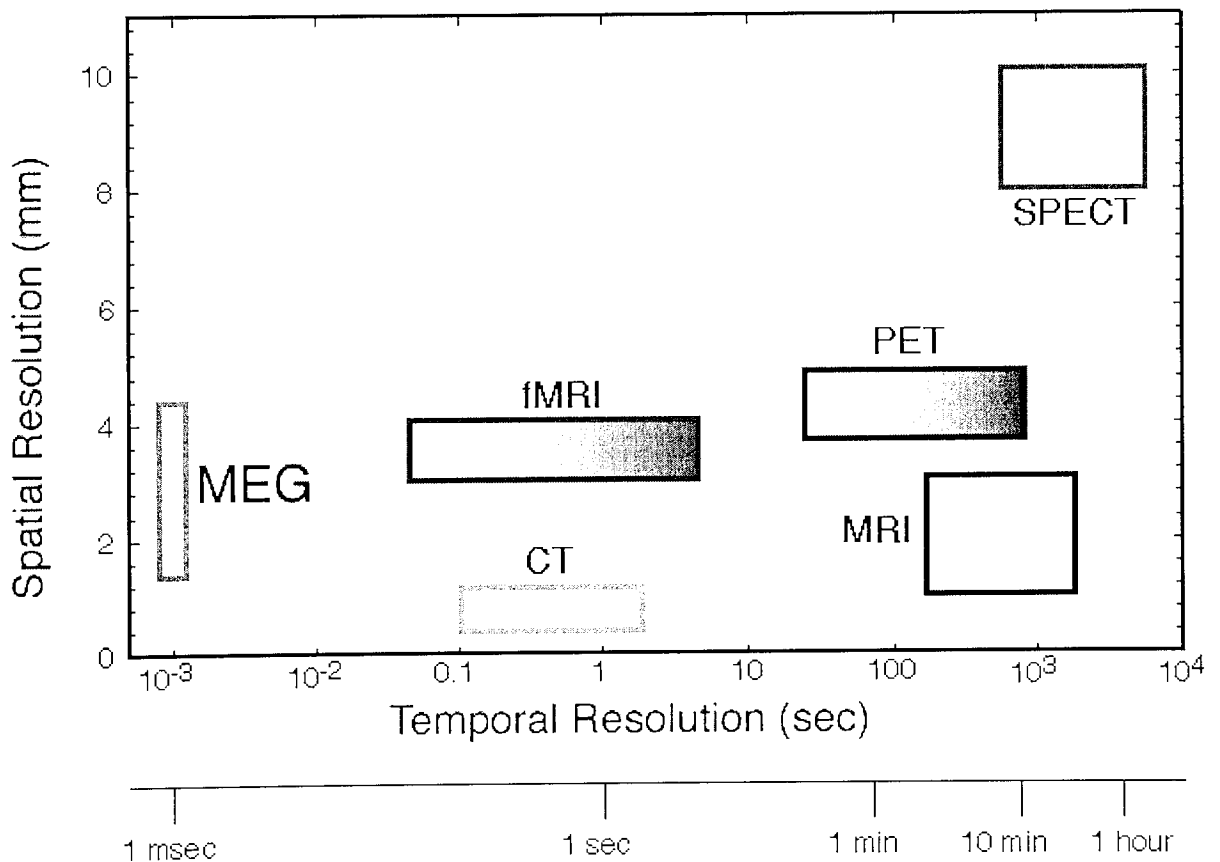


figure 1-1: Comparison of spatial and temporal resolution of different brain imaging techniques. Note the superior temporal resolution of MEG. Source: CTF Systems, Inc.

MEG above EEG is that the magnetic fields are less influenced by the volume conducting properties of the head and that no reference is needed, as MEG measures field strength instead of a potential value. Therefore, the spatial resolution of MEG is higher than that of EEG, whereas the temporal resolution is the same. Another advantage is that there is no contact with the subject or patient (no electrodes have to be attached to the skin, which can take up to 1.5 hours for 124+ electrodes (Gevins *et al.*, 1994), although the use of an electrode cap speeds up the procedure of attaching electrodes), making it an ideal non-invasive method to monitor brain activity for research and clinical purposes.

The measured potential distribution and/or magnetic field can be used to determine the electrical brain activity, using volume conducting theory based on the law of Biot and Savart. In order to solve this so-called bioelectromagnetic inverse problem, one must solve the forward problem. This is the problem of computing the output of the sensors for a certain amount of activity in the brain. This is only possible when the brain activity is modelled by current sources and the head by a volume conductor. It is, in principle, possible to solve the inverse problem when the forward problem is solved, as shown in figure 1-2. A problem is that the solution of the inverse problem is non-unique, which means that the measured potential distribution or magnetic field can be explained by an infinite number of sources. When additional information obtained from anatomy or physiology is introduced into the inverse problem, the problem of non-uniqueness is reduced, resulting in more stable and hopefully more accurate reconstructions of brain activity. Examples of *a priori* information are assumptions about the number of active brain regions or the possible locations of these regions. The main bottleneck is that it is not straightforward to mathematically encode the information from neurophysiology and neuroanatomy so that it can be used with the algorithms that solve the inverse problem.

There are several optimisation algorithms that all try to minimise the difference between the predicted and measured data in some way. The scheme in figure 1-2 shows that it is possible to add *a priori* knowledge to the inverse problem before the optimisation algorithm starts (i.e. possible source locations are added) or to add the information during the optimisation. We will use an optimisation algorithm for which the implementation of the later scheme is relatively easy.

The result of the optimisation algorithm is the reconstructed source, i.e. the source model with optimum parameters, and often an indication of the goodness of the reconstruction. The results should be assessed with care, since all the source reconstruction algorithms are based on several assumptions. Additionally, the source reconstructions are also influenced by all sorts of errors, i.e. measurement errors, noise, modelling errors and computing errors. The validity of the assumptions and therefore of the source reconstructions should be tested against expert knowledge on anatomy and physiology.

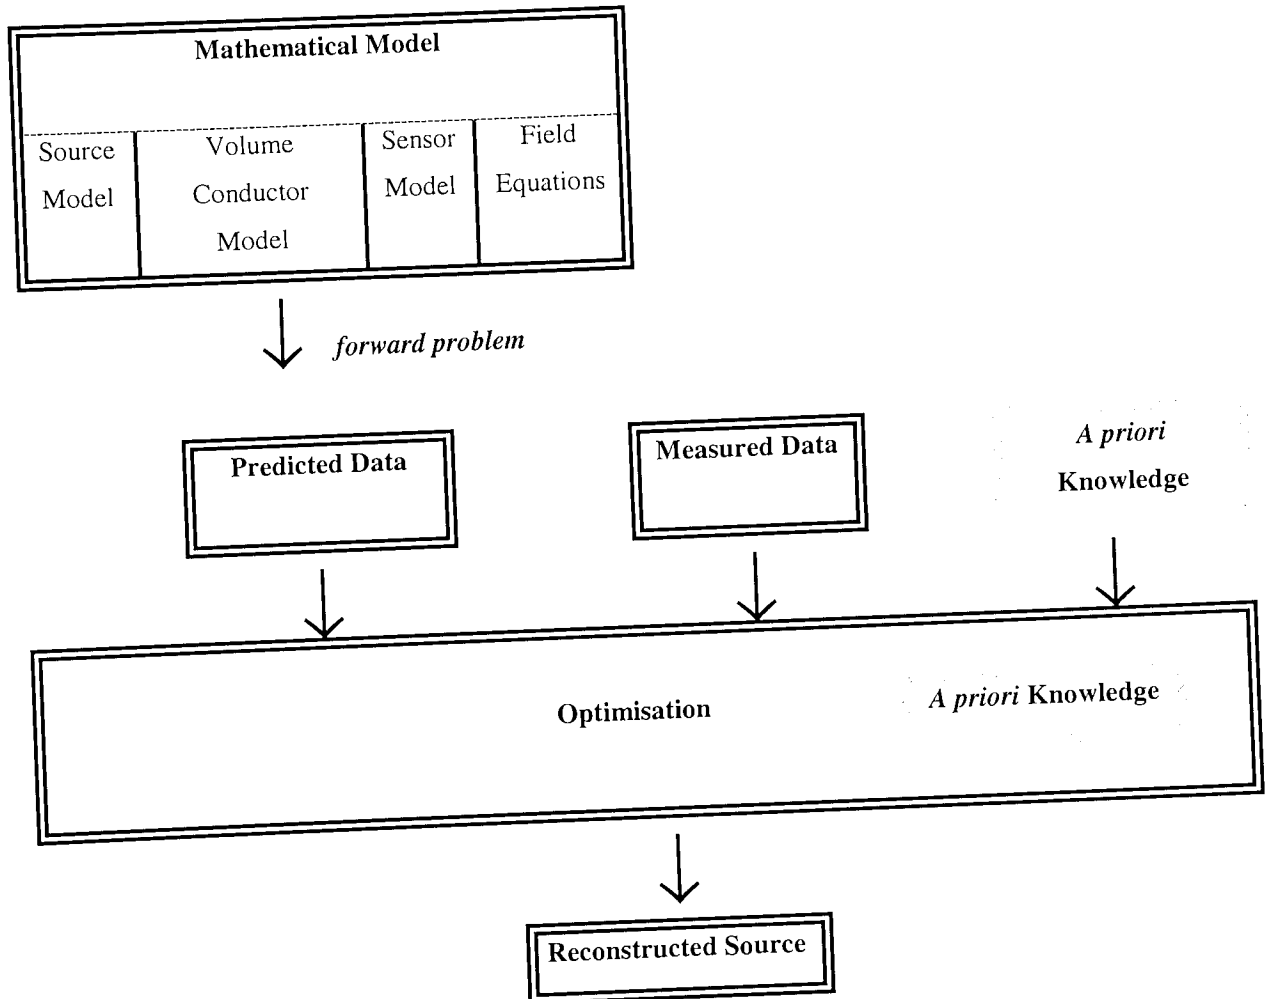


figure 1-2: Basic scheme to solve the inverse problem. Solving the inverse problem relies on the ability to solve the forward problem accurately. The forward problem can be solved when modelling assumptions are made about the sources in the brain and the head as a volume conductor. Furthermore, the sensor configuration must be encoded. Adding *a priori* knowledge at different stages can increase the accuracy of the solution of the inverse problem.

1.2 Objectives

The temporal resolution of MEG allows us to measure the magnetic fields induced by spontaneous brain activity or activity linked to a stimulus on a millisecond by millisecond basis. The theoretical accuracy of the localisation of the underlying sources has only been realised in favourable cases, where the assumptions of a spherical head-shape and dipolar sources are sufficient. Accurate localisation is hindered by the non-uniqueness of the inverse problem. The aim of this work is evaluate the influence of spatial constraints on the accuracy of inverse solutions.

In chapter 2 an overview of the literature is given of the techniques that are currently used to solve the inverse problem. The FOCUSS algorithm is then selected as the favoured approach. This algorithm enables the incorporation of *a priori* information in an easy and intuitive manner. Spatial constraints are included by the selection of a realistically shaped source space, i.e. shaped as the surface of a subject's cortex. Imposing this constraint restricts the sources of electrical activity to lie within the cortex, which is justified in section 1.3.2.

In chapter 3 it is demonstrated how such a realistically shaped source space can be constructed from the anatomical information provided by a subject's MRI. The use of a realistically shaped source space relies on accurate co-registration of the MRI and MEG coordinate systems. Therefore, several co-registration techniques are reviewed in the second part of chapter 3 and the accuracy of an implementation of the surface matching approach is evaluated.

As mentioned before, due to the non-uniqueness of the inverse problem there are many possible solutions that can explain the data. It is therefore necessary to assess the goodness of a solution and the imposed constraints by means of some kind of figure of merit. The FOCUSS algorithm is a hybrid method that starts of with a solution based on linear estimation and converges to a sparse solution, consisting of only a few discrete sources. In chapter 4 figures of merit for linear distributed source reconstructions are evaluated and the applicability to the FOCUSS algorithm is appraised.

In chapter 5 the performance of the FOCUSS algorithm in combination with spatial constraints is compared with the accuracy of an unconstrained dipole fit algorithm and a dipole fit algorithm that constrains the solutions to a small volume around the surface of a simulated sulcal structure. In chapter 6 the same algorithms are compared, using visually evoked magnetic fields recorded with a 151 channel MEG system and a cortical surface constraint, obtained from the subject's MRI.

Finally, a discussion of the main findings of this work follows in chapter 7 and recommendations for further research are given.

1.3 Forward Problem

1.3.1 Introduction

The first recordings of magnetic fields induced by electrical activity within the human brain were performed in 1968 (Cohen, 1968). The magnetic fields produced by alpha activity were detected with a 1 million-turn coil, located in a magnetically shielded room. Signal averaging was needed to increase the signal-to-noise ratio so that the signal could be detected. Non-averaged magnetic recordings of alpha-activity (Cohen, 1972) was only feasible after the introduction of a much more sensitive device, the Superconducting Quantum Interference Device or SQUID. A SQUID consists of a superconducting ring containing two weak links known as Josephson Junctions. Such a configuration has a flux-voltage characteristic that is periodic in applied flux where the period is a flux quantum. By using specialised feedback electronics to maintain the SQUID at a fixed operating point, it is possible to resolve changes in flux of as little as 10^{-5} quanta (see Hämäläinen *et al.* (1993) and Romani (1989) and references therein for more details about the basic principles of SQUIDs).

In order to couple the SQUID to the outside world, and improve its immunity to noise, an external flux gathering circuit known as a flux transformer is used (Romani, 1989). The integrated gradiometer exploits the fact that distant noise sources produce magnetic fields that are almost spatially uniform at the pickup coil, whereas the nearby brain sources produce magnetic fields with (large) spatial gradients. Consider two identical coils that are connected in series but wound in opposite directions (figure 1-3). This arrangement is known as a first-order gradiometer and the distance between the two coils is called baseline. Magnetic fields from distant sources induce virtually identical currents in the two coils giving no net output from the gradiometer. Magnetic fields due to sources near (about a baseline away)




figure 1-3: Schematic of a first order axial gradiometer consisting of two oppositely wound coils inductively coupled to the SQUID sensor.

to the pick-up coil will induce a considerably larger current in the pick-up than in the compensation coil, leading to a net gradiometer output. This arrangement gives noise immunity to distant (noise) sources. Higher order gradiometers can be used to reduce the influence of environmental noise sources even further.

The gradiometer configuration not only reduces the sensitivity to distant noise sources, but also the effective signal that is obtained from deep sources in the brain. Hence, there is an increase of signal-to-noise ratio (SNR) due to reduction of noise and a reduction of SNR due to loss of signal. The optimum gradiometer configuration is the one that maximises the SNR. Details of the estimation of the optimum sensor configuration can be found in Vrba *et al.* (1996), Vrba (1997), Vrba and McKay (1997), Vrba (1998) and Vrba and Robinson (1998). Further details about the general principles of MEG can be found in reviews by Hari and Ilmoniemi (1986) and Hämäläinen *et al* (1993).

The recorded MEG data are usually presented as a set of traces, one for each measurement channel, and/or as a topographic map. These are simply two ways of presenting the same data, where the traces provide information about the timing of events on a millisecond scale and a topographic map provides clues about the spatial information in the data. An insightful way to think about the measured data is that it forms a matrix in which each row is the signal vector of a measurement channel and each column is a sample vector (Hjorth, 1989). Hence a signal vector contains the recorded values for a channel at all latencies, i.e. it gives the timecourse. The sample vector contains the recorded values in all channels at a certain latency. Plotting these values at the locations of the measurement channels and estimating the values for inter-sensor locations via interpolation forms a topographic map (Duffy, 1982). The above-described representations of the data are illustrated in figure 1-4.

Active brain areas can be determined from the measured magnetic field outside the head by solving the inverse problem. In order to solve this inverse problem, it is necessary to solve the forward problem. In order to do so, models of both the electrical activity and the head are needed. It is therefore discussed in section 1.3.2 how the neurons in the brain, which produce the measurable electromagnetic fields, can be described mathematically. Using this formulation together with a mathematical description of the head, described in section 1.3.3, we are then able to give a mathematical formulation of the forward problem in section 1.3.4.

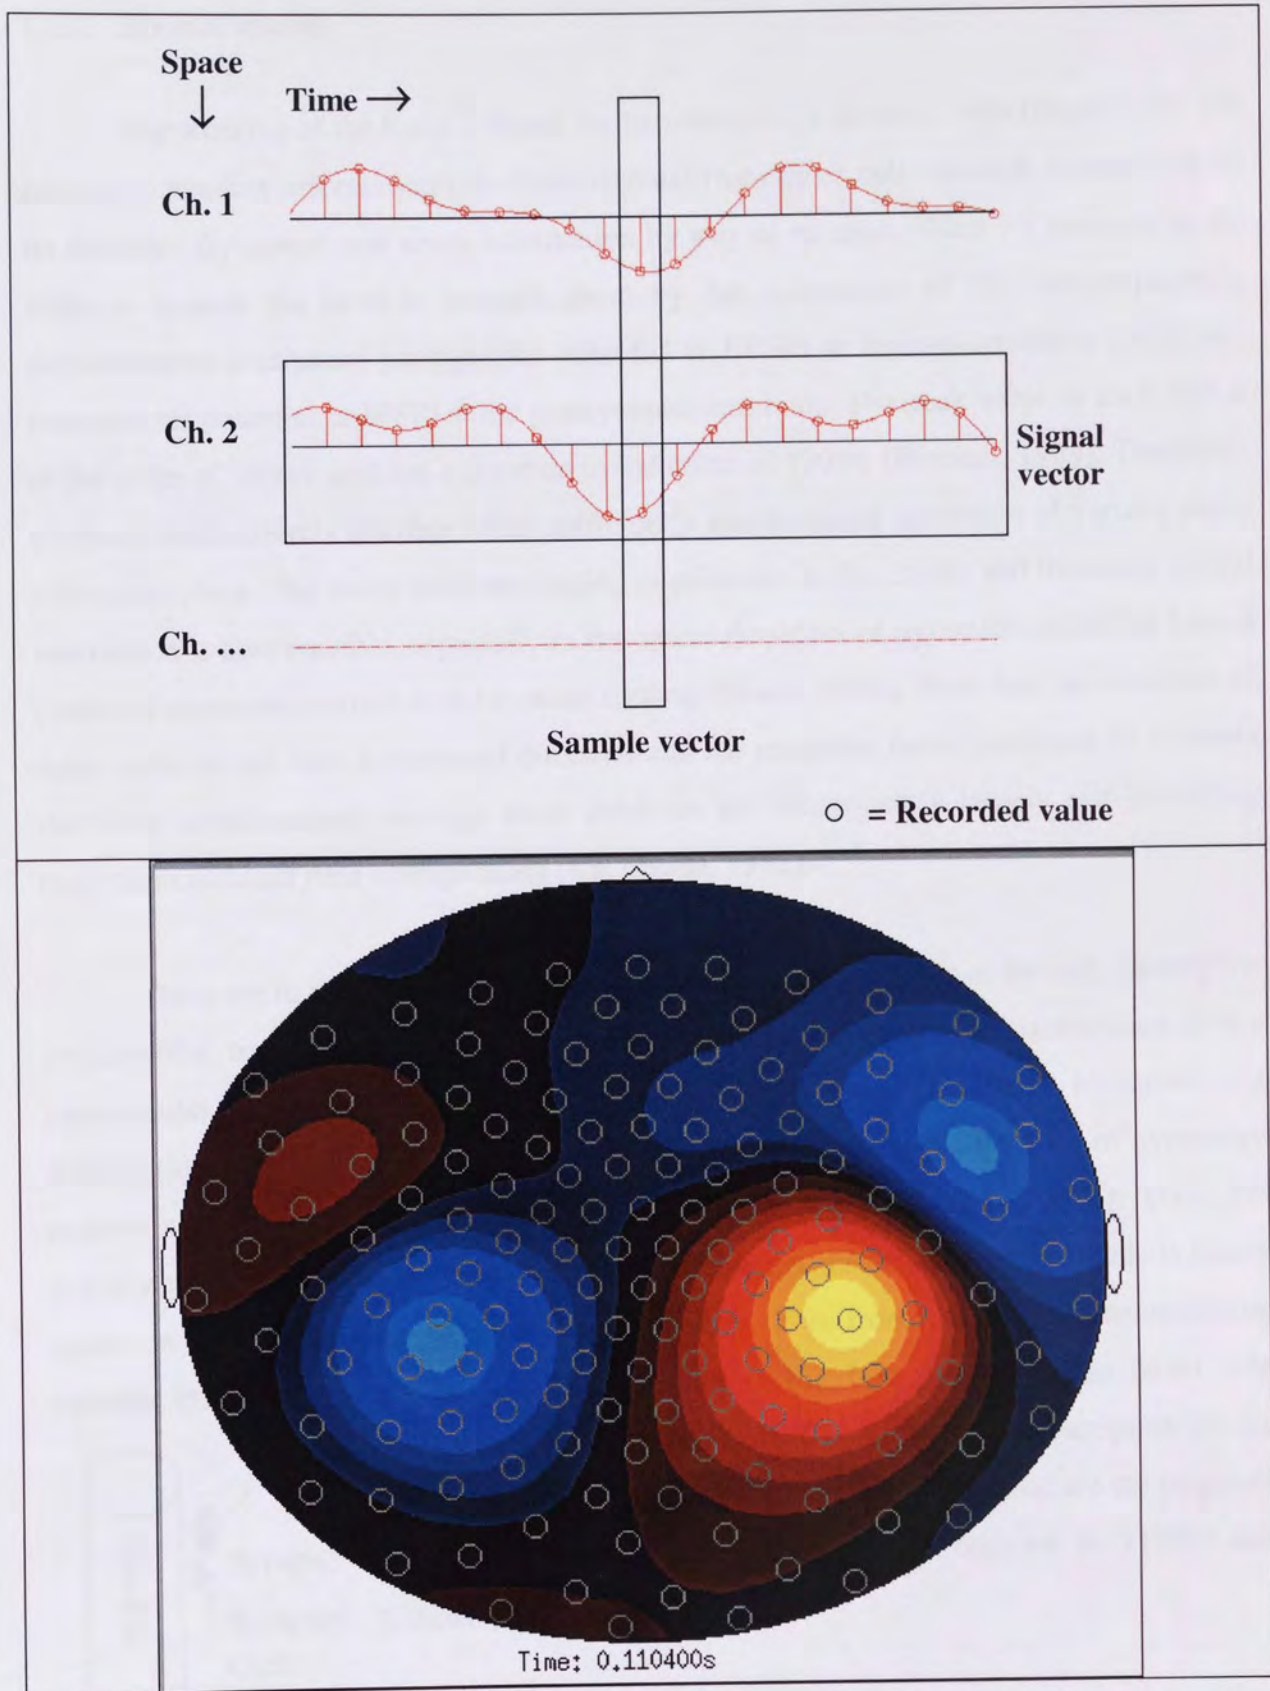


figure 1-4: Matrix representation of the recorded data (top panel). The signal vector contains the recorded values for a channel at all latencies. The sample vector contains the recorded values in all channels at a certain latency (after Hjorth, 1989). In the bottom panel a colour-coded topographic map of a sample vector is shown. In this case the sample vector is the measured field pattern at 110ms after onset of an auditory stimulus. Two dipolar field patterns are clearly visible in the data, suggesting that the magnetic field is induced by two areas of brain activity that can both be modelled by an equivalent current dipole in the left and right auditory cortex, respectively.

1.3.2 Source Model

The working of the brain is based on the interactions of nerve cells (figure 1-5). The nerve cell receives information (electrical signals) from other cells through connections on its dendrites (synapses) and sends information by way of its axon. What we measure on the scalp or outside the head is brought about by the summation of the non-propagating depolarisation (excitatory postsynaptic potential or EPSP) or hyperpolarisation (inhibitory postsynaptic potential or IPSP) at the postsynaptic terminals. The peak value of each PSP is of the order of 10 mV and has a duration in the order of 100ms (Romani, 1989). Therefore, temporal summation is possible when sufficiently synchronised activation of various nerve cells takes place. The nerve cells are aligned in palisades in the cortex and therefore spatial summation is also possible, especially in the apical dendrites of pyramidal cells that have a preferred alignment normal to the cortical surface (Okada, 1982). Note that the dendrites of many cells do not have a preferred direction and the magnetic fields produced by currents that flow simultaneously through these dendrites are consequently largely self-cancelling (they form a *closed field* configuration (e.g. Okada, 1982)).

There are in principle three currents associated with the PSPs at the cell, namely the extracellular, transmembrane and intracellular currents. However, the contribution of the extracellular current can be neglected (Swinney and Wikswo Jr., 1980). Moreover, the contribution of the transmembrane current can also be disregarded because of symmetry reasons and the thinness of the membrane (Swinney and Wikswo Jr., 1980). Thus, the primary current, which is the origin of the detectable neuroelectromagnetic fields, is solely based on the intracellular currents (Okada, 1989). It is ironic that the transmembrane currents that give rise to the cellular excitation in the first place play no direct role




figure 1-5: Schematic drawing of a nerve cell.

(Williamson and Kaufman, 1987). A complete description of the complex electrical processes at a cellular level that are the origin of these currents is given in Swinney and Wikswo Jr. (1980) and Wikswo (1989).

The electrical currents in the dendrites can be modelled by current dipoles in most cases (Okada, 1982). A current dipole is an abstraction that has the dimensions of current times length, though the length is defined as being infinitesimally small. A patch of




figure 1-6: An area of active cortex can be modelled as a layer of current dipoles. This layer of dipoles in turn can be modelled as a single equivalent current dipole when viewed from a large enough distance.

active cortex may be considered as a set of current dipoles distributed in a palisade, as shown in figure 1-6. This forms a current dipole layer, which can be considered as a single equivalent current dipole when observed from a distance much larger than the dimensions of the area of active cortex (e.g. Meijs *et al.*, 1988). The neuroelectromagnetic fields generated by the active cortex area can then be expressed in general terms by a multipole expansion, an infinite sum of multipoles of increasing order (Malmivuo and Plonsey, 1995). When the active area is folded or the measurement distance is not large with respect to the dimensions of the active area, it is necessary to include higher order terms in the source model, such as the quadrupole. It must be stated here that the applicability of the dipole concept follows from the multipole expansion. It does not depend on the fact that the intracellular currents in the neurons have a dipolar structure, as is sometimes mentioned in introductory texts (Snyder, 1991).

The success of the single dipole model follows from the fact that at a typical measurement distance, which is at least 2cm from the source, many current configurations seem 'dipolar' (Mitzdorf, 1985, Hämäläinen and Sarvas, 1989, Hari, 1991). Another aspect is that other source configurations, like the quadrupole, are a superposition of single dipoles. Furthermore, it has been shown that the potential distribution produced by extended sources (dipole disks and dipole annuli) in a spherical head model can be explained by a single dipole without introducing large errors (de Munck *et al.*, 1988, Hara *et al.*, 1999). The deviation between the positions, orientations and strengths of the extended sources and single dipole were small. The results were almost independent of the depth of the sources. This showed that modelling realistic sources by single dipoles does not necessarily yield large errors in location and moment parameters.

A current dipole is completely described by six parameters, three to establish its position within the head and three to define its strength¹. From a mathematical point of view, the maximum number of parameters that can be estimated is equal to the number of independent measurements. This is however an upper limit because noise in the data lowers this number.

1.3.3 Volume Conductor Model

The head can be modelled as a volume conductor and the different biological tissues that constitute the head can be treated macroscopically. This means that inhomogeneities on a cellular level can be neglected. The primary sources described in the previous section set up a passive Ohmic current in the volume conductor, called the secondary or volume current. This volume current completes the loop of ionic flow so that there is no build up of charge (quasi-static approach, see next section). If homogeneous (same value everywhere in space) and isotropic (same value in every direction) conductivity is assumed within the volume, then the effect of the volume currents can equivalently be described by so-called secondary sources at the boundary of the volume, with an orientation perpendicular to the boundary (Geselowitz, 1970).

The geometry of the head is also included in the volume conductor models. The models most commonly used with EEG consist of three concentric spheres representing the brain, the skull and the scalp. In order to give accurate localisation of the current sources it is necessary that the spheres approximate the circumferences of the head compartments well, especially around the source area. This is not always possible, for example when the current sources are located in the temporal area of the brain. In these cases more realistic modelling with the boundary element method (BEM) or finite element method (FEM) is necessary (e.g. Geselowitz, 1970, Hämmäläinen and Sarvas, 1989). With the BEM the different tissues in the head are represented by layers with homogeneous and isotropic conductivity and the boundaries are discretised into triangular elements, so that the effect of the secondary sources can be modelled. With the FEM the volume conductor is divided into

¹ It is not possible to measure the radial component from magnetic measurements when a current dipole is placed in a spherical symmetric volume conductor, as will be shown in section 1.3.4. It is therefore fully

many volume elements of regular shape (tetrahedrons, cubes). The advantages of FEM is that each volume element can be assigned a different conductivity tensor (although such detailed conductivity information is not yet available) and that a detailed model can be constructed. The disadvantages of these numerical methods are that they are time-consuming when used in the inverse problem and that the construction of the models takes a lot of time and computer resources. Additionally, large errors can be introduced in the forward problem if the BEM implementation is numerically unstable due to the large difference in conductivity between the skull and surrounding tissues (Mosher *et al.*, 1997, Zanow, 1997, Mosher *et al.*, 1999). Luckily, with MEG it is usually sufficient to use a spherical model that represents the inner part of the skull near the sources, since the influence of the volume conductor on the magnetic data is largest near the sources (Mondt, 1989). Also, the overall volume conductors effects are small, for reasons described in the next section. For the interested reader, details about the BEM are given in de Munck (1992) and Ferguson *et al.* (1994) and for the FEM in van den Broek (1997).

1.3.4 Theory of Volume Conducting

The aim of neuroelectromagnetism is to obtain information about the primary currents (\mathbf{J}_p), which are directly driven by the neuronal processes. However, there is, in general, no direct relationship between the measured magnetic field or potential distribution and \mathbf{J}_p , since this data is also influenced by the secondary currents (\mathbf{J}_v), which are the Ohmic currents in the volume conductor (Nicholson and Freeman, 1975). The total current density, \mathbf{J} , can be expressed as:

$$\mathbf{J} = \mathbf{J}_p + \mathbf{J}_v = \mathbf{J}_p + \overline{\overline{\sigma}} \mathbf{E}, \quad (1-1)$$

where σ is the conductivity of the volume conductor and \mathbf{E} is the electric field strength in the conductive medium. The double lines above the conductivity indicate that it is a tensor. Throughout this work vector quantities are denoted as bold case symbols.

The mathematical basis that gives the relationships between the total current density and the magnetic field and the potential distribution are the Maxwell equations:

- $curl \mathbf{H} = \mathbf{J} + \frac{\partial \mathbf{D}}{\partial t},$ (1-2)

- $curl \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t},$ (1-3)

- $div \mathbf{B} = 0,$ (1-4)

- $div \mathbf{D} = \rho,$ (1-5)

\mathbf{H} , \mathbf{B} , \mathbf{D} and ρ stand for magnetic field strength, magnetic flux density, electric displacement density² and charge density, respectively. The partial derivative with respect to time is given as $\partial/\partial t$.

We also make use of the following relations, assuming that the medium is linear (Reitz *et al.*, 1992):

$$\mathbf{D} = \bar{\bar{\epsilon}} \mathbf{E}, \quad (1-6)$$

$$\mathbf{B} = \bar{\bar{\mu}} \mathbf{H}, \quad (1-7)$$

Here, ϵ is the electric permittivity and μ is the magnetic permeability. The medium is assumed to be homogeneous and isotropic, turning all tensors into constant scalars. As the magnetic permeability of the brain is close to that of water, which is only weakly diamagnetic, it is assumed that the magnetic permeability is that of free space, $\mu = \mu_0 = 4\pi \cdot 10^{-7}$.

Using the criteria given by Plonsey and Heppner (1967) it can be demonstrated (van den Broek, 1997) that for neuroelectromagnetic signals, which have frequencies far below 1kHz, the time dependencies can be neglected for most of the values of the tissue parameters, σ , ϵ and μ found in the literature. So the effect of capacitances can be neglected

² $\partial \mathbf{D} / \partial t$ is known as the displacement current.

(Nicholson and Freeman, 1975) and the problem reduces to a stationary problem, giving the following reduced Maxwell equations:

- $curl \mathbf{H} = \mathbf{J}$, (1-8)

- $curl \mathbf{E} = 0$, (1-9)

- $div \mathbf{B} = 0$, (1-10)

- $div \mathbf{D} = \rho$, (1-11)

The divergence of the curl of another vector field is zero, therefore equation 1-10 implies that \mathbf{B} can be written as:

$$\mathbf{B} = curl \mathbf{A}, \quad (1-12)$$

where \mathbf{A} is the magnetic vector potential.

This leads with the reduced Maxwell equations to the following Poisson equation:

$$\nabla^2 \mathbf{A} = -\mu_0 \mathbf{J}, \quad (1-13)$$

using $div \mathbf{A} = 0$.

The solution of this equation, found by using the boundary condition that the magnetic flux density is zero at infinity, is called the law of Biot and Savart. It relates the current density in a region Ω of the volume conductor to the magnetic induction at observation point P, located at \mathbf{r} (figure 1-7).




figure 1-7: Definition of the location of the current density and the measurement position with respect to the origin.

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi} \int_{\Omega} \frac{\mathbf{J}(\mathbf{r}') \times (\mathbf{r} - \mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|^3} d^3r', \quad (1-14)$$

The law of Biot and Savart acts as a spatial low-pass filter, because of the integration over space. A consequence of this is the so-called smoothing effect. This can result in a misinterpretation of the measured fields, as for example the complex spatial pattern of several point sources can be smoothed, with a dipolar field pattern as a result (Kullmann, 1991). Additionally, it can be seen that the magnetic induction falls off with the squared distance from the source. Hence, superficial sources induce a much stronger signal in a sensor than deep sources.

It has been shown that the secondary currents in a spherically symmetric volume conductor do not contribute to the radial component of the magnetic flux density (Geselowitz, 1970, Nunez, 1986, Sarvas, 1987). To reiterate: the volume currents are influenced by the conductivity profile, but the influence on the magnetic field is in general small.

If we further assume that the primary current density at \mathbf{r}' can be modelled as a current dipole, we can write:

$$\mathbf{J}(\mathbf{r}') = \mathbf{J}_p(\mathbf{r}') = \delta(\mathbf{r} - \mathbf{r}')\mathbf{M}, \quad (1-15)$$

with δ the Dirac delta function and \mathbf{M} the current dipole moment.

Sarvas (1987) showed that the law of Biot and Savart can then be written as:

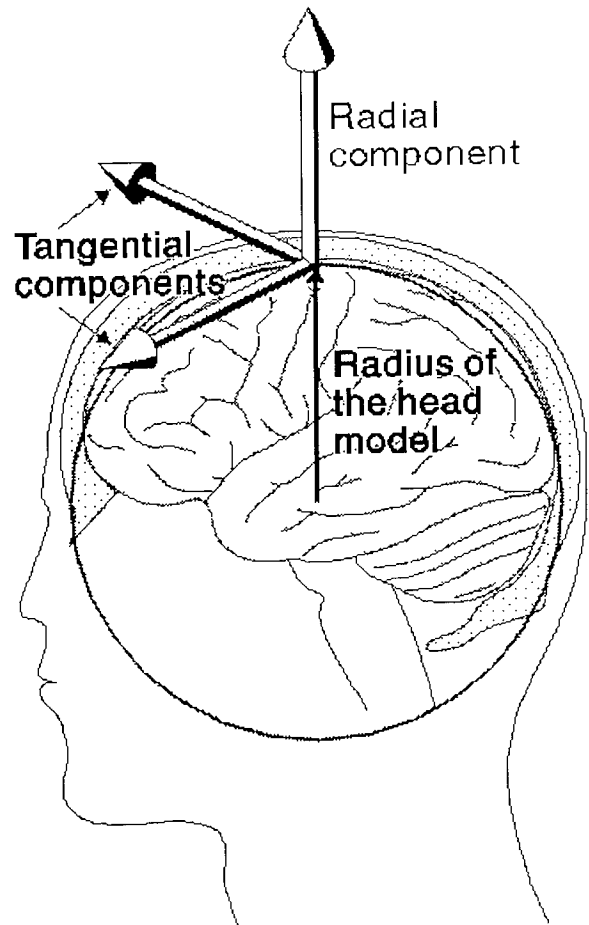
$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi F^2} (\mathbf{F}\mathbf{M} \times \mathbf{r}' - \mathbf{M} \times \mathbf{r}' \cdot \mathbf{r} \nabla F), \quad (1-16)$$

with $\mathbf{a} = \mathbf{r} - \mathbf{r}'$, $F = |\mathbf{a}|(|\mathbf{r}| |\mathbf{a}| + |\mathbf{r}'|^2 - \mathbf{r}' \cdot \mathbf{r})$ and

$$\nabla F = \mathbf{r} \left(\frac{\mathbf{a} \cdot \mathbf{r}}{|\mathbf{a}|} + \frac{|\mathbf{a}|^2}{|\mathbf{r}|} + 2|\mathbf{a}| + 2|\mathbf{r}| \right) - \mathbf{r}' \left(\frac{\mathbf{a} \cdot \mathbf{r}}{|\mathbf{a}|} + |\mathbf{a}| + 2|\mathbf{r}| \right).$$

Four points follow from this formula for a spherically symmetric volume conductor:

- A radially oriented dipole (see figure 1-8) does not produce an external magnetic field, since $\mathbf{M} \times \mathbf{r}'$ is zero in that case. This implies that sources that are oriented perpendicular to the surface of the sphere cannot be detected with MEG (Baule and McFee, 1965).
- There is no accurate knowledge of the conductivity profile of an individual. Luckily, the radial component of the magnetic flux density is independent of the conductivity and the number of layers of the volume conductor. When the assumption of a spherically shaped volume conductor does not hold the volume conducting effects can be taken into account by the BEM or FEM.
- The radius of the sphere is not important, only the centre of the sphere is.
- The magnetic fields sum linearly. Therefore, more complicated source configurations like distributed sources may be composed by simply adding the contributions of several single dipoles.



Forward solutions for different volume conductor models can be found in Moshier *et al* (1997) and Moshier *et al* (1999), as well as forward solutions for the EEG (see also Zhou and van Oosterom, 1992).

figure 1-8: Definition of the radial and tangential components of a dipole (figure from Malmivuo *et al.*, 1997). A dipole with only a radial or tangential component is referred to as a radial or tangential dipole, respectively.

2 Inverse Problem

2.1 Introduction

The inverse problem is the estimation of primary currents from measurements of potentials or a magnetic field. The solution of this problem is non-unique, as there are many different source configurations that produce exactly the same measured potential or magnetic field. In addition, the existence of electrically or magnetically silent sources further complicates the issue. Namely, radially orientated currents do not produce a magnetic field outside a spherically symmetric volume conductor (e.g. Baule and McFee, 1965, Cuffin and Cohen, 1977). This theoretical result has been confirmed by measuring the MEG and EEG simultaneously (Wood *et al.*, 1985). It was shown that the radial sources that were found by measuring the EEG could not be found with MEG, whereas tangential sources were found with both modalities.

Theoretically, a loop of primary current does not produce a potential difference on the surface of a spherically symmetric volume conductor (Roth and Wikswo, 1986). This is however a physiologically implausible situation. Electrically silent magnetic fields can also be produced due to volume effects when a tissue has a 'spiral' or 'helix'-like conductivity (Roth and Wikswo, 1986). Examples of tissues that are likely to exhibit such a conductivity pattern are given in the original paper.

It is possible to reduce the non-uniqueness of the inverse problem by adding *a priori* information. In previous years many methods have been developed to solve the inverse problem, each one with its merits and limitations. Several methods are described in this chapter and special attention is given to methods that can easily incorporate *a priori* information.

Two main types of *a priori* information can be distinguished: constraints based on spatial constraints and temporal information.

Spatial constraints can be obtained from a reconstruction of the cortical surface, using the MRI scan of a subject or patient, or from a standardised model. Such a reconstruction can further be used to set a constraint on the orientation of the sources. The orientation can be restricted to be perpendicular to the cortical surface, as most of the dendrites of the pyramidal cells are oriented perpendicular to this surface (Okada, 1982, Mitzdorf, 1985, Snyder, 1991). Indeed, Williamson and colleagues found that dipoles reconstructed from auditory evoked magnetic fields were orientated in the average direction of the nearest sulcus (Williamson *et al.*, 1991). Other kinds of spatial constraints can be based on symmetry considerations.

Temporal constraints are generally based on a hypothesis about the development of the cortical activity over time (i.e. about the shape of the waveform of the source and the time of onset/offset of the activity). Utilisation of these constraints was first undertaken by Scherg and colleagues (i.e. Scherg and von Cramon, 1985a, Scherg and von Cramon, 1985b) and they are now extensively used.

Additional constraints can be based on source strength. The observed source strengths usually found with measurements lie between 2 and 30nAm (Hämäläinen, 1987). This means that when a source is found with strength larger than 30nAm one can assume that the activity should be modelled by two or more sources. Scherg and Berg (1991) used a constraint based on the source energy (i.e. the squares of the source strength over all time points) in order to distinguish simultaneously active sources. Another type of constraint is based on the covariance matrix of the noise (Sekihara *et al.*, 1994).

It should be emphasised here that it is possible to find a unique solution by introducing *a priori* information, but that this is only a unique solution with respect to the constraints that were used. The general inverse problem is still non-unique. An example of this is given by Scherg and Berg (1991). They analysed electrically measured mismatch negativity, using a spatio-temporal dipole model. Although all solutions, obtained with different constraints, explained the data well, some of the solutions contained sources in remarkably different locations. The choice of the 'correct' solution depends then on the experimenter.

In the next sections four main groups of algorithms are described: algorithms for the localisation of discrete sources, algorithms for the localisation of distributed sources, hybrid methods (methods that fall in between the previous two groups) and an algorithm based on

signal-space projection. The assumptions underlying each method are given, together with its limitations. Furthermore, descriptions of the constraints that can be used and significant results obtained with the methods are given. This chapter is finished with a section in which a comparison of some of the inversion techniques is made. Additionally, the motivation for choosing the FOCUSS algorithm as the technique to use in the rest of this work is given.

2.2 Reconstruction of Discrete Sources

2.2.1 Equivalent Current Dipole Fitting

2.2.1.1 Assumptions

- The activity in each active patch of cortical area can be modelled as a single equivalent current dipole.
- The number of active sources that are needed to explain the measured data must be given explicitly in the algorithm. Several methods are suggested that can be used to obtain an estimate of the number of active brain areas. First of all, a visual inspection of the data can give a skilled experimenter an idea about the number of sources. This soon fails when the number of active areas increases. In addition, a Singular Value Decomposition of the data matrix can give an idea about the number of active sources (this will be explained in section 2.2.3).

2.2.1.2 Finding Dipole Parameters

The method starts with selecting a latency from the recorded waveforms at which a dipole will be fitted. Usually this is the latency at which the total field power, or alternatively the signal-to-noise ratio, over all measurement channels is at a maximum. If the magnetic field map at this latency contains a dipolar pattern (see figure 1-4), the position, orientation and strength of the equivalent current dipole can then found as follows (e.g. Schneider, 1974, Kavanagh *et al.*, 1978, Ary *et al.*, 1981):

1. An initial guess is made of the parameters of the current dipole. This initial guess is usually based on where the experimenter expects the active brain area to be. Alternatively, the initial guess can be based on the separation between the positive and negative field maxima (Kaufman and Williamson, 1986): The dipole position is half way between the two extrema and the depth of the dipole is estimated as the separation between the two extrema divided by the square root of two. The orientation of the dipole is perpendicular to the line connecting the two field extrema.
2. The magnetic induction is computed (solving the forward problem) for this initial guess.

3. The difference between this magnetic induction and the measured magnetic induction is computed, in a least-squares sense (i.e. the value of the cost-function is computed).
4. The source parameters are adjusted and step 2) and 3) are repeated until the computed induction closely matches the measured induction. Expressed differently: the cost-function is minimised until the best-fit source parameters are found, this source is referred to as the localised source.

A visualisation of the process described above is given in figure 2-1.




figure 2-1: Hypothetical example of the minimisation process of the cost-function. Here, the cost-function is only optimised for one of the dipole parameters (r), the other parameters are kept constant. The numbers between brackets denote the iteration number in the optimisation process. The algorithm starts with the initial guess (r_0), traverses through the local (step 2) and global minimum (step 9) and finally (step 21) arrives back at the global minimum (r_s).¹

The values for the parameters of the dipole in the final solution depend on the specific cost-function that is minimised and on the optimisation algorithm that is used.

2.2.1.2.1 Cost-function and Validity of the Model

The final value of the cost-function gives an idea of the goodness of the fit, i.e. whether the chosen model correctly explained the measured data. A cost-function that is still frequently used is based on the percent of variance (or sometimes referred to as the g-value), defined as:

$$\%V = 100\% * \left[1 - \frac{\sum_{i=1}^N (\mathbf{B}_i - \hat{\mathbf{B}}_i(r, \theta, \phi, M_\theta, M_\phi))^2}{\sum_{i=1}^N \mathbf{B}_i^2} \right], \quad (2-1)$$

¹ A global minimum is truly the lowest value; a local minimum is the lowest value in a finite neighbourhood and not on the boundary of that neighbourhood (Press *et al.*, 1992).

where N is the number of measurements, \mathbf{B}_i is the measured magnetic induction in the i^{th} sensor, $\hat{\mathbf{B}}_i(r, \theta, \phi, M_\theta, M_\phi)$ is the computed magnetic induction in this sensor. The computed magnetic induction depends on the five dipole parameters. The dipole position parameters are given here in spherical co-ordinates r, θ and ϕ . The moment in the radial direction, M_r , is not a variable here, as a radial dipole in a spherical symmetric volume conductor produces no radial component of the magnetic field (see section 1.3.4 and Baule and McFee (1965) and Cuffin and Cohen (1977)).

The function that is minimised during the optimisation is the residual variance, defined as 100% - %V. A value of 0% for the residual variance means that the chosen source model exactly reproduces the measured data. In the presence of noise this is obviously not correct. This measure does not contain information about the noise and therefore a high value could be due to noise or due to an incorrect model. For example, it is possible to explain the data perfectly by using many, say a hundred, dipoles, whereas a physiological more adequate solution could be given by, say, two dipoles that explain the data nearly perfectly. Besides, Miltner *et al.* (1994) showed for their simulated data that, although this measure was useful as an indicator for the completeness of a solution, it did not correlate well with the localisation accuracy of identified sources. A measure that takes the noise into account should give a more conclusive understanding of the goodness of the fit. Such a measure is the chi-square function, defined as:

$$\chi^2 = \sum_{i=1}^N \left(\frac{\mathbf{B}_i - \hat{\mathbf{B}}_i(r, \theta, \phi, M_\theta, M_\phi)}{\sigma_i} \right)^2, \quad (2-2)$$

with σ_i the standard deviation of the noise in the i^{th} sensor. The chi-square is based on the assumption that the noise is normally distributed. It can usually still be used successfully as the cost-function when the noise is not normally distributed (Press *et al.*, 1992).

Noise in the data can cause chance fluctuations in the values of the minimised chi-square. The probability that the minimised chi-square exceeds the observed value by chance, assuming the source model and given the data, is a good measure of the goodness-of-fit. In most cases this probability is adequately given by the reduced chi-square (Press *et al.*, 1992). The reduced chi-square is defined as:

$$\chi_r^2 = \frac{\chi^2}{\nu}, \quad (2-3)$$

with χ^2 as in equation 2-2, ν the degrees of freedom (= N - number of source parameters).

An appropriate model is indicated by a reduced chi-square value of $1 \pm \sqrt{2\nu}/\nu$ (Press *et al.*, 1992). However, this value will not be obtained when the noise level is incorrectly estimated. The correct model order can still be evaluated though by estimating the reduced chi-square value for different model orders, since the decrease of the value of the reduced chi-square found by increasing the model order becomes negligible once the adequate model order is reached (Supek and Aine, 1993).

Additionally, it is possible to acquire an impression of the stability of the solution by performing Monte Carlo analysis (Medvick *et al.*, 1989, Press *et al.*, 1992). The measurements that are used for the source localisation are only one particular experimental realisation. A repetition of the experiment would therefore, due to the noise, result in a slightly different measurement and therefore a slightly different localised source. Instead of repeating the experiment several times a simulation can be done, replacing the true, unknown, source parameters by the ones found during the source localisation. The measurement values for this source are then computed and random noise with a Gaussian distribution is added (or any other type of noise, if this is more appropriate). This noisy data set is then used to obtain another source localisation. This is repeated several times and the variability of the fitted solution gives a measure of robustness. The best strategy is to compute both the reduced chi-square value and do a Monte Carlo analysis, since Supek and Aine (1993) showed that under-modelling can give rise to small variability. This could lead to false conclusions about the quality of the solution when only Monte Carlo analysis is done.

2.2.1.2.2 Optimisation Techniques

A non-linear estimation technique has to be used to find the source parameters for which the cost-function is at a minimum, since the magnetic induction is a non-linear

function of the position parameters. The techniques that are most often used are the Levenberg-Marquardt algorithm (Marquardt, 1963), Powell's method (Powell, 1964) and several versions (Achim *et al.* 1988, Huang *et al.* 1998) of the Simplex method developed by Nelder and Mead (1965). The main difference between these algorithms is that the Levenberg-Marquardt algorithm converges fast, as it uses the derivatives with respect to the source parameters, which Powell's method does not. The Simplex method upholds several ($=$ number of parameters + 1) sets of parameter values, called the vertices. These vertices form a geometrical figure, called a simplex. At each iteration, the vertex for which the cost-function is worst is replaced by a better one. Normally, this new one is the vertex found by a reflection through the opposite face of the simplex. Occasionally, other types of steps called expansion or contraction are taken, based on a set of rules. The effect is that the simplex adjusts its shape to the error landscape on which a minimum must be found.

A problem with these minimisation algorithms is that they sometimes find a local minimum of the cost-function function, instead of the global minimum. This can partly be overcome by varying the initial guess or restarting the algorithm when it has converged to a solution (i.e. Nelder and Mead, 1965, Huang *et al.*, 1998).

The relatively new simulated annealing method is less prone to get stuck in local minima and is almost independent of the initial guess. The method uses an analogy of thermodynamics, specifically with the way liquids freeze and solids cool and anneal (Gerson *et al.*, 1994, Khosla *et al.*, 1997). A heated solid that is allowed to cool slowly (so that it reaches thermal equilibrium at each energy stage) reaches its lowest energy state, i.e. the state in which the atoms form a highly regular structure. During the cooling process the energy decreases with temperature, but occasionally the energy increases while the temperature decreases. Minimisation by simulated annealing is based on the same principles in that it minimises a cost-function that depends on the source parameters. It also uses a parameter T that represents the temperature. The algorithm starts at a high temperature T and randomly varies the source parameters, which changes the value of the cost-function. A downhill step in the cost-function is always accepted and uphill steps are sometimes accepted, with a probability that is inversely proportional to the size of the increase in the cost-function and linear proportional to T . The uphill steps ensure that the process can escape local minima and the dependence of the probability on T ensures convergence as T decreases. Equilibrium at each value of T is defined as the condition where the number of uphill steps is equal to the number of downhill steps. When

equilibrium is reached, or alternatively, when a fixed number of random steps are done, the value of T is lowered and the whole process of generating random steps is repeated until the next equilibrium is obtained. The algorithm finishes when it reaches a predetermined convergence criterion. Detailed descriptions of (slightly different) implementations of the algorithm are given in Sekihara *et al.* (1992), Gerson *et al.* (1994), Haneishi *et al.* (1994) Khosla *et al.* (1997) and Uutela *et al.* (1998). Performance could be enhanced by exploiting other analogies with thermodynamics, like entropy, to monitor the progress of the algorithm (Press *et al.*, 1992).

Another global optimisation technique is a genetic algorithm (Uutela *et al.*, 1998) that mimics the evolution of species. It starts by randomly generating, or seeding, an initial population of candidate solutions. Think of each candidate solution as a vector composed of elements representing the parameters of the dipoles. In parlance with genetics, this vector is referred to as the chromosome and the vector elements as genes. The object is now to maximise the fitness of the best solution vector as well as the average population fitness. Fitness implies maximisation and the fitness function is therefore defined as the negative value of the cost-functions described before. Maximum fitness is reached by a combination of three operations: reproduction, crossover and mutation.

In reproduction, only individuals with a high fitness are preserved and unfit individuals are discarded. By doing so, the algorithm focuses on the global maximum.

In crossover, elements of two (or more) individuals are combined to form a new individual. The effect is that the algorithm builds upon the success of the past yet still explores new possibilities (i.e. this enables the algorithm to escape local maxima).

In mutation, individuals are slightly changed. This means that some of the parameters are slightly perturbed with a small probability. This allows local optimisation. Larger perturbations on the other hand enable the escape from local maxima.

By decreasing the number of individuals towards the end of the optimisation, the algorithm can be made to concentrate on good solutions and converge (where convergence means that the individual members in the population start to resemble one another and additional evolution produces little improvement in fitness).

The results from simulated data obtained by Uutela *et al.* (1998) and Nagano *et al.* (1998) show that techniques based on genetics are promising and can outperform the simulated annealing approach.

2.2.1.2.3 Noise Estimation

As stated in section 2.2.1.2.1, it is important that an accurate estimate of the noise in the data is made, both for the least-squares data fitting and in the Monte Carlo simulations. Such an estimate can be obtained in several ways, if we assume no correlation between the channels²:

- If an average signal is computed, the standard errors of the mean can be taken as the standard deviation of the noise, σ_n . (Hämäläinen *et al*, 1993).
- The standard deviation of the noise can be obtained from the baseline if it contains only background activity. There are two different ways. Firstly, the standard deviation of the time-samples of the averaged pre-stimulus interval can be computed. Secondly, if the raw data are available, one can compute the standard deviation of the time-samples in each trial and then average over the number of trials P (only if the noise is Gaussian). The error in the estimation of the noise variance obtained with the latter method is P times smaller than that obtained with the first method (Knösche, 1997). A drawback is that the time-dependence of the standard deviation of the noise is lost.
- Alternating subaverages can be computed by reversing the sign of every second response. The response at a certain time and the following response are likely to be very similar. The subtraction of both responses is therefore an estimate of the noise (Schimmel, 1967).

2.2.1.2.4 Dipole Models Using *a priori* Information

How can we use *a priori* information with the concept of dipole fitting?

As explained above, the global minimum of the chi-square is searched for. So we can add constraints by changing this function:

$$\chi_{\text{new}}^2 = \frac{w_0 \chi^2 + w_1 f_1(r, \theta, \phi, M_\theta, M_\phi) + w_2 f_2(r, \theta, \phi, M_\theta, M_\phi) + \dots}{w_0 + w_1 + \dots}, \quad (2-4)$$

² Usually there is some correlation between the channels (Knösche, 1997). This means that the noise estimation is not completely correct, although pre-whitening can partly correct for correlated noise.

where f_1, f_2 , etc. are the functions that contain the constraints and w_1, w_2 , etc. are weighting factors. When we know, for example, the maximum strength of a source that is possible at a certain location, a constraint-function could be of the form:

$$f_1(r, \theta, \phi, M_\theta, M_\phi) = \begin{cases} \alpha & , q > q_{\max} \text{ at position given by } (r, \theta, \phi) \\ 0 & , q \leq q_{\max} \text{ at position given by } (r, \theta, \phi) \end{cases} \quad (2-5)$$

with q the strength of the source, given by $\sqrt{M_\theta^2 + M_\phi^2}$ and α a positive factor that enlarges the chi-square function. Consequently, the minimum chi-square will not be found at locations where the sources disagree with *a priori* information.

Scherg and Berg (1991) used this method to incorporate constraints with their spatio-temporal dipole model. Their approach is further described in section 2.2.2.4.

Another constraint could be the assumption that the best solution is composed of the smallest number of dipoles that can explain the solution, as was done by Haneishi *et al* (1994). They showed that this solution is found by setting f_1 equal to

$$\sum_{i=1}^N \left| \hat{B}_i(r, \theta, \phi, M_\theta, M_\phi) \right|^\alpha, \text{ with } \alpha \text{ a real number between 0 and 1.}$$

2.2.1.3 Limitations of Equivalent Current Dipole Fitting

- A stable solution is usually found when an equivalent single current dipole is a correct model for the activity. The problem becomes more ambiguous when more dipoles are used, especially when the distances between the individual sources are small.
- The assumption of a single dipole seems anatomically and physiologically reasonable for early cortical activity (Snyder, 1991). There is however little evidence in the literature that the model can be used to localise the activity associated with higher order processes in the brain.
- Optimisation algorithms are sensitive to the initial guess and local minima. Furthermore, the cost-function around the minima can be very shallow, which makes it difficult to find the exact minimum, especially in the presence of noise. Moreover,

Gerson *et al* (1994) showed that closeness of the initial guess to the correct solution does not guarantee a successful fit (i.e. a good guess of the experimenter about the approximate correct location of the source is not necessarily utilised with success). However, this needs further research, as Supek and Aine (1997) found that accurate initial guesses did provide better solutions.

- Even when the global minimum is found (mathematically), the solution might not relate directly to the underlying physiology (due to errors in the computation of the forward problem (wrong head model, errors in location of the sensors) or due to noise). The physiological solution may correspond to a local minimum or to no minimum at all (Scherg and Berg, 1991). This is of course not only a limitation of this technique, but for any optimisation method.
- The number of channels limits the number of sources that can be resolved (i.e. the total number of source parameters should not exceed the number of channels; when noise is present, this number is even smaller).
- Simulated annealing and genetic algorithms use numerical parameters whose values can not be estimated theoretically and they must therefore be determined empirically. In addition, Sekihara *et al* (1992) showed for simulated annealing that a wrong choice for these parameters could trap the solution in one of the local minima. Furthermore, simulated annealing is easy to implement but computationally expensive.
- Genetic optimisation is a stochastic process and, although it is good in finding *an estimate* of the global optimum, the true optimum solution will only be found by chance.

2.2.2 Spatio-Temporal Dipole Models

2.2.2.1 Introduction

Modelling the recorded magnetic fields at a single latency does not give a complete picture of the underlying brain activity. Even worse, it might give a wrong picture. For example, two sources in adjacent areas with dipole moments that are nearly parallel or antiparallel can produce a dipolar field pattern (Wang *et al.*, 1992). A single dipole solution can then be found that explains almost 100% of the variance in the data, but with a location that is far from the location of the original sources (Scherg and Berg, 1991). Taking into account the fluctuation of the amplitude of the activity over time (also called the waveform) of the two sources would enable the localisation of both sources. Besides, the peaks in the timecourses of temporally overlapping sources do not necessarily correspond to peaks in the recorded signals (Scherg and von Cramon, 1985a, Scherg and von Cramon, 1985b). Thus, source localisation based on selecting a single latency that gives a maximum in the recorded magnetic fields, assuming the brain activity is maximum at this point, is not necessarily meaningful. Again, using more latencies reduces the ambiguity and enables localisation of temporally overlapping sources.

There are two different approaches to spatio-temporal dipole modelling. The first one is simply an extension of single latency fitting with equivalent current dipoles, using multiple latencies instead. The second method is based on modelling of the waveform with a smooth parametric function. Both approaches are described next.

2.2.2.2 Assumptions

- Fitting dipoles at multiple latencies is an extension of equivalent current dipole fitting, so the same assumptions as for that method apply. The advantage of using temporal information is that the measurements/parameters ratio increases. Therefore, the inverse problem becomes better defined. Additionally, the influence of noise on the solutions decreases when more latencies are used in the source reconstruction.

- Several dipole models can be used to model brain activity that is developing over time, namely the moving, rotating and fixed dipole models. With the moving dipole model, the location, orientation and strength of the source are all free. This means that the number of parameters that have to be estimated for each source is $N * 5$ (= 3 location parameters, 1 orientation parameter in the tangential plane and the source strength), with N the number of latencies. With the rotating dipole model, the location of each source is fixed over time, giving a total of $3 + 2 * N$ parameters to be estimated for each source. Finally, with the fixed dipole, the location and orientation are fixed over time, only the strength is allowed to vary. The total number of source parameters to be estimated for each source is then $4 + N$. In this case the only time-dependent variable is the strength of the source. The advantage of the latter two dipole models is that the anatomical part and the dynamic part (representing physiological activities) of the solution are completely separated. A general mathematical framework to incorporate moving, fixed and rotating models was developed by Moshier and co-workers (Moshier *et al.*, 1992).

From a mathematical point of view, the fixed dipole model is the best choice, as it gives the largest number of free parameters. Which model describes the underlying brain sources best, physiologically, is still an open question (Huang *et al.*, 1990). It has been shown by Scherg and von Cramon (1985b) however, that two orthogonal dipolar sources, with small separation and partly overlapping waveforms, produce field maps that give the (false) impression of a moving dipole. Therefore, it seems that the moving dipole model does not always represent the underlying brain processes properly. Moreover, it seems more realistic to assume that different parts of the cortex with different cortical functions are activated when they perform their specific task. The 'movement' seen with the moving dipole model would therefore be more realistically modelled as two or more stationary dipoles that are activated at different times (Moshier *et al.*, 1992). The moving dipole model though often shows clusters around its 'path'. This provides an idea about the number of active sources and their locations, which can be utilised with spatio-temporal modelling, using fixed dipoles (Böcker *et al.*, 1994).

A rotating dipole might also be inappropriate, since a dipole model represents a fixed anatomical structure and therefore its orientation should not change. It seems to be an appropriate model only when the centre of activity in an active brain area shifts around a bend in the cortical surface. In general though, brain activity that is stationary, but that changes orientation appears to be physiologically implausible.

An alternative model is the regional dipoles model (Scherg and von Cramon, 1986). An active brain area is modelled by two dipoles with orientations in the tangential plane (and a third dipole in the radial direction in the case of EEG), with shared location and orthogonal directions. The strength of the sources is allowed to vary over time. The number of free parameters is the same as for the rotating dipole, but the difference is that fixing the orientations allows for easy modelling of the source waveforms. As with the rotating dipole model, this model seems physiologically implausible. Since the rotating dipole model and regional dipoles model are mathematically similar, both methods are referred to as the rotating dipole model from now on.

- Assumptions about the number of active sources must be made.

Methods that model the waveform of the sources require additional assumptions:

- The activity in the brain can be modelled by several dipolar sources that maintain their position (and optionally their orientation) throughout the time interval of interest. Therefore, only the rotating and fixed dipole models can be applied.
- The waveform of the sources is thought to be a smooth function of time. Note that the dipole fit method that uses a single latency implicitly assumes that the waveform is a δ -function at that latency. Of course, this is physiologically unrealistic.

2.2.2.3 Selecting the Correct Model

The spatio-temporal modelling methods can only be used when correct values are given for the number of sources, r . A method was described by de Munck (1990) to obtain this number. It is based on the analyses of the Singular Value Decomposition of the measured data (see section 2.2.3 for a description of Singular Value Decomposition). Let's assume a measured data matrix \mathbf{B} , with its singular values, λ_k , and a predicted data matrix $\hat{\mathbf{B}}$, with rank r . The singular values of \mathbf{B} with index larger than r should be the noise singular values (note that the singular values are arranged in decreasing order). Thus, the square of the Frobenius norm of the residual between \mathbf{B} and $\hat{\mathbf{B}}$, denoted as Δ , should obey the following relationship:

$$\Delta \geq \sum_{k=r+1}^{\text{rank}(\mathbf{B})} \lambda_k^2, \quad (2-6)$$

If this lower limit is larger than the noise level, then there is a residual that cannot be explained by the noise. The number of sources would have to be increased, until the lower level becomes smaller than the noise level.

The goodness of fit is then given by:

$$g \leq \frac{\sum_{k=1}^r \lambda_k^2}{\sum_{k=1}^{\text{rank}(\mathbf{B})} \lambda_k^2}, \quad (2-7)$$

These two equations tell to what extent the measured data can be explained by r sources that are linearly independent. A value for r can be obtained by studying the behaviour of the residual or g as a function of r . One problem still exists though, as pointed out by Moshier *et al* (1992). Even if we selected the correct value for r , it is still not clear how to divide the available degrees of freedom between stationary and rotating dipoles. Additionally, this will only work when the distinction between the singular values due to noise and due to active sources is evident. Moreover, SVD decomposes the data in linearly independent components. Consequently, when the sources are linearly dependent, SVD will underestimate the number of active sources. Other ways to select the number of significant singular values, and therefore the number of linearly independent sources, are described in the section about regularisation with the minimum norm method (section 2.3.2.3).

2.2.2.4 Finding Dipole Positions

Fitting dipoles at multiple latencies starts with choosing the dipole model that is used, being either a moving, rotating or fixed dipole model. Subsequently, a range of latencies is selected from the measured data. The latencies may incorporate all of the major peaks of the recorded waveforms, or some subset of these peaks. The larger the latency range, the more information that is available. Next, the dipole locations must be found. The

procedure for doing this is similar to that of the single equivalent current dipole fit method. Instead of the χ^2 , given in equation 2-2, the cost-function to minimise is (Supek and Aine, 1997, Huang *et al*, 1998):

$$\chi^2 = \sum_{i=1}^N \sum_{t=1}^T \left(\frac{\mathbf{B}_{it} - \hat{\mathbf{B}}_{it}(r, \theta, \phi, M_\theta, M_\phi)}{\sigma_{it}} \right)^2, \quad (2-8)$$

where \mathbf{B}_{it} is the measurement at sensor i and latency t , and $\hat{\mathbf{B}}_{it}$ is the corresponding computed measurement as induced by the reconstructed sources.

The non-linear optimisation takes place only in the space of the non-linear position parameters, while after each step in the iterative algorithm the optimal dipoles at the respective positions are calculated. This involves estimating the orientation and magnitude, in case of rotating dipoles, or optimising the strength, in case of fixed dipoles. These dipoles are then called locally optimal (e.g. the best fitting at these particular positions). This procedure is known as the two-step procedure (Supek and Aine, 1997). Details of how to find the orientation and strength of a rotating dipole can be found in Mosher and Leahy (1998). Finding the strength of a fixed dipole with known location is an easy linear problem, as will be shown in section 2.3.2.

Huang and colleagues (Huang *et al*, 1998) used this method to localise fixed dipoles




figure 2-2: Reduced chi-square for solutions obtained with 300 randomly chosen initial guesses, sorted in descending order. Figure modified from Huang *et al* (1998).

for simulated and experimental visual MEG data. Their algorithm automatically restarted with randomly chosen initial guesses to avoid the case that their solutions were dependent on the initial guesses. This resulted in a set of solutions. A graph displaying the reduced chi-square for each solution, sorted in descending order, typically shows a number of 'flat' portions (see figure 2-2). This shows that there are a number of solutions with similar reduced chi-square values, representing local

minima in the cost-function. The lowest flat portion should represent solutions near the global minimum. Apart from the fact that the global minimum can be found, independent of the users initial guess, the method has another advantage. The solutions with similar reduced chi-square form clusters, similar to those obtained with Monte Carlo analysis. This provides information about the uncertainty in the location, orientation and strength of the estimated sources. There are however 2 differences: the multi-start method obtains the uncertainties for a single noisy dataset, whereas Monte Carlo analysis uses datasets with different noise realisations. Additionally, Monte Carlo analysis starts each dataset with the reconstructed solution, whereas the multi-start method tries to find a (new) solution by starting with a new initial guess (Huang *et al*, 1998).

The localisation from simulated data produced by 2 and 4 independent dipoles, 2 linearly dependent dipoles and a case of 2 linearly dependent with 1 independent dipole was accurate, provided the correct model was used (Huang *et al*, 1998). In case of under-modelling of the 2 dependent sources, the source was located midway between the 2 original sources, with an amplitude that was approximately the sum of the amplitudes of the original sources. The reduced chi-square value decreased significantly when a source was added, clearly indicating under-modelling. An interesting result was obtained when the mixed case was over-modelled. The extra source was clearly incorrect, as indicated by the random behaviour of its timecourse and scatter of its location during re-fits. The saturation of the reduced chi-square also indicated over-modelling. Noteworthy was that the timecourses and locations of the clusters of the other 3 sources were correct, although the variance in the location of the sources had increased. This last result seems to be in contradiction with the results that Supek and Aine (1997) found for over-modelled 4-dipole spatio-temporal data. They found with their 2-step procedure that in some cases the addition of an extra source resulted in inaccurate location and timecourse estimation of all sources. These solutions were probably found in local minima. The multi-start algorithm avoids these erroneous results, because it is able to find the global minimum. These results are therefore not contradictory, but simply justify the application of the multi-start algorithm.

A more effective approach than considering each time sample separately is to take the spatio-temporal course of the signals into account as a whole. The method is based on the idea that the individual features of sensory evoked magnetic fields index discrete stages within one or more parallel streams of sensory processing, each stage involving magnetic

fields generated in one or more cortical areas. Assuming that the activity is spatially stable, a spatio-temporal analysis sets out to resolve the location of these areas, guided by the evolution of the responses over time.

This spatio-temporal model has been described by Scherg and colleagues (Scherg and von Cramon, 1985a, Scherg and von Cramon, 1985b). In their model, both the dipole location and orientation was fixed throughout the selected latency range. Furthermore, they tried to disentangle the contribution of different sources by using assumptions about the behaviour of the sources in time. The temporal behaviour was modelled by a smooth parametric function, characterised by the onset, offset, peak latencies and amplitudes of the peaks. The number of parameters that had to be estimated was therefore greatly reduced³. For each source with a monophasic waveform only 8 parameters had to be estimated: three for the location, one for the orientation in the tangential plane and four for the strength over time. Additionally, by using a smooth function as the waveform, the influence of noise was further reduced. Geva and co-workers (Geva *et al.*, 1993, Geva *et al.*, 1995) later generalised this idea, using a set of base functions with monophasic, biphasic or triphasic shape, so called wavelets (see figure 2-3 for an example of a wavelet). The optimisation parameters are then found by a least- squares fit procedure.

The implementation of Scherg's method is commercially available (Brain Electric Source Analysis or BESA). It has been widely used for the analysis of measured data (i.e. Scherg and von Cramon, 1985a, Scherg and von Cramon, 1985b, Scherg and von Cramon, 1986, Böcker *et al.*, 1994, Böcker *et al.*, 1994). The neuronal sources found with the method correspond well with findings from other techniques. This suggests that the approach has significant power and validity. However, BESA requires a lot of user-interaction. The number of sources must be given, waveforms for the sources must be selected and symmetry and other additional constraints can be chosen. Additionally, the user guides the optimisation process by fixing parameters values at certain times during the optimisation process. The solutions that are found are therefore dependent on the choices

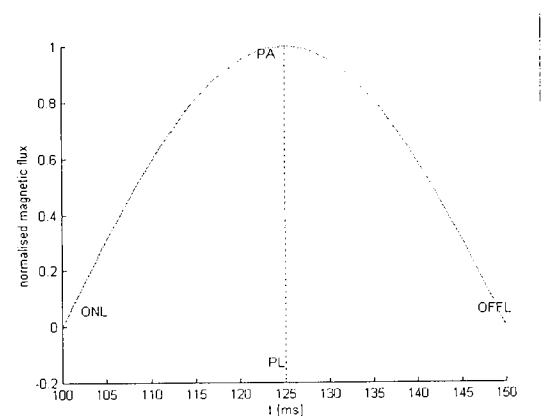


figure 2-3: An example of a monophasic wavelet. It is defined by the four parameters: peak amplitude (PA), peak latency (PL), latency of the onset (ONL) and of the offset (OFFL).

³ even more when symmetry considerations were implemented, i.e. 2 symmetric sources in both hemispheres.

that were made during the optimisation and depend consequently on the user's experience. To test the validity and reproducibility of the BESA approach, a simulation experiment was conducted by Miltner *et al* (1994). They simulated noisy ERP data, produced by ten dipolar sources. The sources were located in the temporal and somatosensory regions of the brain, with three symmetrically located pairs. The waveforms of the sources were partially overlapping. Ten subjects, whose experience with BESA ranged from very experienced to no experience at all, were asked to localise the sources, not knowing the number of active sources nor their locations and waveforms. They only knew that the simulated data represented the response to a somatosensory stimulus in an oddball paradigm by non-painful stimulation of the left middle finger, and that noise had been added to the data. The following observations could be made, showing both the usefulness and limitations of the spatio-temporal modelling approach:

- No subject found the exact solution.
- The number of sources that were localised ranged from 6-10. So, this approach can suffer from under-modelling.
- The grand mean error of all the localised sources by all the subjects was 1.4cm ($n= 75$, $\sigma= 1.0$ cm). The best solution was the localisation of 9 of the sources, with a mean localisation error of 0.81cm ($\sigma= 0.41$ cm). This is promising, given that the estimated best performance that could be expected, in the presence of noise, was a mean localisation error of approximately 0.5cm.
- The waveforms of identified sources were usually accurately reconstructed, even when there was considerable error in the reconstructed location of the source.
- Missing sources in the reconstruction, or errors in the reconstructed locations of some of the sources, did not markedly affect the goodness of the solutions of the correctly identified dipoles. This is in contradiction with the results by Supek and Aine (1997), who found that the reconstruction of all sources was necessary (but not always sufficient) for accurate location and timecourse estimation. The computation of the estimated data matrix involves the linear summation of the fields produced by all reconstructed sources. Assume that there are 9 perfectly reconstructed sources and one erroneous source. Errors in location, orientation and waveform of this source result in errors in the induced fields. These errors have to be compensated for by the other sources, shifting them from their (hypothesised) correct locations. Apparently, the influence of erroneously located sources is small compared with the noise in the data. Another possibility is that the effects are small because the error is compensated for by

many sources, showing only a small effect on individual sources. The difference with the results of Supek and Aine (1997) might be caused by the fact that they used a 2-step procedure and fewer dipoles.

- The peak latencies of the source waveforms were accurately estimated. The onset and offset had relatively large errors. Yet, the sequence of events was usually accurately reconstructed enabling inference about the sequence of events in the brain.

In conclusion, the method provides reasonable estimates of the spatial and temporal characteristics of the sources, but one should keep in mind that the results are not perfect and depend on the strategy chosen during the optimisation process. Also, the results show again that it is possible to obtain a unique solution, but that this solution is unique only with respect to the constraints that were used.

2.2.2.5 Limitations of Spatio-Temporal Dipole Models

- Fitting a series of instantaneous field distributions has several problems. It is time intensive and the number of free parameters is relatively low. Additionally, different model orders must be employed for different latencies, as sources become active or inactive.
- The algorithm can get stuck in a local minimum of the cost-function. The chance of this happening is relatively large, as there are a large number of unknowns. BESA enables escaping from local minima by requiring user interaction. However, this makes the solution depended on the user's experience and the method time-consuming (i.e. unsuitable for clinical use).
- BESA allows the user to solve the inverse problem in parts. That is, a subset of the sources is optimised, followed by the addition of further sources that are then optimised, while the previous sources are fixed. Finally, an optimisation with all sources is done. Although it is an efficient way to obtain solutions, this process favours local minima (Achim *et al.*, 1991).
- A drawback of the method developed by Scherg and von Cramon (Scherg and von Cramon, 1985a, Scherg and von Cramon, 1985b, Scherg and von Cramon, 1986) is that it is difficult to justify a particular choice of temporal dipole activation.

- It is difficult to estimate the number of sources that are active. When the estimate of the number of sources is too high, it is commonly found that two sources are located with large, cancelling, moments (i.e. orientated in opposite directions) (see for example Fernández *et al.*, 1995). Or, a nearly radial source with large moment is found (i.e. a source near the centre of a spherically symmetric volume conductor). An energy constraint could remove these erroneous sources.

2.2.3 Spatio-Temporal Decomposition

In the spatio-temporal modelling approach described in the previous section, the temporal information embedded in the measured data about the underlying brain activity is utilised by either fitting a dipole model to each latency sequentially or by optimising the parameters of a model of the waveforms of the active sources. So, the activity pattern is based on pre-defined time-patterns. These patterns are not directly based on the data, but on the choice of the experimenter (for example activity patterns with monophasic, biphasic or triphasic shape can be chosen (Geva *et al.*, 1993, Geva *et al.*, 1995)). The methods described in this section do not model the activity patterns of the sources but obtain this information 'directly' from the data matrix. This is done by spatio-temporal decomposition of the data matrix. The data matrix is considered as the linear sum of a number of spatio-temporal vectors (the assumption of linearity is reasonable, as was discussed in chapter 1). The decomposition results in a set of temporal and spatial vectors. The temporal vectors that explain most of the variance in the data are viewed as the activity pattern of the underlying sources and the related spatial vectors are then utilised to find the location and orientation parameters of these sources. Note that each spatial vector relates to a single component of correlated activity in the data matrix. However, the underlying source that produced this activity could be a single localised source, a distributed source or even sources in different areas, having correlated activity.

The method most often used to perform the spatio-temporal decomposition is Singular Value Decomposition (SVD). This method is described first, followed by a description of several approaches that utilise the resulting temporal and spatial vectors.

2.2.3.1 Singular Value Decomposition

Singular Value Decomposition is a standard mathematical technique that decomposes the measured data matrix (**B**) of n rows and m columns exactly and reversibly into 3 matrices:

$$\mathbf{B} = \mathbf{U} \mathbf{\Lambda} \mathbf{V}', \quad (2-9)$$

The matrices **U** and **V** are orthogonal matrices with dimensions n x n and m x m respectively, containing the left and right eigenvectors. The diagonal matrix **Λ** is of the same dimensions as **B**. The orthogonality of the matrices **U** and **V** means that each of the temporal and spatial vectors are uncorrelated and linearly independent within their respective matrices.

When the data matrix contains n latencies measured in m sensors, then the **U** matrix represents a series of n temporal vectors. The first vector represents the most prominent, highly correlated and widespread feature of the original data matrix. Subsequent vectors contribute less and less to the original data matrix. The **V** matrix (the transpose of **V** is taken in equation 2-9, necessary for matrix multiplication) contains the spatial distribution across the sensors in a series of spatial vectors. Each spatial vector corresponds to a temporal vector from **U**. More specifically, the first vector in **V** corresponds to the first vector in **U**, the second vector in **V** to the second vector in **U** etc. The diagonal of the **Λ** matrix contains the singular values of **B** in descending order. Each singular value is equal to the square root of the variance (RMS amplitude) of each spatio-temporal feature (Harner, 1990). The number of significant singular values therefore represents the number of significant components in the data. If each component is described by a single source, then the number of significant singular values gives the number of sources necessary to describe the data. However, synchronously active sources will be represented by a single component in the temporal, and therefore spatial, matrix. Therefore, the number of significant singular values gives a lower limit for the minimum number of sources that are needed to describe the data.

Additional operations may be performed on the data matrix before SVD is performed (Harner and Riggio, 1989). Converting the data matrix to a covariance matrix gives more weight to the larger features in the data. The correlation matrix meanwhile gives proportionally more weight to the lower amplitude features. The latter approach is identical

to adjusting the timecourses by subtracting the mean and dividing by the standard deviation of each timecourse before performing SVD on the data matrix. This is called normalisation.

2.2.3.2 Principal Component Analysis

The benefit of decomposing the data matrix in its spatial and temporal components is that relatively few components are needed to describe the data. This concept was first introduced to the evoked potentials community by John *et al.* (1964). They analysed the correlation between changes in evoked potentials in cats and conditioning by means of Principal Component Analysis (PCA). As was described above, computation of the principal components by means of a Singular Value Decomposition of the data matrix provides us with the principal components. The first principal component is the first eigenvector. This eigenvector gives the direction of maximum variance in the data and the first singular value gives the size of the data projection onto this eigenvector. The weight of the component at the recording positions is equivalent to the potential distribution of that component (this is called the factor loading, being equal to the spatial component multiplied with its singular value). A consequence of the assumption of linearity is that the summation of all factor loadings reconstructs the original signals. Note further that the SVD can also be performed on the Fourier transform of the data matrix. The resulting complex matrix, each entry of which has a real and imaginary part, can be decomposed with algorithms for complex SVD. This enables similar analysis in the frequency domain as in the temporal domain (Harner, 1990).

2.2.3.2.1 Limitations

- The method suffers from the phenomenon known as misallocation of variance (Möcks and Verleger, 1986, Koles and Soong, 1998) when there are sources with overlapping timecourses. This means that each original waveform explains a different amount of variance in the data than the associated principal component. The result is a lack of similarity between the principal waveforms and the principal components. Additionally,

a disproportional large amount of variance is normally explained by the first few principal components, leaving less variance to be explained by the other principal components. These small principal components are then wrongly thought off as being noise components. Therefore, PCA is likely to under-model the number of sources that are active (Koles and Soong, 1998).

- The computed principal components are not necessarily meaningful or predictive. For example, one component in one subject and a completely different component in another subject might represent the same waveform. Additionally, the components obtained by SVD are per definition uncorrelated. This means that the components are not necessarily physiologically meaningful (John *et al*, 1964, Maier *et al.*, 1987). This can be understood as follows. Assume that the first principal component is physiologically meaningful. The second principal component is orthogonal to the first principal component. Physiologically, there is no reason why this should be the case (indeed, it is quite unlikely).

An improvement of PCA is a method that transforms the orthogonal set of principal components in a set of physiologically meaningful components (John *et al*, 1964). A manually performed orthogonal rotation was used by Harner and Riggio (1989). They rotated the three largest principal components obtained for individual subjects in such a way that corresponding components were orientated in the same direction. The method works as follows: the values of the three largest principal components (PCA1, PCA2 and PCA3) are assigned to the x, y and z co-ordinate of each 'electrode'. So, the 1st element of PCA1 becomes the x-co-ordinate of 'electrode' 1, the 2nd element of PCA1 becomes the x-co-ordinate of 'electrode' 2, etc.; the same for the y and z-co-ordinates, obtained from PCA2 and PCA3 respectively. This is done for each subject. A 3D-plot of the 'electrodes' for each subject shows general disorder. However, when corresponding 'electrodes' are rotated so that they lie at the same location, a clustering of related areas emerges. An automatic implementation of their orthogonal rotation gave similar results. This enabled a comparison of the flash-evoked potentials across subjects and revealed the similarity in the responses.

The method proposed by Maier *et al* (1987) is based on the assumption that each principal component can be modelled by a single dipole. This spatial component is then rotated so that it best explains the activity produced by a source. So, instead of fitting a single source to each spatial component, a source is fitted to the rotated spatial component (effectively, the rotation parameter is one of the parameters that has to be optimised in the

optimisation process). For details see the appendix in Maier *et al* (1987). Their analysis of visual evoked potentials showed that this method enabled the determination of the dependency of the shape of the responses from a particular visual area on stimulus parameters. Furthermore, it was shown that although the responses varied considerably between subjects, the rotated principal components were remarkably stable. This method therefore enables the comparison of evoked responses between subjects.

The proposed method is not flawless however, as mentioned by Achim *et al.* (1988). The rotation is designed to define the components that are maximally accounted for by current dipoles (criterion 1) and that maximally account for the original principal components (criterion 2). The implementation by Maier *et al* (1987) however, only optimises criterion 2 for single principal components. In general, the optimised solution for criterion 2 will not be found with this method.

Koles and Soong (1998) used the common spatial patterns between two measurements to improve the localisation accuracy of the rotated principal component method. The method can be used when two data matrices (\mathbf{B}_a and \mathbf{B}_b) are given that were produced by the same underlying sources (like brain noise). Furthermore, one of the datasets contains information about additional sources. The two datasets can then be used to extract the extra information from the one dataset. The common spatial patterns are found by a whitening transformation applied to both datasets. This works as follows:

$$\mathbf{B}_{\text{transformed}}^a = \mathbf{W} \mathbf{B}^a \text{ and } \mathbf{B}_{\text{transformed}}^b = \mathbf{W} \mathbf{B}^b, \quad (2-10)$$

where

$$\mathbf{W} = \frac{\mathbf{B}_c}{\sqrt{\lambda_c}}, \quad (2-11)$$

with λ_c and \mathbf{B}_c the eigenvalues and eigenvectors of the composite covariance matrix of the two data matrices. That is:

$$\mathbf{K}_c = \mathbf{K}_a + \mathbf{K}_b, \quad (2-12)$$

where \mathbf{K}_a and \mathbf{K}_b are the covariance matrices of the two datasets.

The transformed data matrices now have the same eigenvectors. Additionally, the eigenvalues of both data matrices sum to one. The consequence of this is that the first eigenvector(s) of the data matrix with the additional source(s) will account for the maximum variance in this data matrix, but for the minimum variance in the other data matrix. This eigenvector is therefore optimal for separating both data matrices.

This method was applied to a simulation in which two common sources were present in both matrices and in one data matrix an additional source was present (Koles and Soong, 1998). Using only rotated principal components did not reveal the third source, but utilising the common spatial patterns unveiled the third source. This shows that spatially correlated noise (i.e. background MEG or active sources other than the sources of interest) can be put to use in order to reveal the sources of interest.

The common spatial patterns are a non-orthogonal set of basis vectors, obtained from the covariance patterns of two data matrices. When two data matrices, such as described above, are not available, a method that decomposes a single data matrix in a set of non-orthogonal components is needed. Independent Component Analysis (ICA) does just that.

2.2.3.3 Independent Component Analysis

Independent Component Analysis (ICA) assumes that the data are a linear summation of signals produced by a number of spatially stationary (either focal or distributed) sources whose timecourses are largely temporally independent. The goal of ICA is then to find temporally independent components, which may have non-orthogonal (even very similar) spatial components. The method is based on a neural network approach, taking into account the 2nd and higher-order dependencies of the timecourses, whereas PCA is based on the 2nd order covariance matrix (Makeig *et al.*, 1996a, Makeig *et al.*, 1997).

In order to appreciate the capabilities of the method, one must realise that the method is designed to extract the individual components from a data vector. In other words, it tries to answer the question: what are the different components? It can potentially be used as a pre-processing stage for a source localisation procedure, which answers the question: where do the components originate?

The performance of the method was tested by means of a set of simulations, designed to test the method in its ability to separate signal components in the presence of noise and

when sources are temporally correlated (Makeig *et al.*, 1996a, Ghahremani *et al.*, 1996).

The results can be summarised as follows:

- The performance of ICA is hardly influenced by the configuration of the underlying sources (i.e. the location and orientation of the sources).
- The performance degrades slowly in the presence of low-level noise sources and depends on the relative amplitudes of the sources and the noise sources and their relative placements and orientations.
- When correlated activity is present, ICA parses the correlation structure into spatially stationary and minimally correlated pieces, and when possible assigns each piece to a different output channel. Therefore, three situations can occur:
 - 1) The correlated activity is distributed over the correlated components.
 - 2) The correlated activity is presented in a single component.
 - 3) The correlated activity is wrongly assigned to a spurious component.

It was concluded that the method is best utilised to compare responses from multiple stimuli, task conditions and subject states. Makeig and colleagues (Makeig *et al.*, 1996b, Makeig *et al.*, 1997) used ICA to analyse the evoked responses to auditory stimuli and confirmed these results. Auditory targets with intensity just above threshold were presented in a 39 HZ click train background. ICA of the recorded data showed three components related to detected targets, 4 components related to undetected targets, 2 components separating out the 39Hz response and 5 components related to noise. Additionally, they demonstrated the intra and inter-subject stability of the principal components.

These results showed that the main limitations of the method are:

- The signals of sources that are strongly temporally correlated cannot be decomposed into separate independent components. Makeig and co-workers (Makeig *et al.*, 1996a) showed though that this could partly be overcome by applying the ICA algorithm concurrently to data from several stimulus and task conditions or adding residual data and brain noise to the data. This enables the algorithm to better separate the partially correlated but functionally independent sources. Also, there is doubt whether correlated sources are expected at all. Scherg and Berg (1991) noted that for physiological reasons the probability is small that two functionally distinct sources should exhibit identical waveforms.

- Interpretation of the physiological and/or psychophysiological significance of each of the independent components is difficult. The number of independent components that the ICA algorithm finds is always equal to the number of channels that are used to measure the data. When the number of underlying sources is less than this number (as will often be the case with the modern, multichannel, systems) the task is to select the independent components that are related to the sources. The remaining components are likely to be related to insignificant (brain) activity, like eye blinking or signal from power lines (Makeig *et al*, 1996a). When more sources are active than there are measurement channels and the 'noise sources' have relatively small independent components, the algorithm is able to assign meaningful independent components to the strongest sources (Makeig *et al*, 1996b and Ghahremani *et al*, 1996). If the number of large independent components exceed the number of channels, the ICA results may be poor (Makeig *et al*, 1996b).
- Localisation of the activity is not feasible directly with this method, although it can serve as a pre-processing stage.

2.2.4 Recursive - Multiple Signal Classification (R-MUSIC)

The multiple signal classification algorithm was described first by Schmidt (1986). It was used to determine the parameters of multiple wavefronts arriving at an antenna array from measurements made on the received signals. It was then used by Mosher *et al* (1992) to reconstruct multiple dipoles from MEG data. They presented promising results from simulations as well as from the analysis of experimental data. An improved version of their algorithm, Recursive MUSIC, was later developed to overcome some of the limitations of the original MUSIC algorithm (Mosher and Leahy, 1998).

Before we can describe the method we have to explain some of the concepts that are used:

Measurement space: The outputs of all sensors that span the measurement space. Each measurement is a point in this space. The line connecting the origin with this point forms a vector. The linear combinations of all the measurement vectors span the measurement space. The measured data not only contains the signals from the actual sources, but it is usually polluted with noise of different origin. The signal part is therefore expected to have

a dimensionality that is lower than the full possible rank of the measurement matrix. We can therefore divide the measurement space in a part that is spanned by the signals from the active sources, i.e. the **signal subspace** and a part that is spanned by the noise, i.e. the **noise subspace**. The noise space is the complement of the signal subspace, when we assume that the signals and the noise are not correlated.

Source space: The area in which the primary current flows, which can be, in its most general form, the whole brain.

Image space: The area in which the primary current is reconstructed. If the primary current is discretised and modelled by a finite set of current dipoles, then the collection of current dipoles forms the image space. In computer simulations the image space is equal to the source space. With real measurements this is in general not the case. This can introduce errors in the parameters of the reconstructed sources.

2.2.4.1 Assumptions

- The measured signal may be attributed to a finite set of dipolar sources, less than the number of detectors.
- A clear separation into noise and signal subspaces is possible.
- Data is polluted with additive, spatially white noise. In the case coloured noise is present, a pre-whitening procedure can be applied before the method is applied (e.g. Koles and Soong, 1998).

2.2.4.2 Finding Dipole Positions

The aim of this method is to find the dipoles in the source space that were active and therefore induced the magnetic signals that span the signal space. One starts by decomposing the measurement space into the signal subspace and the noise subspace. An estimate of this decomposition may be obtained from Singular Value Decomposition (SVD) of the data matrix, or equivalently, from the Singular Value Decomposition of the estimated covariance matrix. The MUSIC algorithm proceeds to identify those sources whose forward

projection onto the signal subspace is maximal (or equivalently, whose projection onto the noise subspace is minimal). Described in a more mathematical way:

Consider the spatio-temporal data matrix \mathbf{B} , containing the measured magnetic field values (or the estimated covariance matrix \mathbf{K}). Decompose \mathbf{B} into the signal and noise subspaces, thereby choosing the rank r of the signal subspace:

$$\mathbf{B} = \mathbf{U}_s \Lambda_s \mathbf{V}_s^T + \mathbf{U}_n \Lambda_n \mathbf{V}_n^T, \quad (2-13)$$

The left and right eigenvectors \mathbf{U}_s and \mathbf{V}_s span the signal subspace. The diagonal matrix Λ_s contains the r largest singular values due to the signals from the sources. The noise subspace is the orthonormal complement of the signal subspace and is spanned by \mathbf{U}_n and \mathbf{V}_n . The smaller singular values, due to the noise, are contained in Λ_n .

Additionally, a grid is defined that forms the image space. In each node of the grid a single equivalent dipole is placed. Let $\hat{\mathbf{B}}$ be the magnetic field produced by all the current dipoles in this grid. The method aims to minimise the difference between the measured magnetic field and the magnetic field produced by the model, in a least-squares sense:

$$\{\hat{\rho}, \hat{\theta}, \hat{\mathbf{q}}\} = \arg \min \|\mathbf{B} - \hat{\mathbf{B}}\|_F^2, \quad (2-14)$$

where F indicates the Frobenius norm. The vectors $\hat{\rho}, \hat{\theta}$ and $\hat{\mathbf{q}}$ contain the estimated location, orientation and strength of the sources, respectively

Using the definition of the signal and noise subspaces (equation 2-13) and the linear relationship between the magnetic field and the source strength, the above leads to the following cost-function (see Mosher and Leahy (1998), for a detailed derivation):

$$\{\hat{\rho}, \hat{\theta}\} = \arg \max \left\{ \sum_{k=1}^r c_k^2 \right\}, \quad (2-15)$$

with c_k the principal correlations between the subspaces spanned by \mathbf{G} and \mathbf{U}_s . \mathbf{G} is the gain matrix (= lead field matrix) that relates the current dipoles in the grid to the magnetic field (see section 2.3.2 for a thorough derivation of this gain matrix).

The above equation shows that in order to find the source parameters we must find the maximum subspace correlations between the two subspaces, which is the same as finding the minimum angle between the two. The conventional MUSIC algorithm proceeds by computing the principal correlations at each grid point (with c_1 the maximum principal component at each grid point). That is, the principal correlations between the gain vector of the dipole at a grid point and all vectors in \mathbf{U}_s are computed. As a graphical aid, $\frac{1}{\sqrt{1-c_1^2}}$ is then plotted for each grid point, showing sharp peaks where the principal correlation is close to unity. At these locations the likelihood that the dipole contributed to the signal as estimated by the data (or covariance matrix) is large. The peaks then define the true dipole locations. Subsequently, the strength of the sources can be computed (Greenblatt, 1993).

Finding the first peak is usually easy, but finding the additional peaks becomes harder. This, and additional problems are solved with the Recursive MUSIC (R-MUSIC) method (Mosher and Leahy, 1998). The limitations of MUSIC and the workarounds of R-MUSIC are as follows:

- MUSIC can only reconstruct linearly independent sources (i.e. Stoica and Nehorai, 1989), because SVD provides us with linearly independent components. When the timecourses are linearly dependent the rank of the spatio-temporal matrix is less than the true number of sources and the complete signal subspace cannot be determined. R-MUSIC introduces the Independent Topographies (IT) model. The data are seen as the sum of a set of ITs, with each IT produced by one dipole or more dipoles that have one single timecourse (i.e. that are linearly dependent). By constructing the source model in this way, the rank of the signal subspace is, by definition, always equal to the number of sources (where a single 'source' can be constructed from more than one dipole with linearly dependent timecourses). The only modification of the original MUSIC method is that, for an IT model with two dipoles, the principal correlations between combinations of two dipoles and \mathbf{U}_s are computed. For higher order IT models, combinations of three dipoles are examined etc. Testing all combinations of two (or more) dipoles in a grid might be too time-consuming. A non-linear optimisation algorithm might be used to find suitable combinations.

In practise, one starts with an IT model of order one. This is used until it becomes impossible to find good principal correlations. The model order is then increased and principal correlations are computed for the new IT model.

- With MUSIC, apart from the largest principal component, finding the other $(r-1)$ large principal correlations can be cumbersome. R-MUSIC masters this by building up a model of all the sources. This works as follows: assume that the first source was found. Scanning the grid for the second source is then done by computing the principal correlation for a grid point plus the first source. The maximum principal correlation is of course due to the first source, and we are therefore looking for the maximum second principal correlation in the grid. Finding the maximum second principal correlation means that the second source is found and it is added to the model. This process is repeated until all r sources are located. In figure 2-4 an illustration is given that shows why R-MUSIC has less difficulty with localising multiple sources than MUSIC.
- R-MUSIC speeds up the algorithm by splitting the localisation problem into two parts. First the non-linear location parameters are found and then the orientation of a source is optimised (see appendix in Mosher and Leahy (1998) for details).
- Instead of maximising the projection on the source subspace, one can also minimise the projection on the noise subspace. The first method is preferred when the number of sensors is larger than two times the rank of the signal subspace (Stoica and Nehorai, 1989). Knösche (1997) stated that projecting on the signal subspace gives less sharp peaks than projecting on the noise subspace. This was probably because the number of sensors was smaller than two times the rank of the signal subspace.
- The image space is formed by a grid and consequently has a limited sample frequency. Therefore, true sources might be located away from a grid point. Therefore, R-MUSIC uses a non-linear optimisation algorithm after a source is located in a grid point in order to improve the estimation of the location parameters.
- As mentioned in previous sections, finding the rank of the data matrix is no trivial issue. Underspecifying the rank for R-MUSIC reduces the performance considerably. Fortunately, overspecifying the rank by a couple of dimensions has little effect on performance. The reason behind this is that the additional subspace vectors span an arbitrary subspace of the noise subspace. The probability that that these vectors correlate with the source model is small. Consequently, the sources that R-MUSIC finds for these vectors can be neglected on bases of their small principal correlations. Therefore, moderate overspecification of the perceived rank of the data matrix is advisable (Mosher and Leahy, 1998).





figure 2-4: Plot of hypothetical principal correlations. The top panel shows the maximum principal correlation for each grid point ($= c_1$). It is clearly seen that with MUSIC and R-MUSIC it would be easy to localise peak A. Finding the other peaks in this figure (as is done with MUSIC) would be more cumbersome, because they are broader and lower. This could be due to noise or modelling errors. The bottom panel shows the second largest principal correlations ($= c_2$) that would be obtained with R-MUSIC after the first source is found from the top panel, using the first localised source as part of the source model. Finding the maximum in this plot is easy. Finding additional sources is therefore much easier with R-MUSIC than with MUSIC.

An additional way to improve the performance of MUSIC is by modification of the MUSIC estimator itself (i.e. the way the principal correlations are computed is modified). More details can be found in Stoica and Nehoria (1989) and references therein.

Finally, information from the frequency domain can be exploited to improve the performance of the MUSIC algorithm (Sekihara *et al.*, 1999). Decomposition of an average time-frequency distribution matrix, computed over a target time-frequency region, results in a set of orthogonal vectors. Again, the vectors related to the largest eigenvalues form the signal subspace and those related to the smaller ones form the noise subspace. These vectors can then be utilised as with the standard MUSIC algorithm in order to localise sources. So, the only difference with the normal MUSIC algorithm is that a time-frequency matrix is decomposed, instead of the covariance matrix of the data. Their simulations showed that incorporating the frequency information facilitates the separation of three closely located sources, not distinguishable with the normal MUSIC algorithm (Sekihara *et al.*, 1999). Application of the time-frequency method (TF-MUSIC) to localise somatic evoked high-frequency magnetic oscillations showed that the method has similar localisation accuracy as a single moving dipole approach when both methods could be applied. However, localisation was possible within some subjects with the TF-MUSIC method, whereas the single moving dipole method failed (Sakuma *et al.*, 1999).

2.2.4.3 Limitations

- The technique requires a clear-cut separation between signal subspace and noise subspace. It is not possible to perfectly define the signal subspace, meaning that the extracted signal subspace will be at some angle to the true one. As a consequence the MUSIC peaks will not be found at the exact source locations. Furthermore, during measurements the noise is not completely random, spatially independent, and zero mean. It might therefore be that some of the signal eigenvalues are indistinguishable from the noise eigenvalues or that the largest noise eigenvalues even become bigger than the smallest signal eigenvalues. As a consequence the signal and noise subspaces will be identified incorrectly.

- When some linear combination of true sources produces the same output as a single source, a false peak occurs at the location of this single source. This can only be resolved by adding more measurements.
- The maximum number of dipoles that can be reconstructed is equal to the rank of the spatio-temporal data matrix, which is in turn equal to the number of measurements. In general the rank is even smaller than this number due to noise in the measured data. This should not be a restriction with the latest multi-channel MEG systems carrying about 150 sensors.
- It remains to be proven that:
 - the method works with distributed sources. Mosher and Leahy claimed that reconstruction of distributed sources is possible, but no details were given (Mosher and Leahy, 1998).
 - a wrongly reconstructed source does not affect the accuracy of the localisation of the other sources via the model of the earlier localised sources that is constructed for R-MUSIC.

2.3 Reconstruction of Distributed Sources

2.3.1 Introduction

One of the main problems with the methods described in section 2.2 is that they rely heavily on the assumption that the measured magnetic fields were generated by discrete localised sources of activity. This assumption may not always be correct. The algorithms developed for discrete sources can still localise sources for magnetic fields induced by distributed sources. The localised sources then represent the centre of gravity of the distributed activity. However, such a reconstruction will be inaccurate when the distributed source is located on a convoluted part of the cortical surface (which will often be the case,




figure 2-5: Graphical representation of a slice through a part of the cortical surface, showing a sulcus. The distributed source, represented by the arrows, extends across both walls of the sulcus. The centre of gravity of the activity is indicated by the black dot. The result of localisation with a single dipole model, located in this centre of gravity, would give a meaningless result as the location is not even within the cortex.

as the cortical surface is highly convoluted). This is further illustrated in figure 2-5.

Thus, algorithms are needed that can reconstruct the properties of distributed brain activity. A report of such algorithms is given in the next sections. First, the minimum norm method, a method that arises naturally from the discretisation of the forward problem, is described. It is then showed that the minimum norm approach is a special case of a generalised Wiener-Helstrom filter. Additionally, a different class of algorithms, based on the probability of finding sources at certain locations, is outlined.

2.3.2 Minimum Norm Least Squares Approach

2.3.2.1 Assumptions

- The primary current distribution can be approximated by a discrete set of dipoles at fixed locations within a source region. This source space is defined *a priori*. It is further assumed that the predefined source space (also known as the image space) is at the same location as the 'true' source space (containing the currents that produced the magnetic fields). For computer simulations this is of course always true (unless deliberately constructed otherwise), but for real data this is not necessarily true. However, from now on we will assume that the image space and the 'true' source space are the same. Both are referred to as the source space, as this is a more intuitive name, since the space is made up out of predefined sources.
- To pick a solution from the infinite number of solutions to the inverse problem, one should take the solution that best fits the data in the least-squares sense and exhibits the minimum norm property.

2.3.2.2 Formulation

It follows from the Biot and Sarvart law that the magnetic induction perpendicular to the observation surface is linearly related to the source current density. Hence, if we assume a certain source space, the inverse problem reduces to finding only the strength of the sources in this source space. This idea was first introduced to the biomagnetism community, independently, by Hämäläinen and Ilmoniemi (1984) and Singh and colleagues (Singh *et al.*, 1984).

It is convenient to introduce the concept of the lead field of a sensor here (Hämäläinen and Ilmoniemi, 1984). We first repeat the law of Biot and Savart:

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi} \int_{\Omega} \frac{\mathbf{J}(\mathbf{r}') \times (\mathbf{r} - \mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|^3} d^3r', \quad (2-16)$$

As described in chapter 1, this law relates the current density in a region Ω of the volume conductor to the magnetic induction, located at \mathbf{r} (see section 1.3.4 for a description of these parameters).

The definition of the lead field then follows from:

$$\mathbf{B}(\mathbf{r}_i) = \int_{\Omega} \mathbf{L}_i(\mathbf{r}') \cdot \mathbf{J}(\mathbf{r}') d^3r' \quad i = 1, \dots, N \quad (2-17)$$

where N is the number of sensors. The lead field $\mathbf{L}_i(\mathbf{r}')$ accounts for the sensitivity of the i^{th} sensor at measurement position \mathbf{r}_i to the presence of the source current density at \mathbf{r}' . Or, alternatively, the lead fields span a subspace within the infinite space. The detectors measure the projection of the (primary) current onto this lead field subspace. A consequence of this is that radial currents are naturally excluded from the source reconstruction, as their projection is orthogonal to the lead field subspace and there is consequently no information about these sources, as shown in figure 2-6.

The elements of the lead field only depend on the volume conductor model, sensor configuration and the source space (i.e. location relative to the sensors). It follows from the equation above and Biot and Sarvart's law that the lead field decreases with distance from the sensor locations.

It is intuitive that the source space must have an anatomically correct shape in order to give a solution to the inverse problem that is physiologically and anatomically meaningful. It is natural to restrict the source space to a folded surface layer of thickness w , representing the




figure 2-6: Projection of the total current vector \mathbf{J} into the subspace spanned by the leads \mathbf{L}_1 and \mathbf{L}_2 . The minimum norm estimate of the total current is given by $\mathbf{J}_{//}$. It can be seen that \mathbf{J}_{\perp} does not contribute to the minimum norm estimate, as $\langle \mathbf{J}_{\perp}, \Omega \rangle = 0$. Generalisation to more than 2 sensors (leads) is straightforward. After Hämmäläinen *et al* (1993).

cortex. In this way, the volume integral in equation 2-17 is reduced to a surface integral. If we further divide the source space into a grid of n cells of area $\Delta\Omega'_j$ centred at position \mathbf{r}'_j , equation 2-17 becomes:

$$\mathbf{B}(\mathbf{r}_i) = \sum_{j=1}^n \mathbf{L}_i(\mathbf{r}'_j) \cdot \mathbf{J}(\mathbf{r}'_j) w \Delta\Omega'_j \quad i = 1, \dots, N \quad (2-18)$$

The lead field specifies the contribution of the j^{th} unit element of the source space to a given magnetic measurement at position \mathbf{r}_i .

Note that $\mathbf{J}(\mathbf{r}'_j) w \Delta\Omega'_j$ has the dimension of ampere-meter, which corresponds to that of a current dipole with strength q . The primary current density, $\mathbf{J}(\mathbf{r}'_j)$, can therefore be replaced by a δ -function centred around \mathbf{r}'_j . Constructing a source space is now simply a matter of placing many current dipoles at predefined positions in a grid.

The δ -function was used as a description of the primary current density at a location. The lead field $\mathbf{L}_i(\mathbf{r}')$ can now also be viewed as the response of the sensors to a current dipole with unitary strength at location \mathbf{r}' . In physics this is known as Green's function, as the kernel in mathematics and in the optics literature it is referred to as the point spread function. In the biomagnetism literature the term *Gain* is also used for the lead field (e.g. Gorodnitsky *et al.*, 1995).

In matrix notation, with the dimensions of the matrices stated in the exponents, equation 2-18 becomes:

$$\mathbf{B}^{N \times 1} = \mathbf{L}^{N \times n} \mathbf{Q}^{n \times 1}, \quad (2-19)$$

Each row of \mathbf{L} represents the sensitivity of a measurement channel to activity in all elements of the source space. As we use dipoles with a moment in only one direction, the number of columns in \mathbf{L} and the number of rows in \mathbf{Q} is equal to $n^{(1)}$. A more formal way to obtain the transition from the continuous to the discrete form of the law of Biot and Savart can be found in the appendices of Smith *et al.* (1990).

⁽¹⁾ In the more general case in which a dipole has a moment in all three directions, this would be $3n$.

The above equation tells us that the magnetic field is a linear summation of weighted (by Q_i) contributions from all source elements.

It seems easy now to find \mathbf{Q} , by using the inverse of \mathbf{L} (the dimensions of the matrices is not given here, for the ease of notation):

$$\mathbf{Q} = \mathbf{L}^{-1} \mathbf{B}, \quad (2-20)$$

A closer look tells us however that it is not that straightforward to find \mathbf{Q} . Usually a lot of elements in the source space are needed to get a good description of the source current density. With realistic measurements the number of elements will therefore far outnumber the number of sensors (thus $n \gg N$). Thus, more parameters have to be determined than there are independent measurements in equation 2-20, which means that this is an underdetermined system. What we need is a type of generalised inverse \mathbf{L}^+ that gives us an estimation $\hat{\mathbf{Q}}$ for \mathbf{Q} so that the residual error, $\mathbf{L}\hat{\mathbf{Q}} - \mathbf{B}$, is minimised in a least squares sense:

$$\Delta = \|\mathbf{L}\hat{\mathbf{Q}} - \mathbf{B}\|_2, \quad (2-21)$$

where $\|\cdot\|_2$ is the Euclidean norm or l_2 -norm. In this work we will refer to this norm with $\|\cdot\|$, unless stated otherwise).

The system is underdetermined, so there are an infinite number of solutions that minimise this residue, from which we have to choose one. A sensible choice is the estimation $\hat{\mathbf{Q}}$ that not only minimises the Euclidean norm of the residue but also the l_2 -norm of the solutions (Hämäläinen and Ilmoniemi, 1984). Such a minimum-norm least-squares (MNLS) estimation is found when the Moore-Penrose generalised inverse (also known as the pseudoinverse) is used (Penrose, 1955). For underdetermined cases this pseudoinverse is given by:

$$\mathbf{L}^+ = \mathbf{L}^T (\mathbf{L}\mathbf{L}^T)^{-1}, \quad (2-22)$$

The MNLS solution is thus computed as:

$$\hat{\mathbf{Q}} = \mathbf{L}^+ \mathbf{B}, \quad (2-23)$$

The solution can be displayed as a plot of all the dipoles as vectors in the source space, with the length given by the strength of the elements in $\hat{\mathbf{Q}}$ (the locations and orientations are predefined⁴). Alternatively, particularly for 2D source spaces, the strength of each source element can be displayed as a colour coded amplitude (see for example the results in Gençer and Williamson, 1998). This results in 'images' of $\hat{\mathbf{Q}}$.

There are also source configurations that do not induce a measurable magnetic induction, as was shown in chapter 1. The advantage of the minimum norm solution is that it automatically discards those elements that have zero contribution to the measured data, i.e. radial dipoles, and that it is unlikely that very large elements add to the solution. Although we find a unique solution, it must be remembered that the complete problem is still non-unique!

The use of this linear formulation with the fixed dipole localisation approach, described in section 2.2.2.4, can now also be clarified. Once the locations of the fixed dipoles are known, the lead fields are known as well. The only unknown is the strength of each source. These unknowns should be given such a value that the residual norm is minimised. Lets construct a lead field matrix that only contains, as concatenated columns, the lead fields of the few localised fixed dipoles. Thus, equation 2-20 now has more equations than unknowns, which makes it an overdetermined system. The pseudoinverse gives in the overdetermined case the solution that minimises the residual norm (Wang *et al*, 1992), which is exactly what we were looking for.

Back to the underdetermined case.

So far, the influence of noise on the MNLS solution has not been discussed. This is done in the next section, where approaches to reduce the influence of noise are also given.

⁴ The orientation was defined *a priori*. In the general case, the strength of each source element in the two tangential directions can be reconstructed, from which the orientation of the source element follows.

2.3.2.3 Regularisation

The distributed source solutions that the MNLS method finds are an exact fit to the data. This means that noise will also be modelled when this is present in the data. Therefore even small amounts of noise in the measurements or errors in the modelling of the head give rise to unstable solutions. Additionally, matrix \mathbf{L} is nearly singular, i.e. the linear system in equation 2-19 is ill-conditioned. As a result, the influence of noise in the data is disproportionately amplified, as will be demonstrated later. Consequently, a compromise must be made between a close fit to the data and a stable solution. It is clear that the difference between the measured data and the magnetic induction, induced by the reconstructed source, can be allowed to be as large as the amount of noise in the measured data (Bedford, 1995). This is further utilised hereafter, where one of the regularisation techniques, truncated Singular Value Decomposition (TSVD), is described.

In order to avoid the instability associated with the matrix inversion in equation 2-22, the pseudoinverse is calculated as:

$$\mathbf{L}^+ = \mathbf{L}^T (\mathbf{L}\mathbf{L}^T + \gamma \mathbf{I})^{-1}, \quad (2-24)$$

Where \mathbf{I} is the unit matrix and γ is a predetermined regularisation parameter.

The error that is minimised with this regularised pseudoinverse is then (compare with equation 2-21):

$$\Delta = \|\mathbf{L}\hat{\mathbf{Q}} - \mathbf{B}\| + \gamma \|\hat{\mathbf{Q}}\|, \quad (2-25)$$

This equation gives insight in what happens with the regularisation. Evidently, there is a trade-off between the minimisation of the error in the data term ($\|\mathbf{L}\hat{\mathbf{Q}} - \mathbf{B}\|$) and the modelling term ($\|\hat{\mathbf{Q}}\|$). When there is little noise, a small value is chosen for γ , and a close fit to the data is obtained. When there is a lot of noise in the data, a close fit to the data will be meaningless and a large value for γ ensures that the modelling term is minimised, thereby reducing the influence of the noise but also decreasing the spatial resolution of the reconstruction. Moreover, a large value for γ can introduce spurious oscillatory components

noise will be disproportionately amplified (see also equation 2-30). This is avoided by zeroing the elements in the pseudoinverse of Λ that correspond to the small singular values in Λ .

Fortunately, the elements in the source space associated with small singular values contribute little to the magnetic induction and hence the difference between the measured data and the magnetic induction is only slightly enlarged when such singular values are truncated. The fact that some signal is removed from the solution is evident if we write the estimated solution in terms of the truncated SVD (Shim and Cho, 1981):

$$\hat{\mathbf{Q}} = \sum_{i=1}^t (q_i \mathbf{v}_i + \lambda_i^{-1/2} n_i \mathbf{v}_i), \quad (2-30)$$

and the residual error is given by

$$\Delta = \sum_{i=t+1}^r q_i \mathbf{v}_i - \sum_{i=1}^t \lambda_i^{-1/2} n_i \mathbf{v}_i, \quad (2-31)$$

with $q_i = \mathbf{Q}^T \mathbf{v}_i$ and $n_i = \mathbf{N}^T \mathbf{u}_i$, \mathbf{N} the vector containing the noise in the data, t the termination index and r the rank of \mathbf{L} . See Appendix A for a derivation of these two equations.

Equation 2-30 shows that for large t there is a close fit to the data (first term), but also a large projection of noise into the solution. This is a problem in case there is a lot of noise in the data. If t is small, there is little projection of noise into the data, but also removal of detailed information about the solution (loss of resolution). This trade-off between influence of noise and resolution can clearly be seen in equation 2-31. It shows that the error between the ideal and reconstructed solution is caused by removal of the small eigenvectors and therefore the projection of data through them (first term) and projection of noise through the larger eigenvectors (second term). A good example of these effects can be found in the article by Shim and Cho (1981) who simulated a CT-reconstruction on a plane.

A second regularisation technique is Tikhonov regularisation, minimising the following error function (i.e. Brooks *et al.*, 1999):

$$\Delta = \|\mathbf{L}\hat{\mathbf{Q}} - \mathbf{B}\| + \gamma \|\mathbf{R}\hat{\mathbf{Q}}\|, \quad (2-32)$$

('ripple') into the solutions (Biemond *et al.*, 1990). Methods to choose the optimum value for γ are described later on in this section.

In practice, SVD is used to compute the pseudoinverse. This allows for a natural incorporation of regularisation and it also gives insight into why the influence of noise can be large on the MNLS solutions.

The pseudoinverse of the lead field matrix follows from the SVD of the lead field matrix:

$$\mathbf{L} = \mathbf{U}\mathbf{\Lambda}\mathbf{V}^T, \quad (2-26)$$

where \mathbf{U} and \mathbf{V} are orthogonal matrices with dimensions $N \times N$ and $n \times n$ respectively. The diagonal matrix $\mathbf{\Lambda}$ is of the same dimensions as \mathbf{L} and contains the singular values of \mathbf{L} in descending order:

$$\mathbf{\Lambda} = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_j, \lambda_{j+1}, \dots, 0, \dots, 0), \quad \lambda_j \geq \lambda_{j+1} \quad (2-27)$$

The singular values represent the magnitude of the contribution of a source to the measured data.

The pseudoinverse is now computed as:

$$\mathbf{L}^+ = (\mathbf{U}\mathbf{\Lambda}\mathbf{V}^T)^+ = \mathbf{V}\mathbf{\Lambda}^+\mathbf{U}^T, \quad (2-28)$$

with the pseudoinverse of $\mathbf{\Lambda}$ given as:

$$\mathbf{\Lambda}^+ = \text{diag}(\lambda_1^+, \lambda_2^+, \dots, \lambda_j^+, \lambda_{j+1}^+, \dots, 0, \dots, 0), \quad \begin{cases} \lambda_j^+ = 1/\lambda_j & , \lambda_j > \lambda_1 \cdot \gamma' \\ \lambda_j^+ = 0 & , \lambda_j \leq \lambda_1 \cdot \gamma' \end{cases} \quad (2-29)$$

where γ' is the parameter that determines which singular values to truncate.

It is clear now from equation 2-29 why the influence of noise on the unregularised MNLS solution is large. The small singular values in $\mathbf{\Lambda}$, associated with the noise, will give rise to large values in the pseudoinverse of $\mathbf{\Lambda}$ and hence the corresponding components of the

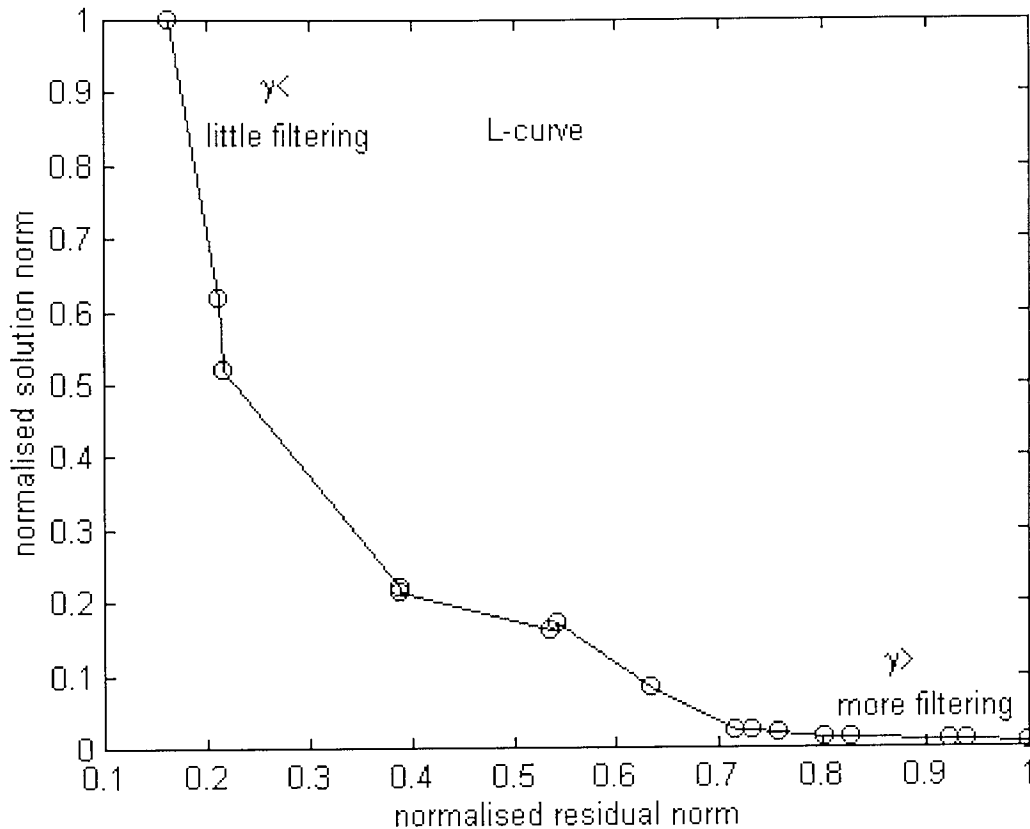


figure 2-7: Example of an L-curve, with on the x-axis the normalised residual norm, $\|\mathbf{B} - \tilde{\mathbf{B}}\|$, and on the y-axis the normalised solution norm, $\|\hat{\mathbf{Q}}\|$. The number of singular values that were used decreases from left to right. One can see that the residual norm decreases when more singular values are used in the computations and that the solution norm increases at the same time. This is expected, as more sources are used to explain the data when the number of singular values that are used increases. The optimum truncation parameter is the singular value that lies near the corner of the L-curve. This example was obtained for a simulation set-up with the 19-channel CES system at Aston University.

with \mathbf{R} a matrix containing additional constraints. If \mathbf{R} is the identity matrix, the solution is the Tikhonov zero-order solution and is constrained in energy (l_2 -norm), as in equation 2-25. If \mathbf{R} is the first or second spatial derivative (Laplacian) of \mathbf{Q} , the Tikhonov first- or second-order solution, respectively, is constrained to have a smooth surface gradient or curvature (Brooks *et al*, 1999).

As illustrated above, a good choice for the truncation parameter is necessary in order to obtain sensible MNLS solutions. A popular method for choosing the truncation parameter is to determine the trade-off curve between the solution norm that we minimise and the residual norm. This curve closely resembles an L-shape and the optimal truncation parameter is found in the ‘corner’ of this curve (e.g. Tu *et al.*, 1996, Johnston and Gulrajani, 1997, Brooks *et al*, 1999). An example is given in figure 2-7.

An L-curve, not to be confused with the gain matrix \mathbf{L} , will only be found if the noise is uncorrelated, random and with zero mean (Johnston and Gulrajani, 1997).

Therefore, if there are large errors in the lead field matrix, due to a shift of the imaging space with respect to the noise space, no L shaped curve will be found. The main disadvantage of this approach though is that the inverse problem has to be solved for a whole range of truncation parameter values. Additionally, the quality of the solution is very dependent on the (correct) choice of the regularisation parameter (Brooks *et al*, 1999). This follows from the shape of the L-curve around the corner: a small deviation from the optimum choice results in a large increase in either the residual error or solution error. Also, Johnston and Gulrajani (1997) showed that the optimum truncation parameter is not necessarily positioned on the corner of the L-curve for high SNR data.

The zero-crossing method described by Johnston and Gulrajani (1997) finds a truncation parameter for which the residual is equal to γ times the squared norm of the solutions (see equation 2-25). This selection is usually close to the corner of the L-curve, but with the additional advantage that the time-consuming computation of the L-curve does not have to be performed.

Another approach is to use information about the noise level to obtain a value for the truncation parameter. When the truncation parameter is based on the SNR, no effort is made to fit the solution to the data more precisely than the noise level (Bedford, 1995). The equation for the optimum truncation index derived by Shim and Cho (1981) was based on this idea:

$$t_{\text{opt}} = \max_k \left\{ k \mid \lambda_k \geq \frac{MK_n}{NK_q} \right\}, \quad (2-33)$$

with M the number of channels and N the number of elements in the source space. It was further assumed that the noise and \mathbf{Q} were white (so that \mathbf{K}_n and \mathbf{K}_q reduced to scalars).

A different regularisation technique was used by Oster and Rudy (1992). They used a modification of the Tikhonov technique, known as Twomey regularisation. Instead of imposing constraints on the magnitude of the solution (zero order Tikhonov regularisation) or its derivatives (first and second order Tikhonov regularisation), the method minimises the difference between the solution, \mathbf{Q} , and some estimate of the solution, \mathbf{P} .

The authors tested several schemes for the calculation of \mathbf{P} , using temporal information. Temporal constraints are, strictly speaking, not regularisers, since the condition of the lead field matrix (containing only spatial information) is not directly improved. But, the temporal information imposes constraints on solutions by restricting the acceptability of candidate solutions based on prior knowledge, and can therefore be treated as regularising (Brooks *et al*, 1999).

Some schemes that were tested incorporated only information from previous time intervals, others incorporated information from both previous and future time intervals. They tested their algorithms with data obtained from a human-torso tank, with a beating canine heart inside. The object was to reconstruct the epicardial potentials from surface potentials, measured in 400 electrodes. Epicardial potentials were measured as well, so that the quality of the reconstruction could be assessed. Although this is not exactly the same inverse problem as described in this work (source localisation from incomplete data), their results will probably hold for this problem as well. What they found was that the methods that use information from the immediate past and future time frames consistently performed better than the other methods that used temporal information and the method that used only Tikhonov regularisation. Another excellent paper (Brooks *et al*, 1999), also reconstructing epicardial potentials, showed that incorporating both spatial and temporal regularisation simultaneously or two types of spatial regularisation simultaneously had several advantages. The authors used an extension of the L-curve, the L-surface, to find the optimum value for the two regularisation parameters. The main findings were:

- Combination of two regularisers provides more robust estimations with respect to the choice of truncation parameters, due to the behaviour around the 'corner' in the L-surface.
- A smaller amount of spatial regularisation is needed when temporal regularisation is used as well. Or a smaller amount of spatial regularisation of one kind is needed when spatial regularisation of another kind is used in addition. This ensures that, although noise is successfully dampened, the solutions are not too much distorted.

See the original paper for further details of the results and further references to regularisation approaches.

2.3.2.4 Application

A first application of the MNLS to simulated and measured magnetic data was reported in 1984 (Hämäläinen and Ilmoniemi, 1984, Singh *et al*, 1984). The paper by Singh *et al* (1984) showed the possibility to obtain source images from magnetic data for a 2D source space. Although no measures were given with which to assess the accuracy of their results, the reported images of the source reconstructions seemed encouraging.

A series of simulations using simple 2D source spaces (Hämäläinen and Ilmoniemi, 1984, Wang *et al*, 1992) evaluated the aptness and limitations of the method more thoroughly. Hämäläinen and Ilmoniemi (1984) reconstructed sources in a plane parallel to the measurement plane that contained point-like sensors. No volume conductor effects were modelled. The effect of an increase in the number of sensors on the reconstruction of current dipoles from, noise contaminated, simulated data was a decrease of the spread in the image. Additionally, the reconstruction of multiple sources and reconstructions for evoked responses were presented.

The source space used by Wang *et al* (1992) was constructed out of two planes, forming an L-shaped 'wall', mimicking the shape of a small portion of a fissure or sulcus. The sources that form the source space were perpendicular to the reconstruction wall. Magnetic fields were computed in a horizontal planar observation plane, located at a distance above the vertical source space. The source space was placed in a conducting half space (i.e. volume conductor effects were not modelled). The following findings were obtained, using noiseless data produced by discrete sources:

- The reconstructed sources were spread around the location of the original source. However, the total strength (i.e. integration of the strength of all elements in the source space) was always within 0.1% of the strength of the original source (i.e. a 'conservation law').
- The spread decreased with smaller grid spacing.
- Side lobes could be observed around the main peak, caused by the ill-conditioning of the lead field matrix (Ramon *et al.*, 1993).
- Shallower sources were better-localised (smaller spread). This is due to the $1/r^2$ dependence in the Biot and Savart law, and therefore in the lead fields. The lead field pattern will vary more sharply with distance near a sensor, and the sensor will therefore be more sensitive to the position of a nearby source.

- Two sources separated by 2cm at the same depth could be correctly reconstructed, not hindered by the fact that they produced a dipolar field pattern. Note that no assumption about the number of sources was necessary to obtain this result.
- For shallow sources, a small amount of noise reduced the side lobes and higher amounts of noise reduced the peak amplitude (the spread was increased) and affected the peak location, although not considerably. For deeper sources the influence of noise was considerable. This could be expected, since the singular values associated with deeper sources are similar in amplitude as the singular values associated with noise. The distinction between noise and signal becomes less clear, resulting in more disturbed reconstructions.
- Noise introduced spurious sources near the edge of the source space. Wang and colleagues explained this as a result of undersampling of the magnetic field. A better explanation follows from the results in the previous section. Small singular values have a large influence on the pseudoinverse when not properly truncated, which explains the large amplitude of these sources. The fact that they are found near the edge of the source space is explained by the fact that the MNLS method reconstructs the solution with minimum norm. A superficial spurious source would need relatively smaller amplitude to produce the same field as a deeper source, due to the $1/r^2$ dependency of the lead fields.

A subsequent article (Wang, 1993) confirmed these results for a slightly more realistic simulation, using again an L-shaped 'wall', this time placed in a spherical volume conductor. The sensors were located on a spherical surface that was located at some distance above the source space. They further found that the side lobes found in reconstructions of discrete sources disappeared when the original source was a distributed source. Another interesting result was found for simulations with the source space and imaging space at slightly different locations. This broadened the peaks and introduced superficial sources near the boundaries. The centre of the pattern was however found at approximately the correct location for shifts smaller than 4mm. They further examined the influence on the solutions of the number of sensors and elements in the source space. In general, detail in the image was reduced with a reduced number of sensors. Additionally, bias towards the sensors increased. The minimum number of sensors needed depended on the number of elements in the source space.

The influence of the number of sensors, source elements and noise on the MNLS solutions was further examined in a simulation study by Ramon *et al* (1993). They used a planar source space with a planar sample measurement array above it. They found an ever-diminishing improvement in resolution for an increase in the number of source elements. This is in agreement with the observation that when the number of source elements is too high, the lead fields of the element will become more alike. This poor linear independence in the lead field matrix causes a greater sensitivity to noise (Okamoto *et al.*, 1983). Additionally, they showed an improvement in resolution when the number of sensors was increased from 81 to 441 (!). Further increase in the number of sensors cause numerical instability, due to the similarity of the rows of \mathbf{L} . Finally, noise did not increase the spread of the main peak by more than 5% for realistic noise levels. This was probably due to the relatively high ratio of sensor to grid elements.

A phantom experiment (Kullmann, 1991) confirmed the results by Wang *et al*, 1992, showing good reconstruction of a single dipole on a plane. The spread increased slightly when the depth of the plane was increased. Two dipoles located in the same plane could be reconstructed when the distance to the sensors was not too large. For deeper planes, only one major peak was reconstructed. A noiseless simulation with a complicated distributed source resulted in the correct reconstruction of the details of the source.

Although these results were promising, Jeffs *et al.* (1987) had already shown in 1987 (!), that the MNLS is unable to accurately reconstruct sources in a 3D source space. They used a sphere of 3cm radius, centred around the origin, as source space. This source space was divided in 1cm^3 elements. First, noiseless data were simulated, induced by 3 sources lying in a single plane. The sensors were lying on a half-sphere, with radius of 4cm. The solutions obtained with the MNLS method were blurred. Moreover, there was a clear bias towards the sensors, and the two deeper lying dipoles were almost completely missed. Similar results were obtained for a noisy dataset. Results for a distributed source in a larger source space (sphere with 6.5cm radius), with sensors lying on a 8cm radius half sphere, showed a complete loss of internal detail of the distributed source. The authors therefore concluded that MNLS is not appropriate as a method for neuromagnetic source reconstructions in a 3D source space.

2.3.2.5 Limitations

- The MNLS estimation introduces a bias that favours solutions closest to the sensors. The lead field is largest for superficial sources. This implies that for a certain magnetic induction, the magnitude of a superficial source can be smaller than that of a deeper source; which makes this the solution that is preferred by the minimum norm method (assuming that the residue is equal for both sources). This effect is seen when insufficient number of sensors are used for 2D reconstructions and in general for 3D reconstructions.
- The solutions found for magnetic data produced by discrete sources are spread out, which makes the results more difficult to interpret. However, the number of active areas does not have to be estimated at the start of the localisation process. The blurred images can therefore be used to estimate the number of sources. Further analysis of the data with one of the algorithms described in section 2.2 might result in more accurate reconstructions.
- The influence of noise is not dramatic for 2D source spaces, when regularisation is applied. However, the MNLS method fails when large external noise sources are present (Sekihara *et al*, 1994).

It will be shown in the next section that the MNLS is a special case of a generalised Wiener-Helstrom filter and that the general form can be used to reduce the influence of noise on the solutions and increase the accuracy of the reconstructions for 3D source spaces.

2.3.3 Generalised Wiener-Helstrom Filter

The incorporation of noise information into the reconstruction method could reduce the influence of noise on the source reconstructions. Assuming that the current density distribution and the noise are random and stationary, Wiener estimation can be applied (Shim and Cho 1981, Kullmann 1991). This estimator provides solutions that agree with the detected magnetic fields in a least squares sense and simultaneously allows to include first- and second- order statistical information about the sources and the noise. Grave de Peralta

Menendez and colleagues showed that the generalised Wiener estimator optimally separates the forward solution of a dipole from the noise in the measured data (Grave de Peralta Menendez *et al.*, 1997).

Second order statistical information is provided with the covariance matrices of the sources and the noise, \mathbf{K}_q and \mathbf{K}_n respectively.

The optimal linear estimator must minimise the mean square error (MSE) between the source current class \mathbf{q} and the reconstructed current class $\hat{\mathbf{q}}$, given by (e.g. Smith and Barrett, 1988, Smith *et al.*, 1990, Kullmann, 1991):

$$\text{MSE} = \text{Trace} \sum_{\mathbf{q}} \sum_{\mathbf{n}} (\hat{\mathbf{q}} - \mathbf{q})(\hat{\mathbf{q}} - \mathbf{q})^T p(\mathbf{n}|\mathbf{q})p(\mathbf{q}), \quad (2-34)$$

with $p(\mathbf{n}|\mathbf{q})$ the probability of occurrence of the noise vector \mathbf{n} given the source current class \mathbf{q} and $p(\mathbf{q})$ the probability of occurrence of \mathbf{q} .

The solution found with the estimator that minimises this error is given by (the interested reader can find a detailed derivation in Smith and Barrett (1988)):

$$\hat{\mathbf{q}} = \mathbf{K}_q \mathbf{L}^T (\mathbf{L} \mathbf{K}_q \mathbf{L}^T + \mathbf{K}_n)^{-1} (\mathbf{B} - \langle \mathbf{B} \rangle) + \langle \mathbf{q} \rangle, \quad (2-35)$$

where $\langle \rangle$ denoted the expectation value. Assuming that we know nothing about the current distribution (rather weak prior knowledge), the expectation values of \mathbf{B} and \mathbf{q} are zero. The generalised Wiener-Helstrom estimator is then given by:

$$\mathbf{W}_{w-H} = \mathbf{K}_q \mathbf{L}^T (\mathbf{L} \mathbf{K}_q \mathbf{L}^T + \mathbf{K}_n)^{-1}, \quad (2-36)$$

It is now clear that the MNLS is a special case of the generalised Wiener-Helstrom filter, for which no *a priori* information about the noise and source covariance matrices is used. Additionally, comparison of equation 2-36 with 2-22 equation shows that incorporating the covariance matrices is a way of regularisation. Instead of inverting the ill-conditioned lead field matrix, the inverse of the better-conditioned matrix $\mathbf{L} \mathbf{K}_q \mathbf{L}^T + \mathbf{K}_n$ is taken.

Utilising the information in the two covariance matrices results in reconstructions that are superior to the MNLS reconstructions with respect to noise. The first application of

the Wiener filter to simulated magnetic data was reported by Smith *et al* (1990). Two different source covariance matrices were used. A strong one, which was the covariance matrix of the actual sources, and a weak one, assuming that the sources were uniformly distributed with no source-source correlation. Additionally, eight different noise covariance matrices were used, assuming different variances for the noise. Sources were reconstructed in a cubic source space. Magnetic fields were computed in a planar and spherical sensor array. The noise added to the data was in correspondence with each *a priori* defined noise covariance matrix. The reconstructions obtained with the spherical array were more accurate than with the planar array. In general, an increase in noise variance reduced the reconstruction accuracy. When weak *a priori* information was incorporated with the planar sensor array, the reconstruction performance decreased with increasing depth, as was already found for the MNLS method. The estimator incorporating strong *a priori* information performed better than the estimator incorporating weak *a priori* information. This showed the improvements that are possible with the Wiener filter, compared with the MNLS method. A resulting question though is how accurate the *a priori* information should be. This was examined for coded-aperture imaging by Smith and Barrett (1988). They showed that using wrong *a priori* information (i.e. inaccurate \mathbf{K}_q) results in poor reconstructions that might contain spurious sources. They found that the results degraded with an increase of incorrect *a priori* information, obtaining almost meaningless reconstructions when little correct *a priori* information was supplied.

The above described simulations use the covariance matrix of the sources. No method was described as how to obtain this information for real data though. A method to obtain the covariance matrix for the sources with real data is given in Smith (1992). A lengthy derivation results in the following expression for an estimation of \mathbf{K}_q at delay time τ :

$$\hat{\mathbf{K}}_q(\tau) = \mathbf{L}^+ (\mathbf{K}_B(\tau) - \mathbf{K}_n(\tau))^{-1} (\mathbf{L}^+)^T, \quad (2-37)$$

The covariance matrix of the data, \mathbf{K}_B , can directly be derived from the measured magnetic data (Sekihara *et al.*, 1999) and the noise covariance matrix can directly be derived from a measurement in the absence of signal or from the unaveraged data (e.g. Srebro, 1996a). So, an estimate of the covariance matrix of the sources is available for use in equation 2-36, using a delay time of $\tau=0$. Consequently, the source reconstructions directly satisfy the

measured magnetic coherences. Unfortunately, no examples of this approach were given in the original paper.

Exploitation of the noise covariance alone also showed superior results to the MNLS method in a simulation with large external noise sources and brain noise (Sekihara *et al*, 1994). The simulation set-up is shown in figure 2-8. The formulation for the Wiener estimator they used is given by:

$$\mathbf{W}_w = \mathbf{L}^T \mathbf{K}_n^{-1} (\mathbf{L} \mathbf{L}^T \mathbf{K}_n^{-1} + \gamma \mathbf{I})^{-1}, \quad (2-38)$$

Minimum norm reconstruction of a single source was correct when no sensor noise was added to the data. Brain noise degraded the reconstruction considerably. When large external noise was added the reconstruction failed completely as was the case when large external noise and brain noise was present simultaneously. Incorporating the noise covariance matrix in the linear estimator gave a good reconstruction, even when both types of noise were present simultaneously. A subsequent paper (Sekihara and Scholz, 1995) modified the estimator in order to incorporate the data covariance matrix as well as the noise covariance matrix. The estimator they used is given by:

$$\mathbf{W}_w = \mathbf{L}^+ (\mathbf{K}_B - \mathbf{K}_n)^{-1} \mathbf{K}_B^{-1}, \quad (2-39)$$

The simulation set-up was similar as in the previous paper, except for the fact that the volume conductor was now an infinite half-space conductor (i.e. boundary effects were not modelled) and there were two original sources in a plane. The results confirmed the findings reported so far, and can be summarised as follows:

- minimum norm solutions were sensitive to the choice of regularisation parameter whereas solutions obtained with the Wiener estimator showed more tolerance to the choice of regularisation parameter and were less blurred. This suggests that there is a tight coupling between weighting the solutions, by using the covariance matrices, and regularisation. This will be confirmed later on.
- For data with very high and very low SNR the MNLS and Wiener estimations were similar. That is, low SNR data resulted in incorrect solutions and for high SNR data the solutions were accurate. For a deep source the MNLS estimation failed whereas the Wiener estimator produced a correct, though blurred, solution. Inaccurate estimation of

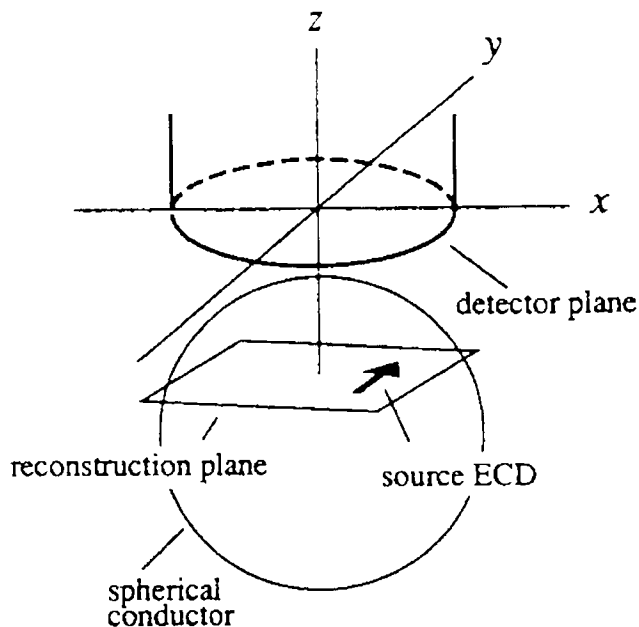


figure 2-8: Detector plane, volume conductor and source space used in the simulations by Sekihara *et al* (1994). The detector plane consisted of 37 sensors.

the data covariance matrix deteriorated the Wiener solutions, and bad estimations of the noise covariance matrix resulted in incorrect solutions.

Dale and Sereno (1993) used this linear estimator to incorporate anatomical and temporal information in order to improve source localisation. They used a different approach to obtain a covariance matrix of the sources. Simulated MEG and EEG data were combined and activity was reconstructed on a model of the cortical surface. This model was

computed from a set of MRI data and about 10000 dipoles were placed evenly on this model. The dipoles were oriented perpendicular to the cortical surface (for reasons described in section 2.1) with a spacing of about 4mm.

The source covariance matrix was used to incorporate the following information:

- Information about the correlation between the dipole strength at different locations. If the correlation between any two cortical patches is known, the source covariance matrix can be given by:

$$K_{s,ij} = \sigma_i \sigma_j \text{Corr}(i, j), \quad (2-40)$$

where σ_i is the standard deviation of the strength of the i^{th} dipole and $\text{Corr}(i, j)$ is the correlation between the strengths of the i^{th} and j^{th} dipoles.

- Information about (a part of) the entire time course, rather than just a single time point. It was described in section 2.2.4.2 how the noise subspace projection of the lead field vector \mathbf{L} could be used to estimate the locations of true dipoles. This measure was incorporated into the linear estimation framework by Dale and Sereno (1993) as something similar to an *a priori* variance estimate of the strength, σ_i^2 , for the i^{th} dipole, as:

$$C_{s,ii} = f\left(\frac{1}{\xi_i}\right), \quad (2-41)$$

with ξ_i as a generalisation of the value used with the MUSIC algorithm and f is a continuous, nondecreasing function. This can be coded into the estimated source covariance matrix as something like the variance estimate. The covariance matrix then becomes:

$$C_{s,ij} = \sqrt{C_{s,ii}C_{s,jj}} \text{Corr}(i,j), \quad (2-42)$$

Both 2-D and 3-D simulation studies were performed, for which a single sphere model (MEG) and 3 sphere model (EEG) was used in the forward computations. No noise was added to the data.

With the 2-D study the EEG and MEG induced by a single source was computed in 32 sensors. It was shown that the use of EEG and MEG simultaneously gave far better results than using one of the techniques alone. Furthermore, the use of temporal information strongly improved the reconstruction of a deep source, that is the solution was not as spread-out as normally found with minimum-norm-like techniques.

The 3-D study, with 61 sensors for both MEG and EEG, showed that the eight original sources, that were well separated, could be distinguished when EEG and MEG was combined and temporal information was used. However, nearby correlated sources were shown to merge if they were too close together.

2.3.4 Probabilistic Methods

In section 2.3.2 it was described how the solution that had the minimum norm property was selected from the infinite number of solutions. Another way to select a solution is by choosing the solution with the highest probability. Each possible (continuous or discrete) source distribution can occur with a certain probability. Therefore, the complete set of possible source distributions can be described by a probability density function.

In order to quantify this idea, the Bayesian technique is introduced.

Bayes' theorem is based on three kinds of probabilities:

1. A *prior* distribution that contains information about the quantity that we are interested in (i.e. *a priori* information about the source distribution).
2. A likelihood function that contains a set of conditional probabilities for the measurements, considering the quantity of interest. This contains information about the measurements when it is conditioned on the value of the quantity of interest (i.e. the magnetic field that would be produced by a certain source distribution).
3. A posterior distribution that contains the beliefs about the quantity of interest following new evidence in the measurements.

It is now possible to define the posterior probability mathematically by the Bayesian technique: The estimator of dipolar activity is the most probable one with regards to measurements and *a priori* considerations. That is:

$$p(\mathbf{Q}|\mathbf{B}) = \frac{p(\mathbf{B}|\mathbf{Q}) \cdot p(\mathbf{Q})}{p(\mathbf{B})} \propto p(\mathbf{B}|\mathbf{Q}) \cdot p(\mathbf{Q}), \quad (2-43)$$

where $p(\mathbf{B}|\mathbf{Q})$ is the probability of observing the data \mathbf{B} , given the sources \mathbf{Q} (this is the forward problem). $P(\mathbf{B})$ is the probability of the observed data that serves as a normalising constant and $p(\mathbf{Q})$ the probability of \mathbf{Q} , which can be used to include *a priori* information.

The most probable source distribution is then found by taking the expectation of this probability density function or by computing its maximum value.

2.3.4.1 Assumptions

- The noise in the measured data is Gaussian with zero mean (Gull and Daniell, 1978, Clarke and Janday, 1989).
- It is assumed that the primary current can be modelled as a patch of single dipole sources (the source space is discretised) when the concept of configurational entropy is used (Gull and Daniell, 1978).

- It is assumed that the primary current is continuous when the concept of probabilistic entropy is used (e.g. Clarke and Janday, 1989).
- Very large currents are not likely to occur. Therefore, the *a priori* probability distribution of current density is assumed to have a finite standard deviation β .

2.3.4.2 Technique

The first probabilistic method is based on maximising the configurational entropy. It was developed by Gull and Daniell (1978) to solve an inverse problem in radio astronomy. The aim was to reconstruct a map of the incident radiation intensity from an incomplete set of noisy data. The method selects a single map of the radiation intensity from all the maps that are compatible with the measured data. This is analogue to the inverse problem in neuroelectromagnetism, which is the problem of selecting a source current distribution from all the possible source current distributions that are compatible with the measured data.

The aim here is to find an estimator of the current dipole magnitudes $\hat{\mathbf{Q}}$ that maximises the posterior probability distribution of \mathbf{Q} , given the measurements \mathbf{B} . This is done by maximising the configurational entropy under the constraint that the chi-square is equal to the number of measurements. In other words, the most likely solution must also satisfy:

$$\chi^2 = N, \tag{2-44}$$

with χ^2 as defined in equation 2-2.

This is a means of regularisation that extracts the meaningful data from the measurements while overcoming the noise. The resulting solutions are always smooth.

A different, but related maximum entropy method has been described by Jeffs *et al* (1987). Their MEG simulations show that the source reconstructions found with the maximum entropy method exhibit the same smoothing and biasing problems as the minimum norm method.

The mathematical framework in which the expectation (or mean) of the probability density function is computed was developed by Clarke and Janday (Clarke, 1989, Clarke and Janday, 1989). The probabilistic entropy is maximised instead of the configurational entropy, again with the constraint on χ^2 that its expectation value is N. A further constraint is that the probability measure is normalised to unity. The advantage of taking the expectation value, instead of the maximum value, is that the method is less sensitive to errors in assigning the probability distribution.

This framework was used to develop a probabilistic inversion algorithm, better known as Magnetic Field Tomography (MFT). This algorithm was tested with magnetic simulations (Clarke and Janday, 1989, Ioannides *et al.*, 1990), evoked magnetic fields (Ioannides *et al.*, 1989, Ioannides *et al.*, 1995), evoked magnetic fields with single trial analysis (Ioannides *et al.*, 1995) and localisation of magnetic epileptic activity (Bamidis *et al.*, 1995). See the original papers for a full description of the implementation of the method. There are three things worth mentioning here:

- Firstly, the spatial prior probability density for $\mathbf{J}(\mathbf{r}'_j)$, $p(\mathbf{r}'_j)$, can be used to bring *a priori* information into the solution. The magnetic induction is then written as:

$$\mathbf{B}(\mathbf{r}_i) = \sum_{j=1}^n \mathbf{L}_i(\mathbf{r}'_j) \cdot \mathbf{J}(\mathbf{r}'_j) w p(\mathbf{r}'_j) \Delta S_j \quad i = 1, \dots, N \quad (2-45)$$

For the case of $p(\mathbf{r}'_j) = 1$ throughout the source space, this reduces to equation 2-18, proving that the minimum norm and probabilistic approaches are similar. It is noteworthy that radial currents are naturally excluded from the source reconstruction with the maximum entropy method, as was seen earlier with the minimum norm method.

- Secondly, the method uses a regularisation parameter that allows a fit to the data that is no more accurate than the standard deviation of the noise. The parameter is furthermore based on the standard deviation β of the prior probability distribution of current density. This is based on the fact that very large currents are not likely to occur. Therefore, the *a priori* probability distribution of current density is required to have a finite standard deviation β .

- Thirdly, the algorithm can be used iteratively to obtain localised solutions. This aspect of MFT will be further described in section 2.4.3.2, as well as the results of several computer simulations and experiments.

2.3.4.3 Limitations

Jeffs and co-workers (Jeffs *et al*, 1987) found that the maximum entropy suffers from the same drawbacks as the minimum norm method. That is, biasing of the solution towards the sensors and smearing of the solution over the source space. A correct implementation of the MFT algorithm can partly overcome these problems by using the spatial prior probability density to correct for the biasing and an iterative procedure to obtain more localised solution. However, the latter solution increases the computing time of the already very time-consuming algorithm. Furthermore, as stated by Ioannides *et al* (1990), the method is most naturally adapted to distributed sources.

2.4 From Distributed Sources to Discrete Sources

2.4.1 Introduction

The methods described in the previous section result in a distributed estimate of the primary current in the source space. The continuous current paths seen in MNLS are difficult to interpret for a neurophysiologist (Hari, 1991), as it is difficult to pinpoint the source locations in these images. The primary current is known to originate in a relatively small area of the brain and we are therefore looking for the more localised solutions that are found with current dipole fitting. In other words, we are looking for a method that gives us a solution for which most of the elements in the source space are zero. Thus, the only assumption made in these methods is the sparsity of the solution (defined by Gorodnitsky and Rao (1997) as a solution for which the number of non-zero elements is equal to or less than the number of measurements). Recently, many groups have developed such algorithms, with most exotic names, like LORETA (Pasqual-Marqui *et al.*, 1994), FOCUSS (Gorodnitsky *et al.*, 1995) and MFT (e.g. Ioannides *et al.*, 1990, Bamidis *et al.*, 1995). These methods are described in the next sections, in such a way that it becomes clear that all these methods utilise the same idea. That is, they use (mostly in an iterative procedure) a weight on the distributed sources in order to obtain localised sources. The recursive weight matrix is simply multiplied with a matrix that contains the *a priori* information and the resulting combined weight matrix is used in the computations. By doing so the algorithm can be forced to give a solution which is a member of a certain class of solutions, for example a solution with non-zero elements in a certain part of the brain. These algorithms therefore find the localised solutions, **without** making any prior assumptions about the number of sources. The weights are further used to correct for the biasing towards the sensors, from which the MNLS suffers. The use of different weights all comes down to: adding a stabilising term that regularises the estimator. This extra term measures something like the smoothness of the desired solution, or a quantity that parameterises the stability of the solution with respect to variations in the data, or a quantity reflecting *a priori* judgement about the likelihood of a solution. Minimising this term provides a "smooth" or "stable" or "likely" solution, which has nothing to do with the measured data itself (Press *et al.*, 1992).

It must be said that although the methods are very similar (equation wise), they are developed independently and with completely different philosophical backgrounds and

assumptions. For each method the kind of *a priori* information that is used to construct the weighting matrices will be described.

2.4.2 Iterative Methods

2.4.2.1 Minimum Dipole Solutions

Matsuura and Okabe developed an iterative method that provides solutions with only a few sources, starting with a distributed source. They tested the method first with an 1D and a cubic source space (Matsuura and Okabe, 1995). They later modified the method so that it could better deal with noisy data and tested the method with a more realistic source space and phantom data (Matsuura and Okabe, 1997).

Assume that there are N sensors. Firstly, a source space is constructed, as would be done for the minimum norm method. Then, exactly N sources in this source space are randomly chosen and the amplitude of the other sources is set to zero. This ensures that the lead field matrix in equation 2-19 is square and therefore a unique solution for equation 2-20 can be found. However, this solution is highly dependent on the N sources that were selected. Therefore, only repeating the random selection of N sources until all combinations are tested would provide us with the best selection of N sources. Fortunately, under the assumption that the solution should exhibit the minimum current property (i.e. the l_1 -norm of the solution is minimised), the problem can be solved by Simplex Linear Programming (Press *et al*, 1992). For the revised simplex method there are some selection rules that give a unique solution for which the l_1 -norm is minimum. Hence, an exhaustive search, impractical for realistically shaped source spaces, is not necessary.

Note that, although in each step N sources are selected, the number of sources in the final solution does not depend on any *a priori* information about the number of sources (and is in general not N).

To avoid biasing towards the sensors, the lead field matrix was multiplied with a weight matrix that punishes superficial sources. The solutions obtained for a 1D source space, with noiseless data produced by one, two or three original sources, were located near the original sources, without side lobes or significant spread. The same was true for a cubic source space, when no noise was present or when the SNR was high. For low SNR large

spurious sources were found in the solutions (Matsuura and Okabe, 1995). By introducing an inequality constraint that could easily be incorporated in the revised simplex algorithm, the influence of noise on the solutions could be reduced (Matsuura and Okabe, 1997). That is, instead of finding a solution for equation 2-20, the following equation is used:

$$\begin{cases} \mathbf{Q} \leq \mathbf{L}^{-1} (\mathbf{B} + \gamma \boldsymbol{\Sigma}) \\ \mathbf{Q} \geq \mathbf{L}^{-1} (\mathbf{B} - \gamma \boldsymbol{\Sigma}) \end{cases} \quad (2-46)$$

where $\boldsymbol{\Sigma}$ is a vector that contains the standard deviation of the noise in each channel and γ is the regularisation parameter.

A proper value for γ was found by performing Monte Carlo simulations with several values for γ .

The aptness of the method was illustrated by the accurate localisation of an artificial dipole in a spherical phantom under noisy conditions.

Another approach was described by Tilg and colleagues (Tilg and Wach, 1995, Tilg *et al.*, 1995, Tilg *et al.*, 1995). They iteratively applied a linear estimator, whereby in each iteration the smallest reconstructed source was removed from the source space, thereby reducing the underdeterminedness of the system. The algorithm was halted when a pre-determined number of sources were left in the source space. The main imperfection of this method is the arbitrariness of the determination of the number of sources with which to stop. With a simulated MagnetoCardioGram (MCG) the minimum norm estimator was used to obtain solutions (Tilg *et al.*, 1995, Tilg *et al.*, 1995). The Wiener estimator (Smith, 1992) or equation 2-36) was used for simulated MEG data, incorporating the covariance matrices of the data and the noise (using equation 2-37). Regularisation was used to deal with the Gaussian noise, using the value of the regularisation parameter obtained with the L-curve method. The results showed that the method is capable of reconstructing focal sources.

A method that removes more than one source in each iteration, and therefore converges faster, was developed by Srebro and colleagues (Srebro, 1994). This method is described in the next section.

2.4.2.2 Shrinking Ellipsoid Inverse

Srebro and colleagues also realised that the distributed solutions obtained with the MNLS method were not physiologically realistic. Their first effort to obtain sparse solutions was based on a search through a pre-defined realistically shaped source space, searching for clusters of sources that could explain the simulated data (Srebro, 1994).

They used a realistic model of a part of the cortex, based on nine digitised contours of a cadaver head. A total of 1271 dipoles were placed on these contours. They used the physiological constraint that cortical sources are mainly the currents in the pyramidal sources that lie perpendicular to the cortical surface (section 1.3.2). The dipoles were therefore given an orientation perpendicular to the surface. The strength of each source was either 0 or 1. A search for all possible combinations of elements that could explain the measurements would be impractical. Therefore, clusters of sources were constructed. The size of the cluster was systematically varied, and all source elements were tested as the centre of these clusters. From the resulting set of solutions the one that best explained the data was chosen. Although the localised sources were close to the original sources (constructed out of clusters of about 12 dipoles), the method has a major drawback. Namely, it will fail for two sources that are located far apart, because a single meaningful cluster cannot be formed.

A more sophisticated approach, using a shrinking ellipsoid, was later reported (Srebro, 1996a, Srebro, 1996b). The first paper (Srebro, 1996a) describes the method and shows the main characteristics of the method. Some small, but important, modifications are described in the follow up paper (Srebro, 1996b).

The method starts off with a MNLS solution. The (normalised) squared amplitude of the sources is treated as if it were the probability of finding sources at these locations. An ellipsoid could then be defined around the source with highest probability. A new MNLS solution was then computed, using only the sources within the ellipsoid. This process was repeated, resulting in a smaller ellipsoid in each step, until sources with a probability larger than 1 occurred (which is obviously not possible). This excluded the need for an extensive search through a source space. Additionally, reducing the size of the source space results in a less underdetermined problem. This in itself results in more accurate source reconstructions (Wang, 1993).

The results of the simulations were promising. Srebro used the realistically shaped source space described above, placed in a 3-shell spherical volume conductor. Correlated

noise was added to the simulated potentials in 111 electrodes. Regularisation was applied, using zero-order Tikhonov regularisation. The original sources were constructed out of 7-30 dipoles that were located in a cluster. Such a source was looked upon as a single macroscopic source. They showed that the minimum norm solution for a single original source had large values around the correct location, but that the solution was very blurred and that it was difficult to tell whether the activity originated in the left or right hemisphere. The final iteration of the shrinking ellipsoid however provided solutions that were near the original sources. The solutions found when the sources were not constraint to be orthogonal to the cortical surface were less accurate. Additionally, reducing the number of electrodes to a more realistic number, 39, spread the solutions over both hemispheres. Two widely spread sources could be reconstructed as well, whereas failure of the method was expected because a single ellipsoid was used. Some spurious sources between the two sources were found with alternating positive and negative amplitudes, and therefore producing largely cancelling magnetic fields. Incorporating the noise covariance matrix in the Wiener-Helstrom estimator reduced the spread around the original sources and the number of spurious sources. The number of spurious sources was greatly reduced when some modifications were implemented (Srebro, 1996b). The solution obtained in each iteration was added to a vector containing the sum of the solutions of the previous iterations. After convergence of the algorithm, a final solution was obtained for the ellipsoid constructed from the large elements in this vector. This had 3 effects:

- 1) Sources that are consistently present in the ellipsoid will have a significant contribution to the final solution.
- 2) Sources that change sign during the iterations will not be present in the final solution. This is why the solution for the two spread sources was better with this method.
- 3) Sources present in the final iterations will have a relatively large contribution to the final solution as the amplitude of the sources increases when the ellipsoid shrinks. This follows from the 'conservation law' found by Wang *et al* (1992). When the ellipsoid shrinks the 'total source strength' is spread over less sources and therefore the contribution of each individual source increases.

Furthermore, a more practical stopping criterion was given. Instead of looking for sources with probability larger than one (a situation that can occur with real measurements), the algorithm was stopped when the number of sources in the ellipsoid became too small to obtain a value for χ_r^2 that satisfied $1 \pm \sqrt{2m}/m$, with m the number of sensors.

The method was extended to localise sources from multiple latencies and used to analyse the VEPs recorded from two subjects (Srebro and Oguz, 1997). The solution obtained from multi-latency data during an iteration consisted of a source reconstruction at each latency. The solutions were then weighted by the SNR at each latency. For each dipole, the maximum value of this weighted solution over time was used to construct a new ellipsoid. Using the maximum value instead of the sum ensures that the method does not favour large sources (that have relatively large amplitude at all latencies). The results obtained with this method agreed largely with the known locations of visual cortical areas V1 and V5 obtained with other methods.

2.4.2.3 Focal Underdetermined System Solver (FOCUSS)

The Focal Underdetermined System Solver is also an algorithm that obtains sparse solutions by iteratively updating previous obtained solutions. Effectively, this reduces the magnitude of many sources to zero, but without the need to create a (shrinking) ellipsoid. The algorithm developed by Gorodnitsky and co-workers (Gorodnitsky *et al*, 1995, Gorodnitsky and Rao, 1997) is based on a MNLS, with an additional weight on the solution. So, the solution can be stated as (Gorodnitsky *et al*, 1995):

$$\hat{\mathbf{Q}}_{\text{FOCUSS}} = \mathbf{W} (\mathbf{L} \mathbf{W})^+ \mathbf{B} = \mathbf{W} \mathbf{W}^T \mathbf{L}^T (\mathbf{L} \mathbf{W} \mathbf{W}^T \mathbf{L}^T)^{-1} \mathbf{B}, \quad (2-47)$$

FOCUSS is an iterative algorithm meaning that the values of the previous step solution are taken as the elements of the weighting matrix \mathbf{W} in order to enhance large elements and reduce small elements in the solution. This is repeated until the change in \mathbf{Q} is not significant anymore, indicating that the algorithm has converged. The simplest version of the algorithm is the non-compound version, for which \mathbf{W} is defined by:

$$\mathbf{W}_k = \begin{bmatrix} \hat{\mathbf{Q}}_{1,k-1} & & 0 \\ & \mathbf{O} & \\ 0 & & \hat{\mathbf{Q}}_{n,k-1} \end{bmatrix}, \quad (2-48)$$

with k the iteration number. So $\hat{\mathbf{Q}}_{n,k-1}$ represents the n^{th} element of $\hat{\mathbf{Q}}$ at iteration $k-1$.

It is also possible to take account of all the previous solutions, which makes it likely that the algorithm converges quicker. This is done in the compound version and \mathbf{W} becomes:

$$\mathbf{W}_k = \mathbf{W}_{k-1} \begin{bmatrix} \hat{\mathbf{Q}}_{1,k-1} & & 0 \\ & \mathbf{O} & \\ 0 & & \hat{\mathbf{Q}}_{n,k-1} \end{bmatrix}, \quad (2-49)$$

The algorithm always converges to a localised solution with at least a quadratic rate of convergence (Gorodnitsky *et al*, 1995, Gorodnitsky and Rao, 1997). When a more general form for the weighting matrix is used, namely \mathbf{W}^p , the rate of convergence is even $2p$, at least. The only drawback is that the initialisation has to be more accurate when $p \neq 1$ is used in order to converge to the correct solution (Gorodnitsky and Rao, 1997). The number of non-zero elements in the solution is never larger than the number of measurements N . It must be stated here that the algorithm is always able to find a localised solution, but that this solution is not necessarily the correct solution. A small perturbation in the initialisation makes sure that a localised solution is found when the algorithm gets stuck in a distributed solution. The solution that is found depends on the initialisation of the algorithm as well as on the choice of the weighting matrix. Furthermore, the number of possible solutions depends on the dimension of the source space. Therefore a more accurate initialisation is needed with a source space that has large dimensions, as only then the algorithm will converge to the correct solution (Gorodnitsky and Rao, 1997).

2.4.2.3.1 Bias Correction

The FOCUSS algorithm can partly overcome the biasing problem associated with MNLS, though the initialisation has to be very accurate (Gorodnitsky *et al*, 1995). Even if this was possible, this is not good enough in the presence of noise, and superficial sources are still favoured. The bias can however be adjusted by introducing extra scaling factors through an extra weighting matrix. It is not possible to estimate a scaling factor for each element in the lead field matrix separately, but approximate bias adjustment is possible by weighting each column in the lead field matrix separately. This corresponds to adjusting the

strength of each current dipole in the source space during the initialisation (i.e. during the computation of \mathbf{L}).

A type of bias adjustment that works well for relatively compact source spaces or when the sensors are closely spaced is called normalisation (Gorodnitsky *et al*, 1995). The weighting matrix is now multiplied by \mathbf{W}_{Bias} , the matrix that contains the scaling factors:

$$\mathbf{W}_{\text{Bias}} = \begin{bmatrix} 1/\|\mathbf{L}_1\| & & 0 \\ & \text{O} & \\ 0 & & 1/\|\mathbf{L}_n\| \end{bmatrix}, \quad (2-50)$$

where \mathbf{L}_n is the n^{th} column of \mathbf{L} .

However, the algorithm favours deep dipoles now. This is due to the fact that there is no compensation for the differences in size between elements within a column. This represents the differences in distance between an element of the source space and the distinct sensors. Extra factors in \mathbf{W}_{Bias} , based on the reciprocal of the average or maximum distance between an element of the source space and the distinct sensors, can partly reduce this problem. The exact values of the factors, which are different for each source space and sensor configuration, are found by trial and error.

It was suggested by Bedford (1995) that it is not justifiable to use a biasing matrix to overcome the depth bias. His conclusion is based on the following reasoning: suppose that a set of simulated magnetic measurements is produced from a deep current dipole source of known location. When an unweighted minimum norm method is used to solve the inverse problem, it reconstructs an erroneous superficial source. With an unbiased algorithm, the original source is correctly reconstructed. So there are two source reconstructions which are both compatible with the data. Suppose that the location of the true source was not known, which is the case with real data, then it would be impossible to know which of the two reconstructions was the correct one. On this basis Bedford concluded that it is not justifiable to use the unbiased algorithm. However, one could come to the opposite conclusion. As mentioned several times before, the inverse problem has no unique solution, so the two source reconstructions are only two representatives from an infinity of possible solutions. It is our task to find the solution that best explains the

measured data and satisfies constraints set by anatomy and physiology. When this is possible by using an unbiased algorithm, why not use it?

The unbiasing matrix can be chosen in such a way that a superficial source is reconstructed by the FOCUSS algorithm when the known source was superficial and a deep source is reconstructed when the known source was a deep source. This can be tested by performing computer simulations.

Recently, Brauer and colleagues (Brauer *et al.*, 1996) described a modified linear estimation approach, tested with a simulation study. They used an iterative algorithm to localise sources in a realistically shaped model of a human torso. Although the authors did not mention this, the algorithm they used worked in the same way as the FOCUSS algorithm. The weighting matrices they suggested was the normalisation matrix and a matrix with the field error, defined as:

$$w_{ij} = \mathbf{B}_{m,i} - \hat{\mathbf{B}}_{j,i}, \quad (2-51)$$

where $\mathbf{B}_{m,i}$ is the measured magnetic field in channel i and $\hat{\mathbf{B}}_{j,i}$ is the magnetic field in this channel due to the source at location j .

They further removed sources with a magnitude below a certain threshold and randomly shifted all the source dipoles with a small amount in each iteration in order to obtain localised sources. Their results were best when both weighting matrices were used, enabling (although not perfectly) the reconstruction of two antiparallel dipoles, separated by 2cm.

2.4.2.3.2 Relation with Statistical Methods

The expression in equation 2-47 is a special form of a formula obtained when the Moore-Penrose pseudoinverse is replaced by the linear estimator known as Wiener-Helstrom filter (see section 2.3.3). This filter is also capable of finding the least squares minimum norm solution. We repeat equation 2-36 here:

$$\mathbf{W}_{\text{w-H}} = \mathbf{K}_q \mathbf{L}^T (\mathbf{L} \mathbf{K}_q \mathbf{L}^T + \mathbf{K}_n)^{-1}, \quad (2-52)$$

with \mathbf{K}_q and \mathbf{K}_n the covariance matrices of the sources and the noise, respectively.

The source covariance matrix could include prior knowledge about the strength of the sources in different regions and correlations between various active areas. When we write $\mathbf{K}_q \mathbf{L}^T$ as \mathbf{W} and \mathbf{K}_n as \mathbf{I} , we obtain equation 2-47 again, with the inverse instead of the generalised inverse. To elucidate, a certain weight matrix in the weighted minimum norm algorithm has its counterpart in a matrix that contains statistical knowledge in the probabilistic method.

2.4.3 Other Methods

2.4.3.1 Low Resolution Electromagnetic Tomography (LORETA)

LORETA (Pasqual-Marqui *et al*, 1994) is another attempt to obtain a more localised solution than the minimum norm solution, although the result is still a distributed source. The method does not assume a limited number of dipolar points but it assumes that the current distribution is smooth. The smoothest of all possible current distributions is found by using the minimum norm formalism with a minimum spatial Laplacian instead of the minimum l_2 -norm. So, instead of minimising $\|\hat{\mathbf{Q}}\|_1$, they minimise $\|\mathbf{W}^{-1}\hat{\mathbf{Q}}\|_2$. The solution is then given by:

$$\hat{\mathbf{Q}} = \mathbf{W}(\mathbf{L}\mathbf{W})^+ \mathbf{B}, \quad (2-54)$$

where \mathbf{W} is an $n \times n$ matrix that provides the weighting on the solution.

\mathbf{W} is composed of two matrices, \mathbf{W}_1 and \mathbf{W}_2 , that are multiplied with each other. \mathbf{W}_1 corrects for the biasing problem, associated with MNLS, through normalisation. This will be more thoroughly discussed in the next section. \mathbf{W}_2 is the inverse of the discrete Laplacian operator matrix.

As a result of using the Laplacian operator, the method has a relatively low spatial resolution. This means that two adjacent, simultaneously active sources are reconstructed with LORETA as one blurred source.

The method has so far only been used with EEG data, either simulated data, evoked responses (Pasqual-Marqui *et al*, 1994) or epileptiform activity (e.g. Lantz *et al.*, 1997). A 3-shell spherical head model was used for the forward computations and a regular 3D grid as source space, lying inside the innermost shell.

By using simulated EEG data produced by two dipoles, it was shown that LORETA was able to separate two dipoles that were 3.1cm apart, whereas a **single** dipole fit resulted in a fit with the dipole located between the two original sources. A fairer case would be the comparison between LORETA and a two dipole fit. When the two dipoles were separated

1.5cm apart, LORETA could not distinguish them anymore and the result was one blurred source, located between the two original sources.

Two types of evoked responses were analysed. One was visual evoked potentials during checkerboard stimulation. The other one was auditory-visual evoked responses. The auditory signal served as a warning signal and the visual stimulus was either a black cross or a black square. The subjects were asked to press a button as soon as possible after presentation of the cross (target). Only the brain responses to targets were analysed. With the first type of stimulation the grand mean average over all 21 subjects was computed and the point of maximum global field power (GFP) between 80 and 130ms (corresponding to the visual P100 component) was used with LORETA and the single and 2 dipole fit method. With the second type of stimulation the grand mean average of 10 subjects was used and three time points of maximum GFP were selected. These time points corresponded with:

- the auditory N100 component
- a time point between 500 and 700ms before the onset of the imperative stimulus (representing the Contingent Negative Variation, CNV)
- a time point between 200 and 400ms after the onset of the imperative stimulus, corresponding to the P300 component.

The results of the visual evoked responses showed that LORETA generates physiologically meaningful results. The results were in fair qualitative agreement with the single dipole fit method for one of the sources revealed with LORETA. The second source was however not found with the single dipole fits (of course), nor with two dipole fits.

The results with the dipole fit methods for the auditory-visual evoked responses were not in agreement with the expected locations for the N100 component and resulted in unlikely localisations for the cognitive responses. However, LORETA gave results that were in agreement with the findings of Scherg and von Cramon (Scherg and von Cramon, 1985a, Scherg and von Cramon, 1985b, Scherg and von Cramon, 1986) for the N100 component, without introducing the *a priori* constraints used by Scherg and von Cramon. The method provided new suggestions about the locations of the cognitive components.

The above described results showed the ability of the method to localise electrical sources in the brain, where the dipole fit methods failed. The results were however still

distributed sources and the spatial resolution was low. The methods described in the following section try to overcome these problems by using an iterative weight (or re-weighting).

2.4.3.2 Magnetic Field Tomography (MFT)

The mathematical framework that is the basis of this probabilistic approach has been described in section 2.3.4. In this section we focus on the tests that were carried out with this algorithm and especially on the iterative version of the algorithm.

Tests

Firstly, the non-iterative version of the algorithm was tested with simulated magnetic fields, induced in 121 sensors by two anti-parallel dipoles located in a disk (Clarke and Janday, 1989, Ioannides *et al*, 1990). No *a priori* information was used, which means that a uniform spatial prior probability density for $\mathbf{J}(\mathbf{r}_j')$ was used and 10% maximum-signal noise was added to the data (see Ioannides *et al* (1990) for the definition of maximum signal noise). The results were more stable when the regularisation parameter was given a value unequal to zero. They also found that the algorithm performed better when more distributed sources were used to generate the magnetic field. The reconstructions with noisy data were good when the regularised algorithm was used. Furthermore, the iterative version of the algorithm was tested with the two dipoles. In each iteration the probability function was defined in terms of the current density found in the previous iteration. This is comparable with the uncompound version of the FOCUSS algorithm. Three versions were tested:

- 1) $p(\mathbf{r}_j') = \mathbf{J}(\mathbf{r}_j')$
- 2) $p(\mathbf{r}_j') = \sqrt{\mathbf{J}(\mathbf{r}_j')}$
- 3) $p(\mathbf{r}_j') = \mathbf{J}^2(\mathbf{r}_j')$

The best choice was found to be version 1, as the second version converged too slowly and the third version converged too strongly towards the strongest source. The results with all three versions were rather localised solutions.

In a 3-D simulation Ioannides and colleagues (Ioannides *et al*, 1989) showed that it is possible to reconstruct two anti-parallel sources in a cylinder, although the magnitude of the lower source was difficult to recover. They used a prior probability weight that favoured deep sources (anti-biasing).

Secondly, the algorithm was tested with measured evoked fields, following the reversal of a checkerboard pattern. A single channel was placed in 42 different positions above the head. The regularised version of the algorithm was used, as well as a uniform *a priori* probability. The results are difficult to interpret, but clearly show the arrival and development of activity in the visual cortex. Such an ‘animation’ is not possible with a single current dipole fit method, since the results of such a method are unstable when the signal to noise ratio is small, i.e. at the onset of the activity sequence.

In another experiment, MFT was used to reconstruct sources from averaged and single trial MEG signals in an auditory “odd-ball” paradigm and an experiment with simple auditory tones (Ioannides *et al*, 1995). The results demonstrated how the technique could be applied to real data and seemed physiologically reasonable. Either the result of the instantaneous MFT or the integration of activity over time was shown. To avoid biasing towards the sensors, the spatial *a priori* probability function was set to be a Gaussian function, centred around the origin of the conducting sphere. A one step iterative algorithm was used, in which the computations were repeated with a spatial *a priori* probability distribution that was equal to the original Gaussian function multiplied with the modulus of the source current density of the first iteration. In this way, more (but not highly) localised solutions were obtained. No reason is given in the paper why the algorithm was not used with more than one step.

The single trial analyses showed a small variability in the solutions from trial to trial. They suggested that in different trials, the brain handles the incoming signals differently. Another explanation however could be that the differences were due to differences in the noise. This could not be tested as no information was given about the noise in the signals.

It was further shown that the algorithm was capable of localising deep sources in the presence of superficial activity, using recorded data.

Thirdly, MFT was used to localise interictal epileptic activity (Bamidis *et al.*, 1995, Ioannides *et al.*, 1995). The same algorithm was used as described above, except that Bamidis *et al.*, 1995) used only one iteration. A spatio-temporal cross-correlation matching technique was used to identify similar epileptiform events before averaging. It was shown that MFT is able to localise single epileptic spikes as well as deep activity, although the physiological plausibility of their results remained unclear.

2.4.3.3 A Bayesian Approach

Two algorithms based on a Bayesian approach have recently been developed independently by Baillet and Garnero (1997) and Phillips and co-workers (Phillips *et al.*, 1997). We describe here the basic idea of their methods, the assumptions they made and the tests they performed. For a full description of their algorithms we refer to the original papers.

Worth mentioning here is another interesting implementation of the Bayesian formalism was published by Schmidt and colleagues (Schmidt *et al.*, 1999). Their implementation differs mainly in that, instead of finding a 'single' best solution, a probability distribution of solutions is obtained. This distribution is then used to make inferences about the properties of active brain areas and the likelihood of these properties.

Method

The Bayesian approach is based on the idea of finding the solution to the inverse problem that maximises the a posteriori probability of \mathbf{Q} . The mathematical framework for this was described in section 2.3.4. This maximum a posteriori (MAP) estimator can be written as:

$$\hat{\mathbf{Q}} = \max_{\mathbf{Q}} (p(\mathbf{Q}|\mathbf{B})), \quad (2-55)$$

with $p(\mathbf{Q}|\mathbf{B})$ as given in equation 2-43.

When \mathbf{Q} and the noise ($= n$) are independent and normally distributed with zero mean and covariance \mathbf{K}_q and \mathbf{K}_n , the MAP of \mathbf{Q} gives as an estimate for \mathbf{Q} :

$$\hat{\mathbf{Q}} = \mathbf{K}_q \mathbf{L}^T (\mathbf{L} \mathbf{K}_q \mathbf{L}^T + \mathbf{K}_n)^{-1} \mathbf{B}, \quad (2-56)$$

This is equal to the solution obtained with the Wiener-Helstrom estimator, equation 2-36. It shows the equivalence between the regularised minimum norm method and the Bayesian approach. With the Bayesian estimate the weighting matrices are equivalent to a generally arbitrary selection of the covariance matrices. Baillet and Garnero (1997) and also Phillips and colleagues (Phillips *et al*, 1997) describe how they select the weighting on the solution. It is based on the observation that $p(\mathbf{Q}|\mathbf{B})$, the MAP estimator, can be defined as a Gibbs distribution. That is:

$$p(\mathbf{Q}|\mathbf{B}) = \frac{1}{Z} e^{-U(\mathbf{Q})}, \quad (2-57)$$

where Z is a normalisation constant, the so called partition function. $U(\mathbf{Q}) = U_1(\mathbf{Q}) + \lambda U_2(\mathbf{Q})$ where $U_1(\mathbf{Q})$ and $U_2(\mathbf{Q})$ are associated with $p(\mathbf{B}|\mathbf{Q})$ and $p(\mathbf{Q})$, respectively.

The MAP estimator is therefore found by minimising the energy function $U(\mathbf{Q})$. Baillet and Garnero (1997) associate all the prior information with $U_2(\mathbf{Q})$ and therefore their method only influences the regularisation, whereas Phillips and co-workers colleagues (Phillips *et al*, 1997) influence the whole energy function.

Assumptions

The assumptions Baillet and Garnero (1997) make to construct their energy functions are:

- Anatomical constraint on source locations and orientations, i.e. sources perpendicular to the cortical surface obtained from an MRI.
- Dipole magnitude evolution is smooth in time, since the frequency of bio-electromagnetic signals is smaller than 100Hz and the measurements are over sampled.
- Sources might or might not have correlated activities, i.e. discontinuities may exist in the source space.

- Information about active sites, obtained from Positron Emission Tomography (PET) or Functional Magnetic Resonance Imaging (fMRI), can be used. This assumes that metabolically active areas are closely related to active neuron groups.
- The MAP estimator can be defined as a Gibbs distribution.

Phillips and colleagues (Phillips *et al*, 1997) made the following assumptions:

- Current sources are sparse and focal. Therefore a binary indicator process is used to model whether each source dipole is on or off. Furthermore, active sources are assumed to have Gaussian amplitude in space.
- The MAP estimator can be defined as a Gibbs distribution.
- Assumption of additive white Gaussian noise.
- Sources are oriented perpendicular to the cortical surface.

Tests

Baillet and Garnero (1997) compared their method with the minimum norm method and LORETA, using simulated EEG data from 16 active areas (8 dipoles in each area) in a three-shell spherical head model. A total of 65 electrodes were used. Four of the patches were activated following a simple activation sequence for dipole magnitudes, using different scaling ratios and delays for each patch. White Gaussian noise was added to the data with a standard deviation that was 10 times smaller than the standard deviation of the simulated data.

Using the minimum norm method and LORETA the reconstructed activity was oversmoothed, although the latter method performed better. Both methods were sensitive to noise. The MAP method reconstructed the activity more accurately, with the best result obtained when temporal information was used. Furthermore the algorithm was more stable when this kind of prior information was used.

Phillips and co-workers (Phillips *et al*, 1997) compared their method with the minimum norm, normalised minimum norm, LORETA and FOCUSS. They simulated MEG data in 80 sensors in a 1-D sensor array. The 2-D source space was a 90° annular segment containing 1098 dipoles, with a distance of 1.5mm between them. A total of 8

sources were activated in 3 clusters. Simulations were carried out with noiseless data and with added white Gaussian noise, resulting in an SNR of 20dB and 8dB.

The minimum norm, normalised minimum norm and LORETA reconstructed the activity oversmoothed. The results of the FOCUSS algorithm and the MAP estimator were comparable, resulting in sparse solutions (3 clusters were reconstructed). The performance of both methods deteriorated when the amount of noise was increased, resulting in the mislocalisation of at least one cluster.

They further tested the methods with experimental phantom data, showing that their algorithm was able to reconstruct the three original sources, although not all three locations were correctly found. The other methods performed comparable to the simulations described above.

A disadvantage of both methods is that they need some well-chosen tuning parameters, called hyperparameters. The values of these parameters can only be defined on an empirical basis. Baillet and Garnero (1997) use two parameter to tune the contribution of spatial and temporal prior information and two scaling parameters. Philips and colleagues (Phillips *et al*, 1997) use three parameters that determine the relative contribution of sparse and clustered sources and the strength of the clustering.

2.5 Synthetic Aperture Magnetometry (SAM)

An effective means of reducing the influence of noise sources on the recorded magnetic fields in the MEG sensors is to form a gradiometer configuration (section 1.3.1). The noise reducing effect is based on the spatial selectivity of the gradiometer, being more sensitive to nearby sources than to distant sources. In effect, the incoming signals are spatially filtered by the gradiometer.

Synthetic Aperture Magnetometry (SAM) also utilises the spatial selectivity of sensors; virtual sensors that is, located in a volume of interest within the brain. The output of each virtual sensor is a linear combination of the signals recorded in the physical sensors, just as the output of a gradiometer is the linear combination of the magnetic fields induced in the individual coils. The spatial selectivity of each virtual sensor is adjusted by optimising the set of weights used for this linear combination. SAM is a so-called statistically optimum beamformer, meaning that the weights are optimised based on the statistics of the data (Van Veen and Buckley, 1988).

Repeating equation 2-17 here:

$$B_i = B(\mathbf{r}_i) = \int_{\Omega} \mathbf{L}_i(\mathbf{r}') \cdot \mathbf{J}(\mathbf{r}') d^3r' \quad i = 1, \dots, N, \quad (2-58)$$

with B_i the measurement in the i^{th} sensor and \mathbf{L}_i the lead field of this sensor. N is the total number of sensors and \mathbf{J} is the current density that produces the magnetic field.

As stated above, a virtual sensor, B_{virtual} is formed by a linear combination of the measurements (Robinson and Rose, 1992):

$$B_{\text{virtual}} = \sum_{i=1}^N B_i w_i^{\text{virtual}}, \quad (2-59)$$

With w_i^{virtual} the weighting coefficients for the virtual sensor that determine the spatial selectivity of the sensor.

Combining the above two equations gives the following for the lead field of the virtual sensor:

$$\mathbf{L}_{\text{virtual}} = \sum_{i=1}^N \mathbf{L}_i w_i^{\text{virtual}}, \quad (2-60)$$

The ideal virtual sensor would be a sensor that only responds to signals that were produced by a source at the location of the virtual sensor. In other words, the ideal lead field of a virtual sensor is the delta-function. Hence, the weights are optimised so that the difference between the virtual lead field and the delta-function is minimal. Additionally, statistical information, in the form of the covariance matrix of the measurements, is utilised in the optimisation of the weights (see Robinson and Rose (1992) and Robinson and Vrba (1998) for details).

The influence of noise sources within the brain decreases when the spatial selectivity of a sensor is increased. The signal from brain sources is however not necessarily orthogonal to the signal from the target source. Hence, increasing the spatial selectivity also decreases the projection of the target signal. In order to maintain a strong response to the target signal, the magnitude of the weights has to be increased. This increases the amount of noise projected into the virtual sensor (Robinson, 1989), hence there is a trade-off between the spatial selectivity and influence of noise of the virtual sensor. This trade-off was seen before for linear estimation methods in the form of the L-curve from which an optimum regularisation parameter can be chosen (see section 2.3.2.3). In this context the trade-off curve between spatial selectivity and variance of a virtual sensor, also L-shaped, is described as the Backus-Gilbert curve and the output of the virtual sensor is regularised by the Backus-Gilbert multiplier μ .

A tomographic image is obtained by placing virtual sensors on a grid placed in the brain volume. The squared output, i.e. the source power, of each virtual sensor is then displayed as a colour coded map on this grid. Activity in areas that produce a signal in the physical sensors that is sufficiently orthogonal to signals from other brain sources are then found as bright areas in the tomographic image. Alternatively, the measured data can be run through a virtual sensor (Robinson, 1989, Robinson, 1999). By doing so, it is possible to reveal brain responses in single epochs, because the influence of other brain sources on the output of the virtual sensor is reduced.

The spatial selectivity in the tomographic SAM maps can be further enhanced by Significance Probability Mapping (SPM), first described by Duffy and colleagues to enhance differences in topographic maps for electrical recordings (Duffy *et al.*, 1981). Two types of SPM can be used:

- Z-transform statistic. This is used to highlight regions in a map that are statistically different from a set of reference maps. The result is a Z-value (the number of standard deviations by which an individual observation differs from the mean of a reference set) for each gridpoint so that another tomographic map can be formed from the Z-values.
- Student's *t*-test statistic. This is used to highlight regions in a set of maps that are statistically different from a set of reference maps. The mean and variance is computed for both sets of maps and Student's *t*-statistic is computed. The result is a *t*-value for each gridpoint so that another tomographic map can be formed from the *t*-values.

The main advantages of SAM are:

- The spatial filtering enhances the SNR so that active brain regions can be found that underlie unaveraged data.
- SPM is performed in the projected signal space, not in the signal space itself. Statistically relevant differences between SAM maps are therefore likely to be due to differences in brain state and not due to differences in recording parameters (like distance from the sensors).
- The statistical comparison of SAM maps can be used to analyse MEG data obtained with paradigms similar to those used in fMRI and PET (for example active versus control state paradigms, see Robinson, 1997).

2.6 Choosing the Inversion Technique

The aim of this section is not to select the 'best' technique. Instead a technique is chosen that can incorporate *a priori* information in a insightful and straightforward manner.

The merits and limitations of the techniques described in this chapter have been outlined in detail for each approach separately. Each technique can provide useful information about the brain activity underlying recorded magnetic fields, provided the results are interpreted with care and in the context of the constraints that have (implicitly) been applied. Additionally, an algorithm that is superior under all conceivable conditions does not exist, due to the very nature of the biomagnetic problem (Lütkenhöner *et al.*, 1996). That is, there are an infinite number of solutions for the unconstrained problem. A single solution can only be found when constraints are applied. Consequently, an algorithm can only be optimal under the constraints that were applied. The price for having advantages under certain conditions is having disadvantages under other conditions (Lütkenhöner *et al.*, 1996). Moreover, comparison of techniques is troubled by the variety of conditions under which they have been tested, i.e. different measurement set-ups, noise realisations, number of active sources etc. A few attempts to compare techniques objectively can be found in Lütkenhöner *et al* (1996) and in Grave de Peralta Menendez *et al.* (1996), Grave de Peralta Menendez *et al* (1997) and Grave de Peralta-Menendez and Gonzales-Andino (1998).

The technique that is going to be used must be able to select one solution out of the infinite possible solutions to the inverse problem. This solution must account for the data and must be stable with respect to the noise in the data. Such a solution should also be able to incorporate *a priori* information in a straightforward manner.

Most of the algorithms that are used to reconstruct a number of discrete sources (section 2.2) minimise a cost-function in order to find the solution that best accounts for the data and any constraints that are applied. *A priori* information can be readily incorporated by altering the cost-function, as discussed in section 2.2. However, the main drawback of these techniques is that the number of sources must be known beforehand, which is no trivial issue (section 2.2.2.3). Additionally, modelling activity of a distributed nature as a discrete source can introduce large errors.

The techniques that are used for the reconstruction of distributed sources (section 2.3) are well suited to incorporate *a priori* information in the form of spatial constraints, as it is necessary to define a source space for these techniques. Hence, the shape and location of this source space exploits *a priori* information of a spatial nature, i.e. anatomical constraints. Furthermore, temporal constraints can be used with these techniques as well. A big disadvantage of these techniques is however the biasing of the solution towards the sensors, although it is possible to counteract the biasing effect by using appropriate weights in the solution. Another disadvantage is that the reconstructed source images are often difficult to interpret, especially when multiple areas are active simultaneously.

The hybrid methods described in section 2.4 combine the advantages of the discrete and distributed techniques. The number of sources can be unknown at the start of these algorithms, yet the final solution contains only a few discrete sources. Also, a source space is used hence the incorporation of anatomical constraints is straightforward. Moreover, the use of weighting enables the incorporation of additional *a priori* information. For example, the weights can be adjusted so that sources in a certain region are favoured, in case it is known from other modalities like PET or fMRI that these regions in the brain were active (Liu *et al.*, 1998). The incorporation of information about the noise in the data is straightforward as well (Sekihara *et al.*, 1994, Sekihara and Scholz, 1995, Sekihara *et al.*, 1999).

The task is now to select one of the iteratively weighted techniques described in section 2.4. Note that although the different methods are based on different principles, the final equations to solve are similar. This is clearly seen when equations 2-47, 2-54 and 2-56 are compared with the formulation of the generalised Wiener filter (equation 2-35) (see also Phillips *et al.* (1997), Grave de Peralta-Menendez and Gonzales-Andino (1998) and Michel *et al.* (1999)). The similarity between MFT and FOCUSS has also been discussed (Taylor *et al.*, 1999). A limitation for the practical use of MFT is that it is a very time consuming technique. A disadvantage of all the techniques described in section 2.4 is that they make use of hyperparameters, whose values can only be found empirically.

The method of choice for this work is the FOCUSS algorithm, because the principles of the FOCUSS algorithm are mathematically straightforward and well described (Gorodnitsky *et al.*, 1992, Gorodnitsky and Rao, 1992, Gorodnitsky, 1994, Gorodnitsky and Rao, 1994, Gorodnitsky *et al.*, 1995, Gorodnitsky and Rao, 1997). Moreover, the

formulation of the FOCUSS algorithm allows for the easy incorporation of *a priori* information in the form of a weighting matrix and by using a source space that can be used to incorporate spatial constraints.

3 Using a Realistically Shaped Source Space

3.1 Introduction

Functional information on its own does not give a complete understanding of the organisation of the brain. Merging the functional information with anatomical data provides us with a more complete account of its functioning and interconnections. It is therefore common practise nowadays to combine the results of functional imaging studies (EEG/MEG, fMRI and PET) with anatomical data obtained with MRI or CT. Normally, this means that the functional data (for example a fitted equivalent current dipole from an MEG study) is overlaid on an image that contains the anatomical structures or on a 3D representation of the cortical surface (see for example Hämäläinen (1991), Fuchs *et al.* (1994), Lütkenhöner *et al.* (1995), Anderson *et al.* (1996)). This can only be done when the co-ordinate system in which the functional data is obtained is linked to the co-ordinate system in which the anatomical data is measured. This is called co-registration of both modalities. This co-registration must be accurate in order for the integration of the functional and anatomical information to be meaningful. Accurate co-registration is even more essential if anatomical information is used to constrain source localisation from MEG or EEG data. Inaccurate co-registration and therefore inaccurate anatomical constraints bias the localisation algorithms in a less predictable way. The applicability of the FOCUSS algorithm, using a realistically shaped source space, is thus dependent on the accuracy of the co-registration method that is used.

In section 3.3 methods of co-registration are examined, followed by an analysis of the accuracy of the proposed co-registration method in section 3.4. The most accurate method for co-registration of two datasets in different co-ordinate systems is based on surface matching (Schwartz *et al.*, 1996). With this method the distance between a set of points in the 'functional' co-ordinate system and a corresponding surface in the anatomical datasets is minimised, best illustrated as fitting a 'hat' onto a 'head'. For this purpose a surface has to be defined in the anatomical dataset from a set of 2D images. This utilises the same techniques

as used for the construction of a realistically shaped source space that can be used as an anatomical constraint in source localisation algorithms. These techniques are therefore first described in section 3.2.

First of all, one has to define the voxels in the brain that form part of the same structures, a process called segmentation. This can be done based on the properties of each voxel and its neighbours. There is a huge interest in this subject as it is an essential part of many research areas, ranging from pattern recognition to assessment of the effect of treatment of tumors. A brief review of the main segmentation approaches is given in section 3.2.2. The result obtained from the segmentation can then be put to use by selecting only the regions that form the grey matter. In section 3.2.3 it is described how these regions are connected so that a description is formed of the cortical surface. This cortical surface is then used to create a realistically shaped source space that is applicable as an anatomical constraint for source localisation algorithms (section 3.2.4). These methods are also used to form the 'head' surface that is used by the surface matching algorithm.

3.2 Constructing a Realistically Shaped Source Space

3.2.1 Introduction

The technique most often used to obtain anatomical information for MEG/EEG studies is Magnetic Resonance Imaging (MRI). It is also possible to obtain anatomical data from other techniques like Computed Tomography (CT) or Positron Emission Tomography (PET). However, PET is less accurate than MRI because of the larger voxel sizes that are used and more invasive because it uses a radioactive tracer (e.g. Aine, 1995). A drawback of CT scans is that the brain boundaries are not clearly defined, although bone is easily recognised.

In this section the focus is therefore on MRI and the technique is briefly described. More details of the technique, its possibilities and its limitations are given in an excellent review of MRI and PET by Aine (1995).

Magnetic Resonance Imaging

During an MRI scan a subject is placed within a static magnetic field of, typically, 1.5T. This field aligns the magnetic vectors (spins) of atomic nuclei. The spins can either be orientated 'parallel' or 'anti-parallel' with the direction of the external magnetic field. Since the parallel orientation is a state with lower energy, there is a net magnetisation in the parallel longitudinal direction. In fact, the spins precess around the direction of the external magnetic field with a frequency that is known as the Larmor frequency. This frequency is directly proportional to the strength of the magnetic field. The spins are out of phase when they are aligned by the external magnetic field and the signal that is produced is therefore zero.

A measurable magnetic field is produced when net magnetisation is created along the transverse axis, in such a way that the spins are in phase. This is achieved by applying a small RF field that is transverse to the steady magnetic field and resonates with the Larmor frequency. The energy delivered by the RF pulse brings spins from the lower energy state in the longitudinal direction into the higher energy state in the transverse direction, with the spins in relative phase. When these spins return to their original state, they release energy.

This signal is picked up by the receiver coil, which is sensitive to the transverse magnetic field only.

The strength of the MR signal that is produced is determined by the number of spins that were brought in the transverse direction and is therefore determined by the strength and duration of the RF pulse. A maximum signal is obtained when the net magnetisation is completely brought from the longitudinal direction into the transverse direction, i.e. the *flip angle* is 90° . Additionally, a larger separation between the lower and higher energy states increases the MR signal. This energy separation is increased by increasing the steady magnetic field, hence the quest for high Tesla MR scanners.

The above discussion explains how a MR signal is obtained from the body. The following reveals how these signals can be utilised to obtain anatomical information.

The signal emitted by spins that return to their low energy state, after the RF field is switched off, decays with a certain rate. The time in which 63% of the spins return to their initial, lowest energy, state is known as T1. Furthermore, the MR signal decays because the spins in the transverse direction start rotating out of phase. The time in which the signal is




figure 3-1: Part of an MRI dataset displayed as a stack of images. Each image shows a coronal view of a cross section of a subject's head. Only every 10th slice is displayed for the ease of interpretation of the figure (starting with the 70th and stopping with the 110th slice). In each slice only some contours are displayed, also for the ease of interpretation. Note that these contours are not part of the original MRI dataset, but are obtained after further image processing (see section 3.2.2).

reduced to 37% of its maximum is known as T2. Thus, T1 and T2 refer to the time that is required to restore 63% of the original longitudinal magnetisation and lose 63% of the recorded signal strength, respectively. The T1 and T2 values depend on the biological tissue and its surroundings. Additionally, the initial signal strength depends on the number of spins that are excited within a voxel, which in turn is mainly determined by the amount of hydrogen in the voxel. The amount of hydrogen that is contained within a certain tissue varies for different tissue types. Thus, it can be determined what the tissue type is on basis of the T1 and T2 values and the relative signal strength.

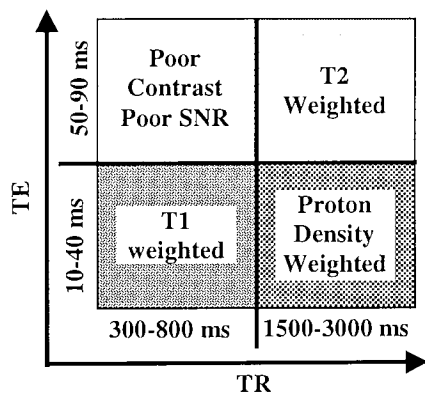


figure 3-2: Diagram that shows how the relative contribution of different tissues to the recorded MR signal is affected by the RF pulse sequence (TE and TR). For example, long TR and TEs highlight differences in T2 between tissues (figure modified from Aine, 1995).

A gradient is applied to the steady magnetic field to determine the origin of the signal that is picked up by the receiver coil. The Larmor frequency is directly proportional to the external magnetic field hence the frequency of the measured MR signal is a function of location. Thus a range of Larmor frequencies is created along one axis. The distinction between different voxels in the other direction is based on phase encoding. For this purpose a gradient in the other direction is applied to the external magnetic field. This gradient is however switched off before the MR response is recorded. The gradient causes frequency differences between the spins

in different voxels. The spins with higher frequency precess faster and get ahead of the spins at lower frequencies. When the gradient is switched off, the spins all precess with the same Larmor frequency again, but the phase difference will remain. The frequency and phase information so encoded in the recorded signals can be used to determine the spatial origin of the signals by means of Fourier imaging.

The RF pulse sequence is determined by two timing parameters, TE and TR. The interval between the excitatory RF pulse and the detection of the response of the tissue to this signal is defined as TE. The time between two RF pulses is known as TR.

The values of TE and TR determine whether the recorded signal is mainly due to T1- or T2-relaxation. As mentioned above, different biological tissue asserts varying T1- and T2-relaxation values. Thus, by varying the values of TE and TR, the contrast between tissues can be manipulated. For example, a T1-weighted image (i.e. TE and TR are chosen to

highlight differences in T1 values) shows a clear distinction between CSF and grey matter (Aine, 1995). It follows from figure 3-2 that an MRI scan with short TE and TR can best be used if one wants to perform segmentation of the cortical surface.

The result of an MRI scan is a dataset that consists of a set of cross section slices. Each slice consists of, typically, 256x256 pixels. The pixel value is the coded response of the tissue to a RF pulse. In figure 3-1 a set of coronal cross sections is displayed as a stack of images. The set of cross section slices form a 3D array. Each element in this 3D volume is then referred to as a voxel. Normally, the in slice resolution is higher than the inter-slice resolution (i.e. pixel dimensions are smaller than the slice thickness). It is however possible to obtain the same resolution in all three directions by combining three orthogonal series of 2D scans (Dale and Sereno, 1993). A volume scan also gives high resolution in all 3 directions, but the scanning protocols are less flexible.

The MRI co-ordinate system is based on the pixels in each slice. Each pixel has a x and y co-ordinate. For the MRI dataset used in this study, the origin is defined as being at the right side, above and in front of the head and with the x-axis running from right to left and the y-axis from top to bottom. The z co-ordinate is the number of the slice, with slice number one at the front of the head. The x-, y- and z-axes are scaled by a scalar in order to obtain the proper physical dimensions.

3.2.2 Segmentation

Different tissue types contain different amounts of hydrogen and consequently their relaxation times vary when measured in a MR scan. The voxels that represent the tissues are usually displayed as 8-bit grey levels. Thus, each tissue is represented by a grey level or intensity value, or in fact by a range of values. Segmentation is the process of finding the intensity values that correspond to the tissues of interest. This is not trivial, due to noise in the scanner and the sheer number of voxels (generally 256x256x128 voxels). Moreover, the differences in intensity levels between various tissue types can be small and variations in intensity level of a tissue type throughout the whole volume can arise due to RF-field inhomogeneities (Wieringa, 1993, Bafra *et al.*, 1994, Dale *et al.*, 1999). Additionally, the

partial volume effect near the boundaries of tissues (i.e. one voxel contains two or more different tissues and the grey value is therefore due to the contribution of all these tissues (Aine, 1995) further troubles the classification (e.g. Bafra *et al*, 1994, Taxt and Lundervold, 1994, Dale *et al*, 1999).

Segmentation has many important applications, like (see also Heinonen *et al.*, 1999):

- 3D visualisation of anatomical structures.
- Co-registration of functional and anatomical information.
- Construction of realistically shaped head models for the forward problem in MEG /EEG.
- Volumetric analysis in order to assess the effect of treatment on tumors, the progression of Multiple Sclerosis and identification of the affected side in temporal lobe epilepsy.

Consequently, a lot of research effort has been put into the development of segmentation algorithms. An overview of different techniques is given in Pal and Pal (1993), Suetens *et al.* (1993) and Clarke *et al* (1995).

The different techniques can be classified as manual, semi-automatic and automatic. Additionally, a classification based on the operation principle is possible, like boundary- or region-based techniques or statistical techniques. Different techniques can be integrated as well.

The simplest segmentation technique, thresholding, is described first, followed by a short description of some of the more complex approaches. A more complete overview of segmentation techniques can be found in the review articles mentioned above.

Thresholding

As described above, in an MRI each tissue is represented by a range of intensity values. The aim is then to choose an intensity range that represents the tissue of interest, for example the grey matter. The segmentation result can then be displayed and examined. If necessary, the range can be adjusted. This task is labour intensive, but easy to perform. A more automatic approach is based on a histogram of the pixel values in an MRI slice. Such a histogram usually contains a few peaks that represent tissues that are well represented in the slice, and a peak due to the background. The range of intensity values that represents




figure 3-3: Histogram of slice 100 from the MRI of the head of subject IEH. The peak on the left side is due to the background. The other two peaks represent mainly the grey and white matter.

the tissue of interest can then be selected from the histogram, with the lower bound based on the beginning of the peak and the upper bound based on the end of the peak. This selection can be done either manually or automatically (e.g. Wieringa, 1993, Ziolkowski and Brauer, 1996, Itti *et al.*, 1997). In figure 3-3 a histogram of a slice in the MRI of a subject's head is shown. The first peak represents the background and the second and third peak represent the grey and white

matter, respectively. The values of other tissues (like the skin) fall within these ranges as well.

The intensity values of a tissue type can vary throughout the volume due to field inhomogeneities in the MR-scanner. This can be accounted for by selecting a separate range of values for each slice (Wieringa, 1993) or by applying intensity normalisation, based on shifting the mean white matter intensity value of each slice to a desired value (Dale *et al.*, 1999). Normalisation is based on the white matter voxels as there is a large number of them and they form large objects (i.e. the influence of the partial volume effect is negligible).

Thresholding works as well as any sophisticated technique if the amplitude features completely characterise the desired tissue (Bafra *et al.*, 1994).

Region Growing

Segmentation based on intensity values alone does not utilise any spatial information. Region growing methods construct regions by grouping spatially proximate voxels that conform to certain constraints (e.g. George *et al.*, 1989, Wieringa, 1993). A number of voxels are chosen in the tissue of interest. Around the boundaries of these seeds, voxels are added that conform to the constraints. These constraints are based on intensity

values and/or morphometric information. The process stops when there are no more voxels left that conform to the constraints.

Segmentation by means of region growing is usually succeeded by morphological operations (Höhne and Hanson, 1992, Wieringa, 1993, Wieringa and Peters, 1993, Psaltikidou *et al.*):

After initial segmentation, based on thresholding and/or region growing, a combination of the following operators is applied.

- Erosion: Objects are reduced in size, holes in objects are enlarged and thin passages between objects are removed by moving a small mask over the image. The centre of the mask is removed from the set if at least one voxel that is not part of the set intersects with the mask.
- Dilation: Enlarges the volume and removes small holes and indentations by moving a small mask over the image. The centre of the mask becomes part of the segmented volume if at least one pixel of the segmented image intersects with the mask.
- Region filling: Internal regions are assigned to the internal volume.
- Boolean masks between 3D regions: Combination of volumes obtained with the above operations can be made following AND, OR, SUB, XOR and NOT operations.

These operations have intuitive effects and in an interactive setting it is possible to guide the segmentation process successfully (Höhne and Hanson, 1992). For example, it is clear that two parts that are incorrectly connected, can be separated by applying an 'erosion'. When this is followed by a 'dilation', the original shape is grossly conserved. Such a sequence is known as an 'opening' operation. The reverse operation ('dilation' followed by 'erosion'), known as 'closing', has a smoothing effect. This operation however is likely to 'fill' small sulci (e.g. Bomans *et al.*, 1990).

This approach fails for datasets with image inhomogeneities or incorrect use of the morphological operators may remove significant structures or connect tissue other than the tissue of interest (Höhne and Hanson, 1992). Particularly, the contours of small sulci are not clearly defined when morphological operators are used (e.g. Wieringa and Peters, 1993).

The fact that the grey matter is a laminar sheet that covers the white matter was utilised by Teo and colleagues (Teo *et al.*, 1997) and later by Dale and co-workers (Dale *et al.*, 1999). Segmentation of grey matter is troublesome because it is a small layer between

the white matter and the CSF, which gives problems due to partial volume effects. Segmentation of white matter, however, is relatively straightforward. It is a large volume and is therefore less influenced by partial volume effects and the variability of intensity of the white matter voxels is relatively small. After segmentation of the white matter, a 3mm thick layer of grey matter is grown onto the boundary of the white matter. This results in anatomically consistent segmentation of the grey matter, if the segmentation of the white matter was performed correctly. The algorithm fails for very thin folds.

Dale and colleagues (Dale *et al*, 1999) used the *a priori* knowledge that the grey matter is a smooth laminar sheet with finite curvature in order to better define the surface of the white matter. This knowledge was utilised to design a filter with a priori known (laminar) shape that was used to enhance the discrimination between white and grey matter.

Edge Detection

The local gradient of intensity values is utilised by edge detection algorithms in order to define edges of regions. The edges are then combined, forming contours on the boundaries between different tissues. An example is the Marr-Hildreth algorithm (Hildreth, 1983). The convolution of a Laplacian of a Gaussian filter, known as a 'Mexican hat', with each slice is taken. A pixel that lies on an edge is then detected by median thresholding. An extension to three dimensions that detects surfaces instead of edges was used by Srebro and Oguz (Srebro and Oguz, 1997). Unfortunately, no details of the segmentation algorithm were given. Bomans and colleagues (Bomans *et al*, 1990) also extended the Marr-Hildreth filter to three dimensions. They found that the contours of the cortex are displaced at corners and curves and tried to correct for this with morphological filters. This resulted in the loss of fine detail though, especially in deep sulci.

A segmentation method based on edge detection does not always succeed. Often there are only small intensity gradients between regions in the MR images. To make matters worse, the local gradients can be reduced due to noise, movement and the partial volume effect, resulting in contours that include spurious edge elements. This approach therefore provides good representations of the major parts of the brain, but not for smaller and more complicated parts as deep sulci (Bomans *et al*, 1990).

Active contours/surfaces

An approach that is the opposite of edge detection is based on active contours. Instead of finding edges based on information within the image, an edge is defined first and then deformed to bring it into alignment with the image (Dale and Sereno, 1993, Dale *et al*, 1999 and references therein). Therefore, not only local information is utilised, but the global shape is also taken into account.

A predefined, deformable, continuous curve is defined around the region of interest. Such an elastic curve is often referred to as a snake or active contour. The energy of this active contour is then minimised so that it converges to the boundary of the region of interest. The energy function consists of two terms, based on the shape of the contour and additional constraints. The first term is based on the MRI data and is made to attract the contour to the brain surface. The difference between the target grey matter value and the intensity values of the voxels on the contour determines whether the contour is attracted (CSF) or repelled (white matter). The second term is a smoothing term that is influenced by the average distance between the contour points and their neighbours (Dale and Sereno, 1993). Note that the computations of the energy terms are based on local information about the MRI data and neighbouring contour points. These computations can therefore be performed by parallel processing. The more global constraint on the shape of the contour is implicitly enforced by the connectivity of all contour points.

Both energy terms can be scaled in order to increase the relative importance of one term over the other during the segmentation. Moreover, when the first term in the energy function is set to zero, the surface is released from the data, resulting in a flattened or unfolded cortex. When functional information (e.g. from MEG or fMRI) is related to the segmented cortical surface, it can be displayed on this flattened surface (Dale and Sereno, 1993). The functional information is then easier to examine, as it is not buried in the sulci, yet the relative topology is still intact.

Opposite sides of a sulcus are spatially close. However, they should not be connected to each other, as this would result in an anatomically inconsistent cortical surface. The advantage of using active contours is that the segmentation of such regions is automatically anatomically correct (Teo *et al*, 1997). In addition, the results obtained with active contour based segmentation are smooth.

Instead of finding an optimum contour in each image sequentially, the method can also be modified to operate on a complete volume at once, using a deformable, or 'shrink-wrapped', surface (Dale and Sereno, 1993, Dale *et al*, 1999). This has two advantages:

1. Contours that lie almost perpendicular to an image are hard to define (Dale and Sereno, 1993). In the 3D case this is no problem.
2. Reconstructing a surface from 2D contours is not trivial because of the branching problem (different slices containing a different number of contours). This is not an issue when a deformable surface is used for the segmentation. Additionally, the 'wrapped surface' does not contain holes or intersections.

Problems with the active contour approach are related to the minimisation process, i.e. initialisation and the problem of local minima. Local minima are usually found near deep sulci with a small entrance (Dale *et al*, 1999). These structures are therefore not always correctly segmented, unless additional terms are used in the energy function that 'push' the surface into the sulci (Le Goualher *et al.*, 1999). Additionally, the method uses several parameters, like the scaling of the energy terms and the scaling of the strength of the attracting and repelling 'forces'. The optimum values of these parameters need to be found empirically.

Statistical approaches

In statistical approaches, the grey level intensity of each pixel in an image is treated as a random variable. The values are taken from a set of possible values, e.g. 256 values in case of an 8-bit image. Define the true pixel labelling (segmentation) of an image as \mathbf{I}_{true} . The objective now is to find the optimal estimate of \mathbf{I}_{true} , given the image to be segmented. Using Bayesian formulation, the algorithm attempts to estimate the maximum a posteriori probability of the true segmentation, given the original image and any prior information. The *a priori* information is given in the form of a Markov or Gibbs Random Field. Such a field gives a value for each pixel based on its neighbourhood in the original image. This imposes connectivity and smoothness constraints onto the set of random variables (e.g. Bafra *et al*, 1994, Liang *et al.*, 1994). Additionally, the segmentation results of previous slices can be used as *a priori* information in subsequent slices (Bafra *et al*, 1994).

Teo and colleagues (Teo *et al*, 1997) used the likelihood of a particular voxel belonging to a certain class (i.e. grey-, white matter or CSF) as *a priori* information. This likelihood was based on the intensity value of the voxel and the mean and standard deviation of the intensity values in each class. They used this *a priori* information with a Bayesian formulation and a smoothing operation working on the posteriori probabilities. This approach resulted in accurate segmentation of white matter.

A drawback of the Bayesian approach is that the number of classes and the class parameters have to be supplied to the algorithm.

Another statistical approach is based on multispectral analysis. In this method images are acquired with different protocols, e.g. T1-, T2 and proton density -weighted, and combined. These combined images contain more information about each tissue type than images obtained with only one protocol (Liang *et al*, 1994). This information can then be utilised to obtain better segmentation (Liang *et al*, 1994, Taxt and Lundervold, 1994). Taxt and Lundervold used a Bayesian framework, with the posterior probability based on the tissue values obtained with all scanning protocols and assuming that the tissue regions can be modelled by a Markov random Field (Taxt and Lundervold, 1994). Liang and colleagues' approach was similar (Liang *et al*, 1994). They proceeded by assuming that the intensity value at each location is a multivariate likelihood function of three weighted signal intensity values. The tissue types are classified by maximum likelihood fitting of the image data. Using information criteria the optimum number of classes is then determined. Using the number of classes and the class parameters obtained with the maximum likelihood fitting, segmentation is again performed using a Bayesian framework with a Markov Random Field prior.

Shareef *et al.* (1999) developed a method based on a biologically inspired oscillatory network. It is thought that the brain groups and segregates visual features on the basis of correlation between neural oscillations (Singer and Gray, 1995). This so-called 'gamma-binding' was demonstrated by Singer *et al.* (Gray *et al.*, 1989, Singer and Gray, 1995). Shareef and co-workers used an oscillatory network in order to define regions with similar properties in medical images. After a few oscillations, oscillators corresponding to a major image region (for example grey matter) will oscillate in synchrony, while any two oscillator blocks corresponding to two different major regions will desynchronise from each other. Oscillators whose corresponding pixels do not belong to a major region will stop

oscillating. Their results are comparable with manual segmented images, if appropriate values for the parameters of the algorithm are chosen. It can deal with small inhomogeneities of the tissue values.

Favoured Segmentation Approach

Most semi-automatic and automatic segmentation methods are sensitive to noise in the image and inhomogeneities of the tissue values over the volume. Additionally, they often fail to segment narrow, deep sulci (Teo *et al*, 1997) and errors in the segmentation remain undetected. For visualisation purposes these methods often suffice, but not for 3D mathematical modelling of anatomical structures.

The results of manual segmentation depend on the knowledge and interpretation of the image data of the operator. This is to a certain extent true for segmentation by interactively applied morphological operators as well. The main drawback of manual applied segmentation however is that it is very time-consuming and tedious. For the purpose of this work this is not a limitation though, as the segmentation of a cortex that is used in a set of computer simulations only has to be done once. Additionally, in order for the results of the simulations to be realistic, an accurate description of the cortical surface was needed, including the sulci. These regions are important to us as two thirds of the total cerebral cortex is buried within the cortical folds (Van Essen and Drury, 1997, Le Goualher *et al*, 1999). Thus, these areas need to be taken into account in order for the results of the simulations to be applicable to real data.

The results of manual segmentation are the most accurate (Shareef *et al*, 1999). Therefore, a combination of amplitude thresholding, contour selection and drawing is chosen as the approach to segment the cortex for construction of a realistically shaped source space. Successful applications of similar approaches are described in the literature. Heinonen and colleagues (Heinonen *et al*, 1999) used thresholding and 2D region growing in combination with free hand drawing in 2D images. Gevins and colleagues (Gevins *et al.*, 1991, Gevins *et al*, 1994) also used 2D threshold based contouring, with manual editing, in order to segment the brain, CSF, skull and scalp. They used the segmented regions to construct finite element models. These models were then utilised for deblurring of the EEG.




figure 3-4: Example of the right side of a coronal slice (left hemisphere) with parts of the CSF marked (yellow). The sagittal and axial slices corresponding to the crosshairs are also shown. The computed contours (red) need modification in order to mark the outer boundary of the grey matter faithfully. The contours marking the inner and outer skull boundaries can easily be discarded.

First, as a visual guide for the selection of the grey matter, parts of the CSF were selected, especially within sulci. Three orthogonal views, containing a coronal, sagittal and axial slice respectively, were displayed (figure 3-4). The CSF voxels could be selected in any of the three images with the cursor. Selected CSF matter was displayed in yellow. It was further possible to flick through the images, in order to give a notion of the 3D structure to the operator (Shareef *et al*, 1999). An example is given in figure 3-4, where most of the CSF in the sulci is selected. As mentioned before, contours that lie almost perpendicular to an image are hard to define (Dale and Sereno, 1993). Using the selected CSF as a guide, defining the grey matter contours in such slices, and in general, was less ambiguous.

Selection of the grey matter was subsequently performed in the coronal slices, using a standard contouring function in MATLAB (The Mathworks). This function uses amplitude thresholding to define contours around objects that consist of pixels with the chosen intensity value. The target intensity value was chosen by trial and error, so that the contour surrounded most of the grey matter. Due to noise and because skin and other tissues have values in the grey matter range, some contours were wrongly placed around tissue other than the grey matter. Noise in

the MRI mainly caused small additional contours. Also, the main contour around the grey matter was interrupted or surrounded by contours wrongly placed around regions containing other tissue types, as is demonstrated in figure 3-4. The main contours were simply selected with the mouse cursor and the faulty contours were discarded. Furthermore, software tools were developed for editing the contours. The mouse cursor was used to



figure 3-5: A manually edited contour where wrongly placed contours were removed and the different parts of other contours were correctly connected. This contour faithfully represents the outer boundary of the grey matter.

select points on the contour for removal, select new points for incorporation in the contour and connecting multiple parts of a contour. It was possible to zoom in and out of an image enabling more comfortable editing of the contours. The final result was a vector-contour through all contour points. This was repeated for each grey matter object in every slice. In figure 3-5 the final result for the left hemisphere in a coronal slice is displayed. This approach was used to select the boundaries of the grey matter in 31 coronal slices of the left occipital lobe of subject IEH.

The utilisation of the contours defined by the MATLAB contouring routine speeded up the segmentation significantly, since large parts of the resulting contours were already correct. Adding and removing parts of contours was a relatively quick and easy process. The inclusion of a 'drawing pen' in the software would further speed up the definition of additional contours, as it is easier than defining a contour by placing points that need to be connected. The drawback of the 'drawing pen' is that a contour will consist of a lot more points.

3.2.3 Surface Reconstruction

Once the contours of the region of interest are known as the result of the segmentation process, a reconstruction of the surface is needed. This is done by connecting contours in adjacent slices with patches, e.g. triangular (Geiger, 1993) or rectangular patches (Odesanya *et al.*, 1993). A straightforward connection of the contours is possible if the number of contours in each slice is constant and when the shape of each contour is approximately convex (i.e. sulci are not modelled, so that the contour of the brain is approximately convex). Triangulation is now simply a matter of connecting the points in

adjacent slices in such a way that obtuse triangles (that is, triangles with one angle larger than 90 degrees) are avoided (see Wieringa (1993) for details). The surface is closed by connecting the triangles in the bottom and top slice. However, keeping the number of contours constant throughout the volume is only possible by artificially connecting grey matter regions that are not connected or by leaving out some isolated grey matter regions. This results in an inaccurate representation of the cortical surface in some regions. For example, near the tip of the occipital lobe, in case coronal slices are used (see figure 3-6), the grey matter boundary is described by two or three contours. Combining these into one contour would result in an inaccurate description of the surface of the cortical lobe.

The surfaces created in this way are accurate enough for the construction of Boundary Element Meshes (Wieringa, 1993, Zanow and Peters, 1995, Ziolkowski and Brauer, 1996), but not for the construction of realistically shaped source spaces.

The problem of connecting a different number of contours in adjacent slices is known as the branching problem.

Not all segmentation algorithms have a set of 2D contours as output. In such cases the boundary of each region has to be determined and a contour must be constructed, based on the boundary points (e.g. Odesanya *et al.*, 1993). Alternatively, a model can be constructed based on the 3D region directly (Lorensen and Cline, 1987, Treece *et al.*, 1998).

Odesanya and colleagues (Odesanya *et al.*, 1993) fitted a smooth spline through the boundary points. This results in data reduction, since only a relatively small number of junction points is needed instead of all boundary points. After subdividing each contour in an equal number of points, pairs of points are connected on contours in adjacent slices. This generates quadrilateral facets between the contours. Alternatively, the cross-section spline curves can be combined to form a spline description of the surface. The cubic B-spline they used is accurate enough to describe the contours of the bones of the hand. However, it would not represent highly irregular shaped contours accurately, like the contour near sulci. A solution for this would be to use a




figure 3-6: Three contours in a coronal slice through the pole of the occipital cortex. Connecting these three contours to a single contour in an adjacent slice poses a problem for some surface reconstruction algorithms.

higher order spline.

The Marching Cubes algorithm (Lorenson and Cline, 1987) is a well-known algorithm that constructs a polyhedral model from a 3D region. The algorithm takes its name from the fact that it marches a cube through the 3D region and produces a set of triangles for each cube. Again, only the voxels on the boundary of each region are used. The basic principle is to reduce the surface reconstruction problem to that of triangulating a single cube, which is intersected by the surface. The whole surface is triangulated by 'marching' this cube through the data and the sum total of all triangles forms an approximation of the surface. The dimensions of the cube determines the resolution with which the surface is reconstructed. The way the cube is triangulated depends on the intersection with the surface, more specifically, it depends on which corner points lie inside or outside the surface. A surface intersects those cube edges where one corner point is outside the surface and the other is inside the surface. It can be shown that, because of symmetry, there are only 15 different ways to triangulate the cube. These cases can be stored in a look-up table, so that fast execution is possible once the case for each cube is known.

The main drawbacks of the Marching Cubes approach are (Treece *et al*, 1998):

- There are many triangles with poor aspect ratio.
- It does not guarantee the surface to be topologically consistent with the data.

One of the approaches to circumvent these problems is to use marching tetrahedra instead of cubes (Treece *et al*, 1998). One advantage is that tetrahedra can be triangulated in only 16 possible ways, which reduces to 3 by symmetry. Thus, only three cases have to be stored in the look-up table. The reconstructed surfaces contain more triangles than the Marching Cubes approach though. However, a surface simplification can be used to reduce the number of triangles and increase the triangle condition. Therefore, the approach results in topologically consistent surfaces with less triangles than obtained with the Marching Cube algorithm (Treece *et al*, 1998).

The segmentation process used in this work results in a set of vector-contours in each slice. It is therefore intelligible to use an approach that connects the contours in adjacent slices. This is not without problems though, since the contours are obtained for a 3D object that is intersected by a number of parallel planes. For concave objects, this can




figure 3-7: Example of a global connection problem. A set of 6 contours in two parallel slices (a). There are several ways to connect the regions in adjacent slices ((b) and (c)) and form a reconstructed object. Intersecting both these reconstructed objects would return the original regions in the cross-section (from Geiger, 1993).

result in several regions in one plane. Unfortunately, there is no unique way to connect these regions, as is illustrated in figure 3-7. Intersecting the reconstructed objects in figure 3-7b and c would both return the original regions in figure 3-7a. Moreover, even when there is only one region per slice there can be situations where the connections are ambiguous (see figure 3-8). Therefore, all automatic methods need to use some criteria to decide on the connections that need to be made, like connecting regions of similar curvature, minimising the surface area, minimising the volume, minimising angles or minimising the edge length. The NUAGES algorithm proposed by Geiger (Geiger, 1996) is based on Delaunay triangulation and therefore based on geometric closeness (see Geiger, 1993, Boissonnat, 1988 or Bowyer (1981) for a description of the principles of Delaunay triangulation). The method starts with the 2D triangulation of the contour vertices, thereby assuring contour containment (re-sectioning the reconstructed surface returns the original contours), removing obtuse angles and adding vertices that are needed for the connections in adjacent slices. Then, overlapping contour parts are connected vertically and non-overlapping parts are connected to their nearest neighbours to form the surface. Addition of extra vertices is done in such a manner that the method can deal with complex contours with multiple branches, the birth and death of holes and complicated splitting lines.

The accuracy of the NUAGES algorithm (Geiger, 1996) was tested with the reconstruction of the surface of a torus and compared with the Marching Cube algorithm (Geiger, 1993).

Pixel, vector, interpolated and hand-drawn contours were obtained for a torus and used as input for the NUAGES algorithm. The results obtained for the torus seem applicable for the reconstruction of the brain, as the brain has no sharp edges either and is relatively smooth as well.

The most important results were that:

- The contour-based reconstructions were more accurate with respect to volume, surface area and surface normals than the reconstructions obtained with the Marching Cube algorithm. Only the maximum surface distance was better with the Marching Cube method.
- The number of triangles obtained with the Marching Cube method was significantly higher than with the contour-based approach, as expected.
- It was found that the surface normals could not be accurately reconstructed for very small triangles.
- The maximum surface distance was smaller for vector and interpolated contours than for hand-drawn contours. The surface normals were more accurate for the hand-drawn contours though, followed by the vector contours.




figure 3-8: Example of a local connection problem. The two contours can be connected in different ways. No discrimination can be made between the choice in the top panel and the one in the bottom panel, unless additional constraints are used (from Geiger, 1993).

It seems that vector contours provide both accurate reconstruction with respect to surface location and orientation. One observation that is relevant to this work is that the largest reconstruction errors are found where the surface is almost parallel to the cross section.

For the construction of the model of the cortical surface used in this work the vector-contours obtained during the segmentation were used as input of the NUAGES algorithm. Internal vertices were added by the algorithm in order to get a more accurate reconstruction of the cortical surface (Geiger, 1993). Surface normals were computed for each triangle vertex, with the normals defined as the weighted sum of the normals of all triangles sharing the vertex v :

$$\mathbf{n}_v = \frac{\mathbf{n}_v'}{\|\mathbf{n}_v'\|} \text{ and } \mathbf{n}_v' = \sum_{i=1}^k \alpha_i \mathbf{n}_i, \quad (3-1)$$

with \mathbf{n}_i the normalised normal of triangle i and α_i the angle between the two triangle edges joining in the vertex v .




figure 3-9: The triangulated cortical surface with the normals in the vertices displayed as arrows pointing in the direction of the normal in each vertex. The left occipital pole is here shown as viewed from the (left) back of the head.

With the original contours as input, the NUAGES algorithm constructed a set of 30702 triangles and 14706 vertices and normals. The triangulated surface with the normals in the vertices displayed as arrows pointing in the direction of the normals is shown in figure 3-9. This reconstructed surface can now be used to construct a realistically shaped source space, as is shown in the next section.

3.2.4 Creating and Displaying the Source Space

The result of the surface reconstruction is a set of points and a connectivity table. Each surface triangle is formed by connecting three points according to the connectivity table. Standard routines in MATLAB (The Mathworks) are used to display all triangles. Phong lighting, using an artificial light source, improves the apparent smoothness of the




figure 3-10: Two active brain areas, displayed in two different ways. The first way is to present the active area as a current dipole, represented by an arrow. The location and orientation of the arrow show the location and orientation of the current dipole. The strength of the dipole is represented by the size of the arrow. The second way to represent the activity is by colour coding each triangle face by the strength of the activity in the vertices of each triangle. In this case strong activity is represented by yellow, and weak activity by dark red. The elements with negligible activity are displayed in pink.

surface. This algorithm interpolates the vertex normals across each triangle face and calculates the reflectance at each pixel in the display.

The triangulated surface can also be utilised to create a flattened representation of the cortex (Van Essen and Drury, 1997). Van Essen and Drury obtained the contours of the cortex with a mouse driven tracing program. A surface reconstruction was then performed with NUAGES. After this the curvature of the surface was computed and the triangles were colour coded for the curvature of that surface area. This surface presentation was then flattened. The information

about the curvature 'stored' as the triangle colours, provided guidance in the flattened cortex.

A source space is created by placing an equivalent current dipole at the vertices of the triangles, with the normal in each vertex as the orientation of each dipole¹ (see section 2.1 for the reasoning behind this orientation constraint). A dipole can then be represented as an arrow, with the length of the arrow expressing the strength of the equivalent current dipole. An alternative way to display the activity in the source space is by colour coding the strength of each source. The colour of the face of each triangle then represents the mean strength of the sources in its vertices. An example is given in figure 3-10 where both methods of displaying activity in the source space are demonstrated. There are two active brain areas, represented both by arrows and colour coding.

¹ The orientation of each dipole can also be unconstrained. In this case a dipole is placed in each vertex with a component in x-, y- and z-directions. In this work the orientation constraint will always be used.

Finally, the information obtained from the segmentation routine can be utilised to combine a surface based representation with images of the interior in the form of cut-away views. Functional information that is located within the volume can then also be visualised. This is achieved by displaying the anatomical data so that the, otherwise invisible, interior is visible at the sliced-away part of the outer surface. The location of brain activity is displayed in the slice that cuts through the MRI volume. The part of the volume that obscures the MRI slice is 'cut-away'. The location of the cortical surface with respect to the outer surface of the head is also shown. Presenting an active brain area in this way shows its location with respect to (sulcal) anatomy and gives a better idea of its 3D location than when the activity is overlaid on an MRI slice alone. An example is given in figure 3-11.




figure 3-11: An active brain area, represented by the red current dipole, is displayed on the cortical surface. The source is visible because a part of the outer surface of the head is 'cut-away'. A presentation like this shows the location of an active brain areas with respect to (sulcal) anatomy and its location with respect to the outer surface of the head.

3.3 Methods of MEG/MRI Co-registration

3.3.1 Introduction

The co-registration of anatomical scans made in two different modalities is a well-examined problem. The datasets represent the same object and co-registration can therefore be volume-based (i.e. based on intensity values). A statistical comparison of the intensity values in, for example, two MRI scans of the same subject, provides the information needed to perform co-registration. The gold standard for co-registration of anatomical scans made in different modalities is based on markers fixed to the bones of a patient. The markers are chosen so that they show clearly in both modalities. This is the most accurate co-registration technique (Fitzpatrick *et al.*, 1998). However, fixing markers to the bones can only be done with patients undergoing surgery.

The techniques available for co-registration of functional data from MEG and anatomical data from MRI are either based on co-registration of landmarks or surface matching. Variations of these approaches are described in the literature. Several implementations of the two techniques are described in the next sections. Then, the accuracy of a co-registering procedure based on surface matching is tested on its accuracy.

3.3.2 Using Landmarks

Method

The methods most often used for co-registration of the functional data from MEG/EEG and anatomical data (usually from MRI) are based on fiducial points (e.g. George *et al.*, 1989, Wieringa, 1993). The markers that are used to define these points are called the fiducial markers. When the locations of these markers is known in both modalities, a transformation can be found that aligns these markers in the two modalities. It can be assumed that the skull is a rigid object² and therefore a rigid-body transformation is

² When care is taken not to deform the scalp tissue during the co-registration, the scalp can be treated as rigid as well.

possible. Thus, a simple translation and rotation is enough to perform the co-registration. Three additional parameters can optionally be used to take scaling factors into account.

Fiducial points are usually defined on anatomical landmarks, like the nasion and the pre-auricular points. The locations of these points can be defined by a 3D digitiser (Polhemus Isotrak). Such a digitiser consists of an emitter, sending out magnetic signals and a receiver. The pen-shaped receiver is used to mark the fiducial points. These points then define a co-ordinate system, which is referred to as the MEG co-ordinate system. These same anatomical landmarks are also defined in the MRI. This can be done by an expert, either by using the MRI slices or by using special software that displays the surface of the skull and a 3D cursor (Fuchs *et al.*, 1995). Prior to the MRI scan, the location of the fiducial points can be marked with oil filled capsules. These oil filled markers are clearly visible in an MRI scan.

With the fiducial markers defined in the MEG and MRI co-ordinate systems, the transformation matrix can be computed. This matrix, $\boldsymbol{\tau}$, follows from:

$$\mathbf{r}_{i,\text{MRI}} = \boldsymbol{\tau} \mathbf{r}_{i,\text{MEG}}, \quad \text{for } i = 1 \dots n, \quad (3-2)$$

with $\mathbf{r}_{i,\text{MRI}}$ and $\mathbf{r}_{i,\text{MEG}}$ the location of fiducial marker i in the MRI co-ordinate system and MEG co-ordinate system, respectively. The number of fiducial markers is n .

As mentioned above, this matrix can be split in a translation and a rotation part, i.e.:

$$\mathbf{r}_{i,\text{MRI}} = \mathbf{R} \mathbf{r}_{i,\text{MEG}} + \mathbf{t}, \quad (3-3)$$

with \mathbf{t} and \mathbf{R} the translation and rotation matrix respectively.

The values for the translation and rotation matrices can be found by a least squares fitting procedure that minimised the Euclidean distance between the two sets of data points (e.g. Singh *et al.*, 1996). Additionally, the contribution of each point in the datasets to the fitting procedure can be weighted. For example, the weight for a marker in a part of the MRI that shows distortions can be lowered. Fuchs and colleagues used a weight for each fiducial marker based on the measured localisation errors of the markers (Fuchs *et al.*, 1995).

Alternatively, it can be shown that rotation and translation are uncorrelated (Fitzpatrick *et al*, 1998). Therefore, one can first determine the optimum translation matrix and then the optimum rotation matrix (or vice versa). First, define the following matrices:

$$\mathbf{C}_{\text{MEG}} = \mathbf{r}_{\text{MEG}} - \mathbf{g}_{\text{MEG}} \quad \text{and} \quad \mathbf{C}_{\text{MRI}} = \mathbf{r}_{\text{MRI}} - \mathbf{g}_{\text{MRI}}, \quad (3-4)$$

with \mathbf{g}_{MEG} and \mathbf{g}_{MRI} the centre of gravity of the fiducial points in the MEG and MRI co-ordinate system respectively.

Additionally, define the matrix \mathbf{P} as the product of \mathbf{C}_{MEG} and the transpose of \mathbf{C}_{MRI} :

$$\mathbf{P} = \mathbf{C}_{\text{MEG}} \mathbf{C}_{\text{MRI}}^T, \quad (3-5)$$

The desired least squares fit follows then from the Singular Value Decomposition of \mathbf{P} (see section 2.2.3.1)

$$\mathbf{P} = \mathbf{U} \mathbf{\Lambda} \mathbf{V}^T, \quad (3-6)$$

The rotation matrix is then given by (Arun *et al.*, 1987, Fuchs *et al*, 1995, Fitzpatrick *et al*, 1998 and references therein):

$$\mathbf{R} = \mathbf{V} \mathbf{U}^T, \quad (3-7)$$

with \mathbf{U} and \mathbf{V} containing the left and right eigenvectors respectively.

The best-fit translation is then obtained from (Fuchs *et al*, 1995, Fitzpatrick *et al*, 1998):

$$\mathbf{t} = \mathbf{C}_{\text{MRI}} - \mathbf{R} \mathbf{C}_{\text{MEG}}, \quad (3-8)$$

Note that the translation and rotation matrices can therefore be determined without making use of an optimisation algorithm.

In practice, the accuracy of a co-registration procedure must be established. One measure of localisation accuracy is the root mean squared Euclidian distance (RMS) between the fiducial markers in the MRI co-ordinate system and the transformed fiducial

markers. This is called the fiducial registration error (FRE). However, this measure is almost independent of the fiducial configuration. Therefore, poor registration caused by a poor fiducial configuration may be unnoticed due to a small value for FRE (Fitzpatrick *et al*, 1998).

The reason for co-registration is to align certain targets (like the scalp or a tumor in the brain). The error measure of interest is therefore the target registration error (TRE). This is the distance between target points in the MRI-co-ordinate system and transformed MEG target points, other than the fiducial points.

Fitzpatrick and colleagues (Fitzpatrick *et al*, 1998) derived the following expression for the relationship between the fiducial registration error and the target registration error at location i :

$$\langle \text{FRE}_i^2 \rangle = \langle \text{FLE}^2 \rangle - \langle \text{TRE}_i^2 \rangle, \quad (3-9)$$

with $\langle \rangle$ denoting the expectation value. FLE is the fiducial localisation error, defined as the distance between the localised location of the fiducial marker and the actual fiducial position.

An interesting and counter-intuitive observation is that a small FRE is indicative of a large TRE. So, near fiducial points that are in close alignment the TRE is largest. A small value for FRE might therefore be a misleading indicator for the accuracy of the co-registration (Fitzpatrick *et al*, 1998).

Fitzpatrick and co-workers (Fitzpatrick *et al*, 1998, Maurer Jr. *et al.*, 1998) derived an expression for TRE that gives insight into some of the issues involved in co-registration:

$$\langle \text{TRE}^2(\mathbf{r}) \rangle \approx \frac{\langle \text{FLE}^2 \rangle}{n} \left(1 + \frac{1}{3} \sum_{k=1}^3 \frac{d_k^2}{f_k^2} \right), \quad (3-10)$$

d_k is the distance of target point \mathbf{r} from the k^{th} principal axis of the fiducial point sets, and f_k is the average distance of the fiducials from the k^{th} axis.

It follows from the above equation that:

- The target registration error increases if the target point is further away from the principal axis, with a squared distance dependency. Singh and colleagues indeed found

a dependency on distance in a Monte Carlo simulation study of MEG/MRI co-registration, although the largest errors were not always found at the most distant locations (Singh *et al.*, 1997).

- The error decreases with a factor proportional to $\sqrt{1/n}$ when the number of fiducials is increased. The graph of the ratio of $\sqrt{1/(n+1)}$ and $\sqrt{1/n}$ in figure 3-12 demonstrates that using more than 5 fiducial points has a relatively small effect on the reconstruction error. A small value in the plot indicates that a lot is gained by adding extra fiducial points and a value of 1 indicates that the error does not decrease due to a change in the number of fiducials (although the addition of a fiducial does change FLE, d_k and f_k and therefore influences the error).
- The error decreases if f_k increases. The fiducial points should therefore be chosen far apart from one another.
- The target registration error is proportional to the localisation error of the fiducial markers.

Because of this last point, it is important that the location of these markers is determined accurately. As described above, normally oil filled markers that light up in the MRI are placed on anatomical landmarks. In Aston a modification of this approach is used (Singh *et al.*, 1996, Singh *et al.*, 1997). A bite-bar system (see figure 3-13) is used, consisting of a rigid plastic holder and dental thermoplastic. Four holes are drilled in the plastic holder.

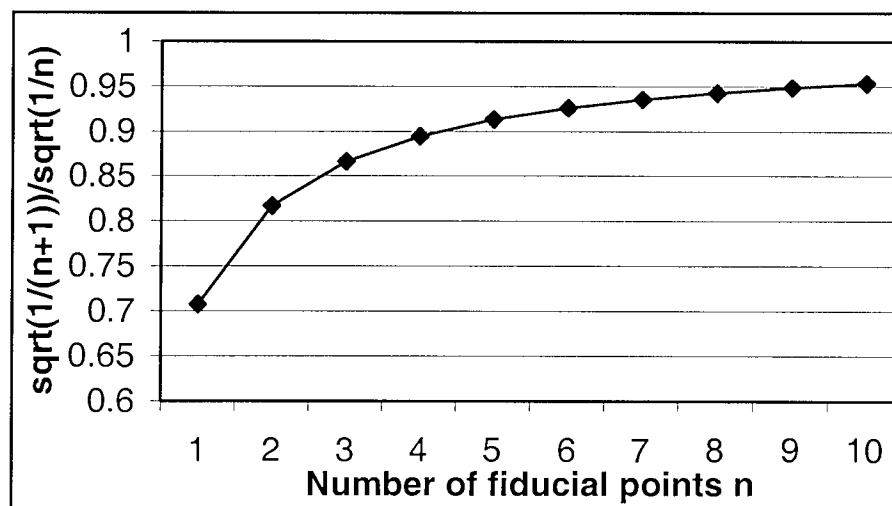


figure 3-12: Plot of the ratio of $\sqrt{1/(n+1)}$ and $\sqrt{1/n}$, with n the number of fiducial points. The minimum number of fiducial points needed for target registration is 3. The graph shows that there is a large decrease of the registration error when the number of fiducial points is increased from 3 to 4 or 5. A further increase of the number of points reduces the registration error with a relatively small amount, as is illustrated by the flattening of the curve.




figure 3-13: Schematic of the bite-bar system used in Aston. It contains four small holes that are digitised in the MEG co-ordinate system. The markers can also be identified in the MRI when filled with oil, enabling the co-registration of the MEG and MRI co-ordinate systems.

The locations of these holes are obtained with a 3D digitiser (Polhemus Isotrak) before and after the MEG measurements are done. The average of these two sets is taken as the location of the fiducial points in MEG co-ordinates. During the MRI scan of the subject, the holes are filled with oil for easy identification in the MRI. Additionally, large oil filled capsules are placed directly above the holes for easy identification of the

smaller holes in the MRI. With the location of fiducial points available in both co-ordinate systems, the co-registration can be performed as described above.

Because more fiducial points are used and the head is stabilised during the co-registration, the co-registration is more accurate than using oil filled markers on anatomical landmarks (Singh *et al.*, 1996, Singh *et al.*, 1997). Also, because the bite-bar is fixed on a rigid arm there is less movement during an MEG recording. Additionally, since no anatomical landmarks are used, there is no problem with identification of the markers. The problem with fiducial points based on anatomical landmarks is that the locations of anatomical landmarks defined by different experimenters or from repeated determination can vary considerably (Huppertz *et al.*, 1998).

Instead of a bite-bar a non-metal helmet (Greitz *et al.*, 1994) or stereotactic headgear (Simpson *et al.*, 1995, Srebro and Oguz, 1997) can be used as well, having the same advantages as a bite-bar.

Van den Elsen and Viergever used triangular markers, with two plastic tubes filled with oil attached to two sides. These tubes light up in several slices in the MRI. The location of the reference point can be taken from the intersection of the two lines through the marker points in the MRI slices, taking into account an offset (van den Elsen and Viergever, 1991). The main advantage of this approach is that there are more points that define one reference location, reducing the probability that a reference location can not be found and reducing the influence of identification errors. A similar approach has been described by Williamson and colleagues (Williamson *et al.*, 1991).

Accuracy

The repeated localisation of a dot on the skin with a 3D Polhemus digitiser has a mean precision of about 1.5mm (Schwartz *et al*, 1996). Brinkmann and colleagues however reported only sub-millimetre variations for repeated inter- and intra-observer marking of 21 marked locations on the scalp. The error for the localisation of an anatomical landmark is at least this high, due to the additional problem of identification in both the MEG and MRI co-ordinate systems. Defining the location of landmarks in the MRI with the help of oil filled capsules also introduces errors, due to the size of the capsules itself.

An issued closely related to the co-registration procedure is how to determine the position of the head with respect to the MEG sensors. When a 3D Polhemus digitiser is used to obtain the location of the fiducial points, the sender is mounted to a pre-defined location on the dewar. The fiducial points that define the MEG co-ordinate system and the location of the sender is therefore known in MEG co-ordinates. The locations of the sensors are then also known in MEG co-ordinates, since the location (and orientation) of the sensors with respect to sender location is known from a calibration procedure.

Alternatively, coils are placed on the subject's head (e.g. Ahlfors and Ilmoniemi, 1989, Fuchs *et al*, 1995). The current fed through these coils produces a magnetic field that is measured by the MEG system. The location of these coils can then be accurately found by a simple least squares optimisation algorithm (i.e. Powell, 1964), as the signal-to-noise ratio of the measured magnetic fields is high and the number of sources (i.e. coils) is known.

The registration errors were reported to be 3mm with a 7 channel system (Ahlfors and Ilmoniemi, 1989). The mean localisation error of three head coils with the CTF 151 channel MEG system is smaller than 2.5mm (CTF Systems Inc.). This value of 2.5mm was found as the mean error for the repeated localisation of a magnetic dipole in a phantom and therefore represents the sum total of all possible errors, including noise in the measured magnetic fields and the actual location of the dipole in the phantom. The errors for the localisation of the head coils are therefore less than 2.5mm.

The head coils can further be used to monitor movement of the subject during the MEG recording, by localising the coils before and after the recordings. In the future, if hardware improvements allow for continuous localisation of the reference coils, it should be possible to correct the recorded data for small movements (CTF Systems Inc.). The data

recorded at the location of a sensor at each instant in time is then interpolated to give a value for the average location of the sensor.

The main drawback of co-registration by fiducial markers is that the identification and accurate localisation of these markers in both modalities is not trivial. Consequently, it happens that surfaces in two modalities are not aligned, even though the fiducial markers are (Bamidis and Ioannides, 1996). An approach that avoids this problem, surface matching, is described in the next section.

3.3.3 Surface Matching

Surface matching was first introduced by Pelizzari and colleagues as a technique to co-register CT, PET and/or MRI (Pelizzari *et al.*, 1989). Several groups have thereafter used the technique for EEG/MEG and MRI co-registration (Wang *et al.*, 1994, Bamidis and Ioannides, 1996, Schwartz *et al.*, 1996, Brinkmann *et al.*, 1998, Huppertz *et al.*, 1998). In an excellent paper by Schwartz and co-workers the accuracy of the method in MEG/EEG and MRI co-registration was examined in great detail (Schwartz *et al.*, 1996).

Co-registration is performed by aligning two surfaces defined in the MRI and MEG co-ordinate system, i.e. a cost-function defining the distance between the two surfaces is minimised. This can be compared with placing a tight fitting 'hat' (i.e. the surface in MEG co-ordinates) onto a 'head' (i.e. the surface in MRI co-ordinates).

Method

Head points

The surface of the skin in MRI co-ordinates is defined with the same techniques as described in section 3.2.2 and 3.2.3 for reconstruction of the brain surface. The definition of the scalp is less prone to errors though, because there is a larger gradient in intensity between the scalp and the background than there is between the cortical surface and the surrounding tissues. Therefore, a simple thresholding technique can be applied to segment the scalp with good results. Small errors due to noise in the MRI scan or interior holes can

be removed by applying dilation, erosion and connection operations (e.g. Schwartz *et al*, 1996) or alternatively by user-editing (Wieringa, 1993). The outer boundary of the scalp can be easily defined by a test on the neighbourhood of each scalp voxel. Each voxel with at least a single neighbour that is part of the background is defined as part of the skin surface (Schwartz *et al*, 1996). A description of the surface can then be made using the surface reconstruction techniques described in section 3.2.3.

Hat points

The surface of the skin in MEG co-ordinates can be defined by a 3D-digitisation process, involving electromagnetic/optical probes, ultrasound (Huppertz *et al*, 1998, Ball *et al.*, 1999) or video cameras (Konyshov *et al.*, 1989). Most often, a 3D electromagnetic digitiser is used to define a set of evenly distributed points on the subject's scalp. The surface that is defined in this way should be as similar to the surface defined from the MRI as possible. Therefore, the digitising pen should not be pressed too hard against the skin to avoid deformation, especially around the face of the subject. For a large part the scalp area is nearly spherical, resulting in a cost-function with flat minima. Hence, points on the nose and around the eyes should be included in the surface in order to get better-defined minima in the cost-function (Wang *et al*, 1994, Huppertz *et al*, 1998).

Cost-function

In each step of the surface matching algorithm, a co-registration matrix is computed. The accuracy of the resulting co-registration should be evaluated. The final co-registration should provide a fit between the 'head' and the 'hat' that minimises the distance between both surfaces. When the MRI surface is densely sampled, the Euclidean distance between the 'hat' points and the nearest 'head' point can be computed for each 'hat' point. The total or mean distance, or its square, can be used as the cost-function. Using the squared distance increases the sensitivity for small displacements (Schwartz *et al*, 1996). Alternatively, the distance between a 'hat' point and the projection of the 'hat' point onto the 'head' surface (i.e. the nearest triangle for triangulated surfaces) can be used (Pelizzari *et al*, 1989).

Fast computation

The evaluation of the cost-function is done in each step of the minimisation algorithm and should therefore be done as fast as possible. The use of a look-up table in which distances to the 'head' surface are stored greatly reduces the computing time, as the table only has to be computed once. Such a table can be computed by the Distance Transform (Borgefors, 1986, Schwartz *et al*, 1996, Kozinska *et al.*, 1997, Huppertz *et al*, 1998). The Distance Transform estimates the distance between voxels that surround the scalp surface in the MRI volume and the scalp surface itself. This is done by passing a small mask (typically 3x3x3, 5x5x5 or 7x7x7) through the volume. During the optimisation, the distance to the 'head' surface for a transformed 'hat' point that 'falls' onto one of the voxels surrounding the surface is then already stored in the look-up table and therefore does not have to be computed again. Borgefors compared several Distance Transforms and presented a new one, a 5x5 mask, that approximates the Euclidean Distance to within 2%.

Optimisation

The cost-function can be minimised with a standard minimisation algorithm (i.e. Powell, 1964, Press *et al*, 1992). Unfortunately, the equations given in the previous section for co-registration of fiducial markers can not be used, since there is not a one to one relationship between all the points in the 'head' and 'hat'. Also, the centre of gravity of both point-sets is used in the equations. One can easily see that, for example, when the 'neck' is part of the 'head' and not of the 'hat' that the centre of gravity of the 'head' is misleading for the co-registration. Also, the value of the residual error should be used with care as an indicator of the goodness of the co-registration. Due to the distribution of points in the two datasets its value can be small, even though the co-registration is inaccurate (Wang *et al*, 1994, Kozinska *et al*, 1997).

Improvements

The method can be improved by the removal of inconsistent points that are only part of the surface in one of the two modalities or by increasing the resolution when the algorithm progresses (Schwartz *et al*, 1996).

Accuracy

The mean registration errors found with the surface matching technique are smaller than the voxel size (0.94mm) and between 2 and 3mm at the boundaries of the MRI volume (Schwartz *et al*, 1996). Huppertz and colleagues reported an average test-retest reliability of 1.6mm (Huppertz *et al*, 1998). An accuracy of 1-3mm was reported by Bamidis and Ioannides (Bamidis and Ioannides, 1996), although it was not specified how this was measured. Wang and co-workers reported a mean error in surface matching of 2.4mm and an even smaller mean error of 1.8mm when scaling was taken into account. The mean error in location of seven scalp points was 2.7mm (Wang *et al*, 1994). Brinkmann and colleagues presented results for a phantom and 5 subjects. The mean reconstruction error was 2.2mm for the phantom and of the order of 4mm for the subjects (Brinkmann *et al*, 1998). Pelizzari reported mean registration errors in the order of 1-2 pixels (typically 1 pixel ~ 1mm for MRI) for any combination of CT, MRI and PET co-registration (Pelizzari *et al*, 1989). We expect the same order of accuracy for MEG and MRI co-registration with their algorithm.

In conclusion, surface based co-registration of MEG and MRI data is more accurate and reliable than co-registration based on fiducial markers. Surface matching is therefore used in this work so that accurate anatomical constraints can be obtained for use with a source localisation algorithm. The reported co-registration errors found in the literature vary between 1 and 8mm due to slightly different approaches and reported error measures. It was therefore necessary to evaluate the accuracy of the algorithm that was used in this work. In the next section the accuracy of the implementation of a surface matching technique was tested.

3.4 Errors in Co-registration

3.4.1 Sources of Errors

In this section the sources of errors in MRI/MEG co-registration by surface matching are described. Errors are introduced by the definition of the scalp in MRI co-ordinates, definition of the skin surface in MEG co-ordinates and by the optimisation algorithm itself (Schwartz *et al*, 1996).

Errors in MRI surface

The first source of errors occurs during the scan itself. Dental fillings and partial volume effects cause local distortions and field inhomogenities cause global distortions. Subject movement may blur the exact scalp boundary.

In order to correct for global distortions in the MRI that was used as a test case, three additional parameters were added to the three translation and three rotation parameters that were optimised by the optimisation algorithm. This was a scaling in both the x- and y direction (i.e. pixel size) and a scaling in the z-direction (slice thickness).

Additional errors could be introduced by the determination of the scalp tissue. An over- or underestimation of the scalp boundary can introduce a scaling error of one or two voxels. As this gives rise to a global error, it can be corrected for by the introduction of the scaling factors in the optimisation procedure.

Local errors exist due to noise in the scanner (i.e. voxels that are part of the background are defined as part of the scalp) or parts where the scalp boundary is not continuous (i.e. holes in the scalp). These local errors were corrected by manually editing of the boundary, using the software tools described in section 3.2.1.

Errors in MEG surface

Errors in the definition of the skin surface with the 3D digitiser (Polhemus Isotrak) are due to the digitising process itself, errors caused by the technician and errors due to movement of the subject.

The error due to the 3D digitiser itself is on the order of $1/10^{\text{th}}$ of a millimetre with the sensor up to 38cm from the source (Polhemus, 1992), and therefore negligible to other sources of error (Wang *et al*, 1994, Schwartz *et al*, 1996). Note that other methods of MEG surface extraction can give larger digitisation errors. Digitisation with ultrasound gave digitisation errors smaller than 2mm (Huppertz *et al*, 1998).

Movement of the subject during digitisation of the head also causes errors. Stabilising the subject with a bite-bar system ensured that movement related errors were negligible. Alternatively, mounting the source for the 3D digitiser to the subject's head ensures that small movements do not influence the digitising accuracy. Large movements disturb the magnetic fields and consequently reduce the accuracy (Wang *et al*, 1994).

Too much pressure on the pen can cause movement of the subject or deformation of the skin. The 'deformation' error varies between 1 and 3mm in areas where no important muscle masses lie between the skin and bone (Schwartz *et al*, 1996). Also, skin irregularities due to hair can cause errors of 2-3mm (Wang *et al*, 1994, Schwartz *et al*, 1996).

Errors due to the optimisation algorithm

Due to the spherical shape of the scalp, the cost-function contains flat local minima. It is possible that the optimisation algorithm gets stuck in one of these minima. Therefore, once the algorithm has converged, it is restarted with the final solution. This enables the algorithm to escape a local minimum (since the parameter steps are larger at the beginning of the optimisation than at the end). Another possibility is using multiple resolutions (Lemoine *et al.*, 1994, Schwartz *et al*, 1996). After convergence of Schwartz and colleagues' algorithm the 'head' was sampled with a smaller sampling interval and restarted. They further removed inconsistent points, defined as points that lie too far away from the surface, as these points might give mathematically good, but anatomically implausible solutions. Also, removal of these outliers increases the slope of the error function around the global minimum and eliminates small local minima (Itti *et al*, 1997).

Finally, there should be enough information available for the optimisation algorithm such that the global minimum of the cost-function approaches the true minimum (i.e. the minimum that corresponds to a perfect co-registration). Thus, there should be enough

points in the 'hat', although using too many points increases the execution time of the algorithm.

In conclusion, global errors are accounted for by the scaling factors in the optimisation algorithm. Local errors are estimated to be on the order of 2 pixels (i.e. about 2mm) due to errors in defining the 'head' and local distortions in the MRI scan (van den Elsen and Viergever, 1991). Additional local errors due to deformation of the skin and hair by defining the 'hat' are estimated to be 2-3mm. In section 3.4.2.3 the influence of these local errors is simulated by adding noise to an ideal 'hat', formed from a sub-set of 'head' points. In a final study, the stability of the co-registration is tested by acquiring several 'hats' and computing the co-registration matrix for each 'hat'.

3.4.2 Accuracy of Implemented Co-registration Procedure

In this section the accuracy of an MEG/MRI co-registration procedure that localises the head in the MEG system by means of head-localisation coils and co-registers the MEG and MRI data by means of surface matching is examined. This procedure is schematically presented in figure 3-14. The accuracy of the surface-matching algorithm is evaluated by means of computer simulations and in a practical situation.

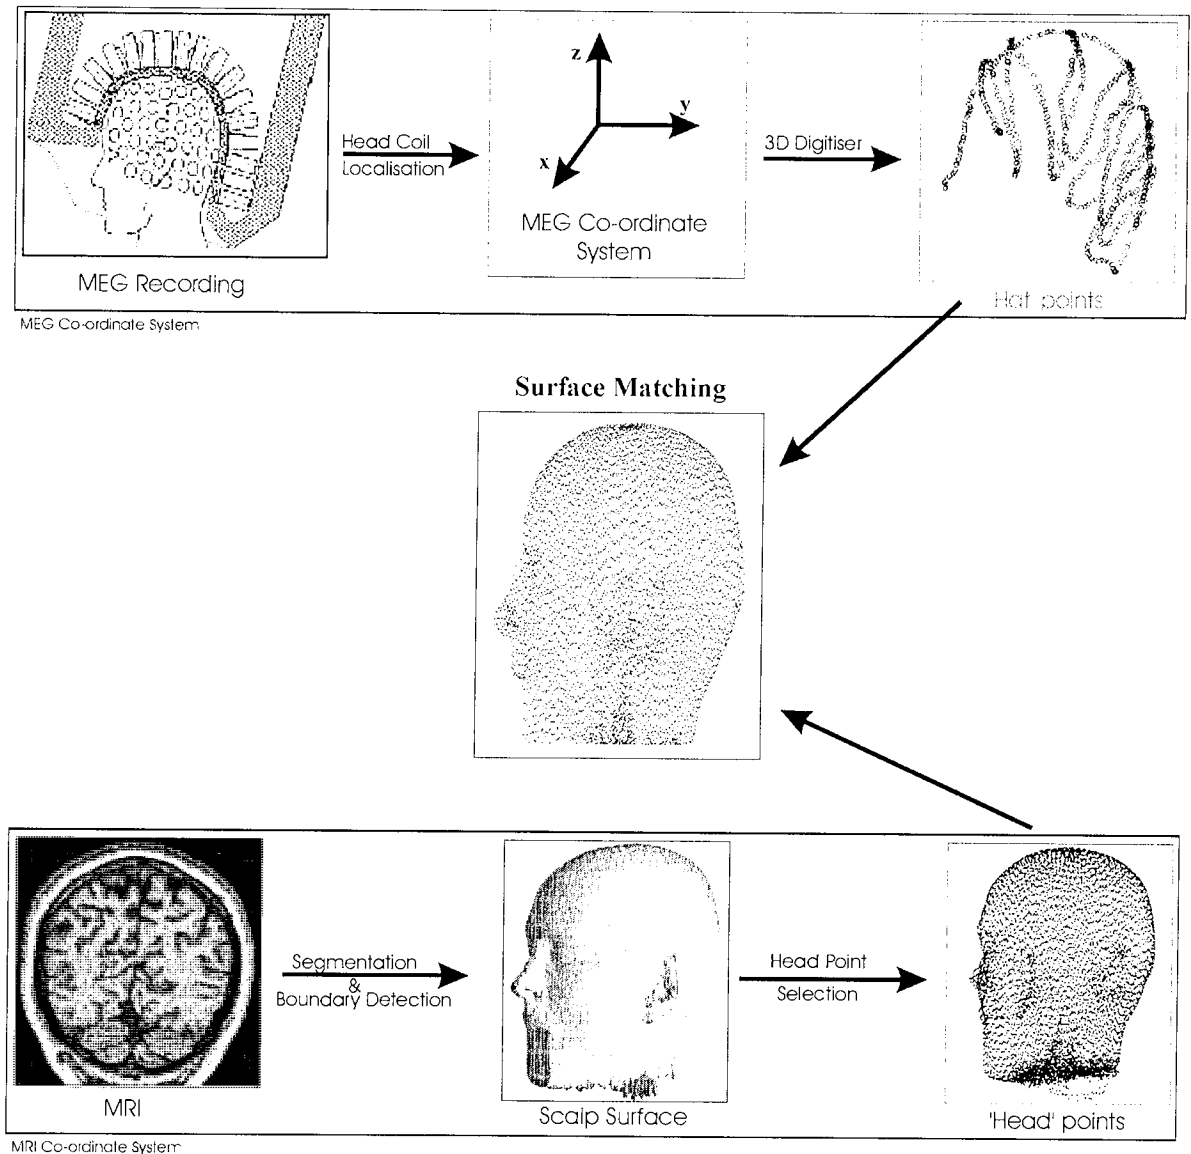


figure 3-14: Schematic of the co-registration procedure. The reference coils determine the location of the MEG sensors with respect to the head. The digitised location of the reference coils (or any other chosen fixed reference system) defines the MEG co-ordinate system. Digitisation of the scalp surface within this reference system gives a set of 'hat' points in MEG co-ordinates. A set of 'head' points in MRI co-ordinates is obtained from the subject's MRI through segmentation and boundary extraction. Surface matching then maps the anatomical information from the subjects MRI to the MEG co-ordinate system. Combining all these steps gives the location of the sensors and the anatomical information in the same co-ordinate system.

Coil localisation

It was already described in section 3.3.2 that the localisation error of three head coils with the CTF 151 channel MEG system is smaller than 2.5mm (CTF Systems Inc.).

Once the locations of the sensors with respect to the coils are known, it is necessary to determine the location of the coils on the subject's head. This is done with a Polhemus 3D digitiser. The location of the coils is determined with respect to some fixed reference system, defined by four points on a bite-bar. This reference system is then transformed to the co-ordinate system defined by the head coils, referred to as the MEG co-ordinate system. This has the benefit that once the MEG co-ordinate system and MRI co-ordinate system are co-registered, the digitisation of the head localisation coils with respect to the bite-bar is sufficient for co-registration in subsequent experiments³. The bite-bar has the additional advantage of keeping the head stable during the digitisation of the head coils.

The test-retest stability of the digitisation of 3 head coils placed on the nasion and pre-auricular points was investigated. The locations of the coils were repeatedly (N=10) digitised with respect to the bite-bar. In each 'run', the bite-bar co-ordinate system was determined by digitising the four holes in the bite-bar. Subsequently, the location of each coil was determined in this co-ordinate system. The location of each coil was digitised twice, moving the digitising pen away from the coil between succeeding measurements. The results are presented in table 3-1. It follows from this table that the smallest errors in

Coil	\bar{x} [mm]	\bar{y} [mm]	\bar{z} [mm]	σ_x [mm]	σ_y [mm]	σ_z [mm]
NA1	22.46	-5.97	74.79	1.42	1.48	0.59
NA2	22.54	-6.06	74.65	1.32	1.19	0.90
PAL1	-71.29	65.45	37.32	0.99	0.83	1.50
PAL2	-71.66	65.45	36.66	1.58	1.00	1.18
PAR1	-65.97	-84.68	30.37	1.41	0.95	1.63
PAR2	-66.47	-84.32	29.69	1.30	1.04	0.81

table 3-1: Mean location and standard deviation of nasion (NA), left pre-auricular (PAL) and right pri-auricular (PAR) as obtained by 10 times digitising the location of a coil at these positions. Each coil was localised twice in a run (NA1 and NA2 etc.).

digitisation of the coils were made in the direction where there was support from the skull. Namely, the standard deviation for the coils at the pre-auriculars was smallest in the y-direction (apart for PAR2). Since the y-axis was defined by the bite-bar as running from the

³ Assuming that the location of the head with respect to the bite-bar does not change over time (i.e. the position of the teeth in the dental impression does not change).

right to the left of the head, this was the direction where there was support from the skull. For the nasion the smallest standard deviation was in the z-direction, which was the direction perpendicular to the plane formed by the 4 bite-bar base points. Thus, for the nasion the support in the z-direction

Coil	ϵ [mm]	σ_ϵ [mm]
NA1	1.89	0.77
NA2	1.85	0.39
PAL1	1.69	0.85
PAL2	1.97	0.76
PAR1	2.03	1.00
PAR2	1.64	0.66

table 3-3: Mean distance from the digitised location of a coil to the mean location as obtained from 10 repeats. Experiment and definitions as in table 3-1.

location of the coils in digitisation 1 and 2 were smaller than 1mm. Note that this digitisation accuracy was not achievable by digitising a coil location once, as follows from table 3-3. This table shows that the digitisation error of a coil, when only digitised once,

Location on bite-bar	σ_x [mm]	σ_y [mm]	σ_z [mm]
Base1	0.32	0.11	0.32
Base2	0.18	0.10	0.17
Base3	0.18	0.08	0.13
Base4	0.13	0.10	0.13

table 3-4: Standard deviation of the repeated (N=10) digitisation of the 4 base points in the bite-bar.

in the digitised location of the coils were not due to variation in the digitised location of the base points, but indeed due to movement of the coils during digitisation. The absolute

	$\Delta\bar{x}$ [mm]	$\Delta\bar{y}$ [mm]	$\Delta\bar{z}$ [mm]	ϵ [mm]
$\overline{NA1-NA2}$	-0.08	0.09	0.14	0.19
$\overline{PAL1-PAL2}$	0.37	0.00	0.66	0.75
$\overline{PAR1-PAR2}$	0.49	-0.37	0.68	0.92

table 3-2: Difference in mean location of the different coils (definitions as in table 3-1) in the x, y and z direction and the distance ϵ between the mean locations. Values were obtained by repeatedly (N=10) digitising the location of a coil at these positions. Each coil was localised twice in a run (NA1 and NA2 etc.).

comes from the ridge of the nose. This suggests that the digitisation errors arose from pressure on the coils by the digitising pen, causing movement in the directions where the coils were not supported by bony structures. This is further supported by the results in table 3-2, where the difference in the mean location and the distance between the mean locations for digitisation 1 and 2 are given. For the pre-auricular points the smallest difference between the mean locations was found in the y-direction, that is the direction where the coils were supported by the skull. The distances between the mean

location of the coils in digitisation 1 and 2 were smaller than 1mm. Note that this digitisation accuracy was not achievable by digitising a coil location once, as follows from table 3-3. This table shows that the digitisation error of a coil, when only digitised once,

was about 2mm.

In a different experiment the 4 base points in the bite-bar were digitised 10 times, removing the digitising pen away from the bite-bar between each digitisation. It follows from table 3-4 that the base points could be repetitively digitised with sub-millimetre variation. This demonstrates that the variations

accuracy was not tested here, i.e. there could be an offset in the locations of the points. A global bias has no consequences for this work, since it is only the relative position of the hat-points with respect to the head coils and the locations of the hat points relative to one another that are important. A local error in absolute location could be caused by distortions of the magnetic fields by a large metal object. Still, away from large metal objects the error in position should be in the sub-millimetre range (Polhemus, 1992).

Summarising, the main errors in the location of the head coils were sub millimetre when each coil location was repetitively digitised. When a coil location was only digitised once, the error was in the order of 2mm. Care should be taken not to press the digitisation pen too hard against the coils in order to avoid movement of the coils in directions where they are not supported by bony structures.

Surface matching

The object of the surface matching algorithm is to co-register a set of 'head' points, obtained from the segmentation of the scalp in a subject's MRI, with a set of 'hat' points obtained by digitising the subjects scalp with a 3D digitiser. This was done by minimising the squared Euclidean distance between the two sets of data points. Using the squared distance as the cost-function instead of the distance itself ensured a larger sensitivity to small changes in the transformation parameters. An example of a set of hat points co-registered with a head surface is shown in figure 3-15.

The computation of the cost-function was speeded up by computing a look-up table that contained the distance from each MRI pixel to the nearest 'head surface' pixel in the MRI volume. Furthermore, it was assumed in this work that the initial guess for the transformation parameters was




figure 3-15: Example of co-registration of a set of hat points obtained with the 3D digitiser (dots) and the head surface obtained from the subjects MRI.

reasonably accurate. Therefore, only a look-up table for MRI pixels lying within a cube of 25x25x25 voxels from the 'head surface' voxels was computed. The distance from the voxels within this cube to the centre of the cube were stored in memory. The centre of the cube was then placed on each 'head' voxel sequentially. The minimum of the distances so obtained for the MRI voxels in the volume around the head surface were stored in the look-up table.

The Simplex algorithm (Nelder and Mead, 1965) was used to minimise the cost-function. All computations were performed in MATLAB (The Mathworks).

It was found during initial test of the algorithm that it got stuck in local minima. These local minima were caused by the discrete nature of the look-up table. That is, a small change in the transformation parameters can result in the distance as obtained from the look-up table not changing. The result is that the value of the cost-function does not change as a result of the change in the transformation parameters and the optimisation gets stuck. This effect was not reported by Schwartz and colleagues (Schwartz *et al*, 1996), who also used a look-up table, because they used fixed steps for the transformation parameters. Huppertz and co-workers (Huppertz *et al*, 1998) neither mentioned problems with local minima, caused by the discrete nature of the look-up table, possibly because they used a different optimisation algorithm.

A way to avoid these local minima is based on the observation that, although the 'hat' points may map to the same MRI pixels in sequential iterations, the 'head' points map to different locations in MEG co-ordinates, even with small changes in the values of the transformation parameters. Therefore, instead of a look-up table containing distances for each MRI voxel, the index of the nearest 'head surface' voxel was stored. The index of an MRI voxel, in an MRI volume with slices that contain 256x256 pixels, was defined as:

$$\text{MRI_index}_{xyz} = x + (y - 1) * 256 + (z - 1) * 65536, \quad (3-11)$$

with x, y and z the MRI co-ordinate of the MRI voxel.

During the optimisation, the hat points were mapped onto the MRI voxels and the indices of the nearest head points were read from the look-up table. Then, the voxels corresponding to these indices were transformed into MEG co-ordinates and the squared

distance between the hat points and the transformed head points was computed. The local minima caused by the discrete nature of the look-up table were therefore avoided at the cost of having to transform some of the head points to MEG co-ordinates and compute the distance between these points and the hat points.

It was further found that convergence to the global minimum was improved by restarting the algorithm at the point at which it last converged.

For the simulations in the next sections the set of hat points was obtained from the set of voxels that form the head surface. Hat points were randomly selected from the head points, so that a set of hat points was obtained that covered the head approximately uniformly. Kozinska and colleagues (Kozinska *et al*, 1997) reported that this is the best way to sample hat points (although no details of other methods were given). Additionally, points were added to the 'nose' when needed in order to ensure better-defined minima in the cost-function.

3.4.2.1 Intrinsic Co-registration Accuracy

In this section the ability of the algorithm to find the global minimum was tested in an ideal situation.

In these simulations 9 parameters were used to perform the transformation from MEG to MRI co-ordinates and vice versa. Three parameters for translation in x, y and z direction in mm (p_x , p_y and p_z , respectively), three parameters for rotation around the x, y, and z axis in degrees ($p_{\theta x}$, $p_{\theta y}$ and $p_{\theta z}$, respectively). Additionally, a parameter to scale the inter-slice distance (p_{sl}) and two parameters to scale the pixel dimensions (p_{pix1} and p_{pix2}) was used (in mm).

A set of 50 hat points was sampled randomly from the head points. Evidently, these hat points were perfectly co-registered with the head points. The hat points were then transformed to MEG co-ordinates with a perfect transformation matrix. This matrix was formed from $p_{perfect} = [p_x, p_y, p_z, p_{\theta x}, p_{\theta y}, p_{\theta z}, p_{sl}, p_{pix1}, p_{pix2}] = [0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1.9 \ 1.1 \ 1.1]$, where the values for p_{sl} , p_{pix} and p_{pix2} were taken as in a typical MRI. The aim of the simulations presented here was to determine the ability of the algorithm to find the correct co-registration parameters for the 'hat' points, given in MEG co-ordinates, when an incorrect guess of the values of the transformation parameters was given at the start ($p_{initial} \neq p_{perfect}$).

Simulations

Firstly, the initial guess of only one transformation parameter was varied at a time, keeping the other parameters constant (at the correct value). By changing only one parameter at a time it was examined whether the algorithm was able to cope with translation, rotation or scaling errors. The initial guess was varied over a wide range to examine how close the initial guess should be to the optimum value. Translation and rotation were given in different units. This complicated the interpretation of the results. To give an idea of the influence of a rotation on the locations of the hat points: a rotation of the hat points by 5° resulted in an average displacement of 7.0mm. Additionally, a scaling of the inter-slice distance by 0.05mm resulted in a change in distance between the most frontal and occipital head points of 5.7mm. A change in pixel scaling of 0.05mm resulted in an increase in distance of 7.7mm between the two outer points on the left and right and

10.7mm between the top and bottom points. The ranges over which the initial parameters were varied were: 1-10mm for the translation parameters, 1-9° for the rotation parameters and 0.1-0.3mm for the scaling parameters.

Secondly, a simulation was carried out in which the 3 translation and 3 rotation parameters were optimised. The initial values were taken from a Gaussian distribution, centred around zero, with a standard deviation of 2 (mm in case of a translation and degrees in case of a rotation). The value of 2 was selected as the standard deviation, since it was found that 100 realisations of such an initial guess gave a mean co-registration error of 4.0 ± 1.6 mm at the initial guess. This mean error was of the same order as the accuracy that can be expected if the initial guess is obtained with the bite-bar system (Singh *et al*, 1997). Each simulation was repeated 50 times.

Finally, the simulations were repeated with all 9 transformation parameters being optimised. The initial value for *psl* was taken from a Gaussian distribution, centred around 1.9, with a standard deviation of 0.01mm. The initial values for *ppix1* and *ppix2* were 1.1 ± 0.01 mm. The initial guess for the translation and rotation parameters was obtained as before. It was expected that the algorithm would have more difficulty converging to the global minimum, because of an increase in the number of local minima when more parameters were optimised.

Results

When only one transformation parameter was optimised at a time, the perfect




figure 3-16: Plot of the cost-function after several rotations of the set of hat-points from the perfect location. The left panel shows the cost-function for rotation in the y-direction and the right panel the cost-function for rotation in the z-direction.

transformation matrix was always found, except when the translation or rotation parameters in the z direction were optimised and when $p\theta_x$ was 3° from the correct value. The results for the translation and the rotation in the z direction are given in table 3-5 and table 3-6.

The perfect solution was not always found, but this was not necessarily due to the fact that the initial guess was far from the perfect solution. For example, when the initial guess was 4° , the perfect solutions were not found, but when the initial guess was worse, i.e. 9° , the perfect solution was found. This was probably due to a local minimum in which the algorithm got stuck, as was indicated by the fact that the same solution was found for $p\theta_z$ equal to 4 or 5° and for $p\theta_z$ equal to 6, 7 or 8° . A comparison of the plots of the cost-function for several values for the rotation in the y- and z-direction shows indeed that the behaviour of the cost-function was more erratic for the rotation in the z-direction (figure 3-16). This problem can be avoided by using a multi-start approach.

Initial Translation [mm] (z direction)	1	2	3	4	5	6	7	8	9	10
d_{end}^2 [mm ²]	0.00	0.00	281.50	0.00	0.00	0.00	644.96	0.00	0.00	0.00

table 3-5: Value of the cost-function (total squared distance, d^2) after optimisation of the translation parameter (z direction). The initial guess for this parameter was off by 1-10mm in the z-direction.

Initial Rotation [°] (z direction)	1	2	3	4	5	6	7	8	9
d_{end}^2 [mm ²]	0.00	0.00	0.00	281.29	281.29	293.13	293.13	293.13	0.17

table 3-6: Value of the cost-function (total squared distance, d^2) after optimisation of the rotation parameter (z direction). The initial guess for this parameter was off by 1-9 degrees in the z-direction.

It is noteworthy that the performance of the algorithm was worse in the direction in which the voxel dimensions were largest. The worse performance in the z-direction could however also be caused by the geometry of the head; the head dimensions were largest in the z-direction.

For the other parameters the optimisation algorithm always found the correct transformation parameters, even when the initial guess was very poor. It can be concluded from the above that the performance of the algorithm was relatively independent of the quality of the initial guess, but that it can get trapped in local minima. In these cases a restart of the algorithm might enable it to find the global minimum.

The results for the simulation where the values of the three translation and three rotation parameters were randomly chosen are displayed in figure 3-17. In three runs perfect co-registration was obtained, and in one run only a small residual error was obtained. In the other six runs the correct solution was not found. The correlation between the mean distance at the start and finish of the surface matching algorithm was 0.89, illustrating that the result of the surface matching was dependent on the initial guess. A correlation of 0.98 between the mean distance and the value of the cost-function after convergence existed indicates that the chosen cost-function was a good indicator of the goodness of the surface matching.




figure 3-17: The mean distance between the head and hat points and the value of the total squared distance at the start and after convergence of the surface matching algorithm. The simulations were repeated 10 times. For each run an initial guess for the three translation and rotation parameters was obtained from a Gaussian distribution with zero mean and a standard deviation of 2mm and 2° respectively. Note the strong correlation between the mean distance and the cost-function.

Summarising:

- The algorithm was able to converge to the global minimum, even when the initial guess was far from the correct solution. For example, in run 8 in the simulation with 9 parameters the global minimum was found, despite an initial mean distance of 4.7mm.
- The number of local minima increased with the number of parameters that were used and thereby the chance that the algorithm got stuck in one of these local minima.
- There was a high correlation between the mean distance obtained with the initial guess and the mean distance obtained after convergence of the algorithm. Hence, the chance of getting stuck in a local minimum increased when the initial guess was further away from the correct solution. The influence of the initial guess on the final result was however not found when only 1 parameter was optimised. This indicates that the quality of the final result depended on a combination of the quality of the initial guess and the number of parameters that were optimised.
- The high correlation between co-registration accuracy and the value of the cost-function after convergence of the algorithm indicates that the value of the cost-function can be used in an experiment as an indicator of the goodness of the match between the head and hat points.
- Local minima could be escaped by restarting the algorithm with the solution found in the local minimum. When this failed, a new initial guess, or alternatively, a new (sub-) set of hat points enabled the algorithm to find the global minimum.

Run	Initial Mean Distance (mm)	Final Mean Distance (mm)	Cost Function Value	Number of Local Minima
1	1.2	1.1	0.15	1
2	1.5	1.4	0.18	1
3	2.0	1.9	0.22	2
4	2.5	2.4	0.28	3
5	3.0	2.9	0.35	4
6	3.5	3.4	0.42	5
7	4.0	3.9	0.50	6
8	4.7	4.6	0.60	7

3.4.2.2 Influence of Number of Registration Points

The performance of the algorithm with only 50 hat points was good when not many parameters were used in the optimisation. However, when the number of parameters was increased, the number of initial guesses that had to be made in order to obtain the global minima increased rapidly, due to the increase in local minima. It was expected that using more hat points would flatten or remove these local minima and hence increase the chance of finding the global minimum. This was examined in a simulation in which the number of hat points was varied. As in the previous section, the hat points were randomly selected from the set of head points, with the exception that it was ensured that there were always 'nose' points in the set. Additionally, only the three translation and three rotation parameters were optimised, with their initial guess obtained from a Gaussian distribution with zero mean and a standard deviation of 2mm and 2° respectively. The simulations with each set of hat points were repeated 5 times.

In table 3-7 the results of these simulations are summarised. As expected, the number of times the global solution was found increased with the number of hat points that were used for the surface matching. The general trend was that the mean distance error per hat point decreased. However, the accuracy decreased when the number of point was increased from 100 to 200 and from 300 to 400. The relative large value for 400 points was caused by a single outlier. It is expected that a smoother dependency on the number of points will be found when the experiment is repeated with more realisations. A faster implementation of the algorithm is needed to do so.

Nr. of Points	100	200	300	400	500
Run					
1	0.08	1.65	0.01	1.18	0.00
2	0.03	0.11	0.01	0.00	0.00
3	1.00	0.07	0.06	0.00	0.00
4	0.04	0.10	0.02	0.02	0.00
5	0.01	2.35	0.05	0.03	0.10
Mean:	0.232	0.856	0.03	0.246	0.02

table 3-7: Influence of the number of head points on co-registration accuracy. The mean distance per hat point after co-registration is given in the table. Note that there was a decrease in accuracy when the number of points was increased from 100 to 200 and from 300 to 400. However, the trend was that the accuracy increased when the number of points was increased.

3.4.2.3 Influence of Noise

In the simulations in the previous sections the hat points that were used were ideal in the fact that there always was a corresponding head point for each hat point, so that the mean distance after surface matching was exactly zero when the correct transformation parameters were found. In a real experiment the head and hat points that are used will contain location errors, as was discussed in section 3.4.1. Due to these errors, the value of the cost-function, and the mean distance between head and hat points, will be larger than zero, even when the correct transformation parameters are used. In the following simulation the influence of such errors on the surface matching algorithm was examined.

The same set of 500 hat points as in the previous section was used. Random noise was added to the location of each hat point, obtained from a Gaussian distribution with zero mean and a standard deviation of 2mm. The value of 2mm was chosen, because this was believed to be the worst case accuracy with which one can digitise a set of hat points. The noiseless hat point set with this many points gave perfect co-registration in 4 out of 5 trials, independent of the initial guess that was used. Any residual distance between the head points and noisy hat points should therefore be due to the noise.

The three translation and rotation parameters were optimised, with their initial guess obtained from a Gaussian distribution with zero mean and a standard deviation of 2mm and 2° respectively. The simulations were repeated 10 times. For each run, the mean distance due to noise alone was also computed, using a perfectly co-registered set of head and hat points.

In figure 3-19 the results of these simulations are summarised. The mean distance over all runs between the head-and hat-points after the algorithm had converged was 1.9mm and the mean distance caused by the noise was 1.5mm. The mean distance after convergence was thus equal to the slice thickness in the MRI scan. The correlation between the mean distance due to the noise and the mean distance after convergence was 0.48. Hence, the accuracy that could be obtained under noisy conditions was limited by the largest dimension of the voxels and correlated to the amount of noise in the head points. Voxel dimensions were found to be unimportant in the noiseless cases, since the perfect solution could be found. However, the voxel dimensions determined the limit of the accuracy that was obtainable when the perfect solution was not attainable due to the noise.

Furthermore, the final result was found to be dependent on the quality of the initial guess (correlation of 0.84).




figure 3-19: The mean distance between the head and hat points due to noise alone and at the start and end of the surface matching algorithm. The simulations were repeated 10 times. For each run an initial guess for the three translation and rotation parameters was obtained from a Gaussian distribution with zero mean and a standard deviation of 2mm and 2° respectively. The noise that was added to the location of each hat point was obtained from a Gaussian distribution with zero mean and a standard deviation of 2mm.

3.4.2.4 Test-Retest Accuracy

The simulations carried out in the previous sections gave valuable clues about the performance of the algorithm in ideal cases and when noise was added to the hat points. The ultimate test for the algorithm is to co-register the scalp surface obtained from a subject's MRI with a set of hat points from the same subject, obtained with the 3D digitiser. In this case the real error can not be determined, as the 'perfect' transformation parameters are unknown. However, the co-registration parameters obtained from a test-retest will vary slightly. The co-registration matrices formed from these varying co-registration parameters transform each hat point to a slightly different location in the MRI co-ordinate system from trial to trial. The spread of each transformed hat point is therefore an indication of the accuracy of the surface matching algorithm. A bias in the algorithm can not be detected in this way, but a visual inspection of the co-registration results can serve as a strategy against incorrect co-registration matrices. Additionally, it was shown in section 3.4.2.1 that the value of the cost-function can be used as an indicator of the accuracy of the co-registration, since there is a high correlation between this value and the mean distance between the head and hat points.

Method

The same set of head points was used as in the previous sections. Additionally, ten different sets of hat points were recorded for the subject for whom the head points were obtained. The head was stabilised with a bite-bar and a set of hat points was obtained with a Polhemus 3D digitiser. The hat points were acquired in a continuous mode, meaning that points were recorded continuously while the experimenter moved the digitisation pen over the subject's scalp. Many points from the nose area and the regions around the eyes were included, as these locations influenced the shape of the cost-function considerably due to the non-spherical shape of these areas. Care was taken not to deform the scalp or the tip of the nose during the digitisation process by applying as little pressure with the pen as possible. A total of 10 hat point sets were obtained in this way.

Due to the continuous mode of recording the number of hat points in each set was quite large, with a small distance between points and their neighbours. The large number of hat points increases the computing time of the algorithm and it was further found in 3.4.2.2

that little is gained by using too many points. Therefore, the hat points were resampled, so that no point had a neighbour located within a sphere with radius ϵ . The areas around the nose and eyes were not resampled. The value of ϵ was chosen so that the total number of hat points was 500.

Each set of hat points was co-registered with the head points, using the same surface matching algorithm as used in the previous section. For the first dataset, the initial guess for the transformation matrix was obtained from manual co-registration. The co-registration parameters were modified on a trial-by-error basis, for which visual feedback was obtained from the display of the head points and transformed hat points. For the additional hat points sets the transformation matrix obtained from surface matching of the previous set of hat points was used as the initial guess.

Transformation of a single point with the 10 different transformation matrices gave a cluster of 10 points. The spread of this cluster was an indication of the variability of the co-registration procedure. The root mean square (RMS) value of the standard deviations in the x, y and z-directions was therefore used as the error measure, ϵ_v .

Results

The transformation parameters as estimated by the surface matching algorithm were used to form 10 transformation matrices. Each vertex in the realistically shaped mesh was then transformed from MRI to MEG co-ordinates with all ten transformation matrices and the resulting error was computed. The mean value found for ϵ_v was 1.6 ± 0.2 mm, with a minimum value of 1.1 and a maximum of 2.1mm. Thus, the minimum variability of the surface matching approach was equal to the smallest voxel dimension.

A map of ϵ_v for each point vertex in the mesh revealed that the largest variability was found in the occipital pole (figure 3-20). This was mainly due to the relatively large variability in the z-direction.

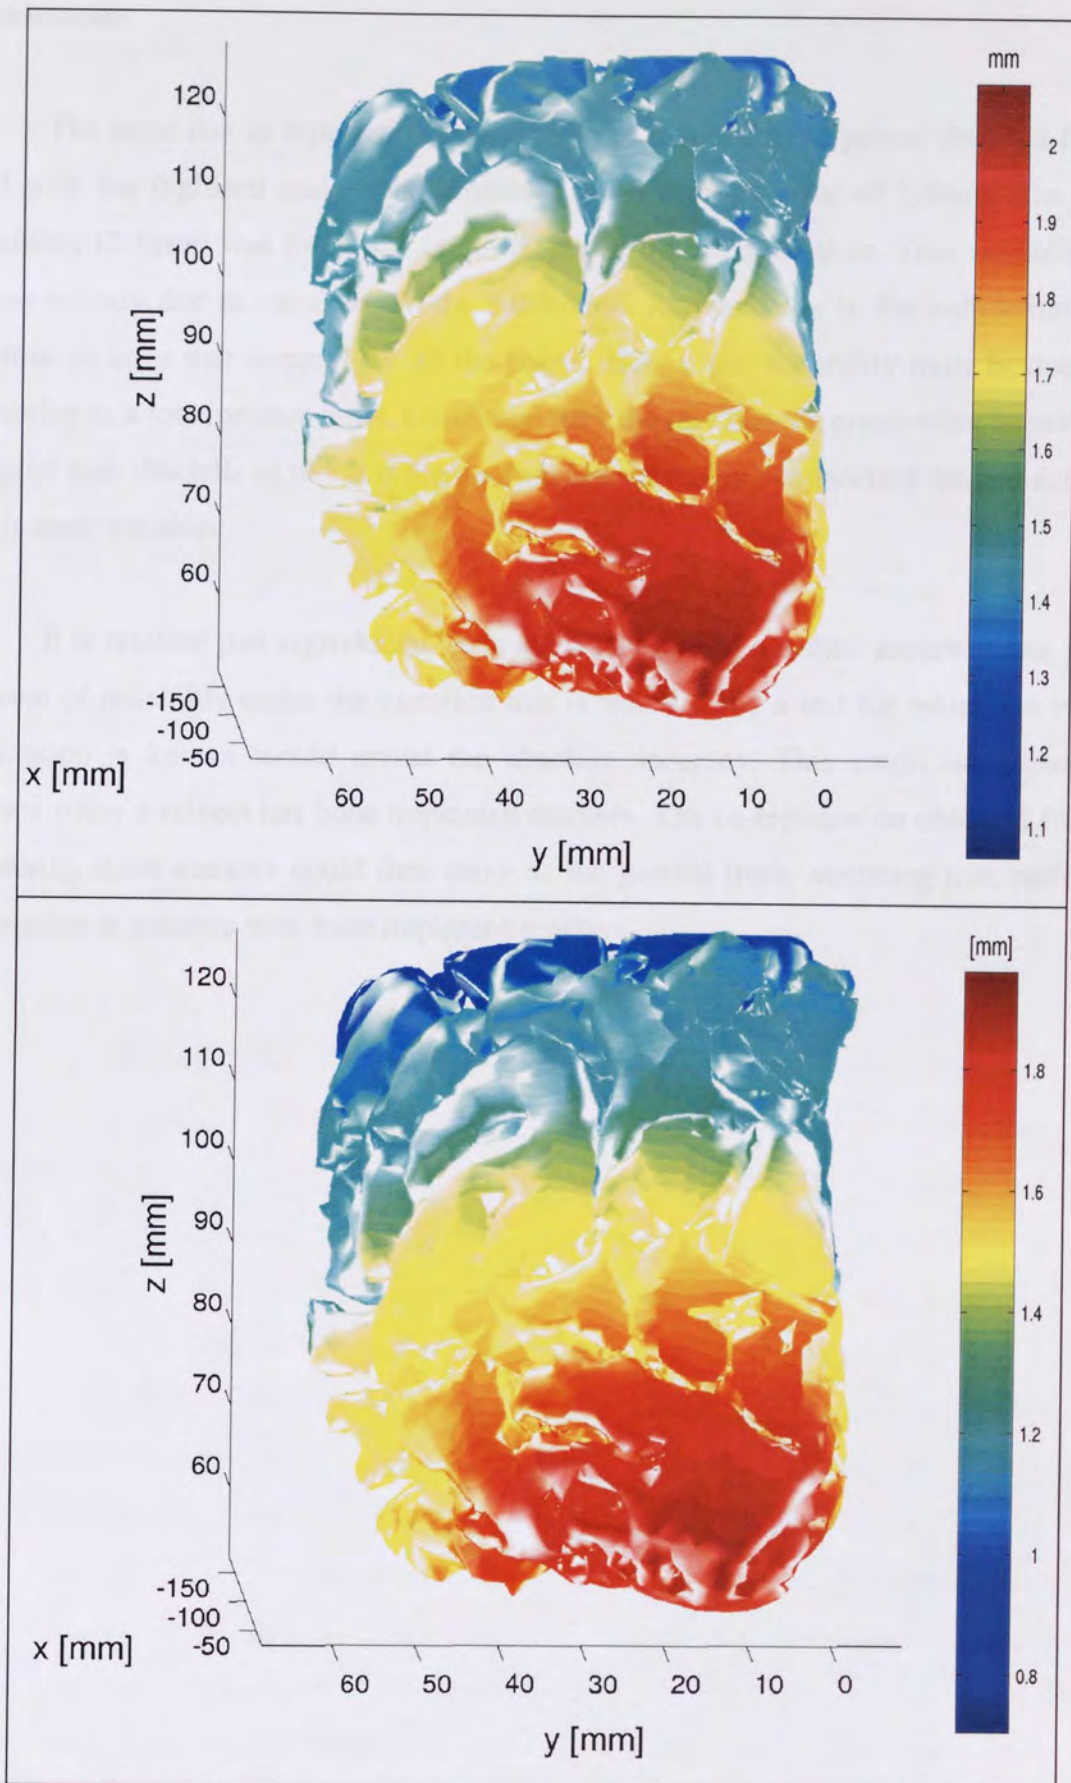


figure 3-20: Variability due to repeated co-registration for different regions of the brain. The total RMS error is given in the top panel and the standard deviation in the z-direction is given in the bottom panel. Note that the largest errors were found towards the back of the head and that these errors were mainly caused by the variability in the z-direction.

Conclusions

The error due to repeated co-registration of a set of head points obtained from an MRI with the digitised scalp points resulted in an average error of 1.6mm. The largest variability (2.1mm) was found for target points in the occipital pole. This variability was almost entirely due to variations in the z-direction. A translation in the z-direction would result in an error that is equal for all the points. Hence, the variability must be due to the variability in a rotation parameter. Combined with the fact that the errors were largest in the occipital pole this tells us that it is the rotation around the line connecting the pre-auriculars that is most variable.

It is realised that reproducibility is not a measure of absolute accuracy, but gives a measure of reliability under the variation that is tested. Only a test for which the true co-registration is known would reveal the absolute accuracy. This could for example be realised when a subject has bone implanted markers. The co-registration obtained from co-registering these markers could then serve as the ground truth, assuming that perfect co-registration is possible with bone implanted markers.

3.5 Discussion

The inverse problem in biomagnetism suffers from the fact that this problem is non-unique, hence constraints are needed in order to obtain a unique reconstruction from the measured magnetic fields. Anatomical constraints are easily incorporated into the formalism of the FOCUSS algorithm. A natural candidate for such a source space is the cortical surface, since most of the brain activity that produces magnetic fields that are strong enough to be measured with an MEG system takes place in the grey matter. A model of the cortical surface can be retrieved from a subject's MRI. The approach to use such a realistically shaped source space as a constraint on the inverse problem depends on two factors:

1. The cortical surface can be accurately recovered from a subject's MRI
2. The co-registration of the MEG co-ordinate system (in which the measurements were taken) and the MRI co-ordinate system (in which the realistically shaped source space is defined) is accurate. Inaccurate co-registration would result in inaccurate constraints and consequently erroneous results.

Both issues have been examined in this chapter.

The accuracy of the description of the cortical surface depends on the segmentation technique that is used to separate the different tissues in the head. Inaccurate segmentation can result in local errors in the definition of the cortical surface. A lot of research effort has been put into the segmentation problem, both because there are many applications for its use and because it is difficult to design techniques that can perform segmentation reliably and automatically. Automatic tissue classification is troubled by noise in the scanner, small differences in intensity levels between various tissue types, variation of the intensity level of a tissue type throughout the whole volume and the partial volume effect near the boundaries of tissues. A closely related problem is finding the boundary of the grey matter tissue, which can also add errors to the cortical surface. There is not yet a technique that can extract the cortical surface reliably from an MRI. Therefore a mainly manual approach was used in this work. The grey matter tissue was defined by a simple thresholding technique. Subsequently, the boundary of the grey matter was defined manually, using a combination of MATLAB's contouring algorithm and software tools that enabled the selection and editing of contours and drawing. This approach resulted in the accurate definition of the cortical boundaries, but was rather time-consuming. For these anatomical constraints to

become available for the day to day use in source reconstructions, a technique is needed that can perform the extraction of the cortical surface (semi-) automatically and reliably. It is expected that with the ongoing research effort into segmentation techniques that such a technique will become available in the near future.

A further cause for error in the construction of the cortical surface is the reconstruction of the surface from the contours defined in each MRI slice. Each slice can contain a different number of, possibly intersecting, contours. The reconstruction of a surface from such slices is non-unique and is therefore dependent on the method that is used to select the best way to connect the contours. The NUAGES algorithm used in this work is based on Delaunay triangulation and therefore based on geometric closeness. It can deal with the branching problem and therefore poses no restrictions on the number of contours that are defined in each slice. It was reported (Geiger, 1993) that vector contours, such as the contours defined around the grey matter, provide both accurate reconstruction with respect to surface location and orientation when used as input for the NUAGES algorithm. Another reported finding was that the largest reconstruction errors are found where the surface is almost parallel to the cross section. This should be taken into account when the orientation of the MRI slices in which the boundaries are defined is chosen. For example, if one wants to use the cortical surface as a source space for source localisation from somatosensory evoked magnetic fields, then the contours should not be selected in axial MRI slices.

A surface matching approach was taken for the co-registration of the MEG and MRI co-ordinate systems. This approach does not suffer from the problem of identification and localisation of a few (anatomical) markers in the MEG and MRI co-ordinate systems, as is the case when fiducial markers are used for the co-registration. Additionally, with fiducial markers, a relatively small digitisation error in the location of one of the markers can have a profound effect on the co-registration accuracy. With the surface matching approach many more points are used and a digitisation error in the location of one of the points will only have a small effect on the overall co-registration accuracy.

It was reported that the accuracy of localisation of the head coils with the CTF 151 channel system, normally used as fiducial markers, is smaller than 2.5mm. The digitisation error of these coils was 2mm when only digitised once, and smaller than 1mm when the location was determined from repeated localisation of the coils. Care must be taken throughout the digitisation not to move the coils in directions where they are not supported

by the underlying bony structures. The total error of determining the location of the coils in the MEG co-ordinate system is of the order of 2mm. However, identifying the location of the coils in the MRI system is prone to larger errors due to the problem of identification, placing oil filled coils at the correct locations and the size of the oil filled capsules. The total error in the co-registration of the fiducial markers is therefore at least double the error of the determination of the location in the MEG co-ordinate system, i.e. 4mm. The co-registration error of points far away from the marker locations, like areas in the visual cortex, is again larger than this error (see equation 3-10). The bite-bar system used by Singh and colleagues for co-registration (Singh *et al*, 1997) stabilises the head during recordings and increases the accuracy with which reference points can be identified. However, an error in location of one of the reference points still causes large co-registration errors due to the small number of reference points that are used.

The surface matching algorithm used in this study used a look-up table in order to speed up the optimisation, since the distances from voxels around the head surface to the surface were already stored in the table. The initial computation of the look-up table was time-consuming, but only had to be performed once and could then be read from disk. Additionally, fast algorithms exist that result in a look-up table containing approximations of the Euclidean distance to the head surface (Borgefors, 1986, Schwartz *et al*, 1996, Kozinska *et al*, 1997, Huppertz *et al*, 1998).

The ability of the algorithm used in this work to avoid local minima depended in a complex way on the number of parameters that were used in the optimisation, the quality of the initial guess and on the number of points that were used. It was found that, in the noiseless case, the accuracy of the co-registration was largest when the most hat points were used. In order to reduce the influence of local minima it was further necessary to use as many points as possible from the non-spherical parts of the head, like the nose and the bony structures around the eyes.

The algorithm was able to find the perfect co-registration parameters when no noise was present in the location of the hat points. When noise was present, the error in co-registration correlated to the error in the hat point locations. Therefore, the accuracy of the co-registration procedure was dependent on the accuracy with which the hat points were digitised. The hat points were obtained in a continuous mode, although it was demonstrated that the digitisation error of a point was smaller when the location of a point was digitised

several times and the mean location was taken as its location. However, this would be too time-consuming due to the large number of hat points that were needed for the surface matching. An approach that could give a large number of hat points in only a few seconds is based on an optical system (Hurley, 1999) A laser-light with low energy is rotated around the subject's head while a digital camera records the reflections on the head. The accuracy of this system has yet to be determined, although it is accurate enough to produce realistically looking 3D images of the head. One problem with this method might be that it also measures (some) reflections from the hair, instead of the scalp.

Additionally, when no noise was present the accuracy of the algorithm was independent of the voxel size. However, under noise conditions, the maximum accuracy that was obtainable was limited by the slice thickness. Therefore, volumetric MRI scans with isotropic voxel dimensions should be used when MRIs are used for co-registration with MEG. Previously, obtaining volumetric MRI scans with small slice thicknesses was time-consuming and therefore prone to movement induced errors. Nowadays, obtaining such MRI sets in a reasonable acquisition time is feasible with fast scanning techniques (Aine, 1995).

The variability due to repeated co-registration was largest (2.1mm) for target points in the occipital pole, mainly due to variations in the z-direction. Hence, it was difficult to find the correct value for the rotation around the line connecting the pre-auriculars. Still, the mean error due to test-retest was only 1.6mm.

4 Figures of Merit for Distributed Source Reconstruction

4.1 Introduction

The solutions found by any algorithm that localises brain activity from measured electric and/or magnetic fields outside the head are non-unique, unless additional constraints are imposed on the solutions. It is therefore necessary to assess the goodness of a solution and the imposed constraints by means of some kind of figure of merit.

A commonly used error measure is the dipole localisation error (DLE) when algorithms that reconstruct dipolar sources are tested with computer simulations. In this case the correct location of a dipolar source is known and the DLE is defined as the Euclidean distance between the correct source location and the location of the reconstructed source. With measured data the correct location is unknown and therefore other error measures are needed. Two commonly used measures are the reduced chi-square and the percent of variance (see section 2.2.1.2 for the definition and a discussion of these measures). A deviation from the expected value for these measures indicates that the measured data are not sufficiently explained by the reconstructed sources. This can either be due to noise in the data or due to the use of constraints that are not satisfactory for the measured data. For example, the imposed constraint of a single dipolar source is invalid when the brain activity can be better explained by a distributed source or by more than one dipolar source.

Unlike for discrete sources, there are no figures of merit that are widely used with tomographic reconstruction algorithms. When tomographic techniques are tested by means of computer simulations, the accuracy is usually assessed by computing the Euclidean distance between the location of the original dipolar source and the location of the centre of activity of the reconstructed distributed source (Wang *et al.*, 1992, Wang, 1993, Wang, 1994). This is not a fair measure, as tomographic techniques were never designed to localise discrete sources. Note that figures of merit like the reduced chi-square or the percent of variance are meaningless with tomographic reconstruction techniques. Due to the

large number of sources it is possible to explain 100% of the data. This means that noise in the data will also be reconstructed as spurious sources in the brain. This unwanted effect can be reduced by means of regularisation (section 2.3.2.3).

There are measures though that can be used to assess the quality of tomographic reconstruction algorithms. Recently, Grave de Peralta Menendez and colleagues have brought the concept of the resolution matrix to the attention of the biomagnetism community (Grave de Peralta Menendez *et al*, 1996, Grave de Peralta Menendez *et al*, 1997, Lütkenhöner and Grave de Peralta Menendez, 1997, Grave de Peralta-Menendez and Gonzales-Andino, 1998, Grave de Peralta Menendez and Gonzalez Andino, 1999). This concept is based on the lead fields and a matrix that represents the inversion procedure, including constraints. Thus, it is only based on the definition of the source space, the volume conductor model, the sensor configuration and any *a priori* information. It is not based on measured data and can therefore not assess the quality of a source reconstruction. On the other hand, it provides the experimenter with a method to judge the applicability of a reconstruction technique for the case in hand, as it can identify regions for which the algorithm gives inaccurate reconstruction results, whatever the quality of the data is.

In the next sections, these figures of merit for distributed source reconstruction algorithms are described in more detail.

Solutions found with the FOCUSS algorithm consist of only a few discrete sources (Gorodnitsky *et al*, 1995). The figures of merit used with distributed source reconstruction algorithms can therefore not be applied to the final solutions obtained with FOCUSS. However, the solutions are of a distributed nature when the algorithm starts. The figures of merit for distributed solutions mentioned above can therefore be used to help with the choice of the initial constraints (i.e. the initial weighting matrix). The usefulness of these measures is examined by means of computer simulations in section 4.3.

4.2 Theory

4.2.1 Decrease of Singular Values

It was described in section 2.2.3.1 how a matrix can be decomposed in orthogonal components by means of Singular Value Decomposition. To repeat briefly:

For a general matrix \mathbf{A} , with dimensions $N \times M$, the SVD results in:

$$\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{V}^T, \quad (4-1)$$

where \mathbf{U} and \mathbf{V} are the orthogonal matrices with dimensions $N \times N$ and $M \times M$, respectively. The matrix \mathbf{U} contains the left singular vectors and \mathbf{V} contains the right singular vectors. The diagonal matrix $\mathbf{\Lambda}$ is of the same dimensions as \mathbf{A} and contains the singular values of \mathbf{A} in descending order:

$$\mathbf{\Lambda} = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_j, \lambda_{j+1}, \dots, 0, \dots, 0), \quad \lambda_j \geq \lambda_{j+1} \quad (4-2)$$

The singular values represent the magnitude of the contribution of each singular vector to \mathbf{A} . The summation of all components, weighted by the singular values, returns the original data matrix.

The decay pattern of the singular values can also be used to assess the quality of the lead field matrix \mathbf{L} that contains the information about the forward problem in linear inverse algorithms. This can be seen as follows:

Each right eigenvector of \mathbf{L} is a basis function with the same dimensions as the number of elements in the source space. A basis function can therefore be displayed in the source space, where it gives insight into the relative contribution from different areas to the basis function. The left eigenvectors are the orthonormal field patterns that are induced in the sensors by each corresponding basis function. The singular values weight the contribution of each left eigenvector to \mathbf{L} .

In general, the low frequency spatial components are represented in the first few eigenvectors and subsequent eigenvectors contain information about finer details. Unfortunately, the singular values corresponding to the latter eigenvectors are small, hence the contribution of these eigenvectors is relatively small. This means that fine detail in the

basis functions is lost when the difference between the larger and smaller singular values is big. Thus, the *Condition Number* of the lead field matrix, defined as the division of the largest singular value by the smallest singular value, provides a single measure that gives an idea of the spatial resolution that can be obtained.

Ideally, the decay pattern of the singular values should be flat, so that the relative contribution from low and high frequency components is equal. In reality though, the decay pattern is almost exponential. Additionally, noise in the measurements ensures that eigenvectors corresponding to small singular values can not be used in source reconstructions. Therefore, some high frequency components need to be discarded, depending on both the decay pattern of the singular values and the SNR.

The quality of the lead field matrix, which is at the heart of any linear inverse algorithm, can be judged by the decay pattern of its singular values. This approach was used to determine the optimum reference electrode for source localisation from EEG (Gençer *et al.*, 1996). They found that the first few basis functions peak near the reference electrode of choice. Therefore, a measurement is more sensitive to activity near the reference electrode. It follows from this that, if one wants to localise activity in, for example, the occipital cortex, one should use an occipital reference. However, the spatial resolution of a source reconstruction that can be obtained depends on the number of basis functions that can be used and hence on the decay pattern of the singular values and the SNR. Using an average reference electrode gave the flattest decay pattern and would be preferable in cases where high resolution is wanted. The flat decay pattern for data with an average reference electrode was caused by the fact that there were no dominant basis functions as the average reference electrode is not near any source areas. However, eye-blinks or other artefacts influence the measurements in all electrodes when an average reference is used. Hence, the choice of reference electrode depends on the type of experiment performed.

In a subsequent paper (Gençer and Williamson, 1998), the decay pattern of the singular values was used to demonstrate that the decay pattern flattens when more sensors are used in an MEG or EEG recording. Therefore, more basis functions may be used to reconstruct activity and consequently the resolution of the source reconstruction increases. However, the amount of extra basis functions that can be used depends on the SNR of the measurements. When high SNR data are used, the number of basis functions that can be used increases significantly when sensors are added. However, when the SNR is low, the increase in the number of basis function is less dramatic. This is illustrated in figure 4-1. It




figure 4-1: Decay pattern of the singular values for several set-ups with different numbers of sensors. With low SNR data only a few singular values are larger than the singular values due to the noise, for set-ups with a few as well as many sensors. The cut-off level is indicated with the line annotated with *Low SNR*. On the other hand, with high SNR data many more singular values are larger than the noise singular values and the difference in the number of singular values that is usable between a set-up with a few or many sensors is large. The cut-off level for high SNR data is indicated with the line annotated with *High SNR* (modified from Gençer and Williamson (1998)).

follows from this figure that for data with low SNR the increase in extra basis functions that can be used is small or even negligible, whereas for high SNR data the increase is significant (compare ~20 basis functions for a system with 25 sensors with ~90 basis functions for a system with 289 sensors for the high SNR data). It was reported that for a realistic SNR the benefit of increasing the number of sensors beyond 100 is small (Gençer and Williamson, 1998).

4.2.2 Resolution Matrix

The concept of resolution kernels (or averaging kernels) was first developed by Backus and Gilbert for the analysis of inverse problems in geophysics (e.g. Backus and Gilbert, 1968). This inverse problem involves the reconstruction of the internal structure of the earth from measurements of properties of the earth (like the mass or frequency of oscillation). They developed the concept of an averaging kernel to quantify the idea of

constructing a desirable estimator that computes a local average of some earth variable from the measured earth data, for example the average density in a region. The idea is that the average at a certain location should emphasise heavy contributions from points near the target location and little contribution from points further away. The ideal averaging kernel is therefore a delta function around the target location. They then chose a function that describes the deviation of the averaging kernels from the delta function. Minimising this function results in an optimal estimator, in the sense that the value of an estimated earth variable at a certain location is the best average that can be obtained at that location (i.e. the average is only based on contributions from a small region around the target location).

Robinson and Rose (Robinson, 1989, Robinson and Rose, 1992) used a similar idea in order to obtain an optimal spatial filter. They minimised the deviation of 'synthesised lead fields' from a delta function for the design of filter coefficients in spatial filter imaging with MEG. With the filter coefficients estimated, an image of source activity is obtained by projecting the MEG measurements through an array of regularly spaced 'synthesised lead fields' (See section 2.5 for more details).

Grave de Peralta Menendez and colleagues described the concept of resolution kernels within the context of the inverse problem in biomagnetism (Grave de Peralta Menendez *et al*, 1996, Grave de Peralta Menendez *et al*, 1997, Lütkenhöner and Grave de Peralta Menendez, 1997, Grave de Peralta-Menendez and Gonzales-Andino, 1998, Grave de Peralta Menendez and Gonzalez Andino, 1999). They presented the resolution matrix, in which each row is the discrete analogue of the resolution kernels described by Backus and Gilbert. Additionally, they derived figures of merit based on the resolution matrix that are easier to visualise than the full resolution matrix. These measures were used for the comparison of different linear estimators.

The derivation of the resolution matrix goes as follows (Grave de Peralta Menendez *et al*, 1996, Grave de Peralta Menendez *et al*, 1997, Grave de Peralta-Menendez and Gonzales-Andino, 1998):

We repeat equation 2-19, the discrete formulation of the forward problem, using the lead field concept:

$$\mathbf{B} = \mathbf{L} \mathbf{Q}, \tag{4-3}$$

with \mathbf{B} the column vector containing the magnetic field measurements in the sensors at a certain instant in time, \mathbf{L} the lead field matrix and \mathbf{Q} the column vector containing the source strengths of all elements in the source space.

It was shown in section 2.3 that several approaches to find a linear estimation of the strength of the sources can be formulated in a general way as:

$$\hat{\mathbf{Q}} = \mathbf{G} \mathbf{B}, \quad (4-4)$$

The matrix \mathbf{G} represents the inversion procedure that is used, including any *a priori* information that is included in the reconstruction. The details of the most commonly used implementations of the linear estimation matrix \mathbf{G} are given in section 2.3.

Substituting equation 4-3 into equation 4-4 gives:

$$\hat{\mathbf{Q}} = \mathbf{G} \mathbf{L} \mathbf{Q} = \mathbf{R} \mathbf{Q}, \quad (4-5)$$

The matrix \mathbf{R} in the above equation is known as the resolution matrix. A few important points follow from this equation:

- The estimated source strength of the k^{th} element in the source space is obtained by matrix multiplication of the correct source strength of all elements with the k^{th} row of \mathbf{R} . Thus, the rows of \mathbf{R} 'filter' the contribution from all other sources to the estimated value of the target element. Let's see what happens in a hypothetical situation where two sources, located far apart, are simultaneously active. Assume that the row in the resolution matrix corresponding to source 1 has its highest value at the location of source 1, but that there are also considerable sidelobes so that the value in the resolution kernel is also significantly larger than zero for source 2. This means that the estimated value of the strength of source 1 is the sum of the (correct) contribution from activity at this location and a (incorrect) contribution from the second source. In the case the element for source 2 is negative and the strength of source 2 is larger than that of source 1, the contributions from source 1 and source 2 to the estimated strength of source 1 could even cancel out. This means that it would be impossible to detect the activity of source 1 when source 2 is active at the same time, no matter how high the SNR of the data is!

- From the example given above it follows that, ideally, each row k should have a value of 1 at the k^{th} element and all other elements should be zero (i.e. the discrete version of a delta function) as this would mean that the estimated strength of a source is independent of activity from other sources. Consequently, the ideal resolution matrix is the identity matrix \mathbf{I} . This requires that \mathbf{L} has a rank equal to the number of source elements¹. This is impossible to occur with biomagnetic measurements. The best we can try to achieve is a resolution kernel with maximum values near the source element associated with that row and small sidelobes. The extent of the area that surrounds an active source and influences its estimated strength can be used as a measure of the resolution of the linear estimator, since fine detail within this area cannot be revealed. Hence the name resolution kernel. An example of the resolution kernels of a source, obtained with different linear estimators, is given in figure 4-2.
- The resolution matrix is based on the linear estimator and the lead field matrix. It is therefore completely defined by the inversion technique that is used, *a priori* information, definition of the source space, volume conductor model and sensor configuration. Hence, it is independent of the data and an analysis of the resolution matrix can be done before any measurements are made. Potential shortcomings of the used approach can therefore be detected before actual measurements, and source reconstruction from these measurements, are performed.
- An important point to make is that the resolution matrix only describes what the possible influence of other sources on the estimated value of the strength of a source is. The exact contribution from other sources depends on the location and strength of these 'disturbing' sources during the measurements. The information obtained from the resolution kernels is still valuable though, since the activity reconstructed with a linear estimator that has large sidelobes in the resolution kernels is more likely to be affected by activity in the brain that is not associated with the stimulus of interest than with an estimator with smaller sidelobes. An estimator with smaller sidelobes would therefore be preferable.

¹ If each source element has an orientation in the x-,y- and z-directions (i.e. the source orientation is not constrained to be normal to the surface), then the rank should be 3 times the number of source elements.

- The columns of the resolution matrix contain useful information as well. They are referred to as the impulse responses of the linear estimator and reveal the ability of the linear estimator \mathbf{G} to localise single point sources. To clarify this: assume that the measurements are exactly those that are produced by a unitary single current dipole located in the k^{th} element of the source space. Note that this is the definition of the lead field \mathbf{L} of the k^{th} element. Hence, in this case $\mathbf{B} = \mathbf{L}$. Using this, it follows from equations 4-4 and 4-5 that the solution provided by the linear estimator \mathbf{G} is identical to the k^{th} column of the resolution matrix. Thus, the columns of \mathbf{R} are the impulse responses of the linear estimator. As with the resolution kernel, the amplitude of the main peak in the impulse responses should occur at the location of the corresponding source elements. The magnitude reflects the amplification or reduction of the activity of



Aston University

Content has been removed for copyright reasons

figure 4-2: Resolution kernels for three different linear estimators for the target point A. The reconstructed activity in A will be influenced by activity from other sources if their value in the resolution kernel is non-zero. For example, activity in B will contribute to the reconstructed activity in A for all three linear estimators. Reconstructed activity in location A could either be due to activity in A or due to slightly stronger activity in B, so the spatial resolution of the linear estimator is related to the width of the main peak in the resolution kernel. When the Minimum Norm estimator is used, activity in C will have a negative contribution to the activity in A. This could cancel the activity of source A when a strong source is present in C. Activity in C has a small positive contribution to the activity in A when the B&G estimator is used and no contribution if the WROP estimator is used. Hence, the optimal linear estimator should have a narrow peak around the target source (influence of source B is minimal and therefore high spatial resolution) and small sidelobes (influence of distant sources like C is small). From Grave de Peralta Menendez *et al* (1997), see original paper for details about the linear estimators.

an actual source due to the estimator \mathbf{G} . The width of the main peak reflects smearing of the activity and sidelobes reflect spurious sources introduced by the estimator.

- It follows from the previous discussion that the resolution matrix can be computed by concatenating the impulse responses that are obtained for all elements in the source space sequentially. This allows for the computation of the resolution matrix for algorithms that can not be formulated as in equation 4-4. For example the resolution matrix for those algorithms that compute the linear inverse by means of an iterative algorithm can be computed this way. It would be especially interesting to compare the resolution matrix at the start of the FOCUSS algorithm (using $\mathbf{R} = \mathbf{G} \mathbf{L}$) and after it has converged (using the concatenation of impulse responses). This should demonstrate a significant improvement of the resolution matrix.
- The impulse response, or measures derived from it, should be used with care. They are based on the ability to reconstruct discrete localised sources, which is something linear estimation techniques were not developed for.
- A plot of the resolution kernels for different estimators can help to choose between different linear estimators. Assume that there is an estimator that gives a narrow peak around the location where an active source is expected, and thus a high spatial resolution, but that also gives large sidelobes. Also assume a second estimator has a broader peak, but no significant sidelobes. One can choose to use the first estimator if no significant brain noise is present. If there are a lot of brain sources that are not related to the process of interest, one should use the second estimator. This ensures that these noise sources have negligible influence on the estimated activity of the target sources, at the expense of a reduction in spatial resolution.
- The derivation of the resolution matrix given above is only valid for linear distributed estimators. A derivation that is valid for linear estimators in general can be found in Lütkenhöner and Grave de Peralta Menendez (1997).

A plot of each row in the resolution matrix provides us with useful information about the possibilities and limitations of the linear estimation approach that is used, as was first shown by Backus and Gilbert (Backus and Gilbert, 1968). Yet, in practice, a realistically shaped source space consists of thousands of elements. A plot of each row would be impractical. Therefore, Grave de Peralta Menendez and co-workers described measures derived from the resolution matrix that contain the essential information, but are

easier to display (Grave de Peralta Menendez *et al*, 1996, Grave de Peralta-Menendez and Gonzales-Andino, 1998):

Source Identifiability

As described above, the ideal resolution matrix is equal to the identity matrix. In reality this can not be achieved, but one can try to minimise the deviation of the resolution matrix from the identity matrix. A resolution kernel should therefore be correctly centred around the target element and have a value of 1. This can be quantified for each row k in the resolution matrix as:

$$\mathbf{SI}_k = \frac{(D - d_{ki})\mathbf{R}_{kk}}{D}, \quad (4-6)$$

with D the maximum distance between elements in the source space, and d_{ki} the Euclidean distance between element k and the element where the absolute maximum in the resolution kernel was found. \mathbf{R}_{kk} is the absolute value of the k^{th} diagonal value in the resolution matrix. A small value in \mathbf{SI} means that either the main peak is at the wrong location (giving rise to large contributions from other sources to the estimated source activity) or that the peak value is small (a source at that location is difficult to retrieve). A large value of \mathbf{SI}_k means that the activity in that source element can be correctly estimated, independently of the spatial blurring.

Bias in Dipole Localisation

The *Bias in Dipole Location* is defined as the Euclidean distance between the element in the source space where the maximum in the k^{th} column of \mathbf{R} is obtained and the ideal, k^{th} , element, normalised by the maximum distance:

$$\mathbf{BDL}_k = \frac{d_{ki}}{D}, \quad (4-7)$$

with D and d_{ki} as above.

A large value in **BDL** can be interpreted as the localisation error that can be expected if a unitary source is active in that particular source element.

Crosstalk

Liu and colleagues presented the use of the crosstalk metric for the analysis of a method to integrate MEG with fMRI data (Liu *et al*, 1998). They defined the crosstalk between two sources i and j as:

$$\xi_{ij}^2 = \frac{|(\mathbf{GL})_{ij}|^2}{|(\mathbf{GL})_{ii}|^2} * 100\% = \frac{|\mathbf{G}_i \cdot \mathbf{L}_j|^2}{|\mathbf{G}_i \cdot \mathbf{L}_i|^2} * 100\% , \quad (4-8)$$

It follows from equation 4-5 that this is the same as element j in the i^{th} row of the resolution matrix (which is the i^{th} resolution kernel), apart from weighting by the estimated activity from source i itself. Thus, the crosstalk quantifies the sensitivity of the estimated activity at location i to activity at location j , relative to activity at location i . A crosstalk of 0% means that the estimated activity at location i is completely insensitive to activity at location j , whereas a crosstalk of 100% means that the estimated activity at location i is equally sensitive to activity at locations i and j .

This measure was used to determine how strong fMRI based constraints should be to get optimal localisation from MEG data (optimal in the sense of minimum crosstalk between sources). Their approach was as follows:

- The source covariance matrix was used as a weighting matrix for their linear estimator. No correlation between sources was assumed, so all off diagonal elements were zero. A diagonal element corresponding to a source in the fMRI data was given a value of 1 and the other diagonal elements, corresponding to invisible fMRI sources, were given a value of 0, 0.01, 0.1 and 1. This corresponds to 100%, 99%, 90% and 0% fMRI weighting, respectively. A relative fMRI weighting of 0% means that no fMRI constraints are used and the linear estimator reduces to the minimum norm estimator.
- Visible fMRI sources were randomly selected and the crosstalk between the sources was computed. In order to get rid of bias effects, this was repeated several times, resulting in an average value for the crosstalk.

- Also, the average crosstalk was computed with several different relative fMRI weightings.
- This procedure was repeated to compute the average crosstalk between a visible fMRI source and an invisible fMRI source as well. Also, instead of point sources, sources with a larger spread were used (i.e. elements in the source covariance matrix corresponding to such an extended visible fMRI source were all given a value of 1).

The main conclusions of their simulations are that:

- Relative fMRI weighting provides a way to compromise between the ability to resolve visible sources (high relative fMRI weighting) and reducing the influence of invisible fMRI sources (low relative fMRI weighting) on estimated activity. The optimum weighting was found to be 90%.
- The crosstalk was smaller when orientation constraints were used for the sources.
- Errors in the orientation of the sources influenced the average crosstalk only little. Using orientation constraints, even when they were not really precise, gave more accurate source reconstructions than when no orientation constraints were used.
- The crosstalk was relatively independent of the number and extent of the fMRI sources. Thus, the reconstruction method that was used is a true multisource imaging method.

It should be noted that the optimum relative fMRI weighting of 90% depends on the shape of the resolution kernels and therefore on the linear estimator, source space, head model and sensor configuration that was used. It is expected that this value will vary for different measurement set-ups. Two approaches can be taken to use the crosstalk to determine the optimum weighting for the combination of measured fMRI and MEG data:

- 1) The average crosstalk between the sources within the fMRI data and randomly chosen invisible sources is computed. The optimum relative fMRI weighting is found by repeating this with several values for the weighting. Source localisation from the MEG data is then performed with the fMRI data as constraints, using the optimum relative weighting.
- 2) An initial source localisation is performed from the MEG data, without using any fMRI constraints. The localised sources that do not correspond to a fMRI source then act as invisible sources in the procedure described above. The advantage of using the 'MEG sources' instead of randomly chosen invisible sources is that the weights are tailored to give the lowest crosstalk based on both the fMRI and MEG data, instead of being only

based on the fMRI data. This approach is only beneficial if the initial MEG localisation is reasonably accurate.

A plot of the crosstalk kernel, the homologue of the resolution kernel, demonstrates the spatial spread of estimated activity. Liu and colleagues used these plots to demonstrate that the spatial resolution for sources on a gyrus is higher than within a sulcus. This was only demonstrated for one pair of sources though, and the lower spatial resolution of the source within the sulcus could be due to other factors, like its orientation. Further research is needed to verify whether this effect holds for other sources.

Image Sensitivity Maps

The diagonal entries of the resolution matrix provide information about the maximum performance of the linear estimator. Each diagonal element gives the peak response of the linear estimator for unit dipole activity at the corresponding element in the source space. It is therefore possible to create a colour coded map of the values of all diagonal elements. This reveals the sensitivity of the linear estimator to activity at different locations (Gençer *et al*, 1996). Hence the name image sensitivity map.

Gençer and colleagues computed image sensitivity maps for a simulated EEG set-up. The maps showed that image sensitivity was highest for activity on the gyri and a decreasing sensitivity for activity within the sulci.

Expected squared error due to other sources

Another measure that summarises the information in the resolution kernels was given by Dale and Sereno (Dale and Sereno, 1993). The expected squared error of the i^{th} element due to activity in all other elements is given by:

$$\varepsilon_{se} = \mathbf{R}_i \mathbf{K}_q \mathbf{R}_i^T, \quad (4-9)$$

with \mathbf{K}_q the source covariance matrix. The covariance matrix reduces to the identity matrix if no *a priori* information about the sources is used, and ε_{se} is simply obtained by

multiplying each resolution kernel with its transpose. The result is a single value for each element in the source space and a map, similar to the image sensitivity maps, will reveal the sensitivity of activity at a location to activity at other locations.

Influence of noise

So far, the potential influence of noise has not been considered in this section. If we consider noise, equation 4-3 becomes:

$$\mathbf{B} = \mathbf{L} \mathbf{Q} + \mathbf{N}, \quad (4-10)$$

with \mathbf{N} the matrix that contains the noise and has the same dimensions as \mathbf{B} .

Equation 4-5 then becomes:

$$\hat{\mathbf{Q}} = \mathbf{R} \mathbf{Q} + \mathbf{G} \mathbf{N}, \quad (4-11)$$

It is clear from this equation that the i^{th} row of \mathbf{G} specifies how much a unit of noise at each sensor contributes to the estimated activity of the i^{th} element in the source space (Dale and Sereno, 1993).

As for the *expected squared error due to other sources*, this information can be summarised in a single measure. The *Expected squared error due to noise* for each element is given by:

$$\boldsymbol{\varepsilon}_{ne} = \mathbf{G}_i \mathbf{K}_n \mathbf{G}_i^T, \quad (4-12)$$

with \mathbf{K}_n the noise covariance matrix. The covariance matrix reduces to the identity matrix if no *a priori* information about the noise is used, and the error measure is simply obtained by multiplying each row in \mathbf{G} with its transpose. The single value for each element in the source space can again be displayed as a colour coded map. Such a map reveals the sensitivity of activity at a location to noise in the measurements when linear estimator \mathbf{G} is used to localise the activity.

Another approach to look at the potential influence of noise is based on varying the number of basis functions that are used in the construction of the resolution matrix (Gençer

et al, 1996). Gençer and co-workers used a formula based on the estimated SNR that gives the number of basis functions that should be used. However, this formula will not always give the correct truncation level, as follows from the following example. Assume two situations with exactly equal SNR and that their formula orders to use the first 30 basis functions. Assume further that in case 1 the angle between the 30th spatial component and the noise vector is small and in the 2nd case the angle is close to 90°. Consequently, one should not use the 30th component (and smaller components) in the 1st case but can use the 30th component in the 2nd case. The potential influence of noise can therefore better be examined by means of ϵ_{nc} .

4.3 Usefulness of Figures of Merit in Conjunction With FOCUSS

4.3.1 Introduction

The figures of merit described in the previous section can be used to evaluate linear distributed source reconstruction algorithms. More specifically, they can be used to select the optimum algorithm for a specific inverse problem (e.g. Grave de Peralta Menendez *et al*, 1996) or evaluate the influence of the measurement set-up on the source reconstructions (e.g. Gençer *et al*, 1996). Grave de Peralta-Menendez and colleagues (Grave de Peralta Menendez *et al*, 1996, Grave de Peralta Menendez *et al*, 1997, Lütkenhöner and Grave de Peralta Menendez, 1997, Grave de Peralta-Menendez and Gonzales-Andino, 1998, Grave de Peralta Menendez and Gonzalez Andino, 1999) have strongly promoted the utilisation of these measures for the biomagnetic inverse problem. However, reports of their usage are still scarce. Moreover, their use is only rarely evaluated with a realistically shaped source space (Dale and Sereno, 1993, Gençer *et al*, 1996). Grave de Peralta-Menendez and colleagues only reported results of the use of these figures of merit in simulations with unrealistically shaped source spaces (i.e. 1 dimensional or formed from slices through spherically shaped source spaces). These simulations give valuable insights into the basic limitations and advantages of the linear estimators that were used. However, it is expected that the features of linear estimators vary when a realistically shaped source space is used in the source reconstructions. Therefore, in section 4.3.2 the figures of merit were evaluated with a realistically shaped source space.

The FOCUSS algorithm used in this work is a hybrid method that is neither a classic optimisation algorithm nor a standard linear estimator. The relevance of the figures of merit used here can therefore be questioned. However, a weighted minimum norm solution is taken as the initial guess of the algorithm. Thus, the figures of merit can at least be used to evaluate this initial guess. Moreover, the relationship between the figures of merit for the initial guess and the final solutions of the FOCUSS algorithm was examined in section 4.3.3.

4.3.2 Evaluation With a Realistically Shaped Source Space

4.3.2.1 Groundwork

Before the figures of merit are discussed, a description and analysis of the simulation framework is presented.

The figures of merit were analysed with the realistically shaped source space that was constructed in section 3.2.4. The maximum distance, D , between any pair of sources in the source space was determined to be 103mm. This distance was used for the computation of the *Source Identifiability* and *Bias in Dipole Location*. The sensor configuration, location and orientation was taken as in a typical visual evoked experiment with the 19-channel CES system at Aston University. This means a diameter of each coil of 15mm and the distances between the gradiometer coils from bottom to top: 1.5, 19.5, 5, 5, 5, 19.5 and 1.5 [mm], respectively. The gradiometers were all oriented perpendicular to the surface of the source space in order to measure the radial component of the magnetic field. The sensor array was positioned around the occipital pole of the left hemisphere.




figure 4-3: Configuration of the 19 channel CES system (ideal). The space between two sensors is 29mm.

The head was modelled as a sphere that fitted the skull around the back of the head. Only a single sphere model was used, since the influence of volume currents on the measured magnetic fields is known to be small in most cases (as was shown in section 1.3.4). With this configuration, the forward problem was solved for each dipole in the source space, having unitary strength. That is, the magnetic induction in the sensors was computed for each dipole.

This sensor response for a unitary dipole is equal to the gain or lead field of that dipole. Hence, the concatenation of the lead fields for all dipoles formed the lead field matrix, \mathbf{L} .

Using the formulation of the pseudo-inverse as used in the FOCUSS algorithm (Gorodnitsky *et al*, 1995), the estimated source strength can be written as:

$$\hat{\mathbf{Q}} = \mathbf{W} (\mathbf{L} \mathbf{W})^+ \mathbf{B}, \quad (4-13)$$

with \mathbf{W} the weighting matrix that contains *a priori* information. The resolution matrix can thus be written as:

$$\mathbf{R} = \mathbf{W} (\mathbf{L} \mathbf{W})^+ \mathbf{L}, \quad (4-14)$$

Initially, no *a priori* information was used, resulting in a weighting matrix that was equal to the identity matrix.

The figures of merit that were used resulted in a single value for each element in the source space. These values could thus be displayed as a colour coded map on the surface. Each triangle face was colour coded by the values of the figure of merit for the dipoles in the triangle's vertices. These colour coded maps are referred to as **surface maps**.

Magnetoencephalography is blind to radial sources in a spherically symmetric volume conductor. It is therefore frequently stated that it is nearly impossible to reconstruct activity in gyri from MEG recordings, implicitly assuming that all gyral sources have a radial orientation. Using the orientation of the sources in the realistically shaped model of




figure 4-4: An MRI slice in which the sources from the realistically shape source space are overlaid. Each source was manually classified as either being gyral (red), sulcal (yellow) or 'undefined' (cyan).

the cortex and a classification of the sources, it was possible to examine this hypothesis. For this purpose, sources were manually classified as being either a gyral or sulcal source. Sources that were not clearly on a gyrus or on a sulcal wall were classified as 'undefined'. Classifying only sources that were unambiguously gyral or sulcal validated the comparison of the properties of the two sets of sources. In figure 4-4 an example of an MRI slice in which the sources were classified is shown.

After this classification, the angle between the orientation of each source and the radius of the spherical




figure 4-5: Histogram of the angles between the sphere radius and the orientations of gyral (left panel) and sulcal (right panel) sources.

model was computed. Thus, each source was now classified and the angle it formed with the radius was known. A histogram of the angles for the sulcal and gyral sources shows the distribution of angles for both classes (figure 4-5). The distribution for the sulcal sources approached a normal distribution, with the mean value being 90° . Hence, the majority of sulcal sources had a tangential, or nearly tangential, orientation. Consequently, activity produced by the majority of these sources will be picked up in an MEG recording, provided the distance to the sensors is sufficiently small. The histogram for the gyral sources shows a large number of sources that were nearly radial. Consequently, activity produced by these sources will not produce a significant MEG signal. However, there was still a notable number of sources that had an orientation that was not radial (the mean angle was 57° for the gyral sources). This is explained by the fact that there is such a large amount of cortex folded into the sulci; two thirds of the total cerebral cortex is buried within the cortical folds (Van Essen and Drury, 1997, Le Goualher *et al*, 1999). Were the brain to consist of no sulci at all, then the orientation of the sources would always be radial. The more sulci there are, the more the gyri are 'pulled' inwards, with the effect that the angle with the sphere radius increases. These non-radial gyral sources will produce a measurable magnetic field outside the head, hence MEG is not blind to all gyral sources (see also Leahy *et al.*, 1998).

The histograms of the distribution of angles for the gyral and sulcal sources suggested that the ability to localise activity for both classes of sources differed. This should be reflected in the resolution matrix and the figures of merit. It was therefore examined whether these measures demonstrated significant differences for sulcal and gyral sources in the succeeding evaluation of these measures.

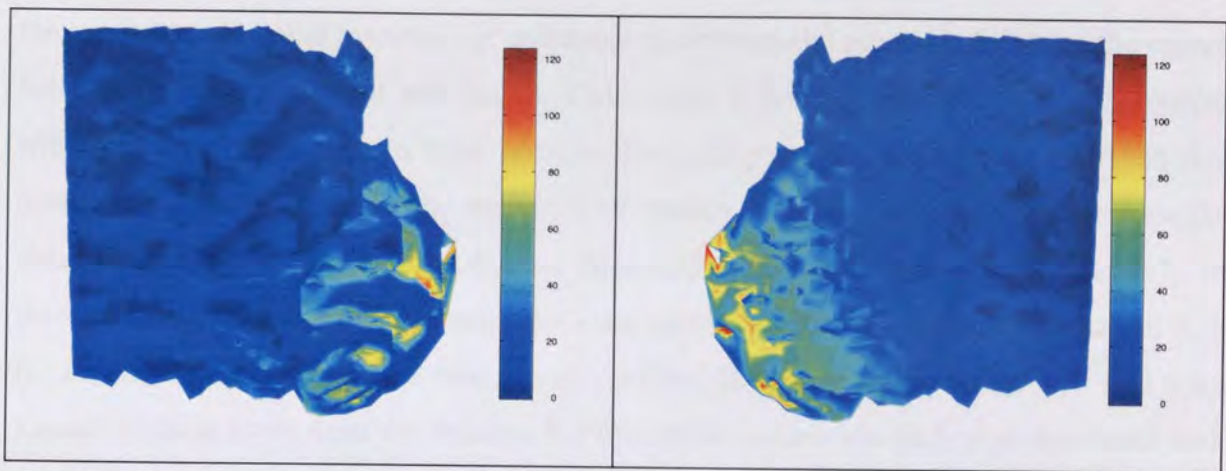


figure 4-6: Surface map of the l_2 -norm of the lead field of each element in the source space, viewed from the left (left panel) and right (right panel). The scale is in fTesla.

l_2 -norm of the lead fields

A source with a 'large' lead field (the l_2 -norm is large) will be easier to reconstruct than a source with a smaller lead field. This can be seen as follows: a source with large lead field but small source strength can still have a significant contribution to the total measured magnetic field, whereas the magnetic field produced by a source with small lead field and small source strength will be relatively small. The ability to reconstruct the activity of a certain source from measured data is therefore to some extent determined by the norm of its lead field. Therefore, the characteristics of the current set-up were first explored by examining the lead field matrix itself.

In figure 4-6 the l_2 -norm of the lead field of each element in the source space is displayed as a surface map. The norm of the lead field was largest for elements closest to

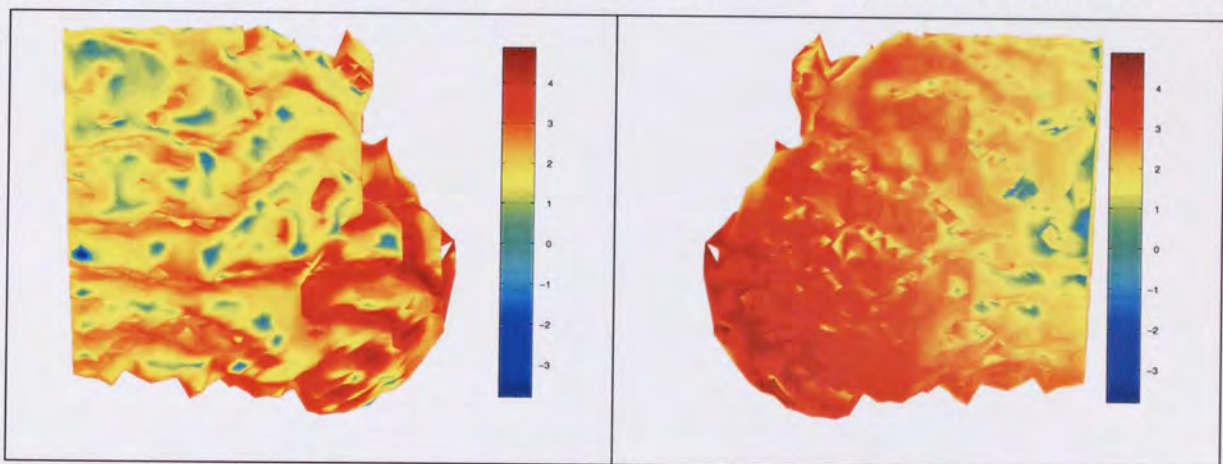


figure 4-7: Surface map of the l_2 -norm of the lead field of each element in the source space, viewed from left (left panel) and right (right panel). The surface maps are displayed on a logarithmic scale.

the sensors as expected from the $1/r^2$ relationship between the magnetic field and the source location in the law of Biot and Savart. There were a few superficial, tangential, sources with very large $\|\mathbf{L}\|_2$. Due to these sources the scaling in the surface map was such that details in the surface map for the sources with smaller $\|\mathbf{L}\|_2$ were not visible. Therefore, the same surface maps as in figure 4-6 were displayed on a logarithmic scale (figure 4-7). In these later surface maps the gyri and sulci were clearly visible due to the differences in $\|\mathbf{L}\|_2$ for sources in these areas. The source with smallest $\|\mathbf{L}\|_2$ were found on the gyri that were located furthest away from the sensors. For the sulcal sources the $\|\mathbf{L}\|_2$ also decreased with distance, but was larger than for a gyral source at similar distance from the sensors. It is noteworthy that even for the gyri there were only a few 'spots' from which it would be impossible to record activity.

In the following section the influence of the differences in $\|\mathbf{L}\|_2$ on the figures of merit for linear inverse algorithms was investigated.

4.3.2.2 Figures of Merit

In this section the figures of merit described in section 4.2 were evaluated. The analysis started with an examination of the resolution kernels and the columns of the resolution matrix. This was cumbersome due to the large number of resolution kernels and columns in the resolution matrix. The inspection of the resolution matrix was therefore followed by an examination of the summarising measures described in 4.2.2, the *Source Identifiability*, *Bias in Dipole Location*, *Image Sensitivity*, *Expected squared error due to other sources* and *Expected squared error due to noise*. The *Crosstalk* measure was not evaluated, since the *Crosstalk* between a source and all other sources is the same as each value in this source's resolution kernel, apart from a scaling factor.

The resolution matrices and figures of merit for the linear estimator were examined for two different weighting matrices. By doing so a distinction could be made between properties of the figures of merit that could be attributed to the linear estimator and the specific set-up or that were due to *a priori* information that was used.

The first weighting matrix used was the identity matrix. With this weighting matrix the linear estimator reduced to the minimum norm estimator. The second weighting matrix that was used counteracted the bias towards the sensors that is characteristic for the minimum norm estimator (see section 2.3.2). Each diagonal element in the weighting matrix was formed by the inverse of the l_2 -norm of the lead field of the corresponding source. The weighting for completely radial sources was set to the mean value for the non-radial sources because otherwise radial sources would get infinite weighting. The weighting matrix thus far obtained also 'punished' tangential sources, since tangential sources have a relatively large $\|\mathbf{L}\|_2$. In order to avoid this, the weightings were sorted according to the average distance to the sensors of the corresponding elements in the source space and a running average of the weightings was taken. These new values were used in the weighting matrix, from now on referred to as **smoothed norm weighting**.

Resolution kernels, using minimum norm weighting

In this case, the weighting matrix was set equal to the identity matrix. This means that no *a priori* information was used and the solution in equation 4-13 reduced to the minimum norm solution. Consequently, the resolution matrix reduced to the one for the

minimum norm solution as well. The resolution matrix for the minimum norm solution is symmetric, hence the resolution kernels were equal to the corresponding columns in the resolution matrix. Also, each column of the resolution matrix can be seen as the impulse response of the linear estimator to the signal induced by unitary source in the corresponding element in the source space. Combining these two facts tells us that each resolution kernel was equal to the minimum norm solution for unitary activity in the corresponding element in the source space.

In figure 4-8 the resolution kernels typically found for a superficial, shallow and deep source are displayed. In figure 4-9 the resolution kernels for the same sources as in figure 4-8 are given, sorted for the distance between each source and the target source.

The surface map for the resolution kernel of the superficial source was close to the ideal one. That is, the values in the resolution kernel were large for sources near the superficial source and reduced in amplitude for more distant sources. This was also illustrated in figure 4-9. For the shallow source the resolution kernel was worse. There were still some elements with a large value in the resolution kernel near the location of the shallow source, but there were also large side lobes at locations far from the target source. These side lobes were found in areas at equal or smaller average distance from the sensors. The resolution kernel for the deep source was the worst case scenario. The reconstructed activity for this deep source was influenced by activity in many places in the source space (figure 4-8). Moreover, figure 4-9 shows that the values in the resolution kernel around the deep source were small and that the main peak was far from this source. Again, this main peak was shifted towards the sensors.

These resolution kernels point out the main drawback associated with the minimum norm estimator. The minimum norm estimator is biased to reconstruct sources too close to the sensors. This was reflected in the resolution kernels; large side lobes were found close to the sensors. Additionally, the finding that the reconstructed activity for the deep source was influenced by activity in many areas of the source space was in agreement with reports that state that minimum norm solutions are more distributed for deep sources than for superficial sources (Wang *et al*, 1992, Wang, 1993).

An interesting observation from the resolution kernels of the superficial and shallow source was that the areas with negative values were close to the target sources. This means

that this estimator might have problems with the reconstruction of activity from these target sources when nearby sources are simultaneously active. Due to the negative sidelobes in the resolution kernels, activity from the nearby sources could reduce or even cancel out the reconstructed activity from the target source, as was also illustrated in figure 4-2.

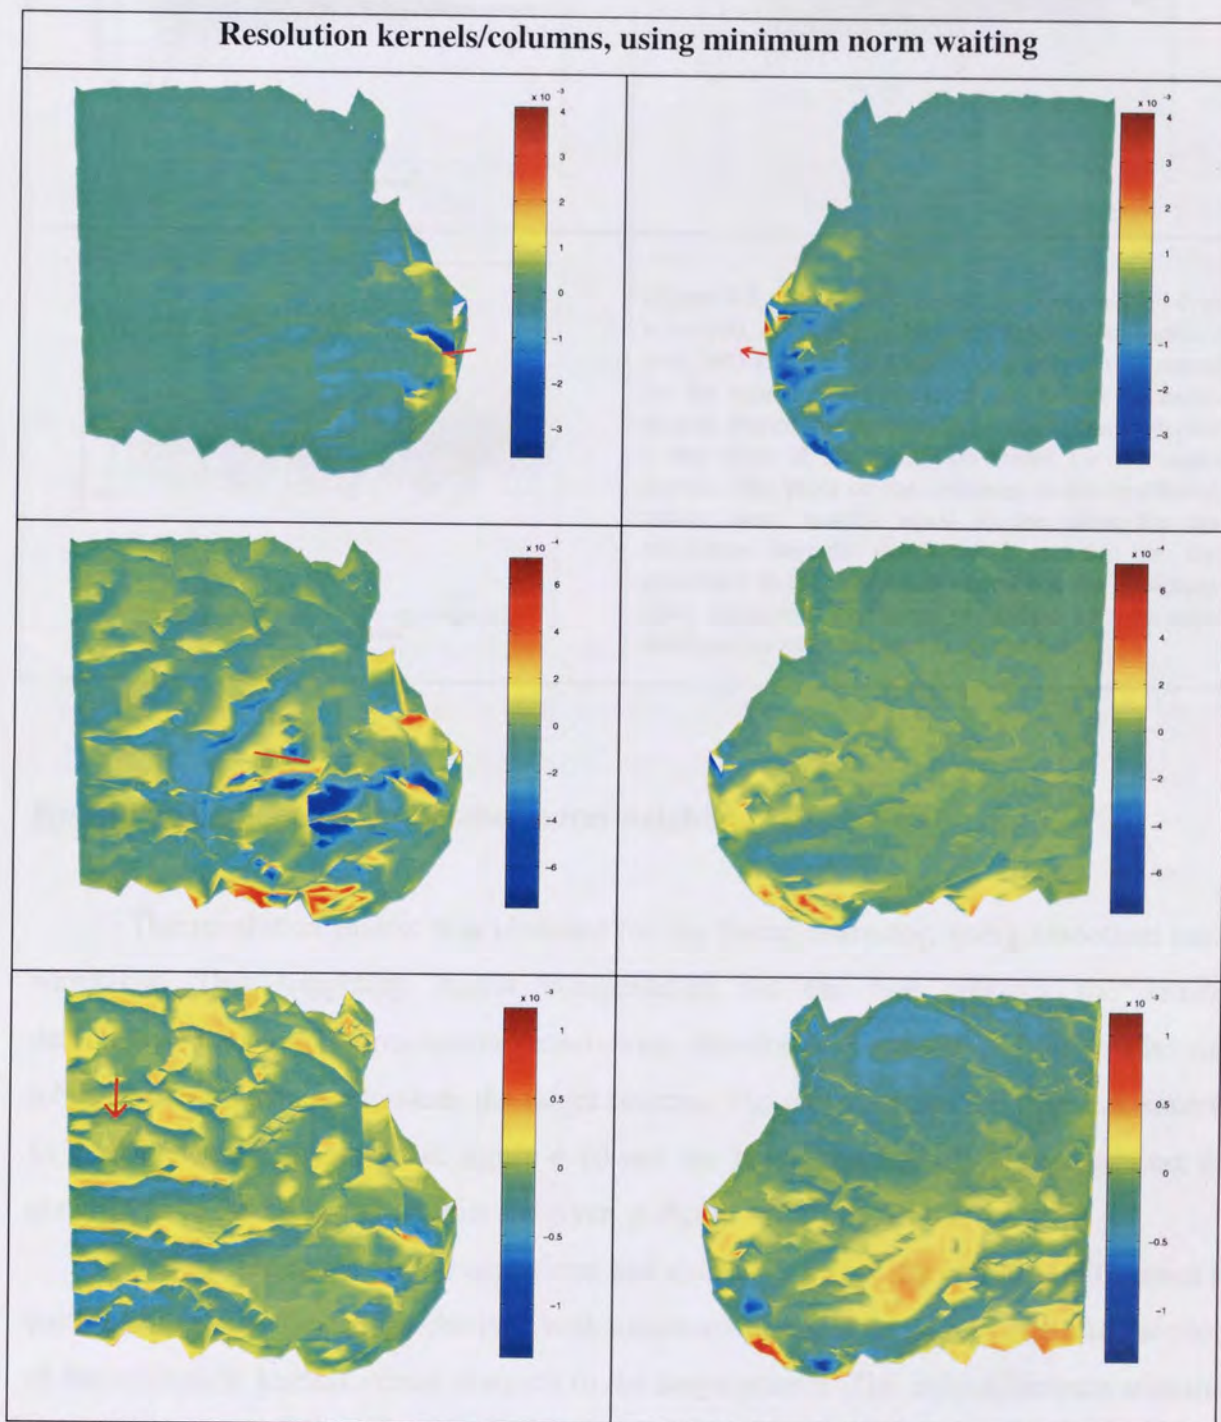



figure 4-8: Surface maps of the resolution kernels for a superficial (top row), shallow (middle row) and deep (bottom row) source. The left column contains views from the left and the right column views from the right. The location of the source for which the resolution kernel was computed is displayed as a red arrow. The surface maps for the columns in the resolution matrix were exactly equal to the resolution kernels displayed here, due to the symmetry in the resolution matrix for the minimum norm estimator. Note that the resolution kernel was closest to the ideal one for the superficial source and was worst for the deep source.



Resolution kernels, using smoothed norm weighting

The resolution matrix was obtained for the linear estimator, using smoothed norm weighting. This weighting matrix compensated for the bias towards the sensors demonstrated above. The resolution kernels were therefore expected to exhibit smaller side lobes, shifted in location towards the target sources. The surface maps for the same sources as used above are displayed in figure 4-10 and the resolution kernels plotted against the distance from the source of interest are given in figure 4-11.

The surface maps for the superficial and shallow sources showed little difference in pattern with the surface maps obtained with minimum norm weighting, neither did the plots of the resolution kernels versus distance to the target source. The only difference was that the plots for smoothed norm weighting were 'smoother' than the corresponding plots for minimum norm weighting. Also, the mean value of the resolution kernels had decreased due to smoothed norm weighting (from $4.5 \cdot 10^{-6} \pm 3.1 \cdot 10^{-4}$ to $4.0 \cdot 10^{-6} \pm 2.4 \cdot 10^{-4}$ for the superficial source and from $2.5 \cdot 10^{-6} \pm 9.9 \cdot 10^{-5}$ to $7.0 \cdot 10^{-7} \pm 3.4 \cdot 10^{-5}$ for the shallow source).

The surface map for the deep source exhibited less areas that had large values but the mean value had not changed ($-3.5 \cdot 10^{-6} \pm 2.1 \cdot 10^{-4}$ for the smoothed weighting compared to $-3.5 \cdot 10^{-6} \pm 1.6 \cdot 10^{-4}$). For the deep source, the plot of the resolution kernel versus distance to the target source showed a side lobe that was not there when minimum norm weighting was used. However, this side lobe was still far from the deep source and the main peak was still found close to the sensors.

The columns of the resolution matrix were different from the rows for the case of smoothed norm weighting. The main features of the resolution columns were however similar to those of the resolution kernels and the maps of the resolution columns were therefore not displayed here. The deeper sources had peaks in the resolution columns at large distances from the target source. Shallow sources also showed peaks away from the target source. The peaks in the resolution columns of the superficial sources were near the target source and in general there was less spread in the resolution columns for superficial sources than in the resolution columns for the shallow and deep sources.

The use of a different weighting matrix influenced the overall performance of the linear estimator, but the main patterns in the resolution kernels were little influenced by the choice of weighting matrix. Additionally, using smoothed norm weighting instead of minimum norm weighting had the largest effect on the resolution kernels of deep sources.

Figures of merit, using minimum norm weighting

The patterns in the surface maps of the *Source Identifiability*, *Image Sensitivity* and *Expected squared error due to other sources* were very similar due to the symmetry of the resolution matrix when the identity matrix was used as a weighting matrix. Therefore, only the results for the *Source Identifiability* are given here. Also, in cases where there were a few elements with a much larger value than the majority of elements in the source space, the surface maps are (only) displayed on a logarithmic scale in order to reveal more detail in the maps. The results for the *Source Identifiability*, *Bias in Dipole Location* and *Expected squared error due to noise* are given in figure 4-12, figure 4-13 and figure 4-14 respectively.

The surface map of the *Source Identifiability* (figure 4-12) showed that superficial sources were easiest to identify, i.e. the reconstructed activity for superficial sources was hardly influenced by activity from other sources. This result was confirmed by the plot of the *Source Identifiability* versus average sensor distance, showing that the *Source Identifiability* was largest when the average distance to the sensors was smaller than approximately 70mm. The surface map further indicated that gyral sources had lower *Source Identifiability* than sulcal sources. The plot showing the relationship between the *Source Identifiability* and the angle with the sphere radius further illustrated this. This plot showed that it was impossible to get a good *Source Identifiability* if a source was largely radial and that the maximum *Source Identifiability* was found for tangential sources (angle with sphere radius is 90°). The *Source Identifiability* for sulcal sources was in general higher than for gyral sources, since the majority of the tangential sources were sulcal sources and the majority of the radial sources were gyral (see figure 4-5). It is noteworthy that also for tangential sources small values for the *Source Identifiability* were found. These tangential sources were probably located far from the sensors, so that the *Source Identifiability* was small due to the large distance between the source and the sensors.

The surface map of the *Bias in Dipole Location* (figure 4-13) exhibits one striking difference with the surface map of the *Source Identifiability* in that there was no distinction in performance between gyral and sulcal sources. This was also demonstrated by the random relationship of the *Bias in Dipole Location* with the angle with the sphere radius (correlation of 0.0). This means that the *Bias in Dipole Location* was independent of the orientation that the source made with the sphere radius.

The relationship between the *Bias in Dipole Location* and the average distance to the sensors was almost linear (correlation of 0.9), with the superficial sources showing the lowest *Bias in Dipole Location*.

Sources with a large $\|\mathbf{L}\|_2$ had a low *Bias in Dipole Location*, whereas when the $\|\mathbf{L}\|_2$ decreased there was more spread in the *Bias in Dipole Location*, i.e. some sources with small $\|\mathbf{L}\|_2$ had a large *Bias in Dipole Location* and some had small *Bias in Dipole Location*.

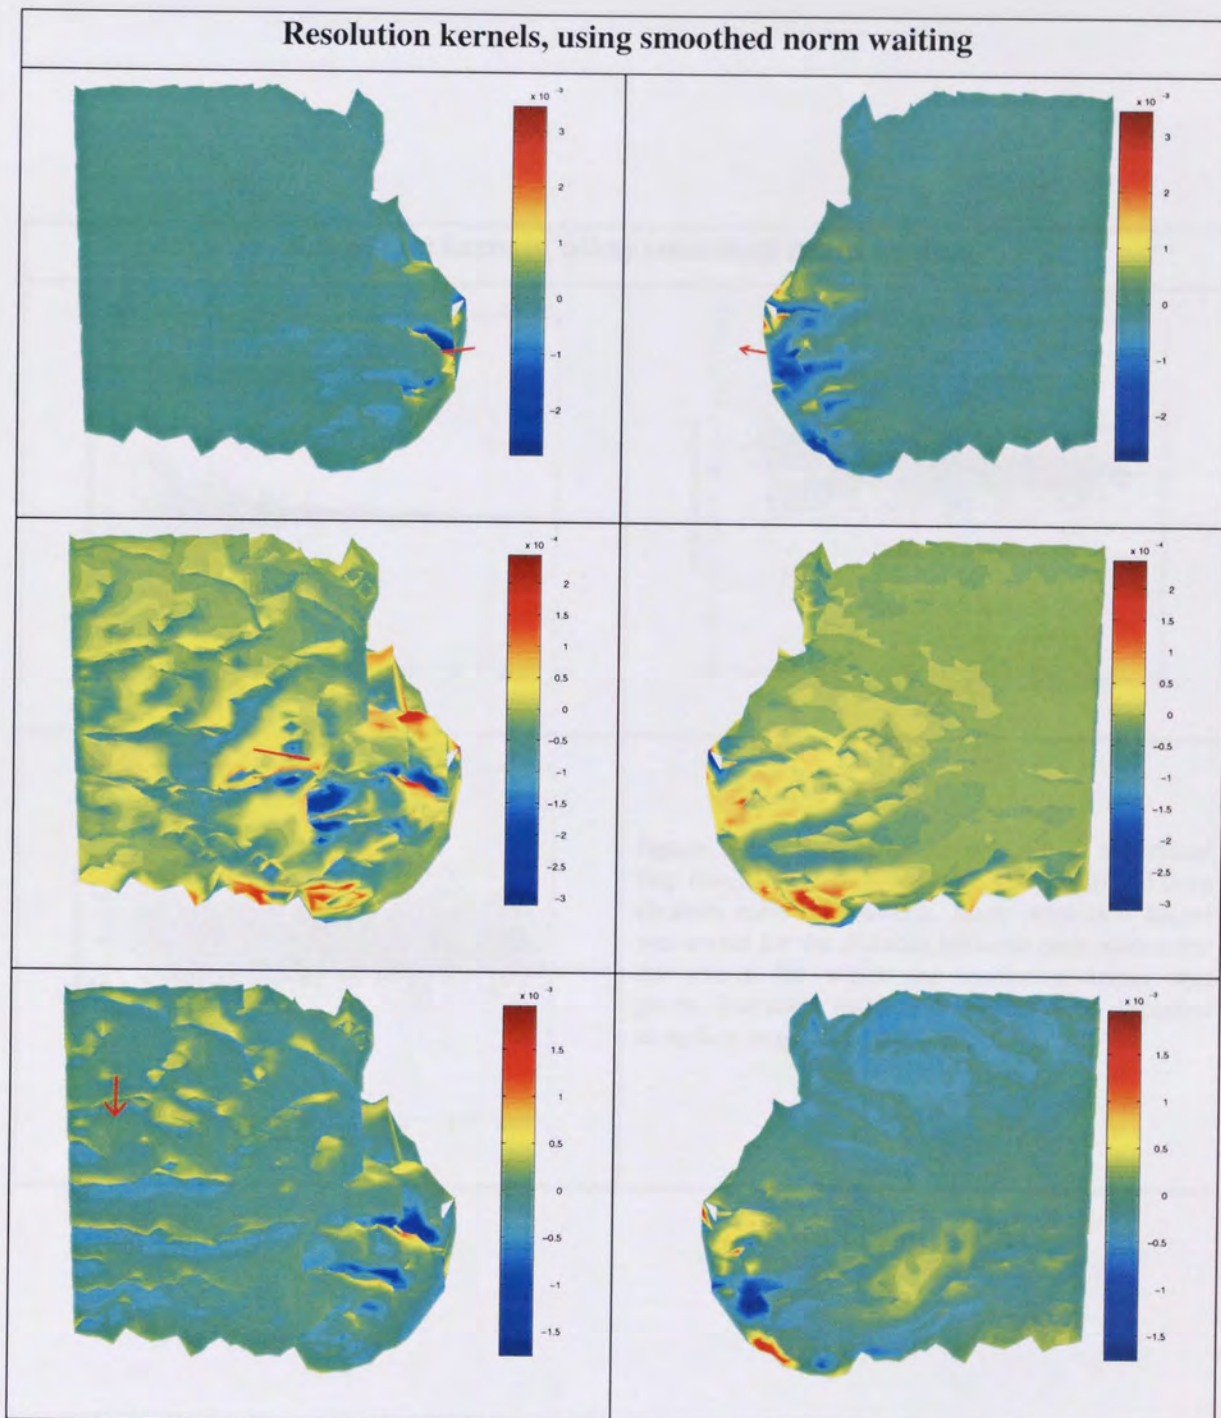


figure 4-10: Surface maps of the resolution kernels for a superficial (top row), shallow (middle row) and deep (bottom row) source. The left column contains views from the left and the right column views from the right. The location of the source for which the resolution kernel was computed is displayed as a red arrow.

Resolution kernels, using smoothed norm waiting

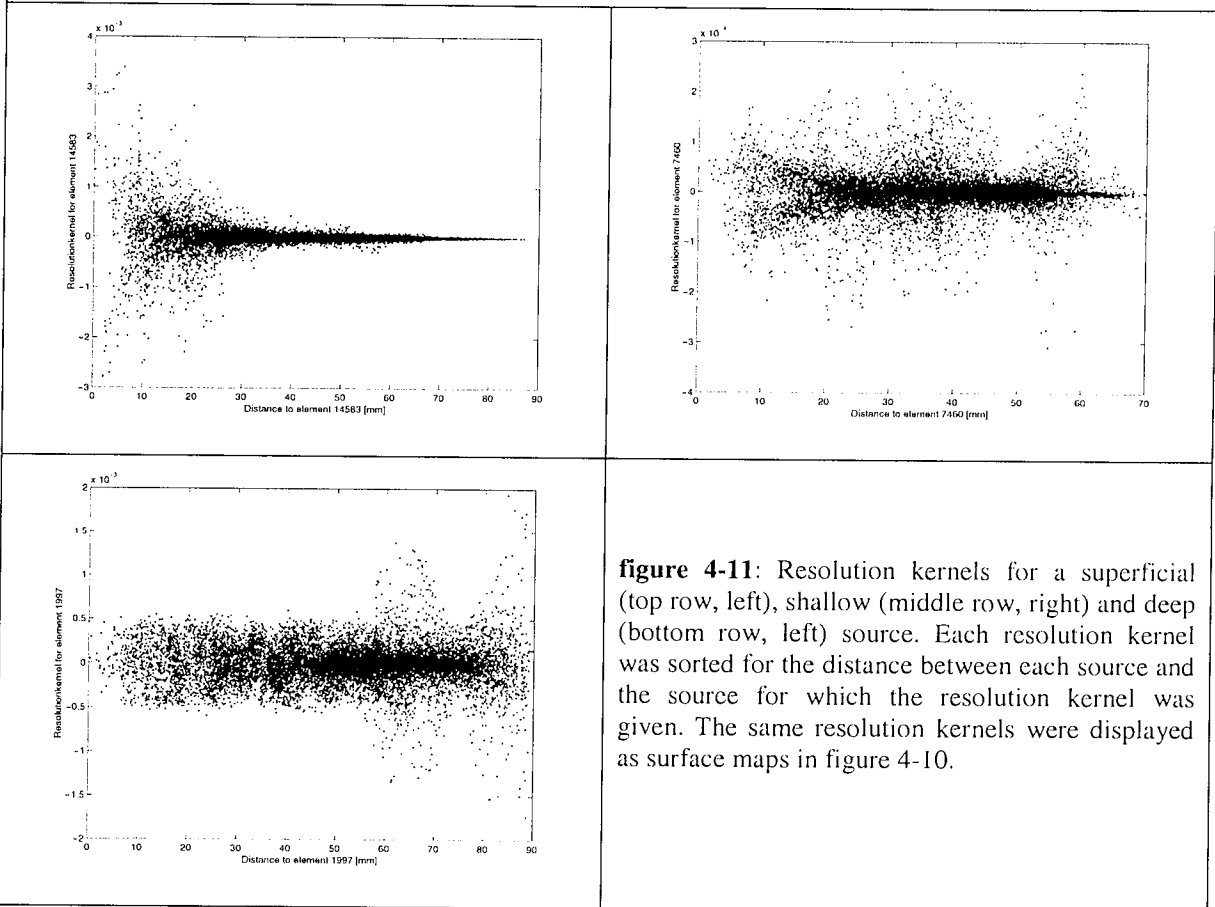


figure 4-11: Resolution kernels for a superficial (top row, left), shallow (middle row, right) and deep (bottom row, left) source. Each resolution kernel was sorted for the distance between each source and the source for which the resolution kernel was given. The same resolution kernels were displayed as surface maps in figure 4-10.

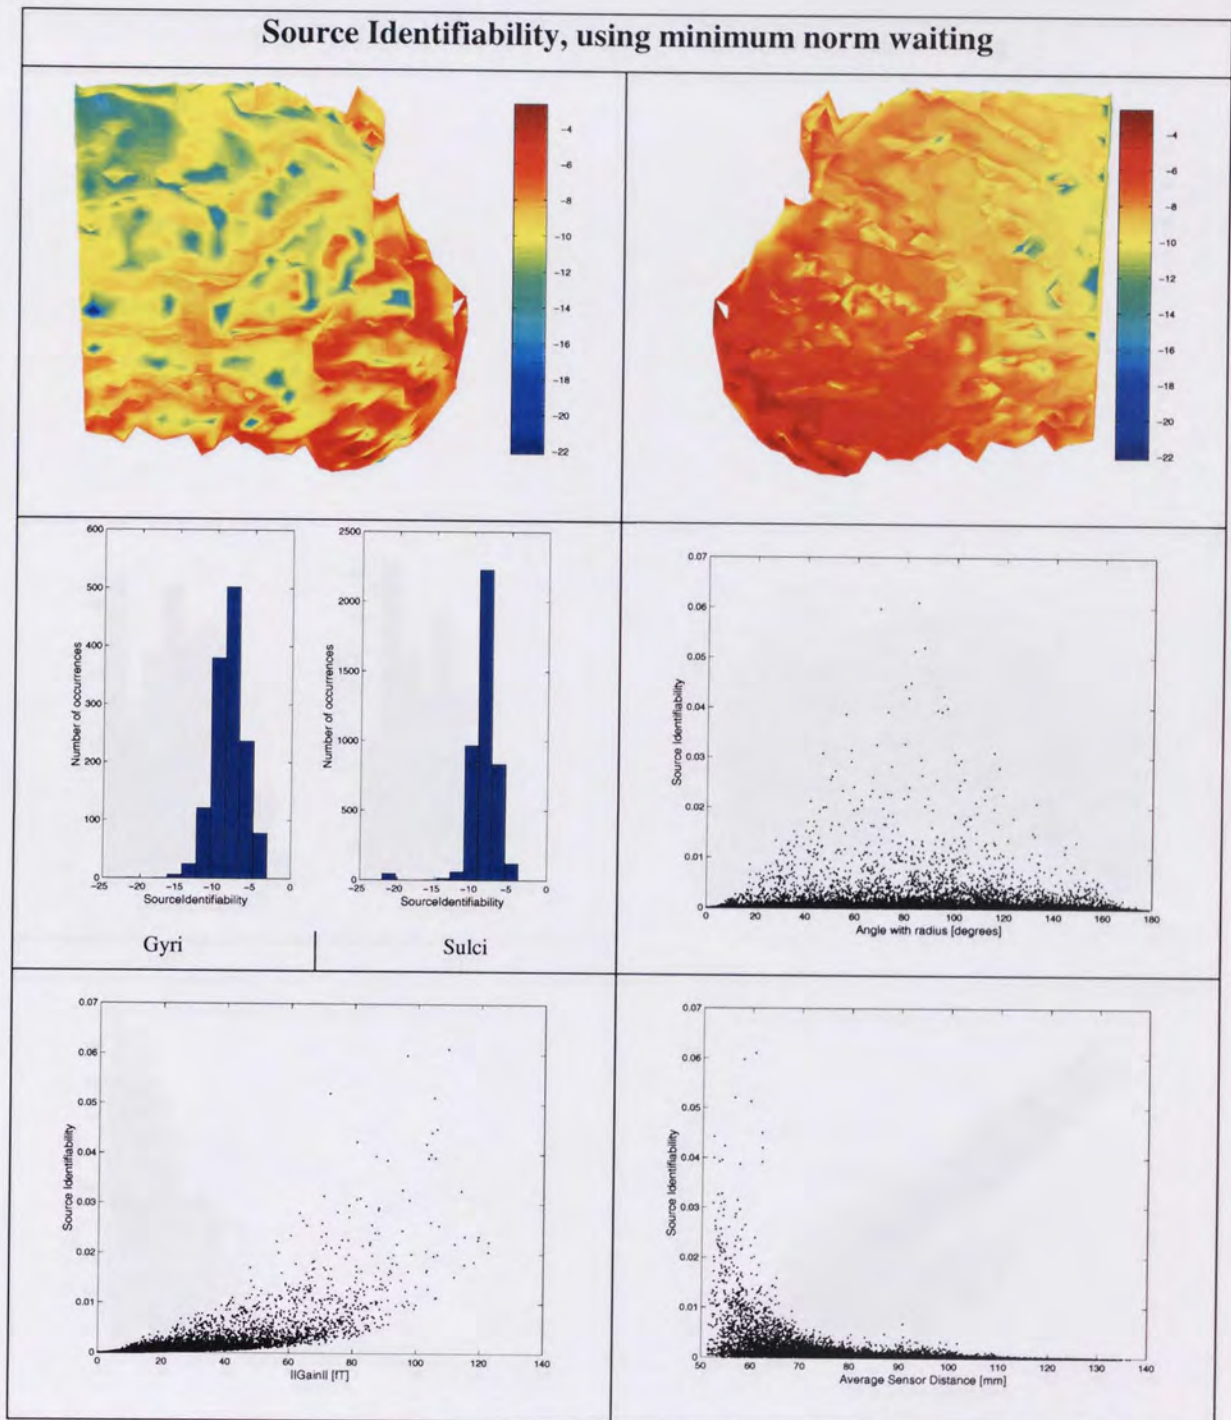


figure 4-12: The top row contains the surface map (logarithmic scale) of the *Source Identifiability* viewed from the left (left panel) and the right (right panel). In the middle row a histogram of the logarithm of *Source Identifiability* is given for both gyral (left in left panel) and sulcal sources (right in left panel) and a plot of the *Source Identifiability* versus the angle that the source orientation formed with the sphere radius (right panel). In the bottom row, the left and right panel show the *Source Identifiability* versus the l_2 -norm of the lead field (Gain) and the average distance to the sensors from each element in the source space, respectively.

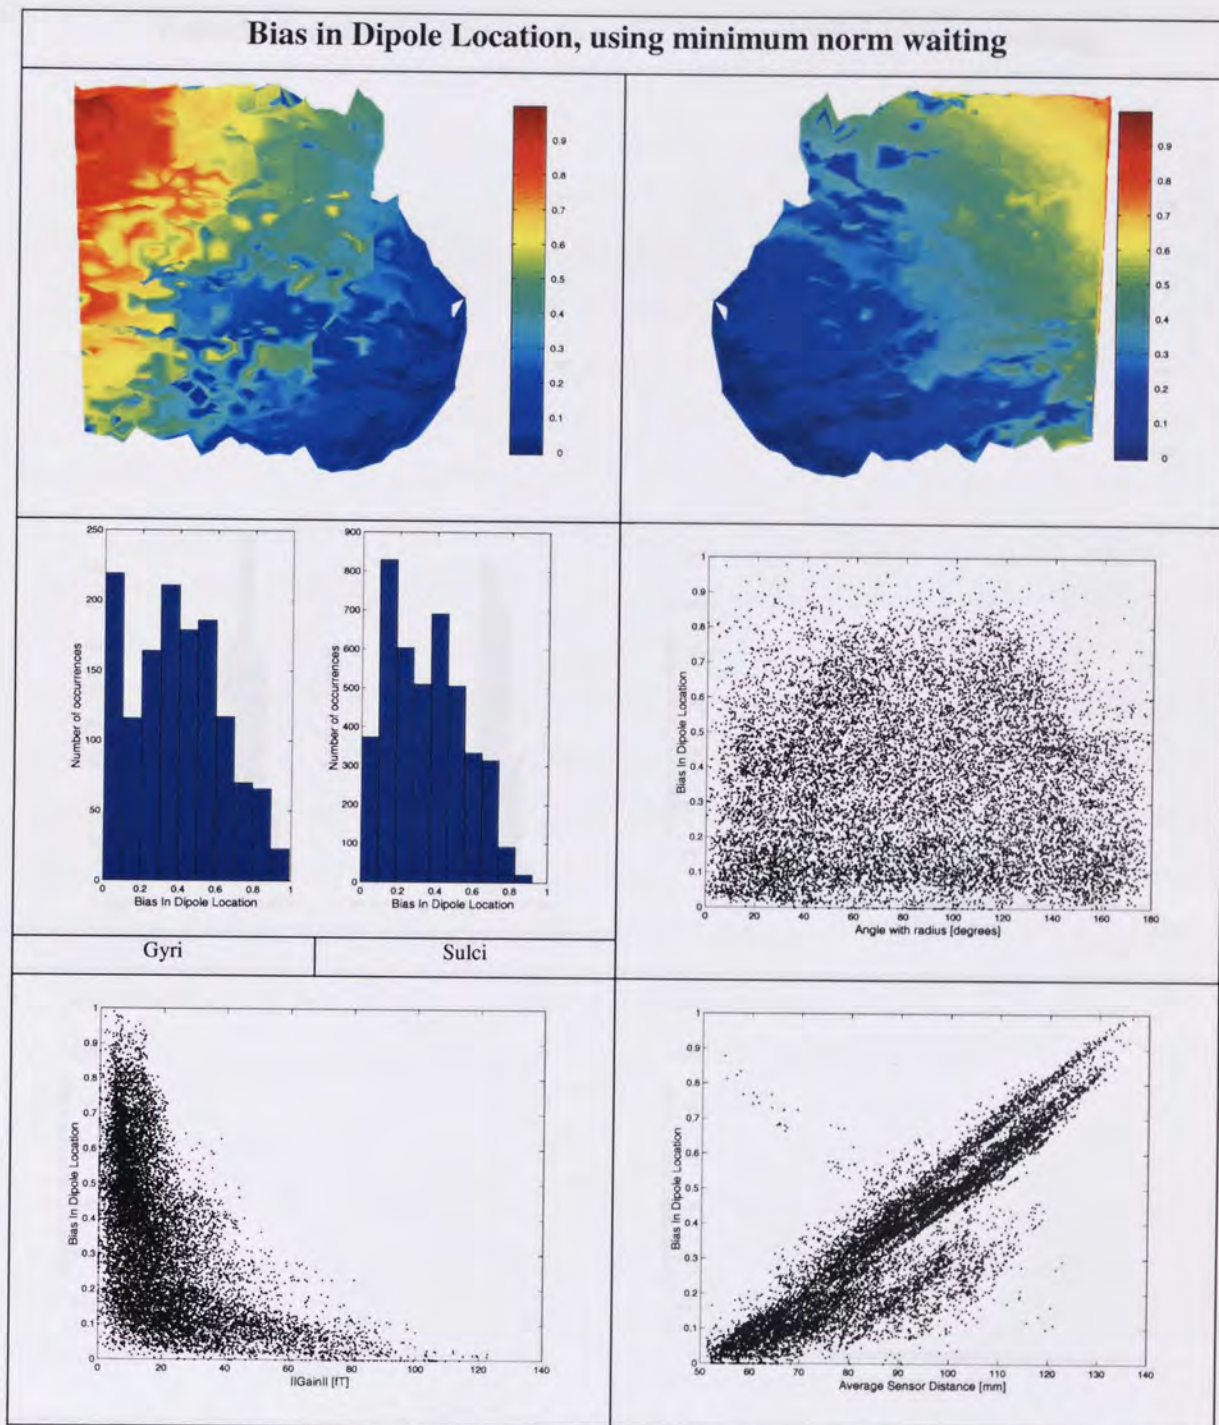


figure 4-13: The top row contains the surface map of the *Bias in Dipole Location* viewed from the left (left panel) and the right (right panel). In the middle row a histogram of the *Bias in Dipole Location* is given for both gyral (left in left panel) and sulcal sources (right in left panel) and a plot of the *Bias in Dipole Location* versus the angle that the source orientation formed with the sphere radius (right panel). In the bottom row, the left and right panel show the *Bias in Dipole Location* versus the l_2 -norm of the lead field (Gain) and the average distance to the sensors from each element in the source space, respectively.

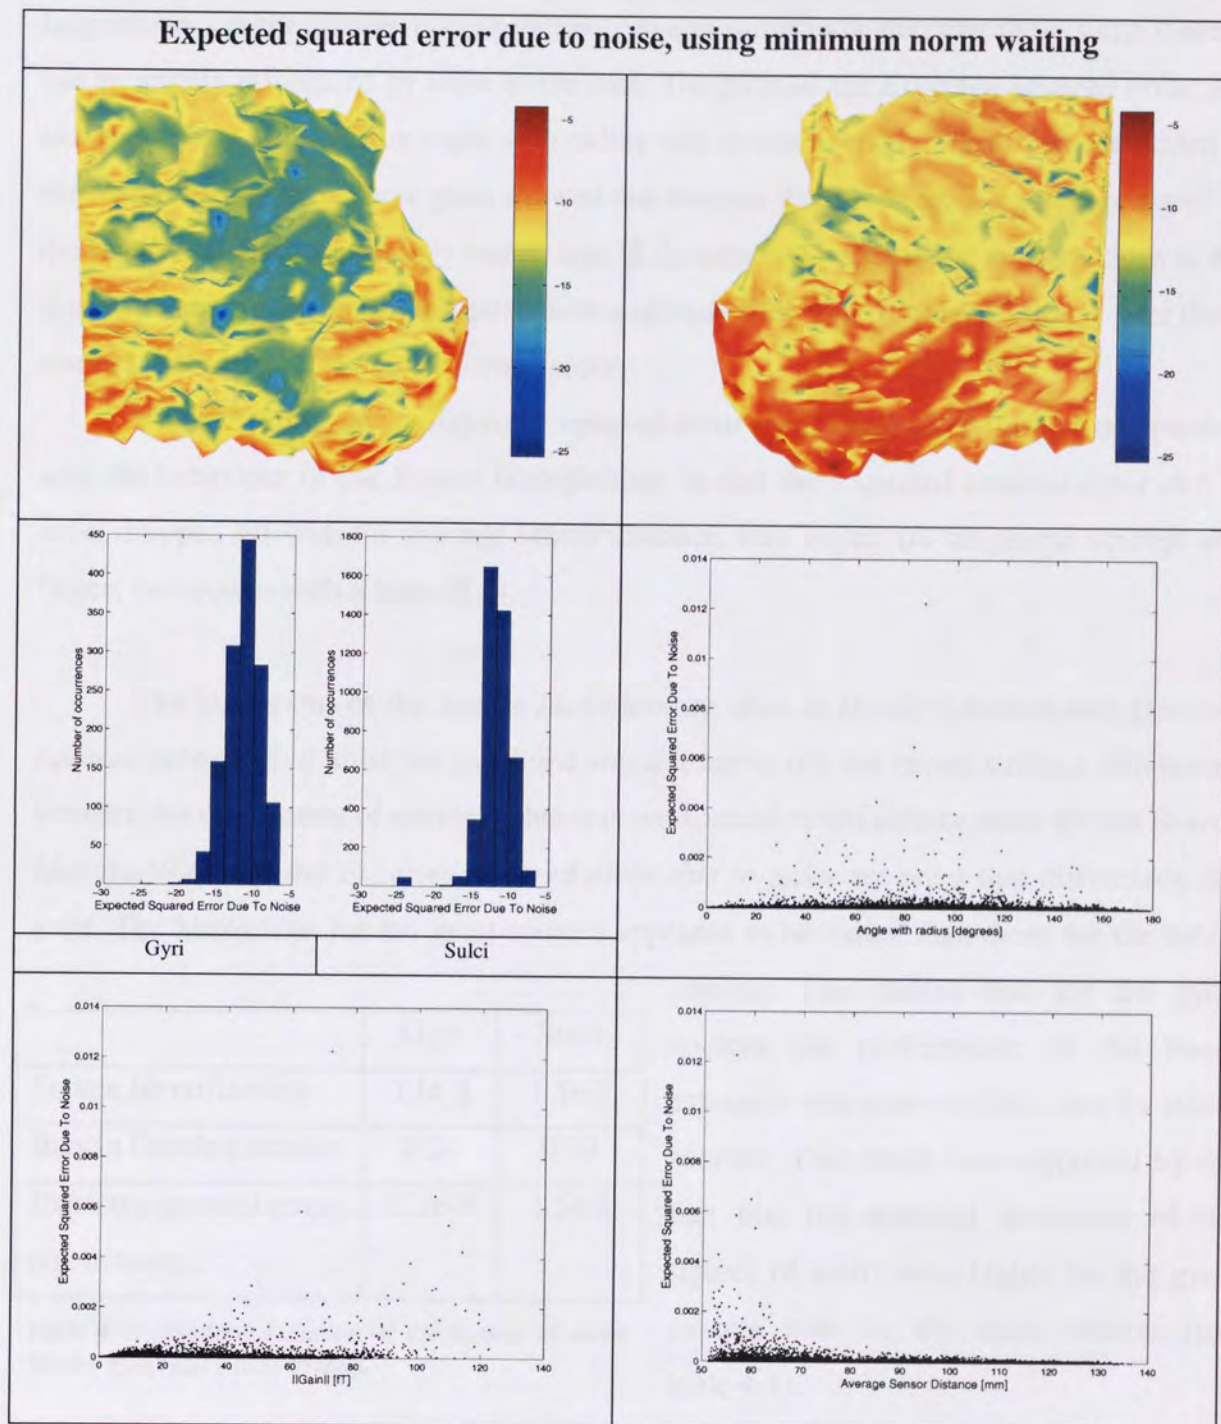


figure 4-14: The top row contains the surface map (logarithmic scale) of the *Expected squared error due to noise* viewed from the left (left panel) and the right (right panel). In the middle row a histogram of the logarithm of the *Expected squared error due to noise* is given for both gyral (left in left panel) and sulcal sources (right in left panel) and a plot of the *Expected squared error due to noise* versus the angle that the source orientation formed with the sphere radius (right panel). In the bottom row, the left and right panel show the *Expected squared error due to noise* versus the l_2 -norm of the lead field (Gain) and the average distance to the sensors from each element in the source space, respectively.

The surface maps for the *Expected squared error due to noise* (figure 4-14) and the *Source Identifiability* (figure 4-12) showed very similar patterns. This was an undesirable result; sources with large *Source Identifiability* were little influenced by activity from other sources hence activity at these locations can be accurately reconstructed. However, the large values for the *Expected squared error due to noise* indicated that these same sources can be greatly influenced by noise in the data. The plots of the *Expected squared error due to noise* versus $\|\mathbf{L}\|_2$, versus angle with radius and versus average sensor distance hinted at the reason behind this. These plots showed that sources that produced a significant signal in the sensors, either because they had a large $\|\mathbf{L}\|_2$, tangential orientation or were close to the sensors, were influenced by the noise. It was already shown in figure 4-12 that it were these sources that had a large *Source Identifiability*.

The behaviour of the *Expected squared error due to noise* was further comparable with the behaviour of the *Source Identifiability* in that the *Expected squared error due to noise* dropped off with the average sensor distance, was largest for tangential sources and largest for sources with a large $\|\mathbf{L}\|_2$.

The histograms of the *Source Identifiability*, *Bias in Dipole Location* and *Expected squared error due to noise* for gyral and sulcal sources did not reveal striking differences between the two classes of sources. This was unexpected as the surface maps for the *Source Identifiability* and the *Expected squared error due to noise* revealed that differences did exist. The histograms for the gyral sources appeared to be 'fatter' than those for the sulcal

	Gyri	Sulci
Source Identifiability	3.1e-3	1.5e-3
Bias in Dipole Location	0.24	0.20
Expected squared error due to noise	3.3e-3	1.5e-3

table 4-1: Standard deviation of the figures of merit for the gyral and sulcal sources.

sources. This means that for the gyral sources the performance of the linear estimator was more variable than for sulcal sources. This result was supported by the fact that the standard deviations of the figures of merit were higher for the gyral sources than for the sulcal sources (see table 4-1).

To briefly summarise the main findings in case the minimum norm weighting was used with the linear estimator:

- The *Source Identifiability* and the *Expected squared error due to noise* of a source were determined by a combination of the $\|\mathbf{L}\|_2$, the average distance to the sensors and the

orientation of the source. The *Source Identifiability* for sulcal sources was in the mean better than for gyral sources.

- The *Bias in Dipole Location* was independent of the orientation of the sources and consequently there was no difference in *Bias in Dipole Location* between gyral and sulcal sources.
- The *Bias in Dipole Location* did depend on the $\|\mathbf{L}\|_2$ and the average distance to the sensors.

Figures of merit, using smoothed norm weighting

The above general conclusions were validated for the case when the linear estimator was used with smoothed norm weighting (i.e. when different *a priori* information was utilised by the linear estimator).

The patterns in the surface maps of the *Source Identifiability* and *Image Sensitivity* were again very similar. The map for the *Expected squared error due to other sources* was also similar to the *Source Identifiability*. Because of the similarity with the map of the *Source Identifiability* the surface maps of the *Image Sensitivity* and the *Expected squared error due to other sources* were not displayed. The surface maps of the *Source Identifiability*, *Bias in Dipole Location*, and *Expected squared error due to noise* are given in figure 4-15. Additionally, in figure 4-16 only the plots of the *Bias in Dipole Location* and *Expected squared error due to noise* versus $\|\mathbf{L}\|_2$ and the *Bias in Dipole Location* versus the average sensor distance and the angle with the radius are given. These figures showed striking differences with the plots for the minimum norm weighting.

The surface map of the *Source Identifiability* (figure 4-15) was very similar to the one obtained with minimum norm weighting. Again, sulcal sources had larger *Source Identifiability* than gyral sources and the *Source Identifiability* fell off with the distance to the sensors. The pattern in the surface map of the *Bias in Dipole Location* (figure 4-15) showed that using smoothed norm weighting instead of minimum norm weighting reduced the bias for sources close to the sensors and increased the bias for sources far from the sensors. Indeed, the relationship between the *Bias in Dipole Location* and the average distance to the sensors was no longer linear, but more uniform (see also the left panel in the bottom row of figure 4-16). Thus, due to the smoothed norm weighting the *Bias in Dipole*

Location was less dependent on the average sensor distance. A similar effect was found for the relationship with the l_2 -norm of the lead field (right panel in the top row of figure 4-16). The plot of the *Bias in Dipole Location* versus angle with the sphere radius demonstrated that the bias was still independent of the orientation of the sources and that the maximum bias had dropped.

The surface map of the *Expected squared error due to noise* indicated that the error had decreased for sources close to the sensors and had increased for sources far from the sensors. Also, deep sources generally had a smaller $\|\mathbf{L}\|_2$. This explains why in figure 4-16 there was a larger number of sources with small $\|\mathbf{L}\|_2$ that had a large *Expected squared error due to noise* than in the corresponding plot in figure 4-14

These results established that the values and patterns of the figures of merit tested here were influenced by the *a priori* information that was used in the weighting matrix. However, the influence of the choice of weighting on the *Source Identifiability*, the *Image Sensitivity* and the *Expected squared error due to other sources* was relatively small. It is hypothesised that these measures are more determined by the measurement set-up than by the *a priori* information that is used, as these figures of merit are based on the resolution kernels. Consequently, the best approach to improve the *Source Identifiability* (or indeed the *Image Sensitivity* or *Expected squared error due to other sources*) would be to change the sensor location or configuration and not the *a priori* information that is used. On the other hand, the *Bias in Dipole Location* strongly depended on the weighting that was used. Using different *a priori* information therefore can significantly improve the *Bias in Dipole Localisation* for sources. The same is true to a lesser extent for the *Expected squared error due to noise*.

It has been mentioned before that the FOCUSS algorithm is based on linear estimation, but that the final solutions consist of only a few sources. In the next session it was therefore examined to what extent the figures of merit for distributed solutions were relevant to FOCUSS.

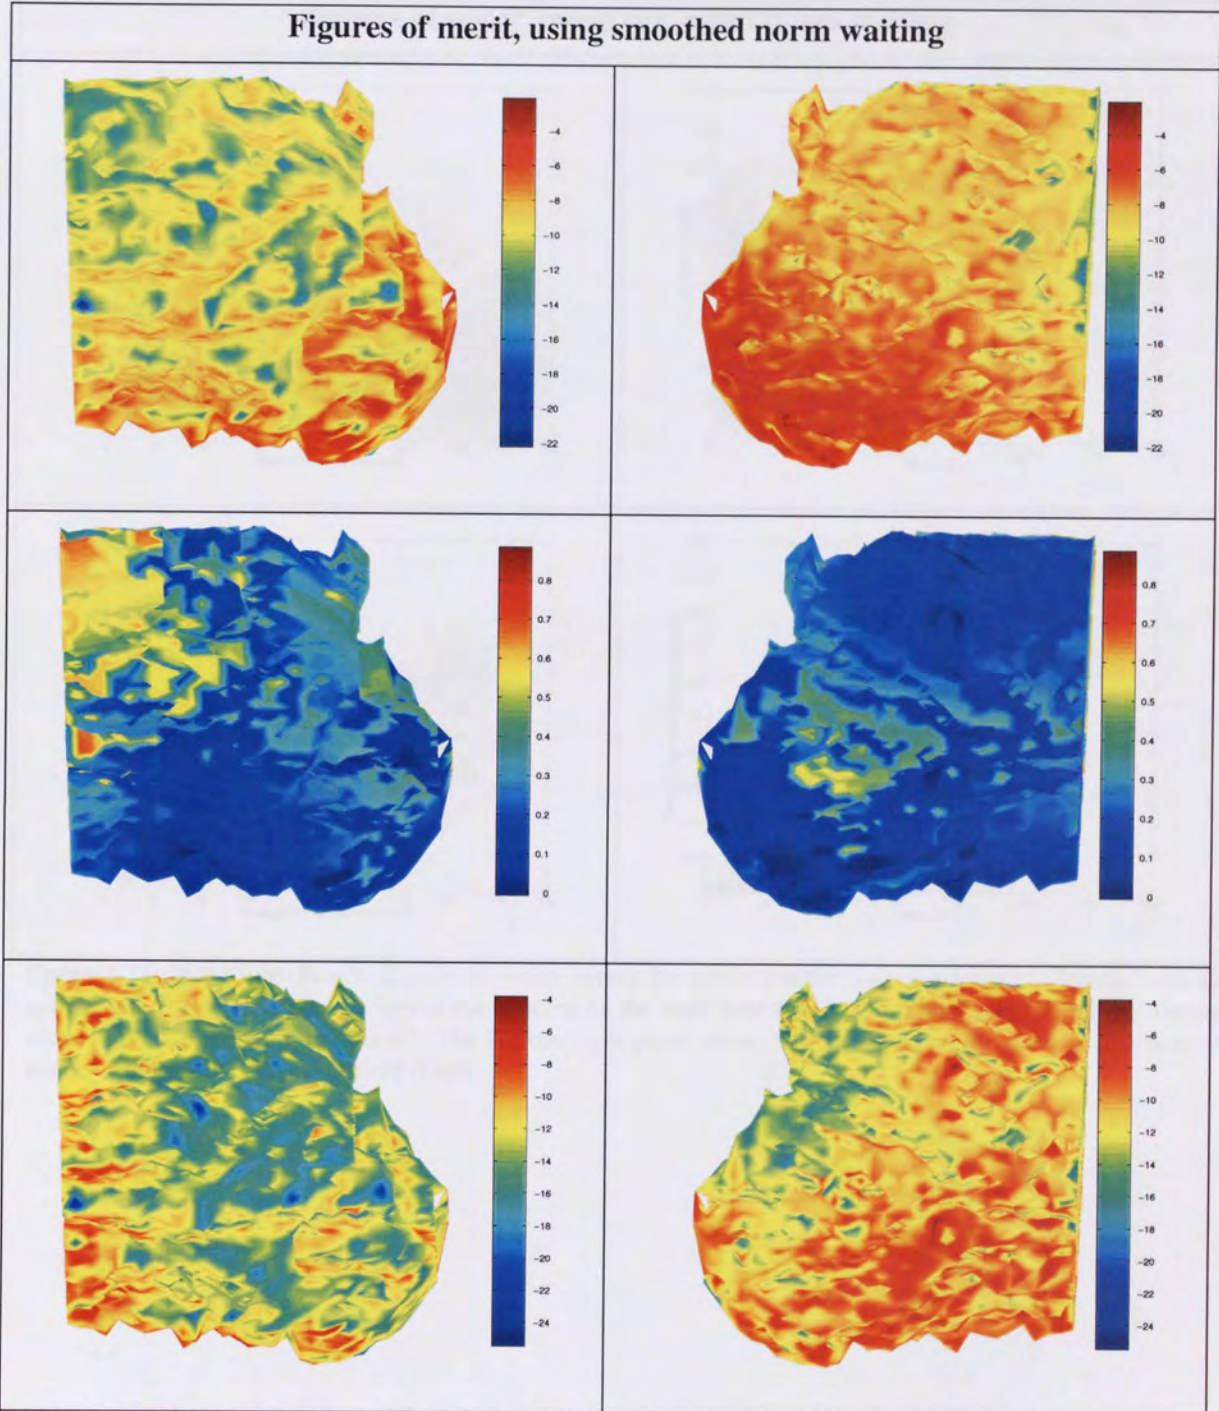


figure 4-15: The surface maps of the *Source Identifiability* (logarithmic scale), *Bias in Dipole Location* and *Expected squared error due to noise* (logarithmic scale) are given in the top, middle and bottom row respectively. The left column contains views from the left and the right column contains views from the right.

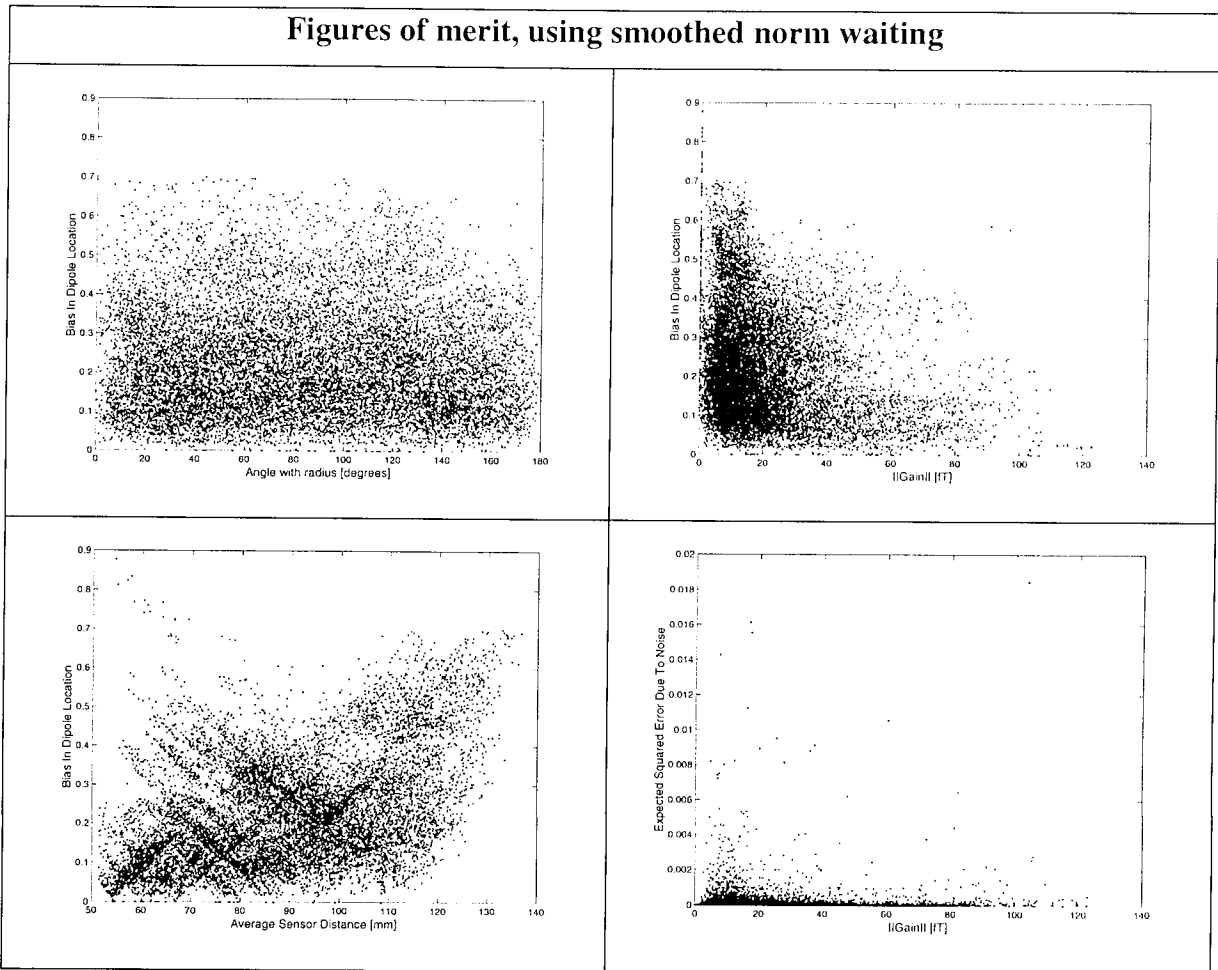


figure 4-16: Plot of the *Bias in Dipole Location* versus the angle that the source orientation formed with the sphere radius (top, left panel), versus the l_2 -norm of the lead field (top, right panel) and versus the average sensor distance (bottom, left panel). The bottom right panel shows the *Expected squared error due to noise* versus the l_2 -norm of the lead field (Gain).

4.3.3 Relevance to FOCUSS

Framework

In this study it was examined to what extent the figures of merit for distributed solutions were relevant to the FOCUSS algorithm. The FOCUSS algorithm starts with a distributed solution obtained from a weighted linear estimator. The figures of merit can thus be used to evaluate the initial guess for FOCUSS. The algorithm then progresses to reconstruct a few discrete sources via an iterative process. This is a non-linear process and it is therefore questionable what the predictive power of the figures of merit is about the quality of the final solutions that can be obtained with FOCUSS.

The applicability of the figures of merit can be examined by applying FOCUSS with different measurement set-ups and/or different *a priori* information (i.e. weighting matrices). A change in measurement set-up or weighting matrix changes both the performance of the FOCUSS algorithm and the figures of merit. It can then be examined to what extent the FOCUSS performance is related to the values of the figures of merit.

The effect of the weighting matrix instead of the measurement set-up was examined. It was shown in the previous section that the *Bias in Dipole Location* was more influenced



figure 4-17: Examples of impulse responses typically found with FOCUSS when the identity matrix is used as initial weighting matrix. The source spaces are viewed from the right-back. The left panel shows an ideal impulse response. The one element with a value of 1 is at the location of the original source (indicated with the red arrow). Note that no colourbar is given for this figure; dark blue represents a value of zero. The right panel shows an impulse response for which the maximum value is not in the same location as the original source. Additionally, there are non-zero elements surrounding the original source. Note however that the vast majority of elements in the impulse response have a value of zero.

by changes in the weighting matrix than the other figures of merit. Therefore the performance of FOCUSS with different weighting matrices was compared with the changes in *Bias in Dipole Localisation* due to these different weights. Also, the impulse responses for only a subset of the sources in the source space were determined with the FOCUSS algorithm. Therefore, the resolution kernels were not 'complete' and consequently neither were the figures of merit based on the resolution kernels.

The following weighting matrices were used:

- Minimum norm weighting, i.e. the weighting matrix was equal to the identity matrix.
- Smoothed norm weighting, described in the previous section.
- Average distance weighting. Bias towards the sensors was suppressed by multiplying the value for each element in the source space by the squared average distance from this element to the sensors.
- Target area weighting. The weighting matrix was focussed on a target area by selecting the values of the weights on the basis of the distance from the element in the source space associated with the element in the weighting matrix to the target source (see figure 4-18). The value of the weights fell off with the distance from the target source,




figure 4-18: Weighting values selected to emphasise a target area. The value of the weights fell off with the distance from the target source, following a 3D-Gaussian distribution.

following a 3D-Gaussian distribution. The standard deviation of the Gaussian was 10mm. This initial weighting matrix was then multiplied by the smoothed norm weighting matrix to counteract the bias towards the sensors.

Target area weighting was used to evaluate whether a weighting matrix could be constructed that reduced the *Bias in Dipole Location* in a region of interest. It was further examined whether such a weighting matrix could be constructed on the basis of the information in the resolution matrix or whether extensive simulations with FOCUSS were needed.

The compound version of the FOCUSS algorithm was used to find the impulse response for a randomly chosen source. Thus, the magnetic field used for the source reconstruction was equal to the lead field of the randomly selected source. A total of one thousand unique impulse responses were computed with each weighting matrix. The resolution matrix was formed from the concatenation of the impulse responses. The *Bias in Dipole Location* was computed for each used source successively. These values were then compared with the *Bias in Dipole Location* for the same sources as obtained from the resolution matrix that was computed at the start of the FOCUSS algorithm.

Results

Two impulse responses typically found with the FOCUSS algorithm, using the identity matrix as initial weighting matrix, are given in figure 4-17. For 139 sources the ideal impulse response was found. That is, the maximum value of 1 was found for the element that contained the original source and all other elements had a value of zero. The other sources had an impulse response with a maximum value in elements near the original source and additional elements with a non-zero value. However, the majority of elements in the source space had a value of zero. A huge improvement was seen when these impulse responses were compared with the impulse responses given in figure 4-8 (remember that a resolution kernel is equal to an impulse response for minimum norm weighting). Consequently, it was expected that the *Bias in Dipole Location* had improved as well. In figure 4-19 the histograms of the *Bias in Dipole Location* obtained at the start and after convergence of FOCUSS were compared. It is evident that FOCUSS reduced the *Bias in*

Dipole Location considerably. The mean *Bias in Dipole Location* at the start was 0.41 ± 0.19 and after convergence it was 0.21 ± 0.19 .

The *Bias in Dipole Location* obtained at the start of FOCUSS and after convergence is given in figure 4-20 for the case when smoothed norm weighting was used. The mean *Bias in Dipole Location* at the start was 0.22 ± 0.13 and 0.18 ± 0.13 after convergence. Comparison with figure 4-19 showed that smoothed norm weighting produced a better initial start for FOCUSS than minimum norm weighting. Equally, the *Bias in Dipole Location* after convergence of FOCUSS had improved as well by using smoothed norm weighting instead of minimum norm weighting. Hence, there was a direct relationship between the 'quality' of the initial guess and the accuracy of the solution after convergence. However, this improvement was qualitative rather than quantitative; the mean value of the *Bias in Dipole Location* in the initial guess improved by 44% whereas the improvement after convergence was (only) 13%. A similar discrepancy between the improvement in performance at the start and after convergence was found for average distance weighting; the mean value of the *Bias in Dipole Location* as obtained from the resolution matrix was 0.36 ± 0.22 and after convergence of FOCUSS it was 0.12 ± 0.09 . Hence, the improvement at the start was 12% and after convergence 44%, with respect to the performance with

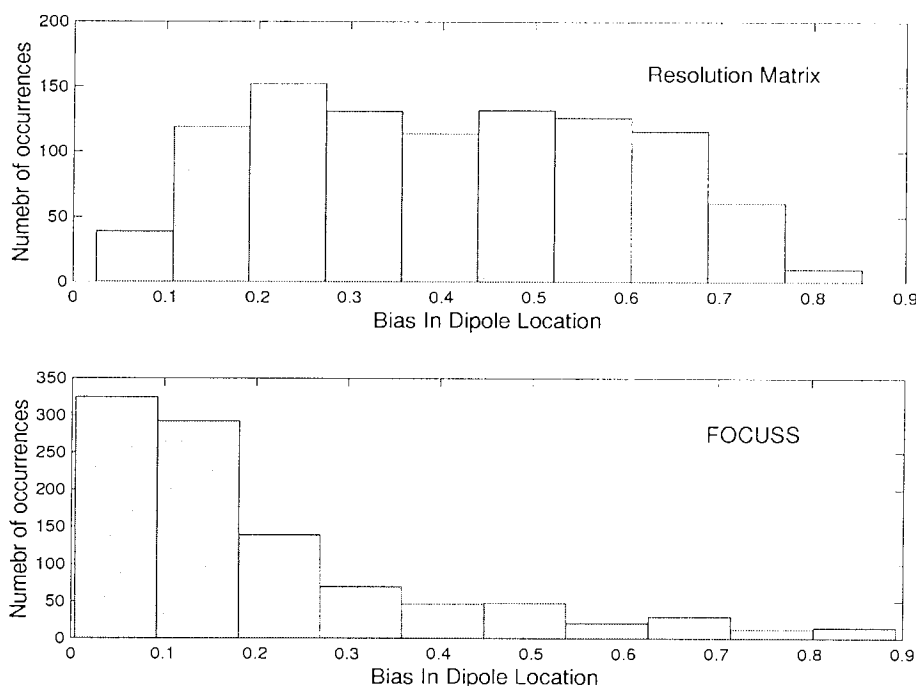


figure 4-19: Minimum norm weighting. Top panel: histogram of the values of the *Bias in Dipole Location* as obtained from the resolution matrix. Bottom panel: histogram of the values of the *Bias in Dipole Location* as obtained from concatenating the impulse responses of FOCUSS. Note the huge reduction of the number of sources with a large *Bias in Dipole Location* for FOCUSS.

minimum norm weighting.

Applying FOCUSS had the largest effect on the *Bias in Dipole Location* for average distance weighting. The mean value of the *Bias in Dipole Location* improved by 68%. The improvement was 49% with minimum norm weighting and only 20% with smoothed norm weighting.

Summarising, a change in the *Bias in Dipole Location* at the start seemed to correlate with an improvement in the accuracy of the solutions obtained with FOCUSS. Also, the improvement of the initial *Bias in Dipole Location* was larger for smoothed norm weighting than for average distance weighting, but the final solutions were more accurate with average distance weighting. Thus, the initial value of the *Bias in Dipole Location* could not be used to favour a certain kind of *a priori* information.

Results for Target Area Weighting

Target area weighing was chosen to improve the accuracy of the source reconstructions in a region of interest. In figure 4-21 it was demonstrated that the *Bias in Dipole Location* was very low in the target area. For other areas in the source space the




figure 4-20: Smoothed norm weighting. Top panel: histogram of the values of the *Bias in Dipole Location* as obtained from the resolution matrix. Bottom panel: histogram of the values of the *Bias in Dipole Location* as obtained from concatenating the impulse responses of FOCUSS.




figure 4-21: Target area weighting. The *Bias in Dipole Location* as obtained from the resolution matrix. Note the small *Bias in Dipole Location* for sources in the target area (the target area was defined in figure 4-18).

range of values was similar as found with minimum norm weighting.

Target weighting was designed to improve the performance for a target area, not the performance for the whole source space. Therefore, the *Bias in Dipole Location* for the 1000 sources nearest to the target source (maximum distance is 16.7mm) was quantitatively compared with the *Bias in Dipole Location* for these sources when smoothed norm weighting was used. The *Bias in Dipole Location* as obtained from the resolution matrix, for these sources, is given in figure 4-22. This showed that using target area weighting improved the *Bias in Dipole Location* for the sources in the target area. The *Bias in Dipole Location* with target weighting was lower than with smoothed norm weighting for all but 3 sources in the target area. To what extent target area weighting resulted in an improvement with FOCUSS follows from figure 4-23. Again, a large improvement was seen. For FOCUSS, 91% of the sources in the target area demonstrated a reduction of the *Bias in Dipole Location* by changing the weighting matrix. Hence, modifying the weighting matrix largely increased the performance of the FOCUSS algorithm for the area of interest.

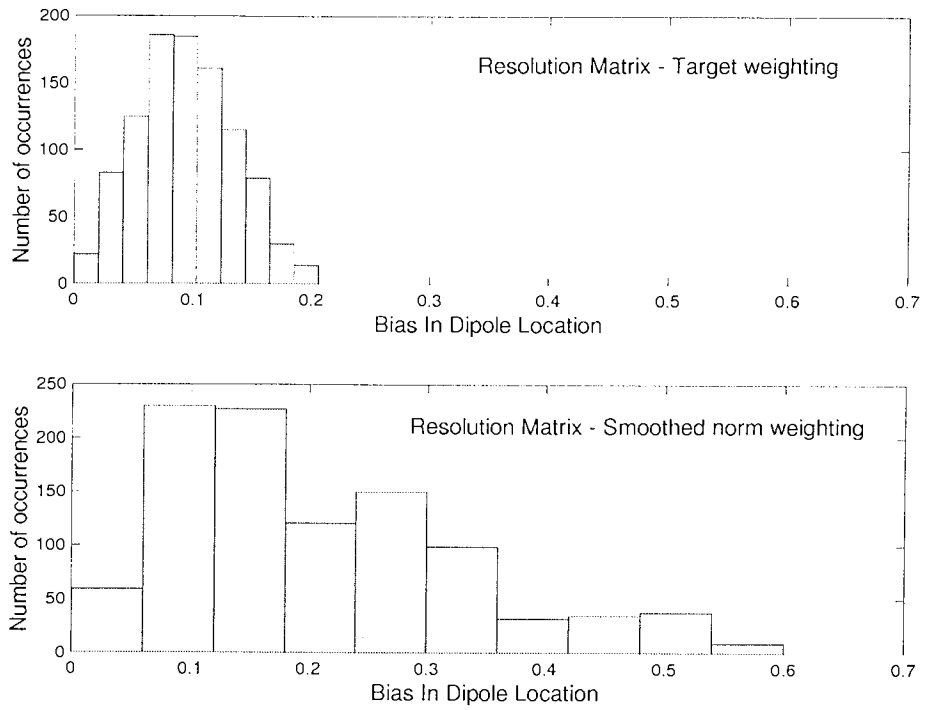


figure 4-22: Comparison of target area weighting with smoothed norm weighting. The *Bias in Dipole Location* was obtained from the **resolution matrix**. Top panel: histogram of the values of the *Bias in Dipole Location* for target area weighting. Bottom panel: histogram of the values of the *Bias in Dipole Location* for smoothed norm weighting. The same scale on the x-axis was used for the ease of comparison.

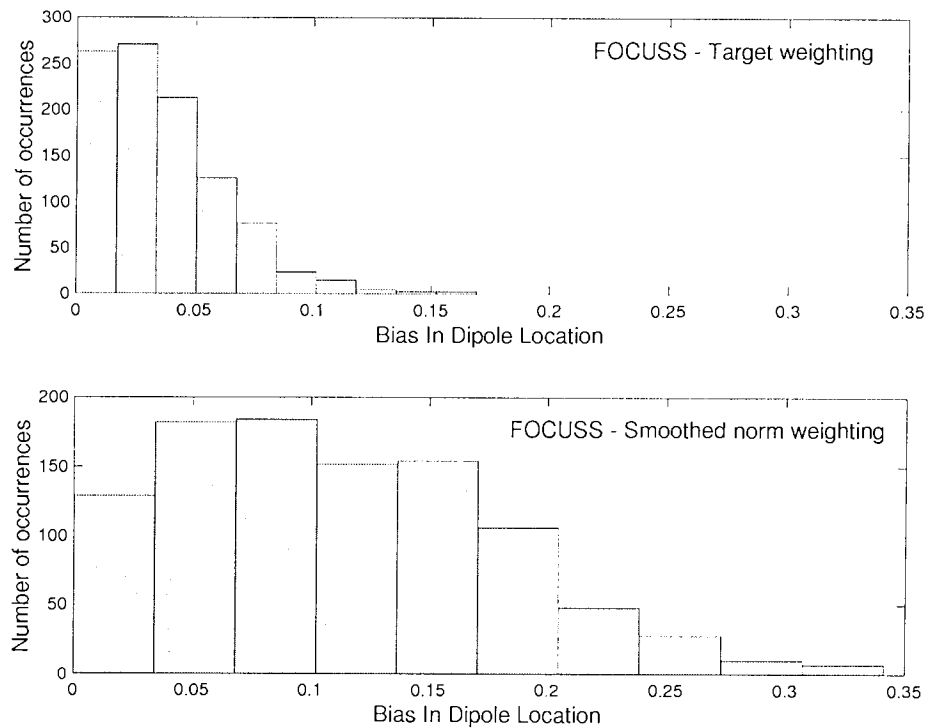


figure 4-23: Comparison of target area weighting with smoothed norm weighting. The *Bias in Dipole Location* was obtained concatenating the impulse responses of **FOCUSS**. Top panel: histogram of the values of the *Bias in Dipole Location* for target area weighting. Bottom panel: histogram of the values of the *Bias in Dipole Location* for smoothed norm weighting. The same scale on the x-axis was used for the ease of comparison. Note however the difference in scale with figure 4-22.

4.4 Discussion

Examination of the resolution kernels revealed that the linear estimator suffered from a bias towards the sensors. This bias was largest for deep sources. The ideal delta function was only approached for superficial sources. For deeper sources the resolution kernels exhibited large peaks that were shifted towards the sensors. Additionally, modifying the weighting matrix had the largest effect on the resolution kernels of deep sources.

These findings were reflected in the surface maps of the figures of merit that were derived from the resolution kernels and the columns in the resolution matrix. That is, the *Source Identifiability* was largest and the *Bias in Dipole Location* was smallest near the sensors. This illustrates how a whole-head MEG system would be advantageous. The *Source Identifiability* of a source was further determined by a combination of the $\|\mathbf{L}\|_2$ and orientation of a source. The *Source Identifiability* for sulcal sources was in the mean better than for gyral sources. However, there was no difference in *Bias in Dipole Location* between gyral and sulcal sources.

The *Image Sensitivity* profile was very similar to the *Source Identifiability* pattern. The behaviour of the *Expected squared error due to other sources* almost mirrored the behaviour of the *Source Identifiability*. These results were not surprising, activity in an area that causes a large signal in the sensors (large *Image Sensitivity*) is easily identified (large *Source Identifiability*) and is less likely to be influenced by activity from other sources (small *Expected squared error due to other sources*). Unfortunately, the *Expected squared error due to noise* was also largest near the sensors and when the $\|\mathbf{L}\|_2$ was large. Hence, sources that are easily identified are also likely to be influenced by noise in the data. However, the fact that the *Expected squared error due to noise* was high for sources with large lead fields needs to be interpreted with care. Consider the extreme case where the noise is completely perpendicular to the space spanned by the lead field vectors. It would then be impossible to project any noise signal into the source space, independent of the magnitude of the lead fields.

The classification of the elements in the source space into sulcal and gyral sources revealed that the distribution of the angle that sulcal sources form with the radius approached a normal distribution with a mean of 90° . Hence, the majority of sulcal sources were tangential sources. Also, a notable number of gyral sources were not radially

orientated, hence it would be possible to detect activity in (most) gyri with MEG. It should be further examined what the angular distributions are for gyral and sulcal sources across a number of realistically shaped source spaces. These findings should then be correlated with the reconstruction accuracy for the two classes of sources.

The use of different *a priori* information had the largest effect on the *Bias in Dipole Location*. The behaviour of the other figures of merit is probably more influenced by factors that modify the lead fields, like the sensor configuration or the shape of the source space. This is explained by the fact that a diagonal matrix was used to incorporate *a priori* information. In effect this modifies the relative strength of each element in the source space. Instead, changing the measurement set-up, or using off-diagonal elements in the weighting matrix (representing for example cross-correlation between different sources), modifies the shape of the lead fields. An analogy can be found in the two methods that can be used to get someone's attention in a crowded room; one can raise his voice (i.e. increase the diagonal elements) or alternatively turn the head into the direction of the other person (change the shape of the lead fields). A source that has a relatively strong activity (shouts hard) is easy to detect and consequently has a smaller *Bias in Dipole Location*. The other figures of merit, like the *Source Identifiability*, are more dependent on the behaviour of other sources. So, although a source can have a strong activity, there might be another source that opposes it. This largely depends on the details of the lead fields and therefore on the measurement set-up.

It was demonstrated that the *Bias in Dipole Location* as obtained from the resolution matrix correlated with the performance of FOCUSS. Thus, an improvement in the *Bias in Dipole Location* due to the use of *a priori* information correctly predicts an improvement in the performance of FOCUSS. However, it was not possible to decide between two different kinds of *a priori* information on the basis of the *Bias in Dipole Location* alone. Thus, it is possible to quickly evaluate whether adding certain *a priori* information to the inverse problem improves the performance, without the need for time-consuming evaluations with FOCUSS. Selection of the *a priori* information most applicable to the specific problem can only be done by a full set of simulations that evaluates the influence on the accuracy of the source reconstructions with FOCUSS.

Finally, it was shown that the *Bias in Dipole Location* was large for many elements in the source space. Taking into account that the maximum distance between elements in the source space was 10.3cm gives that a *Bias in Dipole Location* larger than 0.1 gave a localisation bias larger than 1cm. Fortunately, it was demonstrated that the *Bias in Dipole Location* in an area of interest can be significantly improved by selecting the appropriate weights in the initial bias matrix. The combination of smoothed norm weighting and weighting with a Gaussian centred around a target source proved particularly effective. A weighting matrix like this could be used when *a priori* information about the location of possible sources is known in advance, for example from fMRI studies. The increase of the weights for the target area had little influence on the *Bias in Dipole Location* in the rest of the source space.

5 Evaluation of Source Localisation Algorithms with Computer Simulations

5.1 Aim of This Study

The FOCUSS algorithm has been shown to result in accurate source reconstructions when a simple, circular or spherical, source space was used (Gorodnitsky *et al*, 1995). This study evaluated the performance of the FOCUSS algorithm with a more realistically shaped source space, representing a gyrus and 2 sulci. Additionally, an implementation of the algorithm was used that not only reconstructs sources for a single latency (Gorodnitsky *et al*, 1995), but instead uses the information in the data over a selected time-window.

The source space was positioned at a realistic location in a single sphere and the configuration of the 19 channel CES system at Aston University was used as sensor array (see section 4.3.2.1). Simulated data were used so that the exact location of the target sources was known and the quality of the source reconstructions could therefore be objectively evaluated.

The performance of FOCUSS with respect to its ability to accurately localise the sources underlying the simulated data was compared with the performance of an unconstrained dipole fit algorithm and a constrained dipole fit algorithm.

The simulations were to reveal:

- the efficacy of the implementation of the FOCUSS algorithm
- a method to select the best FOCUSS solution
- the merits of using spatial and orientation constraints with FOCUSS and the dipole fit algorithm
- the optimum spacing of the grid that forms the source space
- the effect of the curvature of the source space on the reconstruction accuracy
- the influence of errors in the spatial and orientation constraints

5.2 Simulation Framework

Source Space

Fitting a polynomial through a set of points that was obtained from a drawing program formed the basis of the contour of the source space. The 3D source space was formed by concatenating the 'slices' constructed in this way. This source space consisted of two parallel sulci, 4cm deep, joined by a gyrus, approximating a true cortical fold (figure 5-1). A source space very similar in overall shape was used in a simulation study by Gençer and Williamson (1998), which was an extension of the model of a single sulcus used by Wang (1994). A slightly different sulcus model was used by Wolters *et al.* (1999). The source space was defined by a 'mesh' of points.

The grid spacing of the original mesh was 0.5mm. Downsampling of this mesh subsequently formed meshes with coarser grids. The source space was formed by placing an equivalent current dipole in each node of the grid. The orientation of the dipoles was constrained to be normal to the surface.

Target Source

A target source was formed by selecting a location in the original mesh (grid spacing 0.5mm). This location formed the centre of an extended source. The amplitude of this extended source fell off with the distance from the centre, following a 2D-Gaussian distribution. The standard deviation of the Gaussian was 10mm. The extended source represented a realistic situation. Modelling such a source by a single dipole should not introduce large errors in the source reconstructions (de Munck *et al.*, 1988, Hara *et al.*, 1999). For example, Hara and co-workers found in EEG simulations that a spherical dipole layer with a radius of 20mm could be approximated by a single dipole without introducing large errors in the produced field patterns. The field pattern produced by a distributed source is more spread than that produced by a single dipole. The influence of the skull boundary spreads this pattern even further in the case of EEG¹ (known as the smearing effect of the skull), producing a less 'dipolar' pattern than for MEG. Hence, the approximation of a dipolar source is certainly valid for MEG if it is valid for EEG.

The activity pattern of the source was given by a sinusoid. The frequency of the sinusoid was chosen so that one cycle fitted in the selected time-window.

Two different target sources were used. One was centred on the gyrus and one was centred on the sulcal wall (figure 5-1). The sensor array was large enough to record the extrema in the dipolar field pattern created by the target sources in these locations (see for example figure 5-3).




figure 5-1: Source space modeled as a cortical fold viewed from the top (left panel) and right (right panel). The channel array is depicted by the filled circles and the centers of the target sources on the gyrus and the sulcal wall are indicated by arrows.

Noise

Noise was added to the simulated data, in order to probe the limits of the tested algorithms. Noise values were obtained from a Gaussian distribution with zero mean and a standard deviation of $20fT$. Similar noise levels are typically found in the literature. For example a noise level of $20.9-27.1fT$ was found for visual evoked data (Aine *et al.*, 1996) and Huang and colleagues used a standard deviation of $14fT$ in their simulations (Huang *et al.*, 1998). The assumption of Gaussianity seems reasonable for MEG/EEG (Dogandzic and Nehorai, 2000).

¹ With a spherically symmetric head model there are no volume conducting effects for MEG.

It is the signal-to-noise ratio (SNR) that ultimately determines the performance of the source reconstruction algorithms. The SNR was varied by varying the amplitude of the target source, and hence the amplitude of the simulated MEG data.

Monte Carlo simulations were performed in order to assess the stability of the source reconstructions. Thus, each simulation was repeated several times, using a different noise realisation. It was found in initial simulations that the mean distance from the solutions to the target source and the standard deviation of this distance stabilised when the number of realisations was in the order of 30. This was found with FOCUSS and with the unconstrained/constrained dipole fit algorithms (figure 5-2). Similar results were obtained for data with different SNRs and version B of the constrained dipole fit algorithm (the algorithms are described below). This means that even when the solutions were influenced by noise, the centroid of the cluster formed by all solutions and its volume became almost independent of further noise realisations.

Each dataset, containing a different noise realisation, was used for the reconstruction of the underlying sources with three distinct algorithms. The implementation of these algorithms is described below.

Implementation of Source Reconstruction Algorithms

Unconstrained Dipole Fitting

A fixed dipole model was used to explain the simulated data. That is, the location and orientation of the dipole was constant throughout the latency range, and the amplitude was allowed to vary. The timecourse of activation was obtained from the data by linear estimation once the location and orientation of the dipole was known (see section 2.3.2). Hence, the parameters to be estimated were: 3 location parameters, 1 orientation parameter in the tangential plane and N activation parameters (with N the number of latencies).

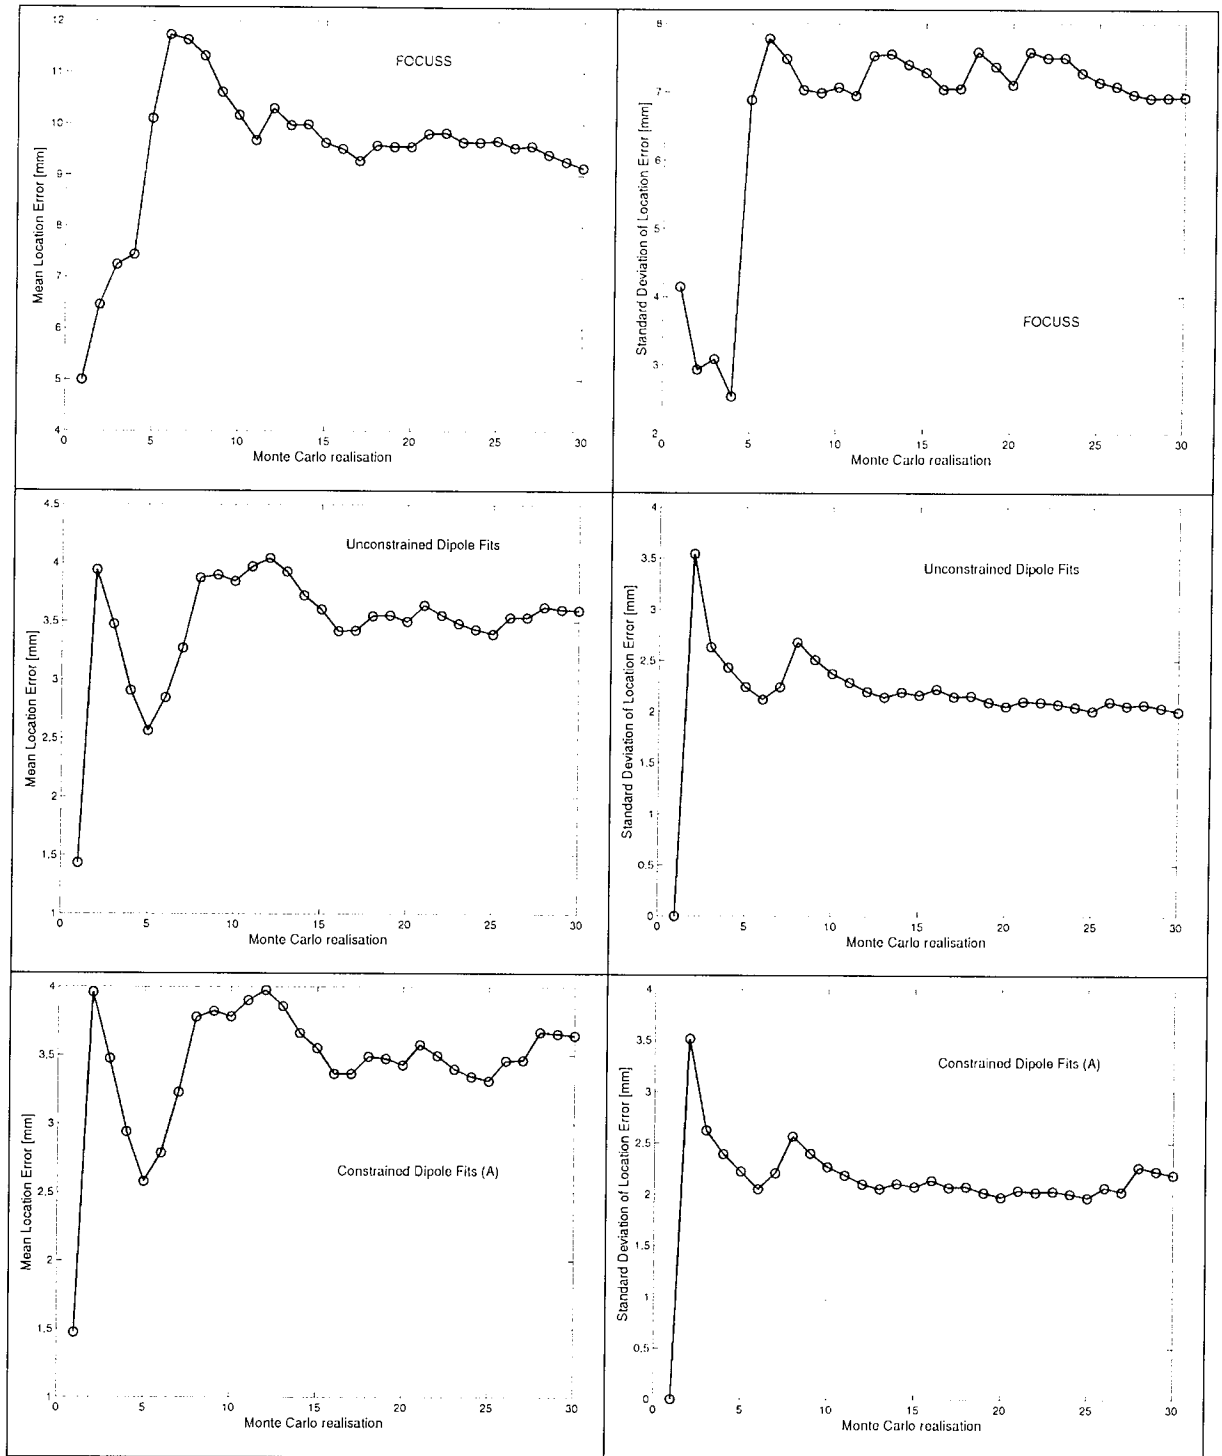


figure 5-2: The mean location error [mm] and the standard deviation of the location error [mm] for different Monte Carlo realisation. The results obtained with FOCUSS, unconstrained dipole fitting and constrained dipole fitting (A) are given in the top, middle and bottom row, respectively (see text for details about the algorithms). These figures were obtained by adding the solution found for a Monte Carlo realisation to the solutions obtained from previous realisations and computing the error measures. The results displayed were obtained for simulated data with an SNR of 1.5 for the target source on the gyrus. Note the stabilisation of the results as the number of realisations increased.

Initial Guess

A single dipole was used in the dipole fit algorithm, i.e. it was assumed that the number of sources underlying the data was known. The location of the initial guess was randomly chosen within a sphere with a radius of 2cm around the target source. The initial guess for the orientation of the source was randomly selected.

Cost-function

The reduced chi-square measure was used as the cost-function (see section 2.2.1.2.1). Its value was estimated for each latency separately and the mean over all latencies was used as the value of the cost-function. Dipoles with a magnitude larger than 100nAm or that moved outside the sphere were 'punished' by squaring the value of the cost-function.

Optimisation algorithm

The implementation of the Simplex algorithm (Nelder and Mead, 1965) in MATLAB (The Mathworks) was used as the optimisation algorithm.

A multi-start approach was used in order to avoid local minima (Huang *et al*, 1998). With each dataset, the algorithm was started with 10 initial guesses (Fernández *et al*, 1995). The solution that resulted in the lowest value of the cost-function was chosen as the solution for the dataset, referred to as *unconstrained dipole solution (UCDIP)*.

Constrained Dipole Fitting

Again a fixed dipole model was used to explain the data. However, this time the location of the dipole was restricted to lie within a distance equal to the grid spacing from the nodes in the source space. Sources that moved further away from the surface were punished by squaring the value of the cost-function. Additionally, the orientation of a dipole was constrained to be the same as the orientation of the normal of the nearest node in

the source space. Hence, only the 3 location and N activation parameters needed to be optimised.

Apart from these constraints, the same implementation as for the unconstrained dipole fitting was used. For each dataset this resulted again in 10 solutions due to the multi-start approach, from which the one with the lowest value for the cost-function was chosen, referred to as *constrained dipole solution (a)* or *CDIPa*. Also, the nearest element in the source space from CDIPa was stored as the *constrained dipole solution (b)* or *CDIPb*. The latter solution was intended to be a fair comparison for solutions obtained with FOCUSS, since the FOCUSS solutions were also obtained by using the spatial and orientation constraints.

Multi-latency FOCUSS algorithm

Implementation

The FOCUSS algorithm was modified so that temporal information could be incorporated. The equation on which the FOCUSS algorithm is based, equation 2-47, is first repeated here:

$$\hat{\mathbf{Q}}_{\text{FOCUSS}} = \mathbf{W} (\mathbf{L} \mathbf{W})^+ \mathbf{B} = \mathbf{W} \mathbf{W}^T \mathbf{L}^T (\mathbf{L} \mathbf{W} \mathbf{W}^T \mathbf{L}^T)^{-1} \mathbf{B}, \quad (5-1)$$

with \mathbf{B} the (simulated) measurements in the MEG sensors, \mathbf{L} the lead field matrix for all elements in the source space and \mathbf{W} a diagonal weighting matrix.

Remember that the FOCUSS algorithm progresses by using the solution from previous iterations as the weights in the current iteration, until the algorithm converges to a final solution. The algorithm was originally described for single latency data only (Gorodnitsky *et al.*, 1995). However, the algorithm can easily be adjusted for multi-latency data (Gorodnitsky, 1998). The solution found in an iteration will consist of the timecourses of activation when multi-latency data is used (i.e. $\hat{\mathbf{Q}}_{\text{FOCUSS}}$ is a matrix instead of a vector). It is obvious that this can not be used directly in the weighting matrix, since the diagonal of the weighting matrix forms a vector, and the solution obtained from the previous iteration forms a matrix. When all the information in the activation curve of each element is

'compressed' to a single value, then this value can be used as the value on the diagonal in the weighting matrix for the element of interest. In formulas:

For single latency data we had for the weighting matrix in iteration k (compound version):

$$\mathbf{W}_k = \mathbf{W}_{k-1} \cdot \begin{bmatrix} \hat{\mathbf{Q}}_{1,k-1} & & 0 \\ & \mathbf{O} & \\ 0 & & \hat{\mathbf{Q}}_{n,k-1} \end{bmatrix}, \quad (5-2)$$

With n the number of the element in the source space. For multi-latency data this becomes:

$$\mathbf{W}_k = \mathbf{W}_{k-1} \cdot \begin{bmatrix} \hat{\mathbf{Q}}'_{1,k-1} & & 0 \\ & \mathbf{O} & \\ 0 & & \hat{\mathbf{Q}}'_{n,k-1} \end{bmatrix}, \quad (5-3)$$

with $\hat{\mathbf{Q}}'_{n,k-1}$ given by:

$$\hat{\mathbf{Q}}'_{n,k-1} = \sum_{t=1}^T \gamma_t \left| \hat{\mathbf{Q}}_{t,n,k-1} \right|, \quad (5-4)$$

$\hat{\mathbf{Q}}_{t,n,k-1}$ represents the solution from the previous iteration, for element n in the source space at latency t. Hence, the combination of equation 5-3 and 5-4 states that the value for each element on the diagonal in the weighting matrix is formed from the weighted contribution of the activity at all latencies for that element in the previous iteration. The weightings γ_t can be adjusted to alter the contribution from different latencies. For example, activity at latencies with peak activity could be given a larger influence on the final solutions by increasing the value of γ for that latency (this could be achieved by basing the value of γ on the SNR). In this study all latencies are given equal weighting. The formulation for the non-compound version of the multiple latency FOCUSS algorithm is now elementary.

Bias Matrix

In this study, a combination of target weighting and distance weighting, as described in the previous chapter, was used. This choice had been empirically shown to give accurate source reconstructions with the given source space. The matrix consisted of a

part that was based on the original distributed target source and a part that corrected for the bias towards superficial sources. This correction was done by multiplying the value for each element in the source space by the squared average distance from this element to the sensors and by an additional factor. This additional factor was equal to the inverse of the difference between the distance to the farthest sensor and the nearest sensor. This further reduced the weighting for nearby sources and enhanced the weighting for deep sources.

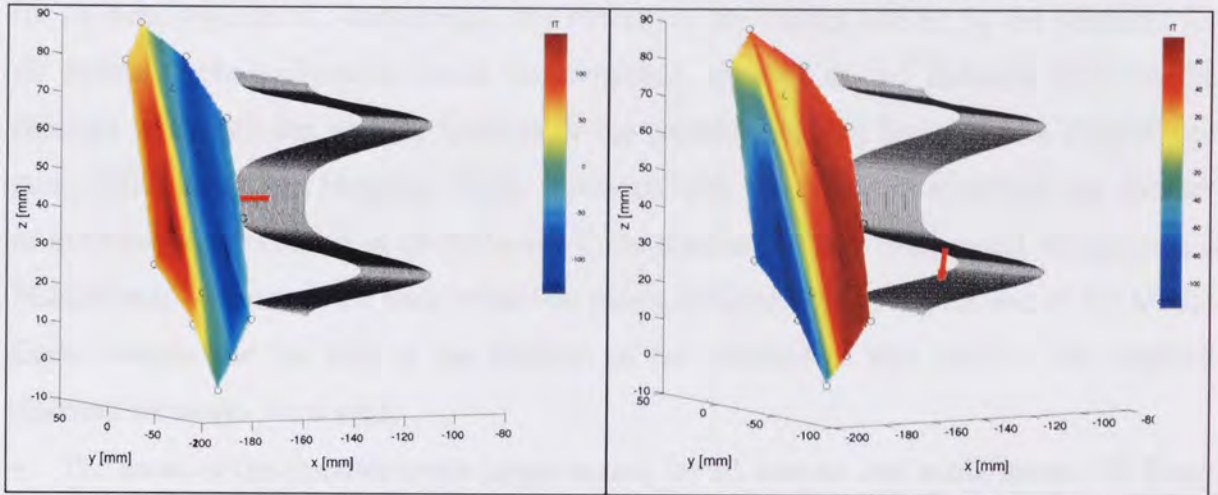


figure 5-3: Fieldmap produced by the target source on the gyrus (left panel) and sulcal wall (right panel), respectively. The dataset has an SNR of 3 and the source activity is near its maximum. Note that the spatial coverage of the sensor array is big enough to incorporate the extrema in the dipolar magnetic field pattern.

Choosing the Solution

Noise regularisation was performed by selecting a truncation value. For each dataset, numerous truncation levels were explored. With each selected truncation level a source reconstruction was obtained. The final solution for the dataset was chosen on the basis of:

- 1) The lowest reduced chi-square value.
- 2) The lowest mean distance from the reconstructed source(s) to the target source.

Note that the reduced chi-square value was computed for the final solution, but that the FOCUSS algorithm itself does not use a cost-function to find a solution. Instead, convergence is achieved by iteratively exploiting the results from previous solutions.

Evaluating Results From Monte Carlo Simulations

In order to compare the solutions obtained with the three different source reconstruction algorithms, error measures were required that could be used with all three algorithms. An obvious choice was to compute the mean value of the reduced chi-square values obtained for all Monte Carlo realisations, since the reduced chi-square was computed for all three algorithms. Additionally, the volume of the cluster formed by the solutions for all Monte Carlo realisations could be computed, as well as the distance between the centroid of this cluster and the location of the (centroid of the) target source (Supek and Aine, 1993, Singh and Harding, 1996). However, with the FOCUSS algorithm the number of different sources found in all 30 Monte Carlo realisations was often small. Computing a Monte Carlo volume on the basis of so few points is futile. Therefore, instead of the Monte Carlo volume and the bias in the location of the centroid of this volume, two slightly different measures were used:

- The mean of the distance to the target source for all sources that made up the (30 final) solutions found with the different Monte Carlo realisations, \bar{d} . This measure revealed any bias in the solutions.
- The standard deviation of this distance, $\sigma_{\bar{d}}$. This measure revealed how much the solutions were spread out. The lower the value of $\sigma_{\bar{d}}$ was, the more confined the cluster formed by the solutions was.

Remember that for the constrained dipole fits and the FOCUSS algorithm two different solutions were obtained for each Monte Carlo realisation. Hence, for both algorithms two different values of \bar{d} and $\sigma_{\bar{d}}$ were acquired.

The simulation framework described above was first used in section 5.3 to evaluate the implementation of FOCUSS. The results obtained with the compound and non-compound version of the algorithm were compared and methods for the choice of the regularisation level were evaluated. Additionally, the influence of gridsize on the constrained dipole fits and FOCUSS was examined. In section 5.4 the effect of inaccurate constraints on source reconstructions was examined. Throughout these simulations the solutions were compared with those obtained with the unconstrained dipole fit algorithm.

5.3 Parameter Selection

5.3.1 Compound Versus Non-compound Version

Two versions of the FOCUSS algorithm were originally published (Gorodnitsky *et al*, 1995); a compound and a non-compound version. The difference between the two versions is the way in which the solutions from previous iterations are used in the iteratively updated weighting matrices. In the compound version the solutions obtained from all previous iterations are used to compute the weights, whereas in the non-compound version only the solution obtained in the last iteration is used. The compound version is claimed to converge faster and to be more robust (Gorodnitsky *et al*, 1995). Hence, it was expected that the compound version would result in more accurate solutions than the non-compound version, with the largest difference in performance for data with low SNR.

The difference in performance of the compound and non-compound versions of the FOCUSS algorithm was evaluated. Monte Carlo simulations (N=30) were performed on datasets with a mean SNR of 1.2, 1.5 and 3, respectively. A grid with a spacing of 8mm was used for data induced by the target source on the gyrus and by the target source on the sulcal wall, respectively. Both an implementation of the compound and non-compound version of the algorithm was used to reconstruct the target sources. For comparison, the same data were used by the unconstrained and constrained dipole fit algorithms for the reconstruction of a single dipole.

Results

The results are summarised in figure 5-4. The following nomenclature is used in this figure and the rest of this chapter:

F_C_Loc: Compound version of FOCUSS. Final solutions selected on the basis of the distance to the target sources.

F_C_Chi: Compound version of FOCUSS. Solutions selected on the basis of the value of χ_r^2 .

F_NC_Loc: Non-compound version of FOCUSS. Final solutions selected on the basis of the distance to the target sources.

F_NC_Chi: Non-compound version of FOCUSS. Solutions selected on the basis of the value of χ_r^2 .

UCDIP: Unconstrained dipole fitting.

CDIPa: Constrained dipole fitting.

CDIPb: Constrained dipole fitting. Final solutions constrained to the mesh.

In general, the accuracy with which the sources were reconstructed increased with an increase in SNR, both with respect to the location of the solutions and the spread of the solutions. A comparison of the top and bottom panel in figure 5-4 reveals that the mean distance and the standard deviation of the mean distance were strongly correlated (0.90).

Choosing the final solutions for the FOCUSS algorithm on the basis of the value of χ_r^2 gave inaccurate results. The noise truncation applied with FOCUSS largely removed the noise components from the data (as well as high spatial frequency data components). The formalism of the reduced chi-square was therefore not suitable. Hence, the selection of the final solutions for the FOCUSS algorithm should be based on different criteria.

The performance with the compound version of FOCUSS was better than with the non-compound version (when the solutions were selected on the basis of the distance to the target sources). The smaller spread in the solutions for the compound version is explained by the fact that the compound version converged quickly to a solution whereas the non-compound version converged more slowly. For the non-compound version the final solution was therefore more dependent on the particular realisation of the noise. This is illustrated in figure 5-6, where the final solutions are shown for the compound and non-compound version for data with SNR of 1.5. It is clear that there were many more different solutions found with the non-compound version than with the compound version. The spread of these sources was also larger for the non-compound version than for the compound version.




figure 5-4: Comparison of compound and non-compound version of FOCUSS, unconstrained dipole fit and constrained dipole fit algorithms for data with different SNRs. Target source is the source on the **gyrus**. The top panel shows the mean of the distance to the target source for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note the increase in accuracy and the decrease of the spread in the solutions with increasing SNR and the relatively poor performance of FOCUSS.

F_C_Loc: Compound FOCUSS. Final solutions based on distance from target; *F_C_Chi*: Compound FOCUSS. Final solutions based on χ_r^2 ; *F_NC_Loc*: Non-compound FOCUSS. Final solutions based on distance from target; *F_NC_Chi*: Non-compound FOCUSS. Final solutions based on χ_r^2 ; *UCDIP*: Unconstrained dipole fitting; *CDIPa*: Constrained dipole fitting; *CDIPb*: Constrained dipole fitting. Final solutions constrained to mesh.

The dipole fit algorithms outperformed the FOCUSS algorithm for all SNR cases. In turn, the performance of the constrained dipole fit algorithm was better than the performance of the unconstrained algorithm, especially at low SNR. The advantage of using constraints diminished when the SNR of the data was increased. Constraining the final solution to lie within the mesh further increased the accuracy of the constrained dipole fits, due to the fact that noise in the data did not cause 'jitter' in the location of the final solutions. For data with an SNR equal to 3 the same source was found with the constrained dipole fit algorithm (UCDIPb) in all Monte Carlo realisations, hence the standard deviation of 0. The residual distance between this source and the target source was caused by the fact that the target source did not lie in a node of the mesh that was used in this simulation.

The differences in performance between the compound and non-compound versions of FOCUSS were very similar when the target source was located on the sulcal wall (compare figure 5-4 with figure 5-5). Thus, these findings were independent of the location of the target source.

The localisation accuracy was lower for the source on the sulcal wall in comparison with the gyral source, independent of the algorithm that was used. This was due to the larger distance between the sulcal source and the sensors. The poor performance of the unconstrained dipole fit for data with an SNR of 1.2 is especially striking. It is also noteworthy that the spread in the solutions was smaller for the sulcal source than for the gyral source when FOCUSS was used, even though the accuracy in location was lower.





figure 5-5: Comparison of compound and non-compound version of FOCUSS, unconstrained dipole fit and constrained dipole fit algorithms for data with different SNRs. Target source is the source on the **sulcal** wall. The top panel shows the mean distance to the target, gyral, source for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. *F_C_Loc*: Compound FOCUSS. Final solutions based on distance from target; *F_C_Chi*: Compound FOCUSS. Final solutions based on χ_r^2 ; *F_NC_Loc*: Non-compound FOCUSS. Final solutions based on distance from target; *F_NC_Chi*: Non-compound FOCUSS. Final solutions based on χ_r^2 ; *UCIP*: Unconstrained dipole fitting; *CDIPa*: Constrained dipole fitting; *CDIPb*: Constrained dipole fitting. Final solutions constrained to mesh.



Conclusions

The performance of the compound version of FOCUSS was better than that of the non-compound version. For the non-compound version the solutions obtained with different noise realisations were more spread out and the mean distance between the solutions and the target source was larger. With the source on the sulcal wall as target source the non-compound version was unable to distinguish between activity on opposite sides of the gyrus.

For the gyral source, and to some extent for the sulcal source, the difference in distance error for the compound and non-compound version of FOCUSS was largest when the SNR of the data was high. The compound version was reported to be more accurate and robust than the non-compound version (Gorodnitsky *et al*, 1995). It was expected that this would be emphasised with noisy data. The contrary was true. Apparently, the slower convergence of the non-compound version did not introduce more spurious sources with noisy data than the compound version did. However, the slow convergence introduced more spurious sources when the SNR was high.

All three algorithms performed better for the gyral source than for the sulcal source for data with comparable SNRs. The larger average distance between the sulcal source and the sensors probably caused this.

The distance error for the dipole fit algorithms was smaller than for FOCUSS. An exception is the extreme case of noisy data for the sulcal source, where the unconstrained dipole fit algorithm resulted in by far the worst performance. Also, the constrained dipole fit algorithm performed better than the unconstrained dipole fit algorithm. The difference between constrained and unconstrained solutions was largest for noisy data. Thus, spatial constraints reduced the influence of noise on the source reconstructions.

5.3.2 Truncation Parameter

The FOCUSS algorithm is based on an iteratively updated weighted minimum norm solution. The influence of noise on the solutions is potentially large, because the inverse problem is underdetermined (the number of elements used in the source space is much larger than the number of MEG sensors). Therefore, noise regularisation needs to be applied. What follows is a summary of the discussion in section 2.3.2.3.

The noise components in the data are normally associated with small singular values. For the reconstruction of the sources the spatial components in the data are multiplied by the inverse of these singular values. Consequently, each noise component is multiplied by a large number and unstable solutions are found due to this magnification of the projection of the noise components into the solution. The solutions can be stabilised by means of regularisation. In this work regularisation is applied by means of Truncated Singular Value Decomposition. In this approach the small singular values associated with the noise components of the data are discarded and consequently the associated noise components are not projected into the solutions. One can intuitively see that when the truncation level is set too high, one discards signal components as well as noise components. This decreases the spatial resolution of the solutions. On the other hand, setting the truncation level too low creates unstable solutions. This trade-off between the spatial resolution and the influence of noise is demonstrated by the L-shaped curve that is obtained when the solution norm is plotted versus the residual norm for different levels of truncation (Tu *et al*, 1996, Johnston and Gulrajani, 1997, Brooks *et al*, 1999). The optimum truncation level is found in the 'knee' of this curve, where both the residual norm and the solution norm are small.

In this section it will be demonstrated that the L-curve method could not be used with FOCUSS. This was due to the fact that the number of elements with zero activity increased as FOCUSS progressed. This reduced the effective number of elements in the source space and consequently reduced the underdeterminedness of the problem. The concept of the L-curve is not suitable for problems that are not underdetermined.

An alternative approach to select the truncation level was also evaluated in this section. The truncation level was set equal to $\alpha\%$ of the largest singular value. The value of α was varied over a wide range of values. As a result, several different solutions were found (i.e. one for each value of α) for each dataset. For simulated data the final solution could be selected by choosing the solution that contained the sources nearest to the target

source. For real data this would be impossible and other strategies to select the final solution were therefore evaluated (see later).

In this section only the FOCUSS algorithm was used for source reconstructions. It is first demonstrated that the L-curve method was not usable with the FOCUSS algorithm. It is then shown how the final solution could be selected from the set of solutions found when the truncation level was systematically varied.

Results

In figure 5-7 the L-curve is shown for a number of iterations of the FOCUSS algorithm. The gyral target source was reconstructed from a dataset with an SNR of 3. In each iteration the truncation level was manually chosen, based on the L-curve. The results were therefore dependent on the choices the experimenter made. Several issues are illustrated by this figure:

- The 'knee' in the L-curve was not clearly defined at the start of the algorithm (iteration 1). This could be due to the small number of sensors with respect to the number of elements in the source space.
- The L-shape became more clearly defined when the algorithm progressed (iteration 2-4). Due to the rejection of sources with low amplitude by FOCUSS, the effective ratio of the *number of sensors* to the *number of elements in the source space* increased. It was demonstrated by Gençer and Williamson that increasing the number of sensors improves the singular value decay pattern (Gençer and Williamson, 1998). A similar effect was found here.
- In further iterations (= 5-8) the shape of the L-curve deteriorated and disappeared altogether. Due to the removal of all but a few large sources from the source space the system became overdetermined. The concept of the L-curve is defined for underdetermined systems and was therefore not valid anymore.

Ideally, the selection of the truncation level should be performed automatically. This is easily done when the curve is clearly L-shaped, as in iteration 2-4. The truncation level can then be selected by automatic detection of the 'knee' in the L-curve based on the derivative

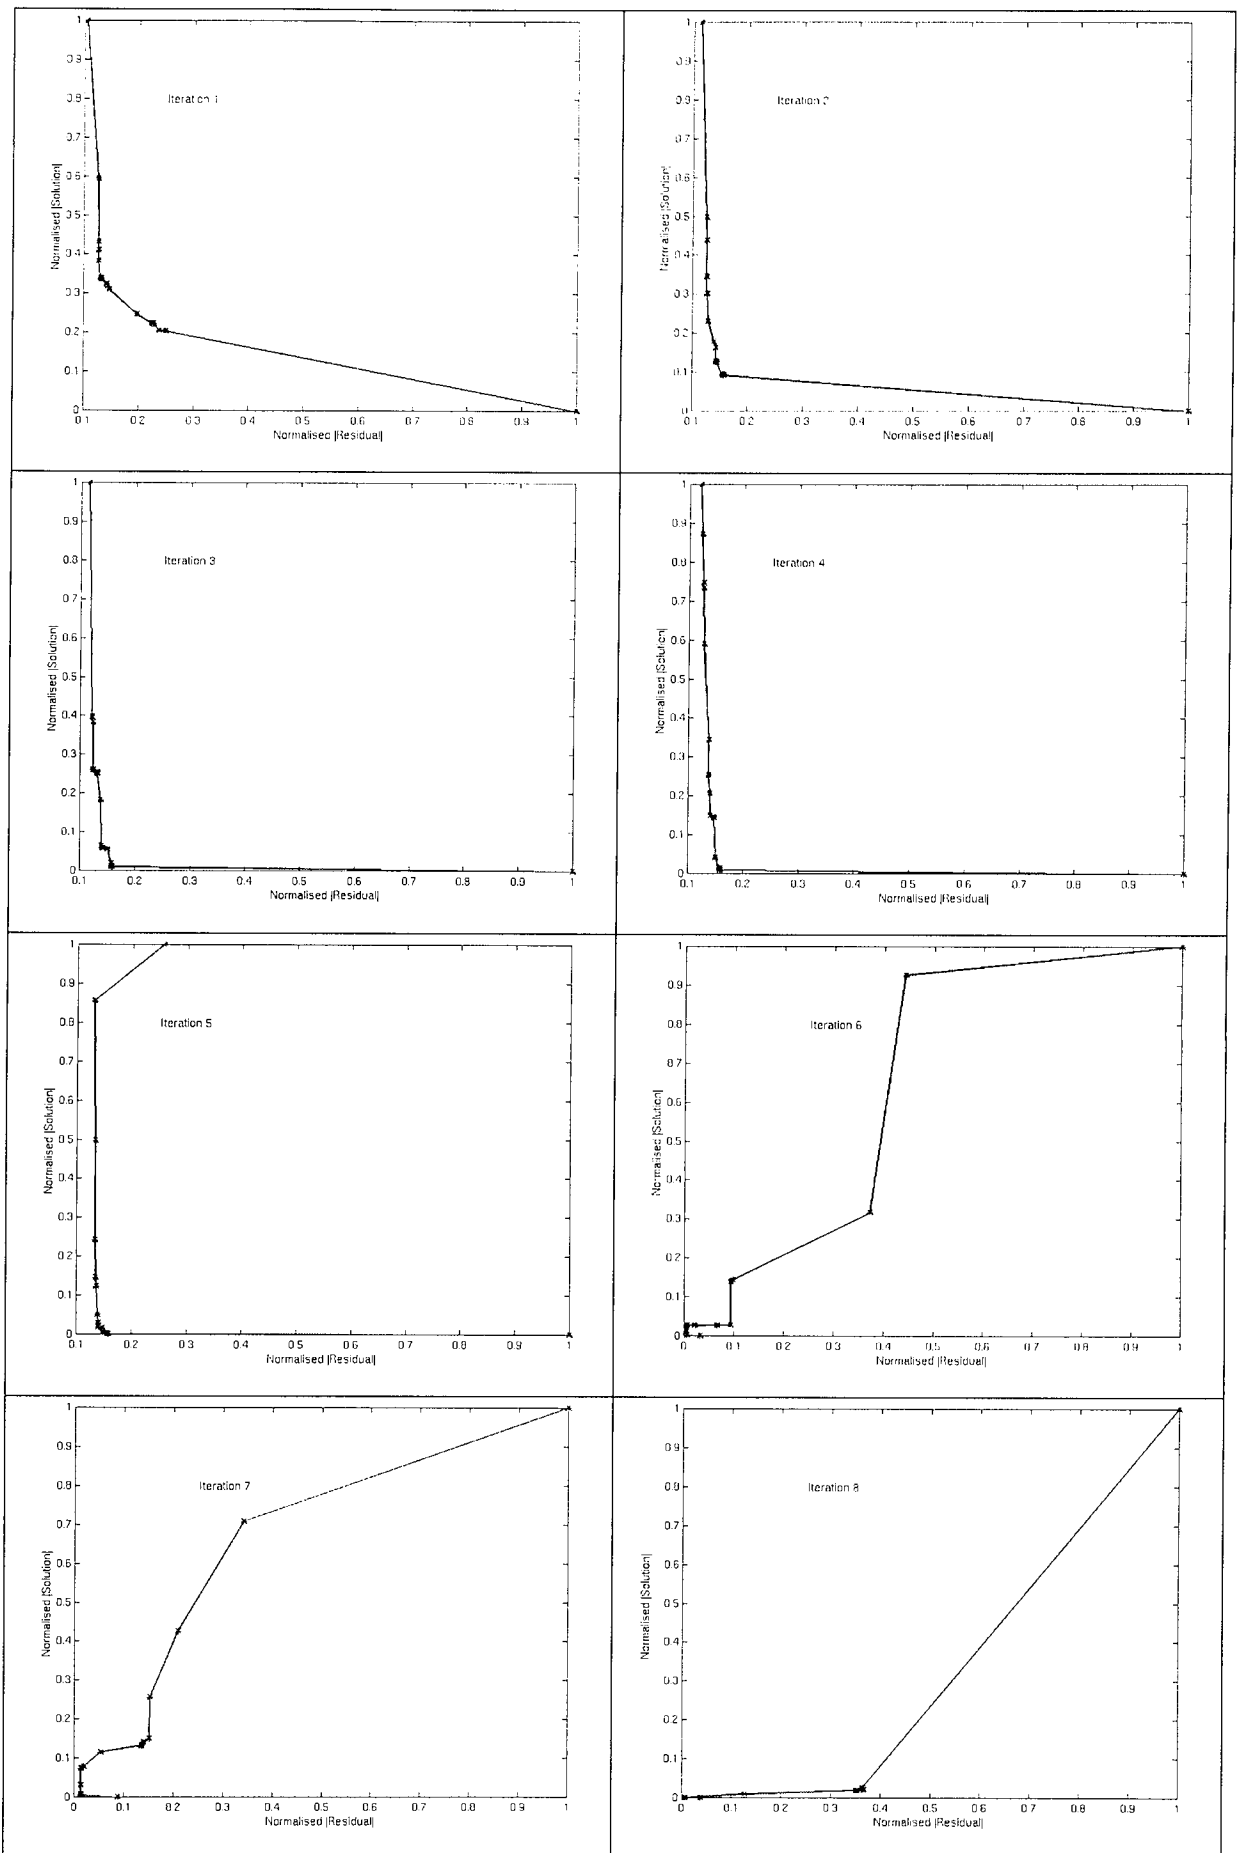


figure 5-7: L-curves for different iterations of the FOCUSS algorithm. Note the improvement in the shape of the L-curve in the first 4 iterations and the loss of the L-shape in the graphs for iterations 5 to 8.

of the L-curve. In the first iteration this is more ambiguous and for the last few iterations this approach would fail altogether. The alternative is a manual approach, where the experimenter selects the truncation level manually, based on the L-curve. In the final iterations the truncation level can not be based on the L-curve. The truncation level can be set at a high level for these iterations, since the number of (noise) sources is already low in these iterations. This is however time-consuming and the results are dependent on the choices the experimenter makes.

An alternative automatic method can be implemented, based on the following observations:

- The graphs with the L-curves for iteration 2-4 showed that the number of points on the horizontal leg of the L-curve decreased when the algorithm progressed. This meant that the number of singular values that should be used in the inversion of the weighted lead field matrix decreased.
- The difference between the first few singular values and subsequent singular values increased when the algorithm progressed.

Consequently, the selection of the truncation level at a fixed percentage ($=\alpha\%$) of the first singular value would ensure a decrease in the number of selected singular values in succeeding iterations. The correct value of α can be found by systematically evaluating the algorithm for a range of values.

Performing source reconstructions with FOCUSS, using a range of values for α , resulted in a set of solutions for each dataset. For simulation purposes the best solution could be selected on the bases of the distance between the reconstructed sources and the target sources. However, for real data this would be impossible. A strategy for real data would be to select the final solution on the basis of the value of χ_r^2 . Alternatively, the solutions obtained for different values of α could be examined manually. The timecourses of the activation of the reconstructed sources provided insight into the validity of the source. This is illustrated in figure 5-8, where examples are given of timecourses typically found for different truncation levels. The left panel displays the timecourse for a source reconstructed near the location of the target source. The sinusoidal activity pattern was clearly visible, even though the pattern showed some distortions due to the noise in the data. The right panel contains the timecourses for two sources located close together. The timecourses were relatively large and almost self-cancelling, indicating that the solution contained too many sources.




figure 5-8: Timecourses typically found when different values of α were used to set the truncation level for noise contaminated data. The location of the constructed sources is given in the figures. The location of the target source on the gyrus was $(-0.137 -0.002 0.043)$ [m].

Solutions exhibiting this kind of behaviour could therefore be rejected on the basis of the timecourses (see also Fernández *et al*, 1995).

In conclusion, the best solution could be selected on the basis of the timecourses for the reconstructed sources. Sources with activity patterns that were very noisy or self-cancelling could be rejected. Note that the final ('best') solution obtained with this strategy is subjectively chosen.




figure 5-9: Normalised error measures for different values of α . The legend denotes the error measure that was used to obtain the corresponding curve. Only the solutions that were different from solutions found for other values of α are displayed. Hence, although only six solutions are displayed here, more values of α were used by FOCUS. Note that the solution with minimum χ_r^2 was at a different α to the solution that gave the minimum localisation error and the solution selected on the basis of the timecourses.

In order to evaluate the accuracy of the method, the mean location error and the standard deviation of the location error were compared for three different methods of picking the best solutions; based on the distance between the reconstructed sources and the target source, based on picking the solution with the minimum χ_r^2 and based on examination of the timecourses. An example of such a comparison is given in figure 5-9. In this case, the solution (subjectively) chosen on the basis of the timecourses was the same solution as selected on the basis of the distance to the target source. The solution selected on

Method	Mean Distance Error [mm]
Distance error	9.2 ± 6.9
χ_r^2 -value	18.6 ± 10.4
Timecourses	13.2 ± 9.4

table 5-1: Mean distance error for solutions found with FOCUSS for Monte Carlo simulations (SNR of 1.5). For each dataset a set of solutions is obtained for different truncation levels. The best solution is then selected with each method. This is repeated for all datasets in the Monte Carlo simulations and the mean location error is computed for each method. Note that the method based on the timecourses outperforms the method based on the value of χ_r^2 .

were used. Hence, also singular values corresponding to noise were used. Consequently, the solutions were more variable for the range of small values of α .

As demonstrated in the example given above, a dataset resulted in three different solutions. This was repeated for all datasets in a Monte Carlo simulation for data with an SNR of 1.5. The mean distance error and the standard deviation of the distance error are given in table 5-1. On examination of the solutions selected on the basis of the value of the χ_r^2 one learns that these solutions often contained a source near the target source and additional spurious sources due to the noise. These sources brought the value of the χ_r^2 down, but increased the overall distance error. These spurious sources could be rejected in the case when the correct number of sources was known. It is illustrated in figure 5-10 that the value of χ_r^2 could not be used to decide on the correct number of sources.

the basis of χ_r^2 was found at a lower value of α . Note that, unfortunately, this solution had a much larger localisation error than the solution at the optimum value for alpha.

It also follows from figure 5-9 that the solutions varied a lot more for smaller values of α than for larger values. This is explained as follows: for small values of α , many singular values

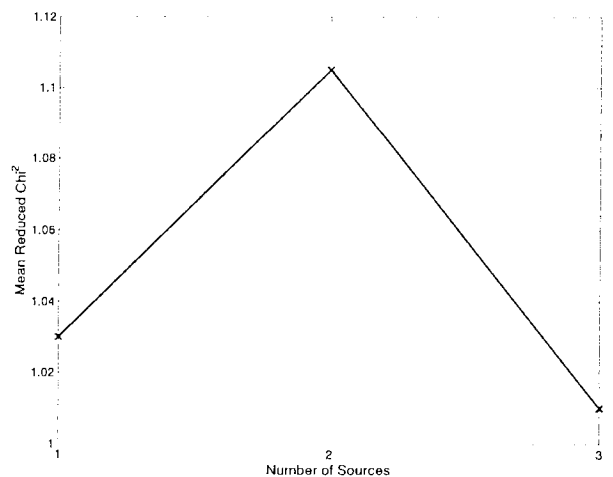


figure 5-10: The mean χ_r^2 for solutions with 1, 2, or 3 sources, respectively. For a dataset the solutions obtained with different truncation levels were grouped on the basis of the number of sources in the solution. Then, the mean value of χ_r^2 was computed for each set of solutions. Note that the correct number of sources was 1, but that the lowest mean χ_r^2 was found for solutions with three sources.

Conclusions

It was demonstrated in this section that selection of the value of the noise regularisation parameter could not be done by means of the L-curve. The shape of the L-curve became better defined in the first few iterations of FOCUSS, due to the increase in the effective ratio of the *number of channels* to the *number of elements in the source space*. However, a further increase of this ratio in subsequent iterations made the system overdetermined, rendering the L-curve approach obsolete.

Setting the truncation level in each iteration at a fixed percentage of the maximum singular value had the potential to improve the localisation accuracy. The exact value of this percentage had to be determined experimentally. This resulted in a set of solutions for a dataset. In simulations, selection of the best solutions on the basis of the minimum distance error is a possible. Picking the final solution on the basis of the minimum reduced chi-square did not give accurate solutions. This strategy can therefore not be used with real data. The failure of the reduced chi-square was due to spurious sources in the solutions. The reason behind the failure of the reduced chi-square approach was already discussed in the previous section; the noise truncation applied with FOCUSS largely removed the noise components from the data (as well as high spatial frequency data components), rendering the formalism of the reduced chi-square obsolete.

It was further demonstrated that the final solutions could be selected on the basis of the activity patterns of the reconstructed sources. The rules for the rejection of solutions were fairly straightforward and can be incorporated in an automatic rule based decision-making algorithm.

For the simulations that follow only the optimum solutions, based on the distance from the target sources, are given.

Different truncation levels were selected for different datasets, and in different iterations of FOCUSS, even though the SNRs were the same for the datasets. Hence, choosing the truncation level on the basis of the SNR of the data alone is not valid for this work. It was found that the shape of the L-curve depended on the effective ratio of *the number of sensors* to *the number of elements in the source space*. Selection of the truncation level based on the formula by Shim and Cho (1981) might therefore be a valid approach:

$$t_{\text{opt}} = \max_k \left\{ k \mid \lambda_k \geq \frac{MK_n}{NK_q} \right\}, \quad (5-5)$$

with N the number of elements in the source space, with non-zero amplitude (see section 2.3.2.3 for details). Further work is necessary to evaluate this.

5.3.3 Gridspacing

The FOCUSS algorithm and the constrained dipole fit algorithm both use a mesh in order to reconstruct the sources of electrical activity that produce the measured magnetic fields. For FOCUSS, the mesh is used to construct a source space, consisting of unitary dipoles located in the nodes in the mesh. The activity pattern for each of these sources is computed in an iterative approach (after convergence many of these sources have an amplitude of zero). For the constrained dipole fits the location of the dipoles is restricted to lie within a distance equal to the grid spacing from the nodes in the mesh and the orientation of the dipoles is equal to the normals of the nearest nodes in the mesh.

It is obvious that the gridspacing of the mesh influences the accuracy with which sources can be reconstructed. Undersampling of the source space results in a decrease in accuracy. On the other hand, oversampling increases the computation time significantly, especially for FOCUSS (due to the increase in the dimensions of the lead field matrix and as a result the time it takes to compute the inverse of the lead field matrix). There are several studies (van den Broek, 1997, Zanow, 1997) where the Boundary Element Method (BEM) or the Finite Element Method (FEM) was used to solve the forward problem. In these studies the grid spacing influenced the accuracy with which the forward problem could be solved. This in turn influenced the accuracy of source reconstructions. However, this did not reveal what the direct influence of the grid spacing is on the accuracy of source reconstructions for 2 reasons:

- 1) The grid was used as part of the forward problem and not to constrain the possible solutions when the inverse problem was solved.
- 2) These studies did not use a constrained source space as the inverse problem was solved using unconstrained dipole fit algorithms.

Other studies that incorporated a source space did not justify the choice of grid spacing that was used (Dale and Sereno, 1993, Lütkenhöner *et al.*, 1995, Fuchs *et al.*, 1998) or the grid spacing was chosen to be the same as the spatial resolution of the MEG system. Using a spatial constraint might however improve this resolution and therefore a smaller gridspacing would be beneficial. The only study found in the literature to examine the influence of gridspacing was performed by Ramon and colleagues (Ramon *et al.*, 1993). In a simulation study with a simple circular 2D source space it was found that decreasing the

grid spacing resulted in an ever diminishing improvement in resolution for their implementation of the minimum norm solution.

In the following simulations the forward problem was accurately solved; it was the influence of the choice of grid spacing on the accuracy of source reconstructions that was under examination.

Define spatial resolution as the distance needed between two sources in order to be able to resolve the activity from each source separately. The ability to resolve activity from two nearby sources depends on the difference in magnetic field they produce in the sensors. Hence, the spatial resolution is determined by the difference in the lead fields of both sources and consequently depends on the surface curvature. It was therefore hypothesised that the optimum grid spacing would be inversely related to surface curvature.

The influence of the grid spacing was evaluated for the target source located on the gyrus and for the target source located on the sulcal wall. Datasets with a mean SNR of 1.2, 1.5 and 3 were used in the Monte Carlo simulations, respectively. Several meshes were used, with the following values for the gridspacing: 2, 4, 6, 8 and 10mm. The results obtained with the unconstrained dipole fit algorithm were also given. It is obvious that the gridspacing does not influence the solutions for unconstrained dipole fitting. Consequently, the unconstrained fits had to be performed only once for each dataset. However, the results were plotted for the meshes with different grid spacing so that they could be compared with the results for the constrained algorithms.

A difference in the influence of the gridspacing on the solutions for the gyrus and sulcal target source could either be due to the fact that the gyrus source was more superficial than the source on the sulcal wall or due to a difference in surface curvature around these target sources. In order to reveal an influence of the surface curvature, the whole mesh was rotated by 90 degrees around the z-axis (figure 5-11). The mesh positioned this way did not represent a realistic situation. However, the target sources were now located at approximately the same average distance from the sensors, so any




figure 5-11: Source space viewed from the back (slightly shifted to the right). The channel array is depicted by the filled circles and the centers of the target sources on the gyrus and the sulcal wall are indicated by arrows. Note that the average distance to the sensors is approximately equal for both target sources.

differences found in the influence of the grid spacing on the solutions could be ascribed to differences in surface curvature around the target source

Results

FOCUSS

The influence of the gridspacing on the solutions for the gyral and sulcal source is shown in figure 5-12 and figure 5-13, respectively. It follows from these figures that:

- The mean distance error and the standard deviation of the mean distance error were almost independent of the gridspacing for the sulcal source. This was found for different SNRs.
- The mean distance error and the standard deviation of the distance error increased with an increase in gridspacing for the gyral source. The largest effect was found when the gridspacing became larger than 6mm. This increase in error took place in a similar way for data with different SNRs.
- For data with high SNR (1.5 or 3) the unconstrained dipole algorithm outperformed FOCUSS. For noisy data (SNR of 1.2) the FOCUSS algorithm outperformed the unconstrained dipole fit algorithm. This advantage disappeared for large gridspacings.

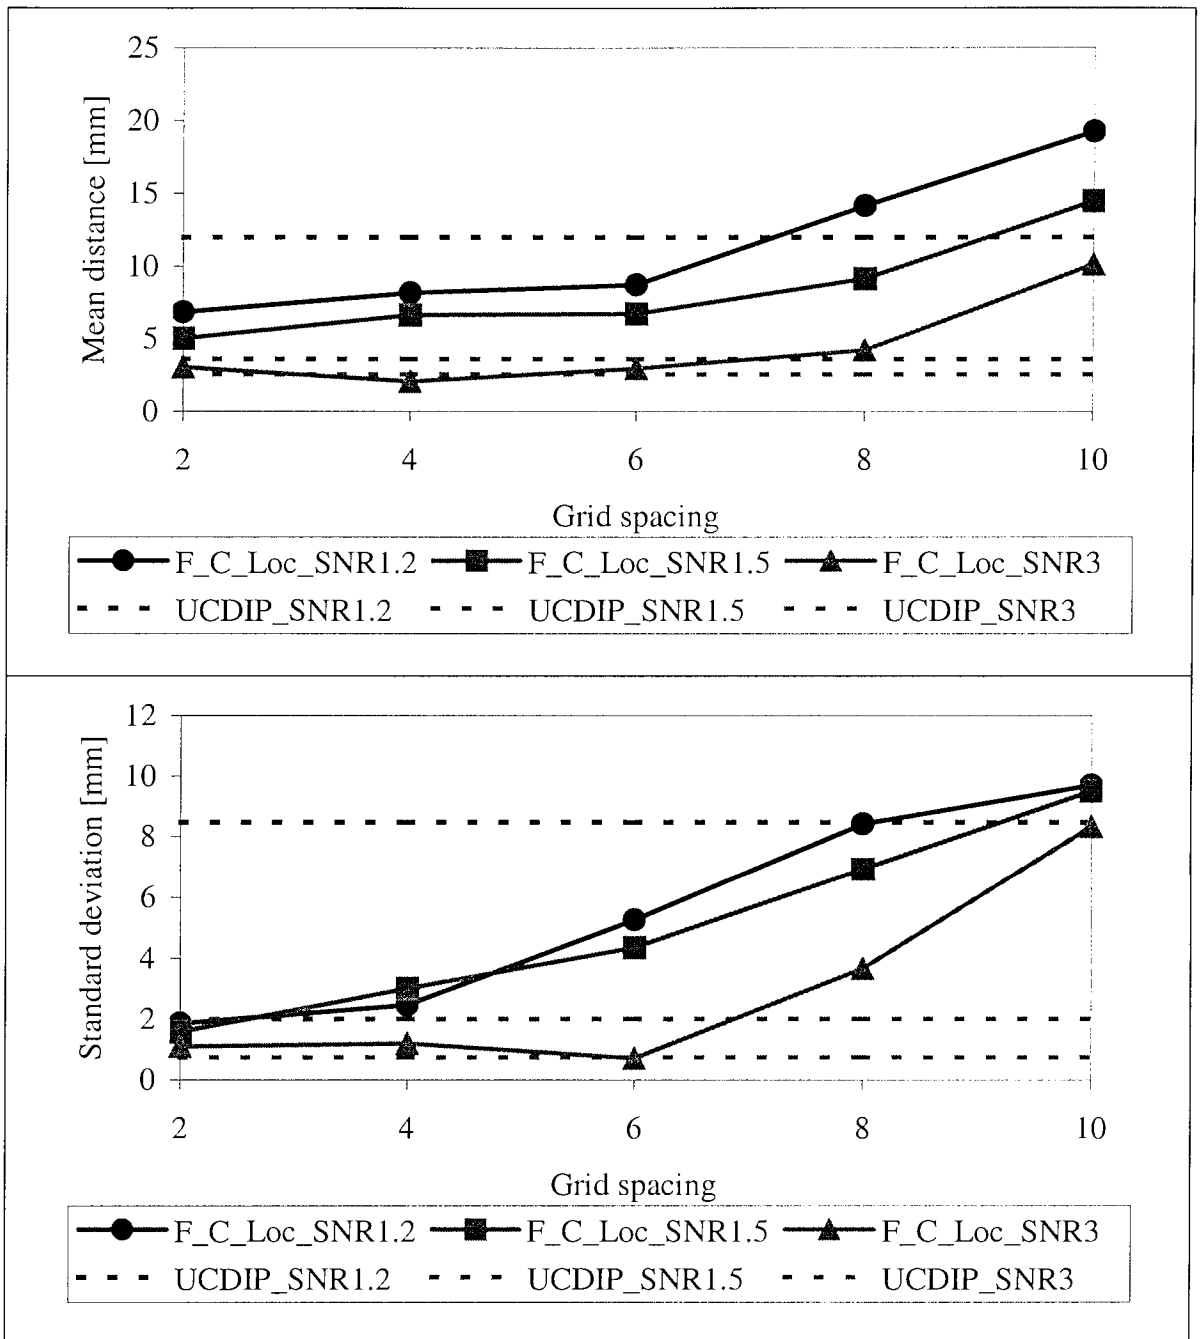


figure 5-12: Influence of grid spacing on source reconstructions with FOCUSS for data with different SNRs. The results for the unconstrained dipole fits are shown for comparison. The top panel shows the mean distance to the target, **gyral**, source for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note the increase in the mean distance error and the standard deviation of the distance error due to an increase in gridspacing.

F_C_Loc_SNR1.2, F_C_Loc_SNR1.5 and F_C_Loc_SNR3: Compound FOCUSS. Final solutions based on distance from target. Data with SNR of 1.2, 1.5 and 3, respectively; UCDIP_SNR1.2, UCDIP_SNR1.5 and UCDIP_SNR3: Unconstrained dipole fitting for data with SNR of 1.2, 1.5 and 3, respectively.

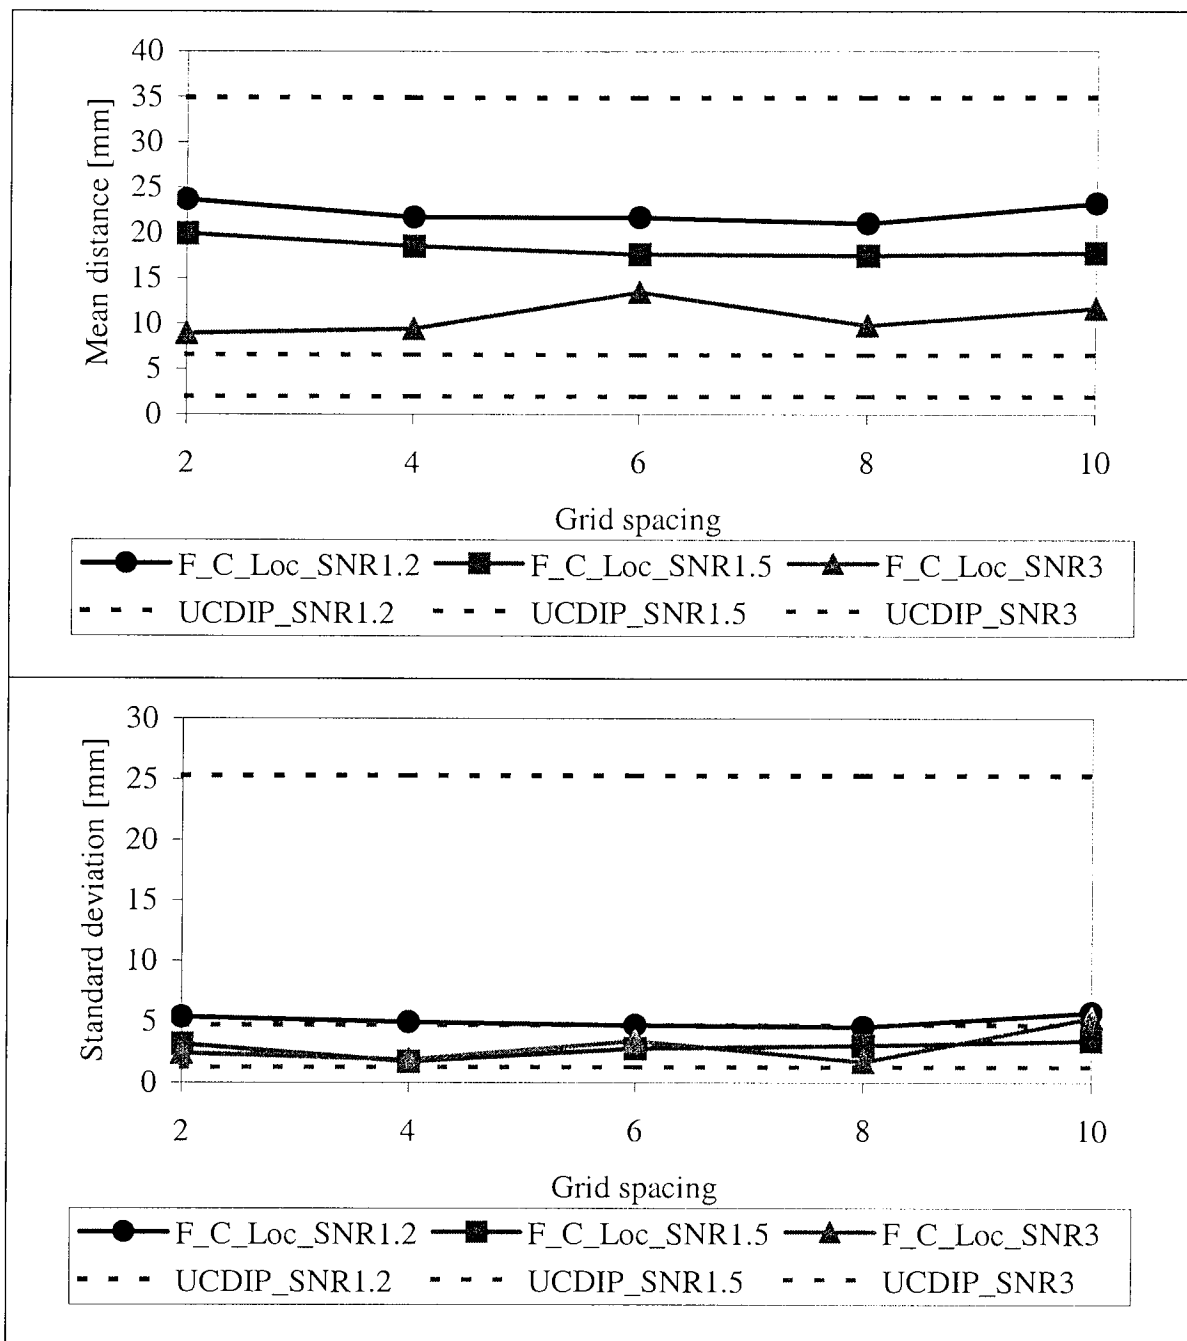


figure 5-13: Influence of grid spacing on source reconstruction with FOCUSS for data with different SNRs. The results for the unconstrained dipole fits are shown for comparison. The top panel shows the mean distance to the target, **sulcal**, source for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note that the results were almost independent of the gridspacing. *F_C_Loc_SNR1.2*, *F_C_Loc_SNR1.5* and *F_C_Loc_SNR3*: Compound FOCUSS. Final solutions based on distance from target. Data with SNR of 1.2, 1.5 and 3, respectively; *UCDIP_SNR1.2*, *UCDIP_SNR1.5* and *UCDIP_SNR3*: Unconstrained dipole fitting for data with SNR of 1.2, 1.5 and 3, respectively.

Constrained Dipole Fits

The influence of the gridspacing on the reconstructions with the constrained dipole fit algorithm are given in figure 5-14 and figure 5-15 for the gyral source and in figure 5-16 and figure 5-17 for the sulcal source.

The following results follow from these figures:

- A small influence of gridspacing on the performance of CDIPA was only found for noisy data (SNR of 1.2) for the gyral source.
- The results with CDIPB did not depend on the gridspacing. Only for a gridspacing of 10, there was an increase in the distance error. However, the standard deviation of the distance error was small for this gridspacing. This illustrates the trade-off between stable solutions due to a large gridspacing and inaccurate gridspacing due to undersampling.
- The constrained dipole fit algorithm gave more accurate and stable solutions than the unconstrained dipole fit algorithm, when noisy data (SNR of 1.2) was used. However, the advantage of using constraints was only marginal for data with an SNR of 1.5 or 3.




figure 5-14: Influence of grid spacing on constrained dipole fits for data with different SNRs. The results for the unconstrained dipole fits are shown for comparison. The top panel shows the mean distance to the target, **gyral**, source for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note that the influence of the gridspace was largest for data with low SNR.

CDIPa_SNR1.2, *CDIPa_SNR1.5* and *CDIPa_SNR3*: Constrained dipole fitting for data with SNR of 1.2, 1.5 and 3, respectively; *UCDIP_SNR1.2*, *UCDIP_SNR1.5* and *UCDIP_SNR3*: Unconstrained dipole fitting for data with SNR of 1.2, 1.5 and 3, respectively.




figure 5-15: Influence of grid spacing on constrained dipole fits (with the final solution constrained to the mesh) for data with different SNRs. The results for the unconstrained dipole fits are shown for comparison. The top panel shows the mean distance to the target, **gyral**, source for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. **CDIPb_SNR1.2**, **CDIPb_SNR1.5** and **CDIPb_SNR3**: Constrained dipole fitting with final solutions constrained to mesh. Data with SNR of 1.2, 1.5 and 3, respectively; **UCDIP_SNR1.2**, **UCDIP_SNR1.5** and **UCDIP_SNR3**: Unconstrained dipole fitting for data with SNR of 1.2, 1.5 and 3, respectively.




figure 5-16: Influence of grid spacing on constrained dipole fits for data with different SNRs. The results for the unconstrained dipole fits are shown for comparison. The top panel shows the mean distance to the target, *sulcal*, source for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note that the results were almost independent of the gridspacing. *CDIPa_SNR1.2*, *CDIPa_SNR1.5* and *CDIPa_SNR3*: Constrained dipole fitting for data with SNR of 1.2, 1.5 and 3, respectively; *UCDIP_SNR1.2*, *UCDIP_SNR1.5* and *UCDIP_SNR3*: Unconstrained dipole fitting for data with SNR of 1.2, 1.5 and 3, respectively.




figure 5-17: Influence of grid spacing on constrained dipole fits (with the final solution constrained to the mesh) for data with different SNRs. The results for the unconstrained dipole fits are shown for comparison. The top panel shows the mean distance to the target, **sulcal**, source for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. *CDIPb_SNR1.2*, *CDIPb_SNR1.5* and *CDIPb_SNR3*: Constrained dipole fitting with final solutions constrained to mesh. Data with SNR of 1.2, 1.5 and 3, respectively; *UCDIP_SNR1.2*, *UCDIP_SNR1.5* and *UCDIP_SNR3*: Unconstrained dipole fitting for data with SNR of 1.2, 1.5 and 3, respectively.

Results for the rotated mesh

FOCUSS

The influence of the gridspacing on the solutions for the gyral and sulcal source is shown in figure 5-18. It follows from this figure that:

- For the gyral source the distance error and the standard deviation of the distance error increased when the gridspacing was increased.
- The accuracy of the reconstructions of the sulcal source did not decrease with increasing gridspacing. Interestingly, an increase in accuracy was found for the range of gridspacings used in these simulations.
- The performances for the gyral source were better than for the sulcal source.
- The dipole fit algorithm outperformed the FOCUSS algorithm.
- The performances for the gyral and sulcal source were very similar for the unconstrained dipole fit algorithm.

Constrained Dipole Fitting

The influence of the gridspacing on the solutions for the gyral and sulcal source is shown in figure 5-19. It follows from this figure that:

- For the gyral source there was an increase of the distance error with an increase in gridspacing. The effects of changes in gridspacing on the standard deviation of the distance error were small. This was found for both versions of the algorithm.
- The influence of the gridspacing was small for localisation of the sulcal source.

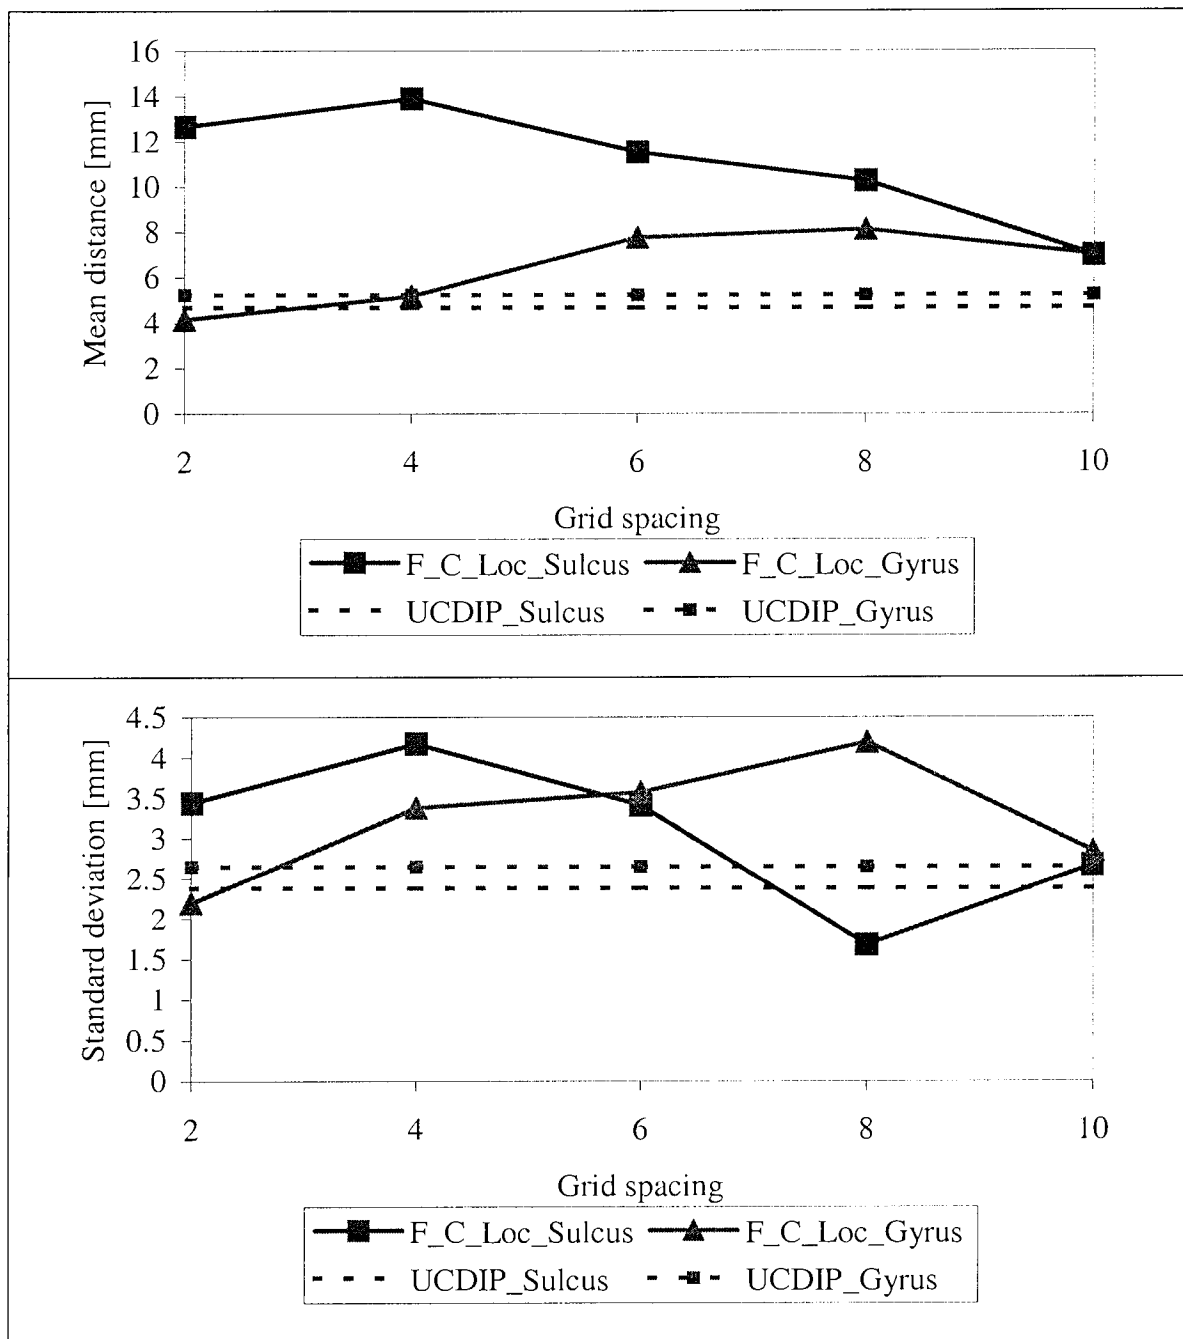


figure 5-18: Rotated mesh with FOCUSS. Influence of grid spacing on the sulcal and gyral target sources for data with an SNR of 1.5. The results for the unconstrained dipole fits are shown for comparison. The top panel shows the mean distance to the target sources for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note that the results for the gyral source were influenced in a different way as the results for the sulcal source. *F_C_Loc_Sulcus* and *F_C_Loc_Gyrus*: Compound FOCUSS. Final solutions based on distance from target. Target source is the sulcal and gyral source, respectively; *UCDIP_Sulcus* and *UCDIP_Gyrus*: Unconstrained dipole fitting for the sulcal and gyral target source, respectively.




figure 5-19: Rotated mesh with constrained dipole fits. Influence of grid spacing on the sulcal and gyral target sources for data with an SNR of 1.5. The results for the unconstrained dipole fits are shown for comparison. The top panel shows the mean distance to the target sources for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. *CDIPa_Sulcus* and *CDIPa_Gyrus*: Constrained dipole fitting for the sulcal and gyral target source, respectively; *CDIPb_Sulcus* and *CDIPb_Gyrus*: Constrained dipole fitting for the sulcal and gyral target source, respectively. Final solutions constrained to mesh; *UCDIP_Sulcus* and *UCDIP_Gyrus*: Unconstrained dipole fitting for the sulcal and gyral target source, respectively.

Discussion and conclusions

FOCUSS

Using the realistically orientated mesh, the influence of the gridspacing was different for the sulcal and gyral source. Possible reasons for this difference are discussed here:

- It is unlikely that the differences were due to SNR, because for the localisation of the gyral source, using data with an SNR of 1.2, the distance errors were of similar amplitude as for localisation of the sulcal source, using data with an SNR of 3. Still, there was only an influence of gridspacing for the localisation of the gyral source.
- It is unlikely that the differences were due to the distance to the sensors, because rotating the mesh preserved the differences in performance for the localisation of the gyral and sulcal source. In the rotated mesh the average distance to the sensors was approximately equal. The fact that the two target sources in the rotated mesh gave very similar results with the unconstrained dipole fit algorithm further supports this.
- It is likely that the differences were due to the differences in surface curvature near the two target sources. The surface curvature was high for the gyral source. Consequently, the lead fields of the sources in this area were changing rapidly. In order to incorporate this information into the source space, a fine gridspacing was needed. This would explain the decrease in accuracy when the gridspacing was increased. The fact that the accuracy did not increase much when the gridspacing was reduced to a value below approximately 6mm could indicate that the angle between the lead field vectors became too small to be detectable by the sensor configuration. Hence, a further decrease in gridspacing was futile, due to the inherent properties of the measurement system. The relationship between gridspacing and surface curvature would also explain the fact that the gridspacing had little influence for the sulcal source in the realistically orientated mesh. The sulcal wall was very 'flat'. As a consequence, the lead fields did not vary much in this area and the accuracy of the source reconstructions was little influenced by the gridspacing. However, it is not yet understood why the accuracy of the reconstruction of the sulcal source in the

rotated mesh increased with an increase in gridspacing. A possible explanation can be sought in the change in the shape of the lead fields of the mesh elements itself. A way to explore this idea further is by use of the concept of the resolution matrix and the figures of merit described in the previous chapter. A change in the shape of the lead fields changes the resolution matrix. These changes might explain the way in which the gridspacing influenced the accuracy of source reconstructions with FOCUSS for different target sources and different orientations of the mesh. Further research in this direction is needed.

The reconstruction of the gyral source was more accurate than for the sulcal source, for both orientations of the mesh.

The influence of the gridspacing was unrelated to the SNR of the data.

Constrained Dipole Fits

The source reconstructions of the sulcal source were hardly influenced by the gridspacing. For the gyral source the performance generally decreased with an increase in gridspacing. This difference in dependency on the gridspacing for gyral and sulcal sources can be explained by the difference in surface curvature for these sources (see the discussion given above).

A trade-off was sometimes found between an increase in stability due to large gridspacing and a decrease in accuracy due to undersampling.

Comparison with Unconstrained Dipole Fits

The FOCUSS algorithm only outperformed the unconstrained dipole fit algorithm for data with an SNR of 1.2. This advantage disappeared for large gridspacings.

It was advantageous to use the constrained dipole fit algorithm when data with low SNR was used. The accuracy and stability gained by incorporating the constraints vanished when the SNR of the data increased. Additionally, the difference in performance between version CDIPA and CDIPB was small.

Optimum Gridspacing

For FOCUSS, the optimum gridspacing depends on the curvature of the surface. The spatial sampling should be high enough to incorporate the information provided by the variations in the lead fields across the surface. The highest spatial frequency that is beneficial is however determined by the shape of the lead fields and therefore by the measurement set-up. Exceeding this limit is disadvantageous because of an increase in computing time. A spatial sampling that is too coarse results in decreased performance due to undersampling. For the mesh and sensor configuration used in these simulations the minimum spatial frequency that was effective was 6mm.

The performance of the constrained dipole fit algorithm was maximal with the smallest gridspacing, i.e. 2mm. However, a large gridspacing can be advantageous in some cases as it can stabilise the solutions.

5.4 Influence of Spatial Errors

5.4.1 Position Errors

A realistically shaped source space is normally obtained from a subject's MRI. Accordingly, the use of spatial constraints in the inverse problems relies on accurate co-registration of the MEG and MRI co-ordinate systems. Errors in co-registration can range from several millimetres to a centimetre or more. In this section the influence of spatial errors on the accuracy of the constrained source localisation algorithms was evaluated. It was hypothesised that large spatial errors result in unacceptably large errors in locations of the reconstructed sources.

In these Monte Carlo simulations, simulated data with an SNR 1.5 was computed for the target source on the gyrus and on the sulcal wall, respectively. The mesh and sphere origin were then shifted over a certain distance in a randomly chosen direction, with the value of this distance taken from a Gaussian distribution with a certain standard deviation and zero mean. The standard deviations used in these simulations ranged from 2, 4, 6, 8 to 10mm. These shifts in the mesh location and the origin represented errors in the location of the mesh due to errors in the co-registration of the MEG and MRI co-ordinate systems. The shifted mesh was subsequently used by the constrained source reconstruction algorithms to find the solutions for the simulated data. As before, the solutions obtained with the unconstrained dipole fit algorithm for the datasets with these SNRs were given for comparison. The results for the unperturbed mesh/origin were also given for comparison.

Results

The influence of the mesh perturbations on the reconstructions of the gyral source and the sulcal source are given in figure 5-20 and figure 5-21, respectively.

FOCUSS

- The source reconstructions for the gyral source were hardly influenced by perturbations of the mesh. Only for perturbations with a standard deviation of 8mm there was an

- increase in distance error of 1.3mm compared to the unperturbed case. However, this increase in error was only 16% of the standard deviation of the perturbation error.
- The mean distance error of the source reconstructions for the sulcal source was also hardly influenced by the mesh perturbations. Additionally, the standard deviation of the distance error due to a mesh movement of 10mm increased only by 23% of this movement.

Constrained Dipole Fits

- The influences of mesh perturbations on the reconstructions of the gyral source with CDIPA was less than a millimetre, even for perturbations as large as 10mm. Similar results were obtained for the sulcal source.
- Reconstructions with the CDIPB version of the constrained dipole fit algorithm decreased in accuracy when the perturbations increased, both for the gyral and the sulcal source. The largest increase for the gyral source was 35mm, found when the mesh was perturbed by 10mm. Note that this was still only 35% of the standard deviation of the perturbations of the mesh. Similarly, for the sulcal source the largest increase in distance error was 53% of the standard deviation of the mesh perturbations.

Conclusions

The influence of mesh perturbations on the source reconstructions with FOCUSS and the CDIPA version of the constrained dipole fit algorithm were in the order of 1mm and therefore negligible compared to the maximum distance by which the mesh was perturbed. Apparently, the algorithms were able to use the elements in the perturbed mesh that had been shifted towards the location of the target source to reconstruct the source activity accurately.

The accuracy of the location and the stability of the sources reconstructed with the CDIPB version of the constrained dipole fit algorithm decreased when the standard deviation of the mesh perturbations was increased. The largest increase in the distance error was found to be 5.3mm for perturbations of the mesh by 10mm. Moreover, the advantage of this version of the algorithm over unconstrained dipole fitting vanished when the standard deviation of the mesh perturbations were larger than 6mm.

The gridspacing used here was not the optimum value for the FOCUSS algorithm. Nevertheless, the influence of mesh perturbations was still small, even when the standard deviation of the perturbations was larger than the gridspacing. It is therefore expected that the results found here for FOCUSS are also valid for meshes with different gridspacings. A similar argument is valid for the CDIPA version of the constrained dipole fit algorithm. For the CDIPB version of the dipole fit algorithm the accuracy of the source reconstructions decreased with an increase in the standard deviation of the mesh perturbations. It was shown in section 5.3.3 that the influence of gridspacing on the accuracy of localisations with CDIPB was small. It is therefore expected that mesh perturbations influence the accuracy of source localisations in a similar way as described above in case different gridspacings were to be used.

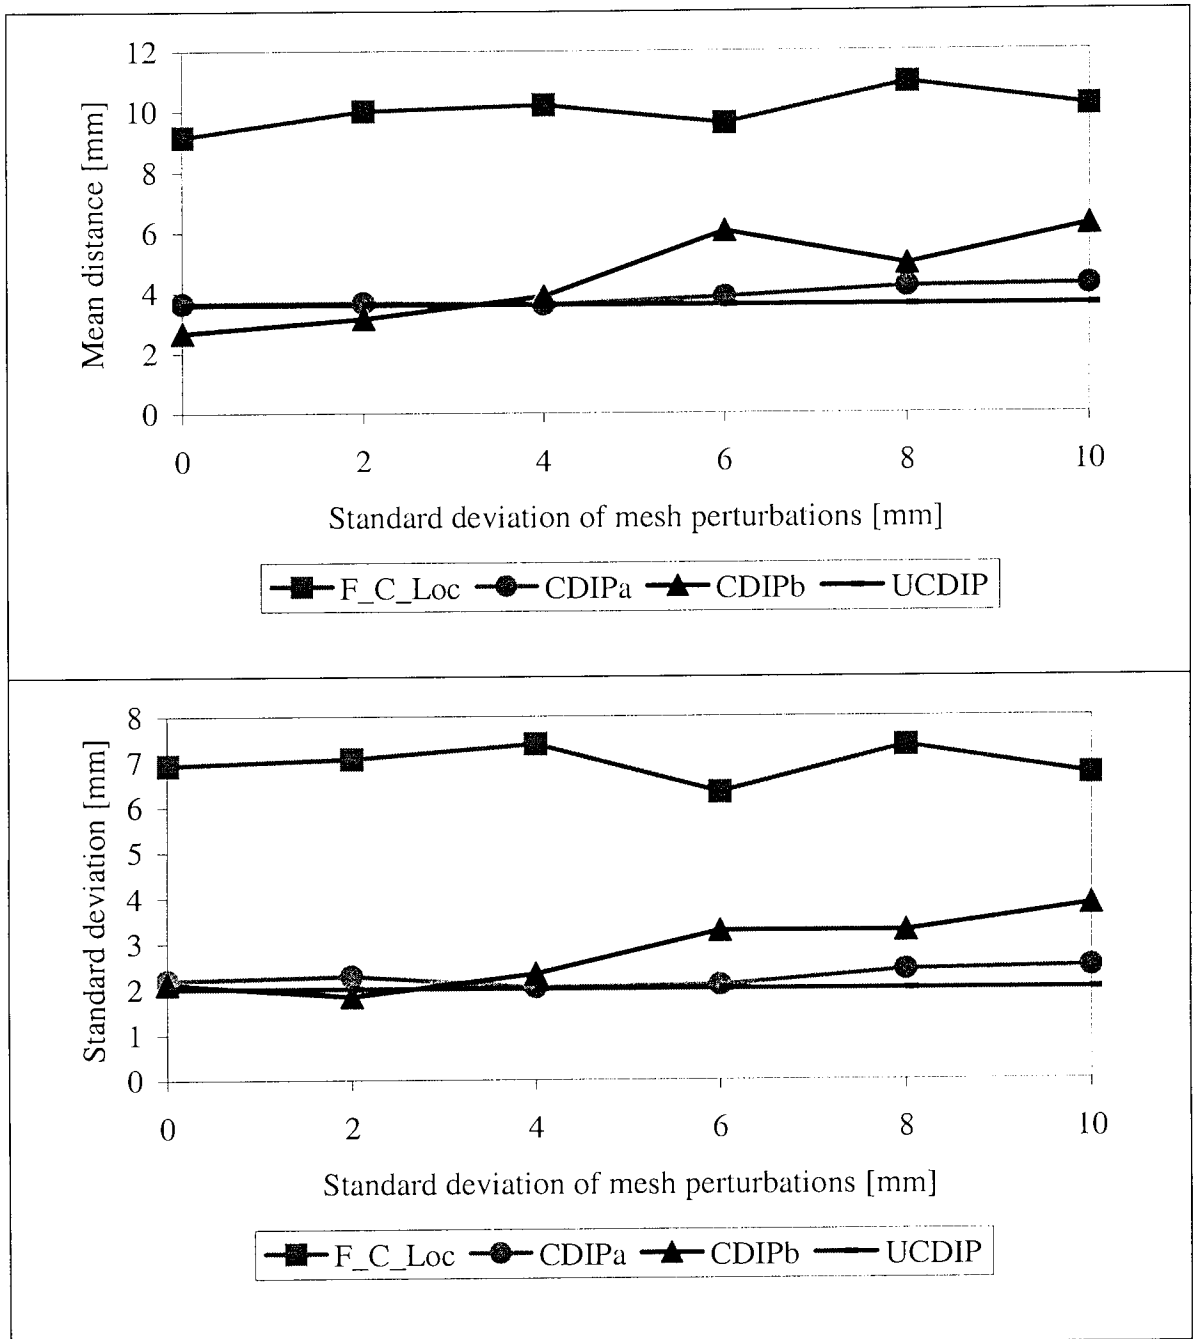


figure 5-20: Influence of displacement errors on the source reconstruction with the three different algorithms. The target source is the **gyral** source. The top panel shows the mean distance to the target sources for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note that the mesh perturbations only had a large influence on CDIPb. . **F_C_Loc:** Compound FOCUSS. Final solutions based on distance from target; **CDIPa:** Constrained dipole fitting; **CDIPb:** Constrained dipole fitting with final solutions constrained to mesh; **UCDIP:** Unconstrained dipole fitting.




figure 5-21: Influence of displacement errors on the source reconstruction with the three different algorithms. The target source is the **sulcal** source. The top panel shows the mean distance to the target sources for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note that the mesh perturbations only had a large influence on CDIPb. *F_C_Loc*: Compound FOCUSS. Final solutions based on distance from target; *CDIPa*: Constrained dipole fitting; *CDIPb*: Constrained dipole fitting with final solutions constrained to mesh; *UCDIP*: Unconstrained dipole fitting.

5.4.2 Orientation Errors

Errors in the construction of the source space can cause local errors in orientation of the surface normals. Consequently, the orientations of the normals of the grid-nodes that are used by the constrained source reconstruction algorithms can be erroneous. For example, the erroneous inclusion of a voxel from the CSF into the model of the cortex causes such a local error. Computing the average orientation of an area of cortex reduces the potential influence of such an error, but does not remove it completely. In this section it was examined how stable the constrained localisation algorithms were with respect to local perturbations in the orientation of the normals of the nodes in the mesh. Based on evidence found in the literature (Liu *et al*, 1998), it was hypothesised that the influence of these perturbations would be relatively small.

The influence of perturbations of the normals in the nodes of the mesh was examined by means of Monte Carlo simulations with data with an SNR of 1.5. The normal in each node was perturbed by an angle about the axis formed by the normal vector. The value of this angle was taken from a Gaussian distribution with a given standard deviation and zero mean. The values of the standard deviation used for this Gaussian were 10, 20, 30 and 40°, respectively. The latter value was an extreme case that was not expected in realistic cases. The FOCUSS algorithm and the constrained dipole fit algorithm were used to reconstruct the sources from simulated data produced by the gyral and sulcal sources. As in the previous sections, the results obtained with the unconstrained dipole fit algorithm were displayed for comparison. The results for the mesh with the unperturbed normals were also given for comparison.

Results

The influence of the perturbations of the surface normals on the reconstructions of the gyral source and the sulcal source are given in figure 5-22 and figure 5-23, respectively.

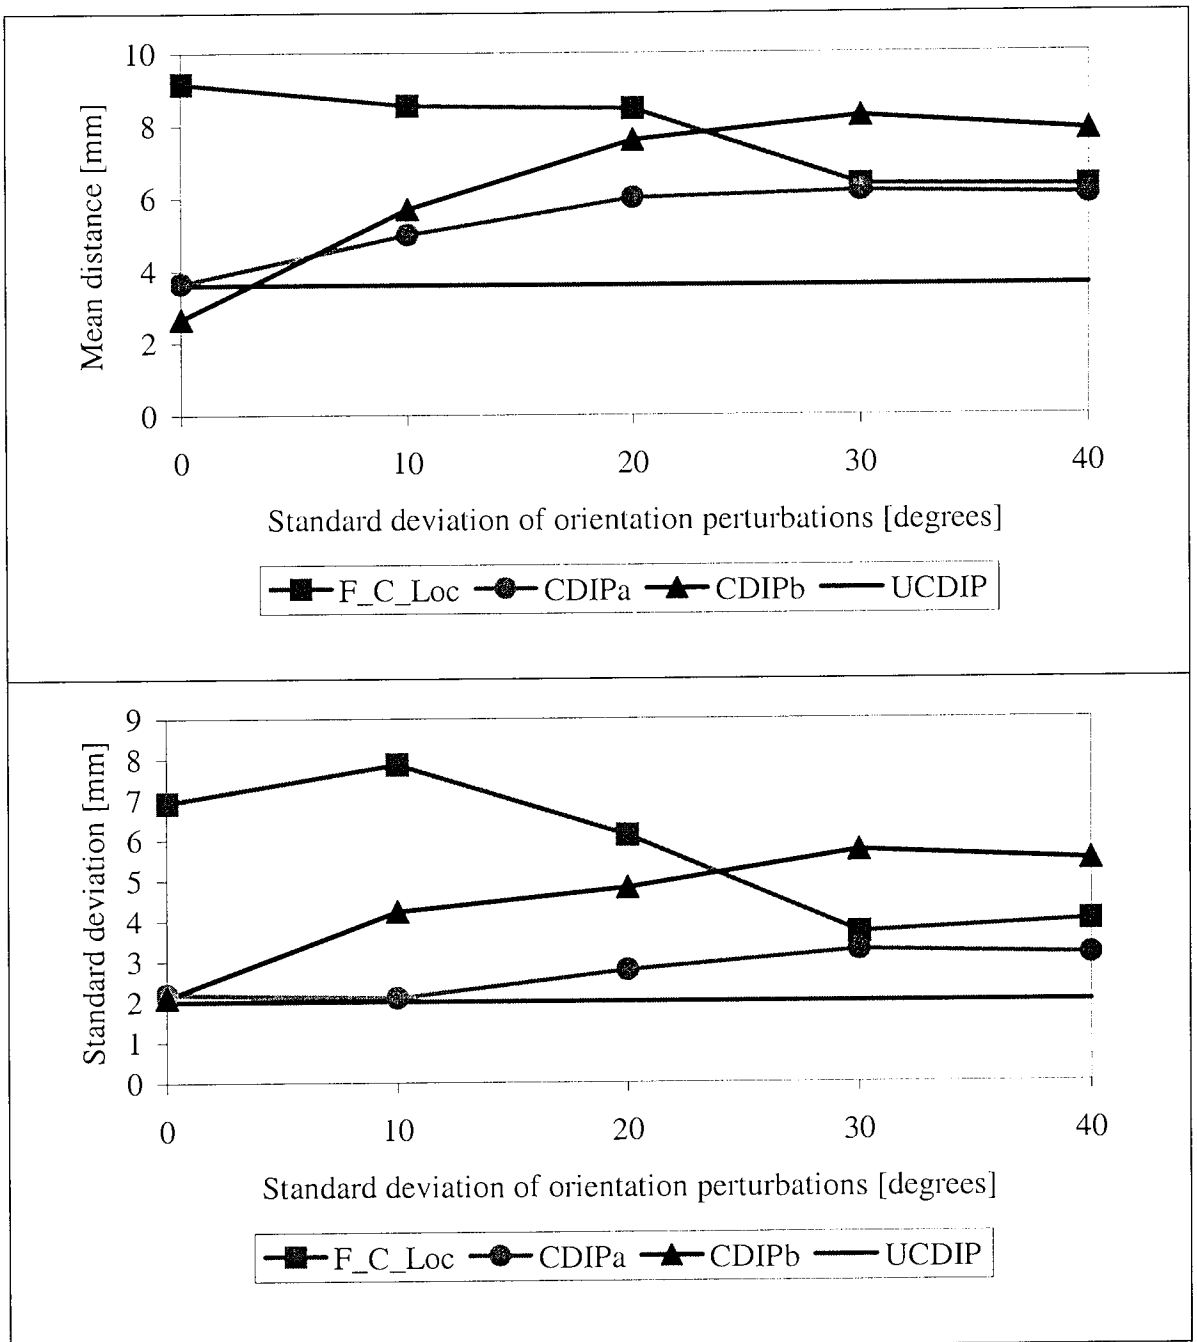


figure 5-22: Influence of errors in the orientation of the surface normals on the source reconstruction with the three different algorithms. The target source is the **gyral** source. The top panel shows the mean distance to the target sources for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance. Note the increase in performance for FOCUSS and the decrease in performance for the constrained dipole fits, due to the perturbations of the surface normals.
F_C_Loc: Compound FOCUSS. Final solutions based on distance from target; *CDIPa*: Constrained dipole fitting; *CDIPb*: Constrained dipole fitting with final solutions constrained to mesh; *UCDIP*: Unconstrained dipole fitting.

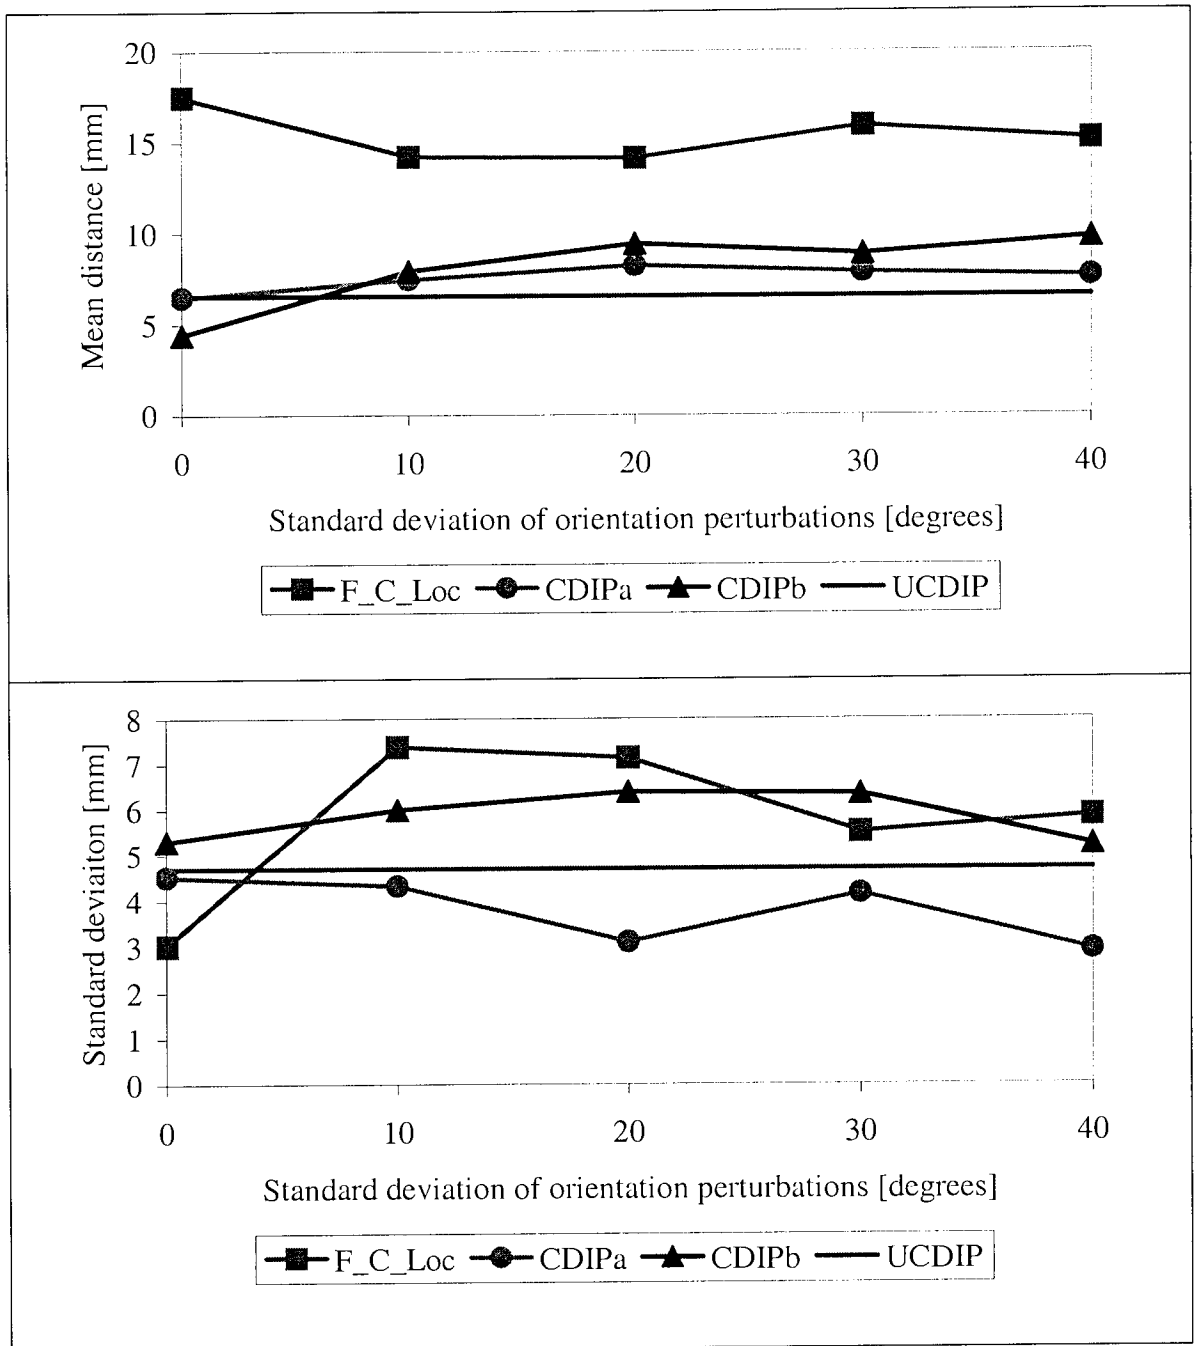


figure 5-23: Influence of errors in the orientation of the surface normals on the source reconstruction with the three different algorithms. The target source is the **sulcal** source. The top panel shows the mean distance to the target sources for all solutions from the Monte Carlo simulations. The bottom panel shows the standard deviation of this distance.

F_C_Loc: Compound FOCUSS. Final solutions based on distance from target; *CDIPa*: Constrained dipole fitting; *CDIPb*: Constrained dipole fitting with final solutions constrained to mesh; *UCDIP*: Unconstrained dipole fitting.

FOCUSS

- There was no degradation in performance when the standard deviation of the perturbations of the surface normals was increased. In fact, for the gyral source a slight improvement was observed. There appeared to be a saturation effect.

Constrained Dipole Fits

- A decrease in performance was found when the standard deviation of the perturbations of the surface normals was increased. There seemed to be a saturation effect.
- The effect of perturbations of the surface normals was more pronounced for the CDIPB version of the dipole fit algorithm than for the CDIPA version.
- Using CDIPB, the maximum increase in the distance error due to the perturbations of the surface normal was 5.6mm for localisation of the gyral source and 5.3mm for localisation of the sulcal source.

Discussion

FOCUSS

Generally, errors in the orientation of the surface normals were found to have no deteriorating effect on the source reconstructions with FOCUSS. Interestingly, a slight improvement in the FOCUSS algorithm was observed. This can be explained on the basis of the concept of the resolution matrix. When the surface normals are unperturbed, many surface elements in the mesh have normals with the same or a similar orientation, due to the way the mesh is formed. The lead fields of the elements with similar orientation are similar as well (if the distance between the elements is small). Consequently, the resolution kernels for these elements consist of many sidelobes. Due to the perturbations of the surface normals, the similarity of the lead fields is reduced and consequently the number of sidelobes in the resolution kernels will be reduced.

There appeared to be a saturation effect. A possible explanation is that the lead fields are not only determined by the orientation of the surface elements, but also by their location and the characteristics of the sensor array. The improvement of the resolution

kernels is thus limited by the measurement set-up. Due to this limit the source reconstructions can not be improved by further randomisation of the surface normals.

These results suggest that the influence of errors in determining the surface normals for a realistically shaped mesh will be small or that they can even increase the accuracy of source reconstructions due to randomisation effect on the lead fields and the improvement of the resolution kernels as a consequence.

Constrained Dipole Fits

The perturbations of the surface normals influenced the accuracy of the source reconstructions considerably, increases in distance errors of up to 5.6mm were found. These increases were more pronounced for the CDIPB version of the dipole fit algorithm than for the CDIPA version, due to the fact that the CDIPB version is stricter bound by the constraints.

The effect of perturbations of the surface normals was so large that the advantage of the constrained dipole fit algorithm over the unconstrained algorithm disappeared even for perturbations with a standard deviation of only 10° . Additionally, the improvement of FOCCUS due to perturbations of the surface normals ensured that for the gyral source its performance was better than for the performance of CDIPB and similar to the performance of CDIPA when the perturbations were larger than 20° .

5.5 Conclusions

The simulations performed in this chapter, using a simulated cortical fold, revealed that the source reconstructions with the compound version of FOCUSS were, compared to the results with the non-compound version, less scattered and that the mean distance between the solutions and the target source was smaller. This is in agreement with reports that the compound version is more accurate and robust than the non-compound version (Gorodnitsky *et al*, 1995).

It was demonstrated that noise regularisation can be applied by setting the truncation level at a fixed percentage of the maximum singular value in each iteration of FOCUSS. The optimum value for this percentage could be determined experimentally. Selection of the final solution for FOCUSS on the basis of the minimum reduced chi-square gave inaccurate solutions, due to spurious sources in the solutions with lowest reduced chi-square. The, subjective, selection of the final solutions on the basis of the activity patterns of the reconstructed sources was proven to be a satisfactory alternative. This approach can be implemented in an automatic rule based decision-making algorithm.

The accuracy of source reconstructions with FOCUSS decreased with increasing gridspacing. The influence was more pronounced for areas with large surface curvature. A smaller gridspacing is needed in order to incorporate the information provided by the changes in the lead field in these areas. The maximum spatial resolution that can be achieved seems to be determined by the shape of the lead fields in an area and therefore by the measurement set-up and the region of the source space where the activity is reconstructed. The influence of gridspacing on the performance of the constrained dipole fit algorithm was generally small. Only for data with low SNR and for an area with large surface curvature a decrease in accuracy was found due to an increase in gridspacing.

MEG-MRI co-registration errors shift the mesh with respect to the sensors. However, it was demonstrated that the influence of errors as large as 10mm was negligible. The increases in localisation errors due to perturbation errors were namely in the order of 1mm, both for FOCUSS and the constrained dipole fit algorithm. On the contrary, perturbations in the orientation of the surface normals increased the localisation error by as much as 5.6mm when the constrained dipole fit algorithm was used.

No increase in error was found for FOCUSS. This is in agreement with the findings of Liu and colleagues (Liu *et al*, 1998) who reported that the effect of perturbations of the surface normals on the accuracy of source reconstructions was only small with their linear estimator.

Interestingly, a small improvement was found for FOCUSS when the surface normals were perturbed. This could be explained by an improvement of the resolution kernels.

It is possible that the localisation accuracy for the constrained dipole fit algorithm could be improved by using a 'soft' constraint for the orientation of the sources (Schmidt *et al*, 1999), i.e. the angle the orientation vector makes with the surface normal is constrained, instead of the orientation vector itself. Further work in this direction is needed.

In conclusion, the comparison of FOCUSS, unconstrained and constrained dipole fit algorithms revealed that adding constraints improved the accuracy of reconstructed sources when noisy data was used. This improvement disappeared for constrained dipole fitting when the SNR of the data was increased or when the surface normals were perturbed. With FOCUSS the advantage disappeared for data with high SNR and when the gridspacing was increased.

6 Source Localisation from Visual Evoked MEG Data

6.1 Introduction

The aim of this study was to evaluate the performance of several source localisation algorithms under realistic circumstances.

The algorithms that were described and evaluated in the previous chapter were applied to averaged evoked MEG recordings obtained with a whole head Omega neuromagnetometer. The use of real data tests the performance of the source localisation algorithms in a situation that is impossible to re-create exactly by simulations.

The chosen experiment concerned the measurement of the response evoked by a brief visual stimulus. The visual modality was used as the modality of choice, because the activated cortical areas can be accurately modelled by equivalent current dipoles (Aine *et al.*, 1989, George *et al.*, 1989, Aine *et al.*, 1996, Barnes *et al.*, 1998). Hence, the use of dipole fit algorithms is feasible. Additionally, the primary visual evoked response is known to involve several closely spaced cortical areas. Resolving the locations and activation sequences of sources in these areas poses a test on the feasibility and limits of the source reconstruction algorithms. The experiment described here is part of an ongoing study to examine cortical magnification by means of MEG (Barnes *et al.*, 1998).

The measurement framework was first examined by an analysis of the figures of merit derived from the resolution matrix. This was followed by source reconstructions by means of unconstrained and constrained dipole fitting, the minimum norm algorithm and FOCUSS.

The next session describes the experimental methods and the analysis strategies that were used. This is followed by a presentation of the results. In the final section the results are discussed.

6.2 Methods

Stimulus Design

The visual stimulus was a Gaussian windowed patch of sinusoidal grating (see figure 6-1), generated using a VSG 2/3 board (Cambridge Research Systems). The spatial frequency of the grating was 2cpd. Two different stimulus conditions were used in the experiment. In one condition the standard deviation of the Gaussian window was 1.67 degrees and in the other condition the standard deviation was 0.83 degrees. These two stimulus conditions were randomly interleaved.

The stimuli were presented in the lower right visual field, at an eccentricity of 1.5 degrees from both the vertical and horizontal meridian. The subject was asked to fixate on a white fixation point with dimensions of 0.25 degrees. The patches were presented in a pattern onset paradigm onto a background of the same mean luminance (42 cd/m^2). The contrast was 70%.

Each stimulus was presented for a period of 500ms and the interstimulus interval was set to 2 ± 0.2 seconds.




figure 6-1: Gaussian windowed patch of sinusoidal grating with spatial frequency of 2cpd and fixation spot in top right corner. The standard deviation of the Gaussian window is 1.67 degrees. Note that this visual stimulus was viewed through a mirror during the MEG experiment, thus it was presented in the lower right visual field.

MEG Recordings

MEG data were recorded using a 151 channel CTF Omega whole head neuromagnetometer. An experienced and co-operating subject (IEH) sat comfortably in an upright position. Three magnetic head coils were attached to the head with a Velcro band that was wrapped around the head. This allowed for the quick and secure placement of the head coils. One coil was positioned on the forehead and the other two coils near the left and

right ear, respectively. An inflatable head-cuff was placed over the subject's head. This secured the head within the MEG helmet once inflated and thereby reduced the movement of the head during the MEG recordings. The head coils were activated before and after the recordings and the locations of the coils with respect to the MEG sensors were estimated. The amount of head movement was determined by comparing the coil locations as estimated before and after the recordings. The total head movement was less than 2mm in this experiment.

A total of 200 epochs were recorded, with a pre-stimulus interval of 190ms and a post-stimulus interval of 800ms. The data were recorded with the 3rd order gradiometer configuration in order to reduce contamination of the brain signals by external noise sources. The data were sampled with a frequency of 1250Hz, using an anti-aliasing filter with cut-off frequency of 400Hz.

After scrutiny of the raw data for artefacts, the data were base-line adjusted (based on the pre-stimulus interval) and averaged. The responses to the two different stimuli were averaged separately, resulting in two sets of averages made up of 100 epochs each. The plus-minus average was also computed to serve as an estimate of the noise in the data. The averaged data were subsequently filtered with a 50Hz comb filter to remove interference from the powerlines (and the higher order harmonics). Furthermore, a low-pass filter of 80Hz was applied.

The locations of the coils with respect to the subject's bitebar were measured with a 3D Polhemus digitiser. The subject's headshape was also digitised in this co-ordinate system in order to form a set of hat points. The positions of the coils and the head points were subsequently transformed to the MEG co-ordinate system that was based on the location of the three head coils see CTF Systems Inc. (1999) for the definition of this co-ordinate system). Also, the scalp surface was extracted from the subject's MRI, using the contouring method described in section 3.2.2. The set of digitised hat points was matched to the scalp surface in order to get the transformation matrix needed for the co-registration of the MEG and MRI co-ordinate systems.

The set of hat points was also used to obtain a best fitting sphere. This sphere was used as the volume conductor model of the head.

***A priori* Information**

The realistically shaped source space used throughout this work, which was incidentally based on IEH's MRI, was used as the source space for the FOCUSS algorithm and provided the spatial and orientation constraints for the constrained dipole fit algorithm (figure 6-2). See section 3.2 for a description of the methods used in order to obtain this source space. The source space was transformed from the MRI co-ordinate system to the MEG co-ordinate system, so that the constraints could be used with the recorded MEG data.

Source Localisation Algorithms

The source localisation algorithms described and evaluated in the previous chapter were used for the reconstruction of stimulus related electrical activity in the brain. See section 5.2 for implementation details. The following modifications were made:




figure 6-2: Vertices of the realistically shaped source space shown in a coronal slice from the subject's MRI

FOCUSS

Three different weighting matrices were used at the start of the algorithm.

- Minimum norm weighting, i.e. the identity matrix was used as the initial weighting matrix.
- Smoothed norm weighting in order to avoid bias towards the sensors (this weighting matrix was introduced in section 4.3.2.2).
- Average distance weighting, as described in the previous chapter. The weighting is based on the squared average distance to the sensors and an additional correction factor. This reduced the bias towards the sensors.

Unconstrained Dipole Fits

In the previous chapter constraints were made effective by squaring the value of the reduced chi-square in case one of the constraints was violated. For measured data this can be ineffective, since the value of the reduced chi-square can fall below 1, for example when the source model is not correct for a certain time-window. Therefore, violations of constraints were penalised by multiplying the value of the reduced chi-square by a factor of 10.

Constrained Dipole Fits

As for the unconstrained dipole fits, violations of constraints were penalised by multiplying the value of the reduced chi-square by a factor of 10.

Unlike for the simulated mesh used in the previous chapter, the realistically shaped mesh had a variable gridspacing. The gridspacing was smaller than 3mm for all but one element, with a mean value of 1.0mm. An uncertainty of 2mm in the mesh location was also taken into account (section 3.4.2.4). The sources were therefore restricted to lie within 5mm from the nodes in the mesh.

The best solutions obtained from the multi-start dipole fit algorithms were used as initial guesses in Monte Carlo simulations. These simulations revealed the stability of the solutions with respect to different noise realisations of the data. For this purpose, zero mean, Gaussian noise was added to the recorded data. The standard deviation of the noise as measured in each channel was used as the standard deviation for the Gaussian. The implementation of the source reconstruction algorithms in MATLAB meant that it was only feasible to use 10 different noise realisations in the Monte Carlo simulations.

In addition to the above-mentioned methods, the (weighted) minimum norm solution was also computed (note that this is the solution obtained in the first iteration of the FOCUSS algorithm). Furthermore, the moving dipole algorithm implemented in CTF's acquisition and analysis software, MEG/4, was used for the reconstruction of a single moving dipole (CTF Systems Inc.). Note that this is not a multi-start implementation.

Analysis of Source Reconstructions

The CTF software was used to display the solutions found with the source reconstruction algorithms. The locations of the sources were shown in three views of a schematic of a human head. The sources were represented by small circles, with the orientation of a source in the tangential direction denoted by a short line extending from the circle in the direction of the orientation of the source. The co-ordinates of the best solution found with each algorithm were also given. For FOCUSS, the best solutions were manually selected on the basis of the timecourses of the sources. For the dipole fit algorithms the solutions with the lowest value for the reduced chi-square were selected from the set of solutions obtained with the multi-start approach.

Moreover, the timecourses of activation were given for the sources in the best solutions, as well as a plot of the change of the reduced chi-square over time.

Wherever applicable, the volume of the 95% confidence interval obtained from the Monte Carlo simulations was also given.

6.3 Results

6.3.1 Initial Inquiry of the Data

The analysis started with an examination of the characteristics of the inverse problem for the sensor locations and configuration and the realistically shaped source space used for the experiment. For this purpose the figures of merit described in chapter 4 were computed and visualised. In figure 6-3 the *Source Identifiability* and the *Expected squared error due to noise* are displayed on a logarithmic scale in order to reveal the patterns in the surface maps. Additionally, the *Bias in Dipole Location* is displayed. The surface maps of the *Image Sensitivity* and the *Expected squared error due to other sources* were omitted, since the patterns in these maps showed great similarities with the surface map of the *Source Identifiability*. The *Source Identifiability* was low for deep sources located towards the midline. Hence, it was difficult to unambiguously reconstruct activity for sources in these areas, due to the influence of activity in other sources. Furthermore, it is clear that towards the left of the source space, the *Source Identifiability* was different for sulcal sources than for gyral sources. This could be due to the fact that the gyral sources were nearer to the sensors than the sulcal sources. To put this into perspective, this result should be compared with 48 in section 4.3.2.2 where the *Source Identifiability* was computed for the 19 channel CES system. The bias towards the sensors was much more pronounced for this smaller sensor array, due to the smaller area of cortex covered by this system. The advantage of using a system that covers the whole head is obvious. The fact that the bias towards the sensors was small suggests that the choice of the initial weighting matrix was not critical with the whole head system. Indeed, the surface maps of the figures of merit computed for smoothed norm weighting were similar to the maps displayed here (not shown). It is noteworthy that the *Source Identifiability* was higher for gyral sources than for sulcal sources, a finding that contradicts the results obtained with the smaller sensor array.

The advantage of a whole head system is also demonstrated by the surface map of the *Bias in Dipole Location*. With the whole head system, the *Bias in Dipole Location* is negligible for most elements in the source space. Hence, similar accuracy in the location of reconstructed activity can be obtained for most sources in the source space. Moreover, the influence of noise in the data shows no systematic relationship with the locations of the sources.




figure 6-3: Figures of merit for the measurement set-up during the recordings. Minimum norm weighting was used in the computations of these figures of merit. Top row: *Source Identifiability* displayed on a logarithmic scale, viewed from the left-back (left panel) and the right-back (right panel). Bottom row, left panel: *Bias in Dipole Location*. Bottom, row right panel: *Expected squared error due to noise*, displayed on a logarithmic scale. The patterns for the figures of merit displayed in the bottom row were very similar when viewed from the right and were therefore omitted. Note the low *Source Identifiability* for deep sources located towards the midline. Also note the difference in *Source Identifiability* between sulcal and gyral sources on the left. Furthermore, it is striking that the *Bias in Dipole Location* was negligible for most elements in the source space. The *Expected squared error due to noise* displayed a random pattern.

The SNR of the averaged data for the larger patch was higher than for the smaller patch. The data recorded for the larger patch were therefore analysed here.

In order to further evaluate the performance of the system the SNR of the recorded data in each measurement channel was estimated. For each channel the SNR was defined as the ratio of the maximum absolute value of the data to the standard deviation of the plus-minus average. The SNR, sorted for amplitude, is shown in figure 6-4. A total of 62 sensors had an SNR larger than 2. Due to the large coverage of the whole of head system there were however many (88) sensors that had a SNR smaller than 2 over the same selected latency range. This illustrates why the utilisation of smaller systems has provided useful insights in the past. If such a smaller array of, say 19 sensors, is placed in a favourable location, then




figure 6-4: Signal to noise ratio over the latency range from 71 to 122ms. The channels were sorted for SNR. Note that almost half of the sensors had an SNR that was larger than 2.

the SNR in most of these channels will be significantly larger than 2 (apart from sensors near the zero-crossings of dipolar field patterns).

The averaged data for the larger stimulus patches were visually examined to find peaks in the evoked responses. Also, fieldmaps for selected latencies were called onto the screen. In figure 6-5 the evoked responses are displayed in a schematic of the subject's head (top panel). In this schematic the nose is pointing upwards and the left and right ear are shown on the left and right, respectively. As expected, a stimulus related response was clearly visible in the sensors located near the visual cortex, i.e. the occipital sensors. A large evoked response was also visible in sensors overlying the left and right temporal lobe. The onset of this response occurred after the response seen in the occipital sensors had reached its maximum. In the bottom panel of the same figure a field map of the magnetic field measured at 87ms after stimulus onset is shown. A dipolar pattern in the field map is clearly visible. This suggests that there was electrical activity in the brain below the zero crossing that can be modelled by an equivalent current dipole.

The occipital channels contained the primary response to the visual stimulus and were therefore examined in more detail in figure 6-6. The initial onset occurred at about 70ms, followed by several peaks in the responses (figure 6-7). At about 120ms a distinguished change in the morphology of the evoked response and the patterns in the

fieldmaps took place, indicating that an essential change in processing occurred towards that latency. In order to reduce the complexity of the analysis, a restricted latency range, from 71-122ms, was chosen. Data recorded for this latency range were used for the source reconstructions.

In figure 6-7 it is noticeable that the maxima in the sensors were shifted in time with respect to one another. This could indicate that the locus of the underlying brain activity shifted in location over this time range or that two sources, separated in space, were active with out-of-phase activity patterns. This is further examined in the next section, where the results obtained with the source localisation algorithms are given.



Figure 6-6. Top panel: Averaged cortical responses recorded for each sensor. The signals are displayed in a clockwise order of the subject's head (nose pointing upwards and the left and right ear on the left and right, respectively). From the associated colored responses (see 6.7) within the recorded latency and the magnitude of response (grey) over the left and right side of the head at a late latency. Bottom panel: Top-down view of the head, displaying an interpolated field map of the primary magnetic field at a latency of 100ms. The color bar shows the strength location. A dipole with positive pole towards the left magnetic pole, suggesting that there was an active source in the visual cortex at this latency.

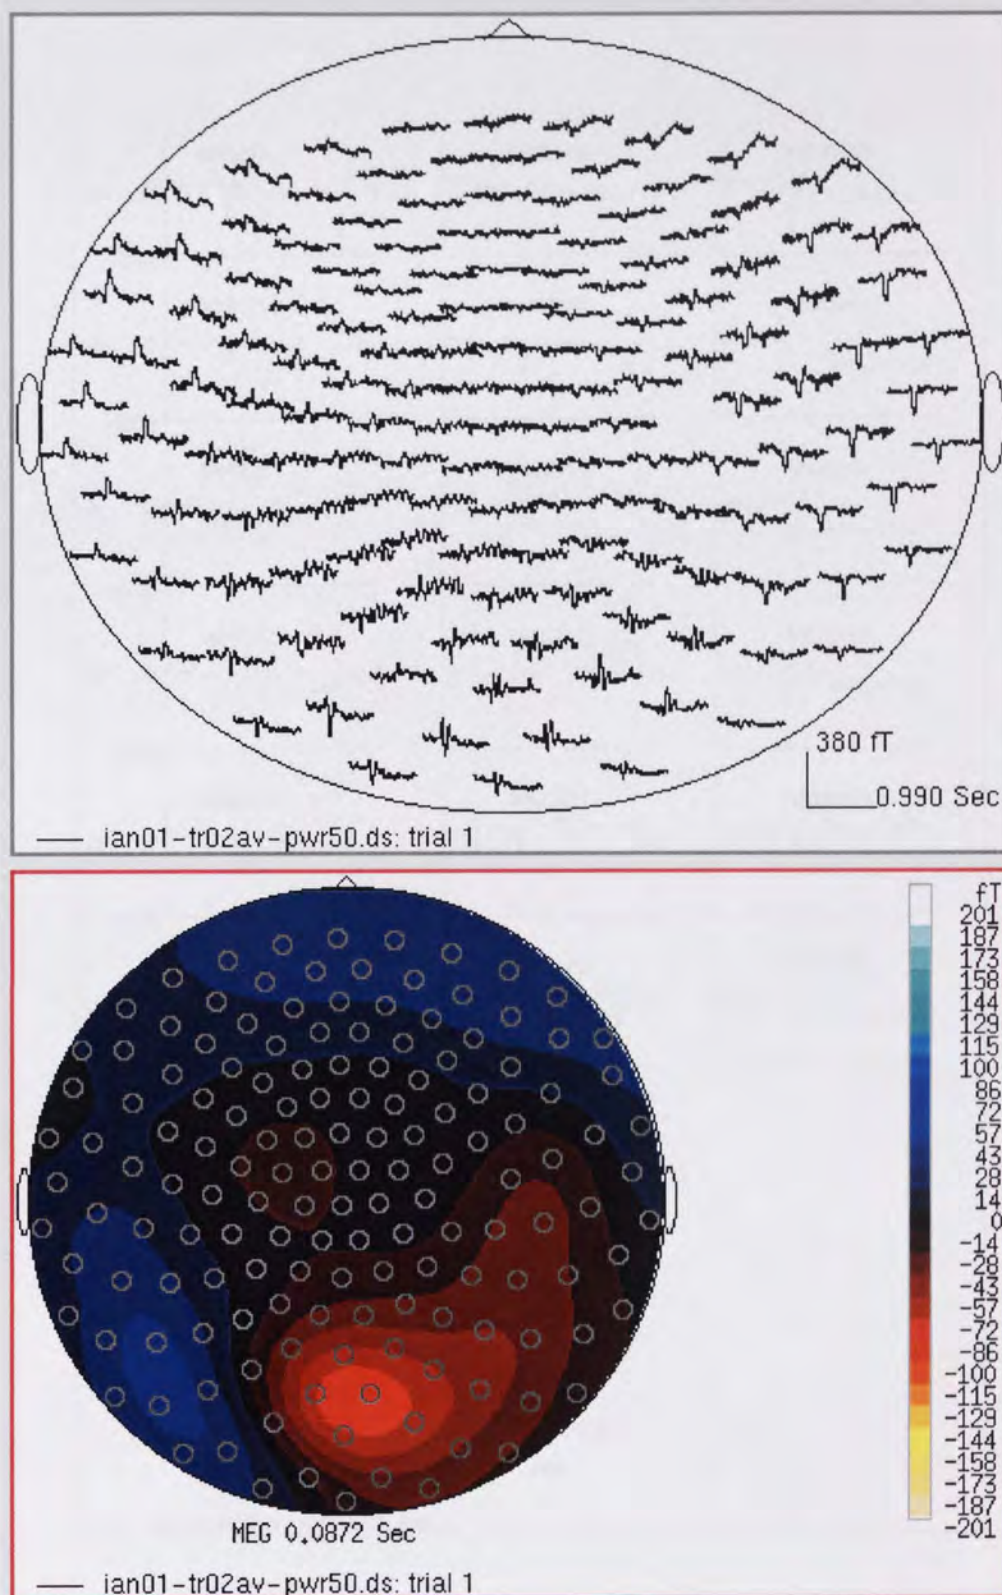


figure 6-5: Top panel: Averaged evoked responses recorded for each sensor. The signals are displayed on a schematic of the subject's head (nose pointing upwards and the left and right ear on the left and right, respectively). Note the stimulus related response that is visible in the occipital sensors and the responses of opposite polarity over the left and right side of the head at a later latency. Bottom panel: The same schematic of the head, displaying an interpolated field map of the recorded magnetic field at a latency of 87ms. The circles denote the sensor locations. A dipolar field pattern is visible over the left occipital pole, suggesting that there was an active source in the visual cortex at this latency.

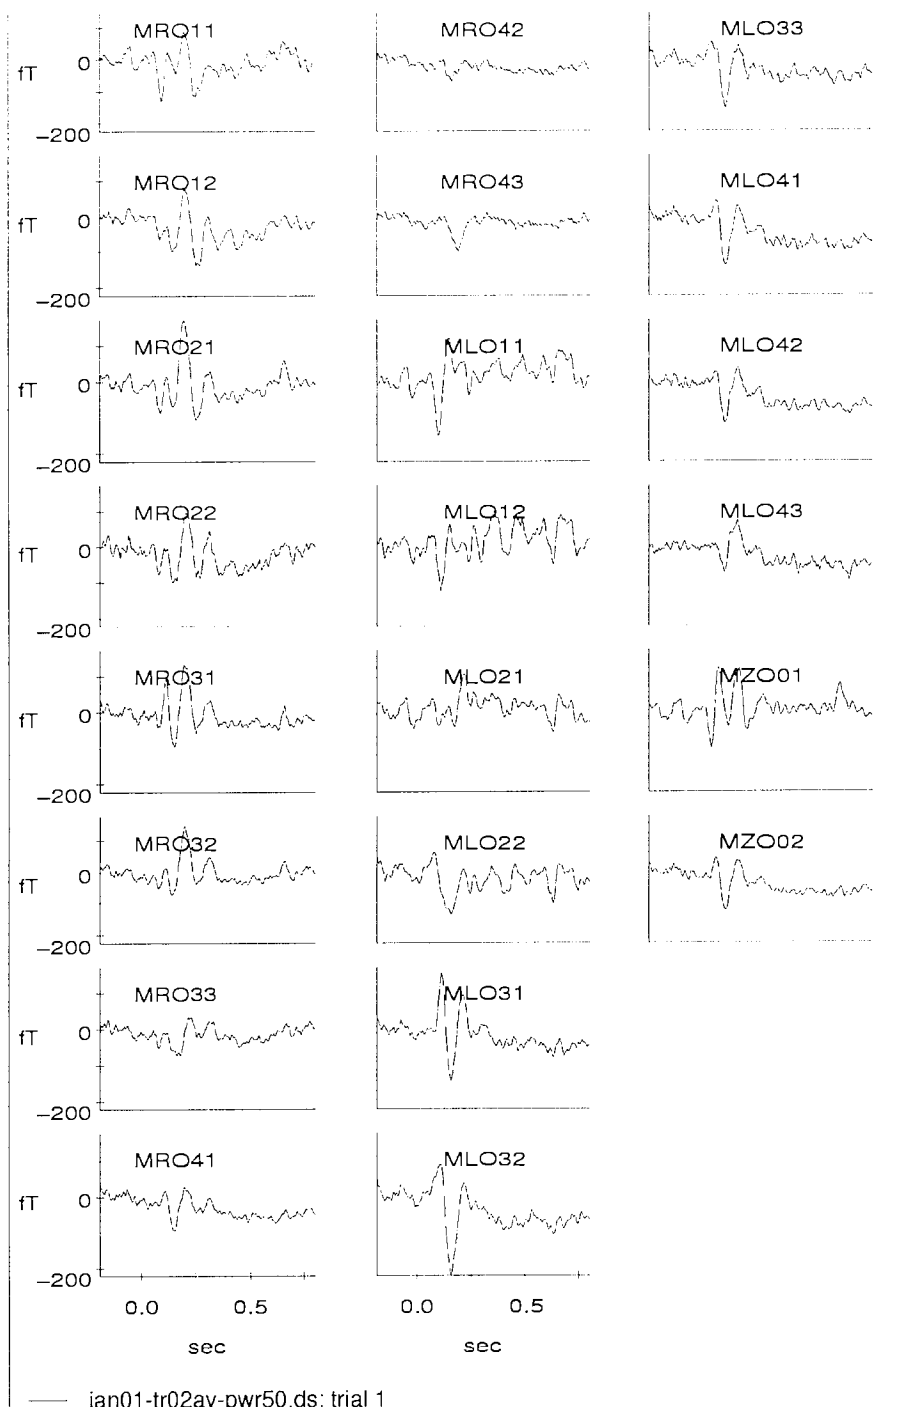


figure 6-6: Averaged evoked responses recorded in the sensors placed over the occipital cortex. An evoked response triggered by the presentation of the visual stimulus is clearly visible.




figure 6-7: The timecourses for the occipital channels overlaid. Only the early visual evoked response is shown (60-130ms). The timecourses show a complicated activity pattern. The sensors with the strongest signals show a positive peak in activity at about 75 and 110ms and a negative peak at about 95ms. Note however that the maxima in each sensor were shifted in time with respect to one another.

6.3.2 Source Reconstructions

Minimum Norm Solutions

Minimum norm solutions were computed for the averaged data for the larger stimulus over the latency range from 71 to 122ms. A range of values for the noise truncation level was explored. The absolute values of the source activity for each element in the source space were summated over the latency range in order to reveal the areas that were most active. In figure 6-8 the summated minimum norm solutions typically found for the different truncation levels are shown. Discarding too many singular values resulted in a loss of detailed information. Consequently, only patterns with low spatial frequency were found for these truncation levels. When more singular values were used for the source reconstructions, patterns with higher spatial frequency became visible. These patterns

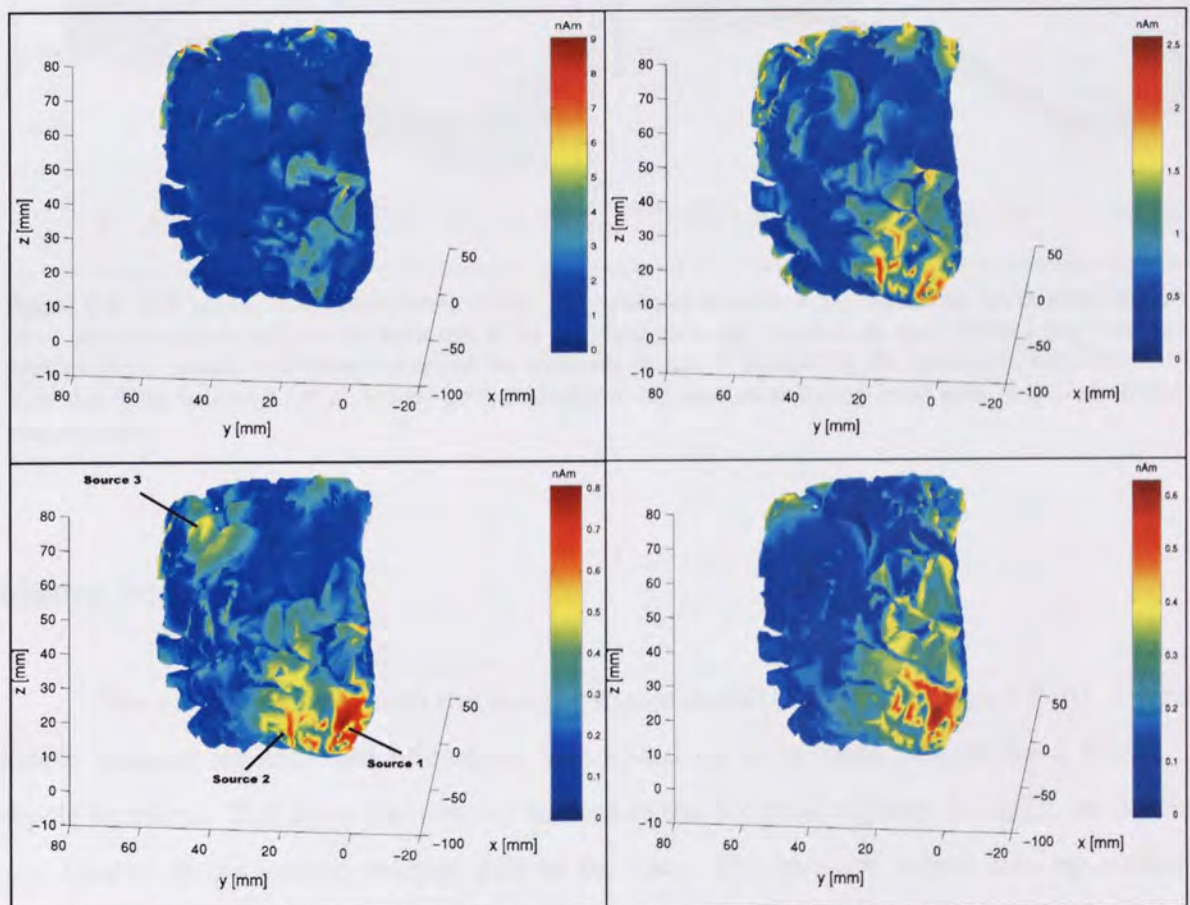


figure 6-8: Summation over time of the [minimum norm solutions], with different levels of singular value truncation. From left to right, top to bottom, the truncation level was set at the 3rd, 10th, 30th and 43rd singular value, respectively. Noise truncation at intermediate levels gave very similar results to those shown here.

When many singular values were discarded (truncation level at the 3rd singular value), mainly patterns with low spatial frequency were found. The decrease of the truncation level resulted in patterns with higher spatial frequency. However, when the truncation level was set at a very low level (43rd singular value), the increase in spatial resolution was impeded by the increase of the influence of noise on the solutions, due to the projection of noise components into the solutions.

suggested three different areas of activity; two areas with a small separation between them in the occipital pole and a 3rd area in a parieto-occipital region with smaller activity. A further increase in the number of singular values degraded the solutions due to an increase in the influence of noise in the measurements.

The timecourses of activation of the elements in the source space located in these three regions of relatively large activity (hotspots) are shown in figure 6-9. The activation patterns for elements in the weaker hotspot were similar to the activation patterns in the two main hotspots. Hence, this activity was likely to be related to the presentation of the stimulus and not due to noise in the data.

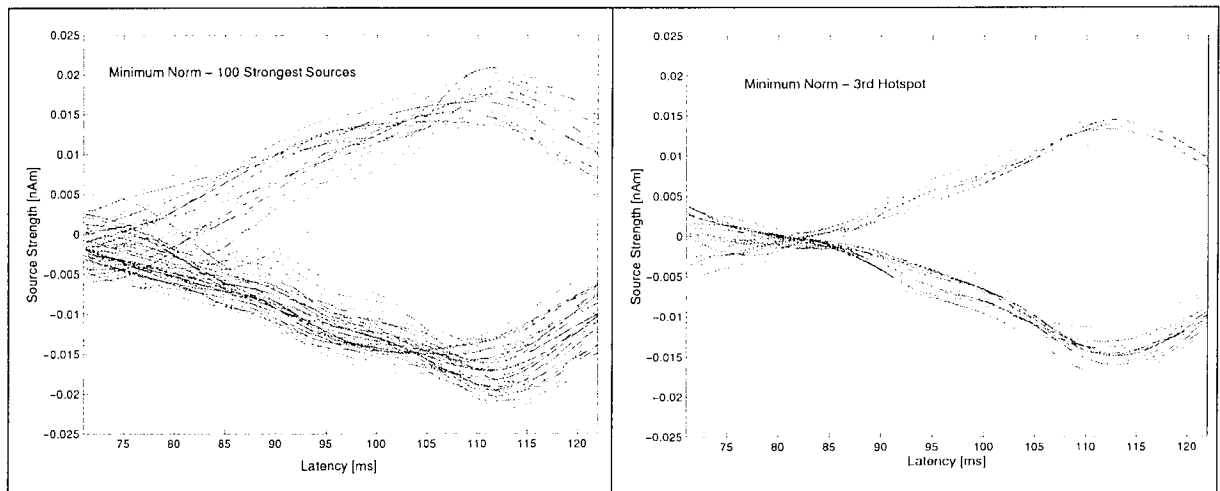


figure 6-9: Left panel: Activation curves of the 100 strongest sources in the minimum norm solution. All these elements are located in two hotspots in the occipital pole (see figure 6-8), apart from a few spurious sources. Right panel: Activation curves of the elements in the 3rd hotspot in the minimum norm solution. Note that these sources have an activity pattern similar to the sources in the occipital pole as well as similar peak strength.

Moving Dipole

The results obtained with the moving dipole model are given in figure 6-10. A new dipole location position and orientation was found for each latency, forming a 'trace' of dipole locations. This trace was entirely located in the occipital regions. Initially, the dipole was located in the anterior-inferior part of the trace. The location varied only by a small amount over the next period, forming a stable cluster. The dipole then moved to a more posterior-superior area in the occipital pole, where a second cluster was formed. At the later stages the dipole moved even more posterior and to the right hemisphere, eventually moving outside the sphere.




figure 6-10: Locations of a single moving dipole fitted to the data over the latency range 72-122ms, forming a 'trace' of dipole locations. The trace is formed by the dipole moving from an anterior-inferior position to a more posterior-superior area in the left hemisphere and then to the right hemisphere and even more posterior.

This 'trace' provided information about the number of active sources and their locations (Böcker *et al*, 1994). It suggested that the activity shifted in location from an initial locus of activity to a second, more posterior locus. This movement of activity in time was already hinted at by the latency shifts of the maxima in the responses in the different sensors, as discussed in the previous section.

The movement of the dipole outside the sphere at the latter stages suggests that the data were not well explained by a single source at these latencies. Hence, for the early period of the latency range a two-dipole model, with different timecourses of activation, seems appropriate. For the later period a more complicated model might be needed to explain the data.

FOCUSS

The results obtained with the FOCUSS algorithm, using the three different weighting matrices, are given in figure 6-11. In this figure all solutions obtained with different truncation levels are plotted in the same schematics of the human head. It should be noted that these figures therefore also contain solutions that consisted of spurious sources that were caused by the noise. However, presenting the results in this manner revealed clusters of sources. These clusters were formed by sources that appeared in many of the solutions and were therefore likely to represent areas of activity that were related to the stimulus presentation. The results obtained with the three different weighting matrices revealed a similar pattern: two separate clusters of sources were found, most clearly defined in the axial and sagittal views. One cluster was located in the occipital pole and the other

cluster was found in a more anterior and superior location. The location of the latter cluster is close to the edge of the source space.

Additionally, the spread in the reconstructed sources was largest in the superior-inferior direction, especially for the anterior cluster. Hence, the spread was largest in the direction away from the sensors. Furthermore, the posterior cluster was also extended laterally. This could indicate the existence of two separate sources in this region.

From each set of solutions displayed in figure 6-11, the best solution was selected on the basis of the courses of activation. Note that this is a subjective method (see section

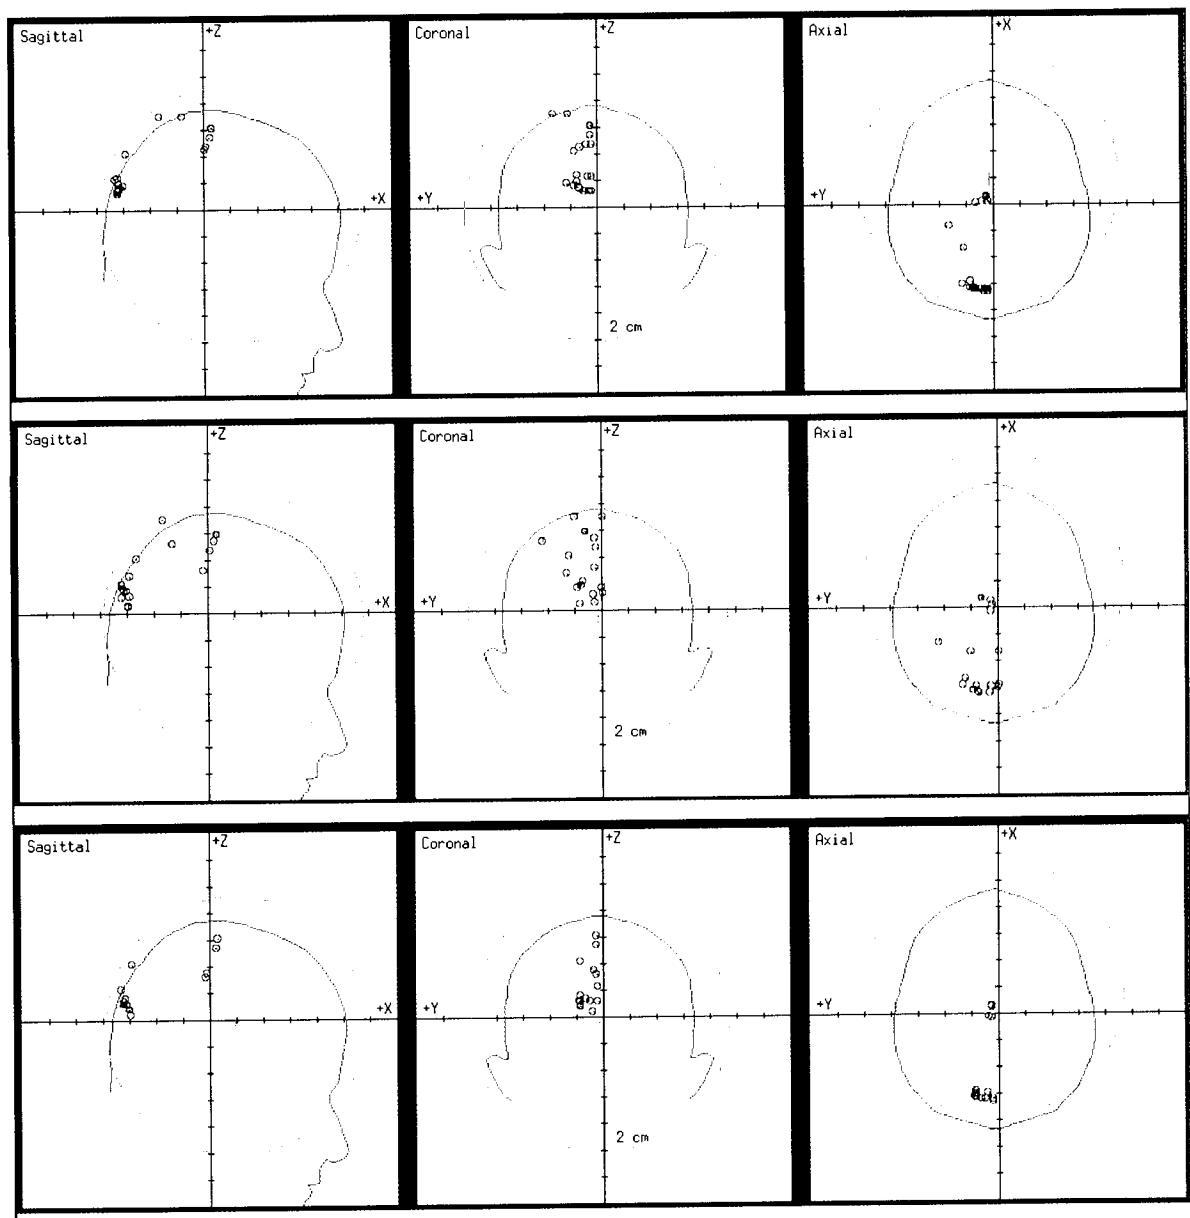


figure 6-11: Sources reconstructed with FOCUSS, using a range of truncation levels and different initial weightings. Top row: minimum norm weighting. Middle row: smoothed norm weighting. Bottom row: average distance weighting. Note that with all three weighting matrices two separate clusters of sources were found, most clearly defined in the axial and sagittal views. The spread in the solutions was largest in the superior-inferior direction, especially for the anterior cluster. The posterior cluster also showed a lateral prolongation.

5.3.2). The best solutions obtained for the different weighting matrices are displayed in figure 6-12, figure 6-13 and figure 6-14. In these figures the activation courses of the sources are also given, along with a plot of the value of the reduced chi-square over time.

FOCUSS solution, using **minimum norm** weighting:

Two sources were reconstructed in the occipital pole (dipole 2 and 3), with dipole 2 located towards the left. This source had the strongest activity, followed by the activity of dipole 3. The peak activity of these two sources was shifted in time by 7ms, with the dipole located towards the midline reaching maximum activity first (82ms). A third source with smaller activity, dipole 1, was located in a more anterior-superior region. This dipole

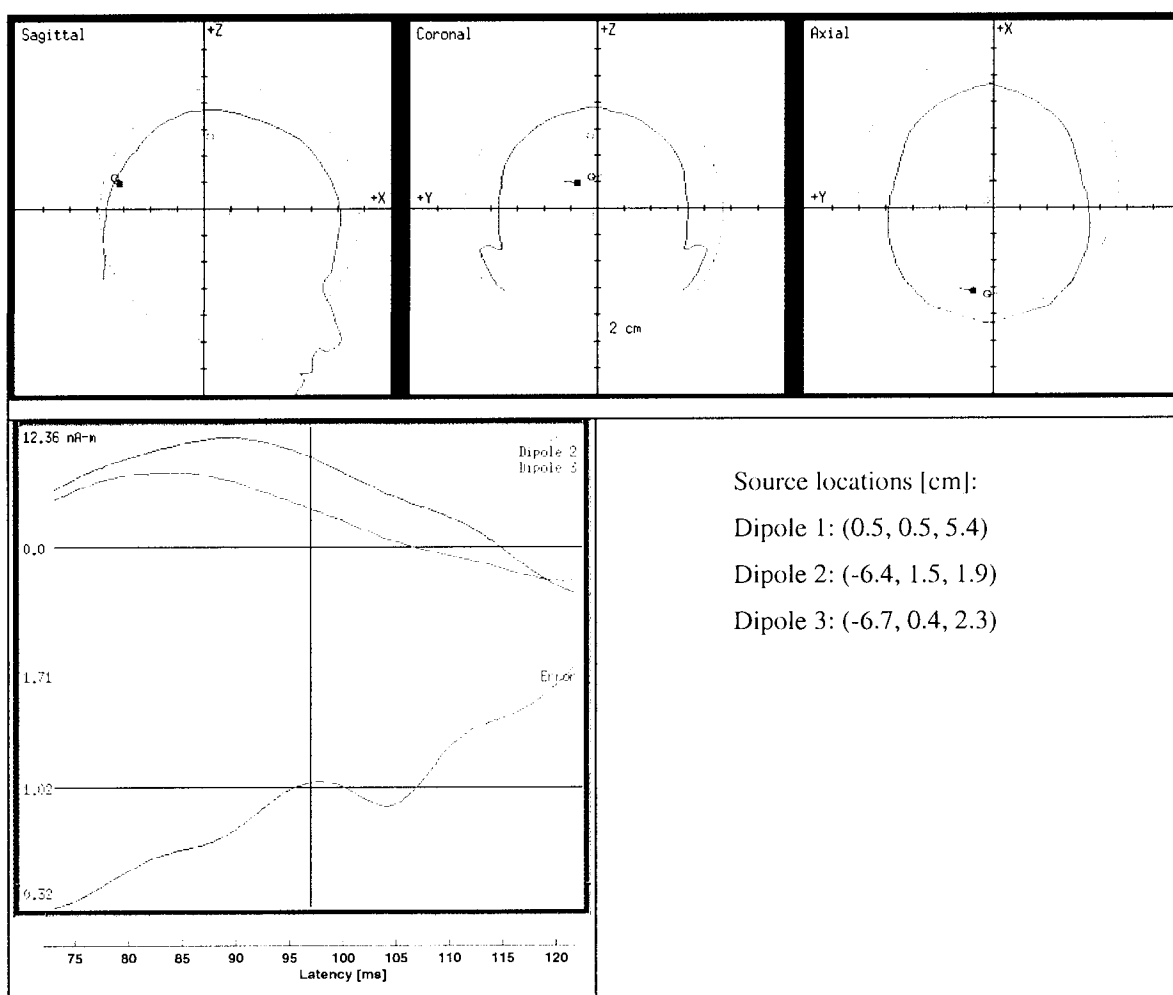


figure 6-12: Best solution found with FOCUSS, using **minimum norm** weighting for source reconstruction over the latency range from 71 to 122ms. The top panel shows the source locations and the bottom panel the activation curves and the reduced chi-square error function. Two sources were reconstructed in the occipital pole (dipole 2 and 3), with dipole 2 located towards the left. This source had the strongest activity, followed by the activity of dipole 3. The peak activity of these two sources was shifted in time, with the dipole located towards the midline reaching maximum activity first. A third source with smaller activity, dipole 1, was located in a more anterior-superior region. This dipole peaked at a later latency.

peaked at a later latency (99ms). The amplitude of all three sources was small in the later part of the latency range (115-122ms). The data at these latencies were therefore not well modelled by these sources, resulting in the large values for the error function in this range. The low value of the error function at the beginning of the latency range could indicate that the data were over-modelled by three dipoles at these latencies. This was supported by the fact that the reconstructed activity for dipole 1 was small for these latencies.

FOCUSS solution, using smoothed norm weighting:

Some of the solutions found with smoothed norm weighting contained three sources. In these solutions the two occipital dipoles were both shifted towards the midline. However, the activation curves were largely self-cancelling and these solutions were

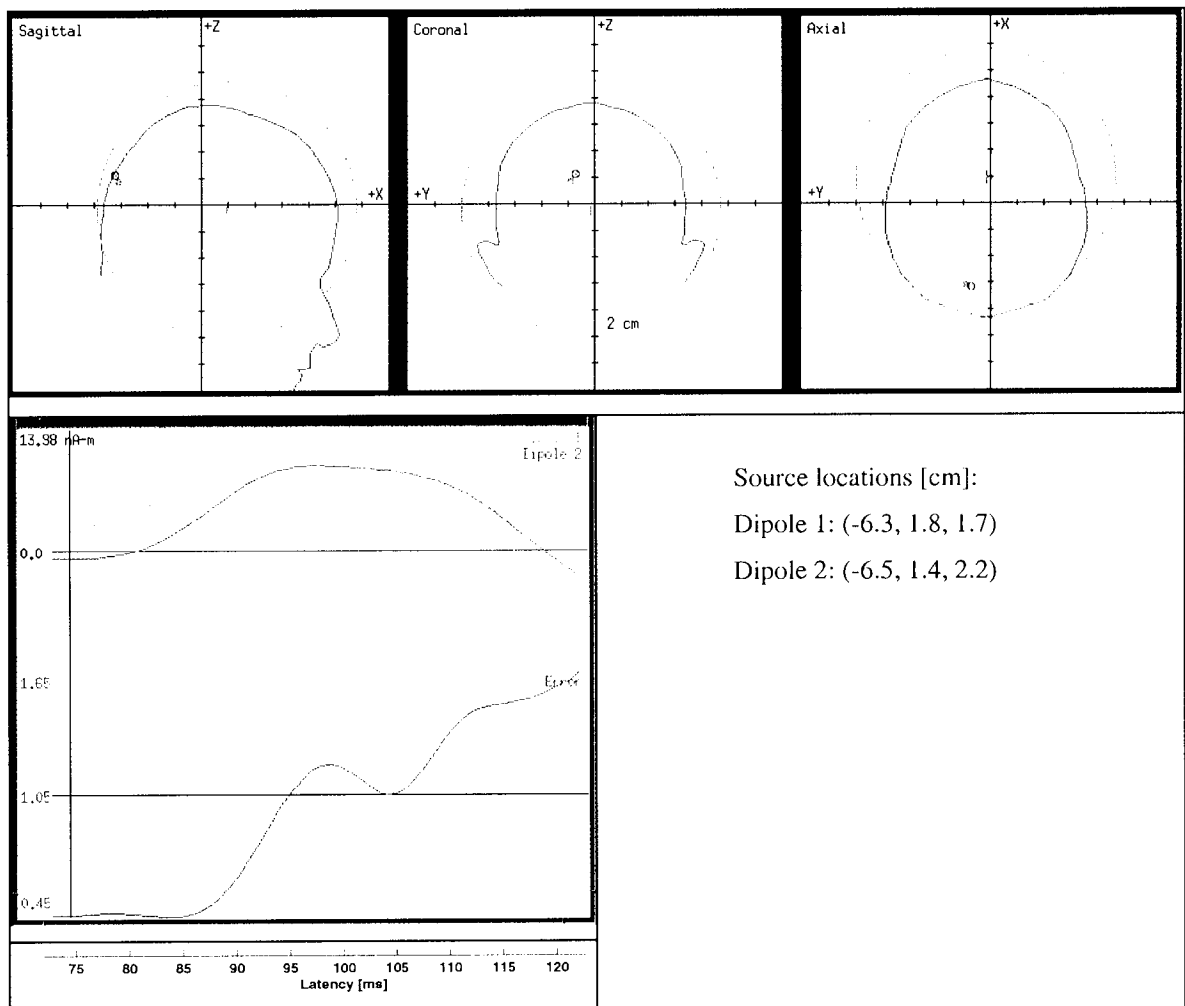


figure 6-13: Best solution found with FOCUSS, using **smoothed norm** weighting, for source reconstruction over the latency range from 71 to 122ms. The top panel shows the source locations and the bottom panel the activation curves and the reduced chi-square error function. Two sources were reconstructed in the occipital pole, with dipole 1 located towards the left. This source had the strongest activity. The peak activity of the two sources was shifted in time.

therefore rejected. The best solution obtained with smoothed norm weighting consisted only of 2 sources both in the occipital pole. The separation in space was only 1mm and 3ms in time. The error function again indicated that the data were not well modelled at the start and end of the latency range.

FOCUSS solution, using average distance weighting:

The results obtained with average distance weighting were similar to the results found for smoothed norm weighting. Again, no well-behaved solution with three sources was obtained. The spatial separation between the two occipital dipoles was larger as for smoothed norm weighting. The dipole closer to the midline was activated 7ms before the more lateral dipole, in accordance with the results for minimum norm weighting.

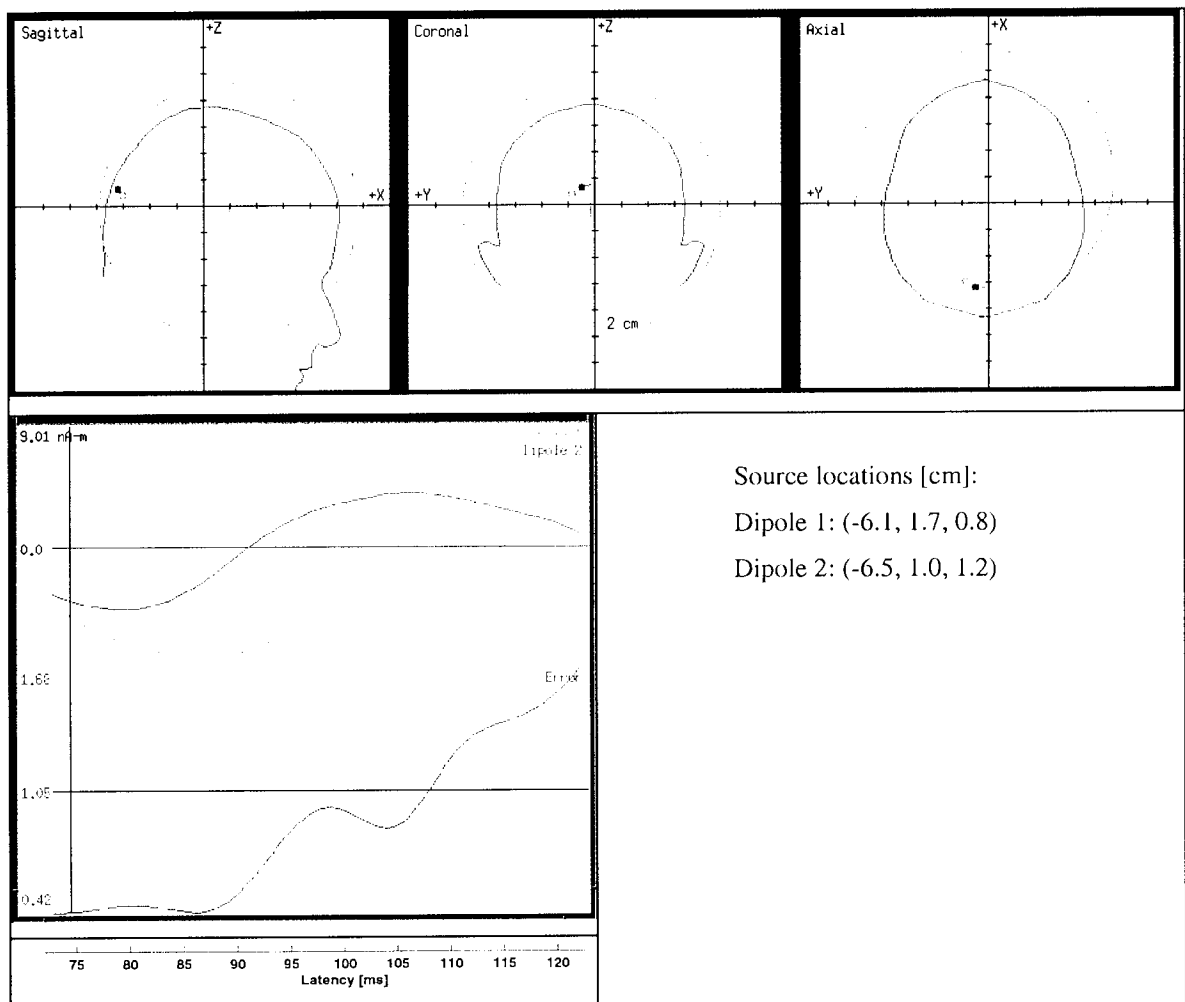


figure 6-14: Best solution found with FOCUSS, using average distance weighting for source reconstruction over the latency range from 71 to 122ms. The top panel shows the source locations and the bottom panel the activation curves and the reduced chi-square error function. Two sources were reconstructed in the occipital pole, with dipole 1 located towards the left. This source had the strongest activity. The peak activity of the two sources was shifted in time, with the dipole located towards the midline reaching maximum activity first.

Dipole Fitting

Single Dipole

The locations of the dipoles reconstructed with the unconstrained dipole fit algorithm are given in figure 6-15. The solutions obtained with the multi-start approach are first shown. These solutions formed a tight cluster of sources in the occipital pole. The solution that resulted in the lowest value for the reduced chi-square ($= 1.31$) was found at the following location: $(-6.5 \ 1.3 \ 0.7)$ [cm]. The overlay of the dipole onto the subject's MRI revealed its location with respect to the anatomy of the head.

The activation curve for this dipole showed a peak approximately in the middle of the latency range (94ms). The error function indicated that the data were well modelled at these latencies, but that the data at the start and end of the latency range was not thoroughly explained by this model. The stability of the solution was confirmed by the value of 9mm^3 for the 95% confidence volume obtained from the Monte Carlo simulations for this best solution (not shown).

The results for the constrained version of the algorithm were similar to those found for the unconstrained algorithm (figure 6-16). The multi-start algorithm gave a broader spread of solutions. The dipole that resulted in the lowest reduced chi-square value ($=1.31$) was positioned in approximately the same location as the unconstrained best solution, namely at $(-6.3, 0.8, 0.4)$ [cm]. The activation curves were also very similar, peaking both at 94ms. The 95% confidence volume obtained for the best solution was only 3mm^3 , illustrating that, although the constraints increase the extent of the cluster that was obtained with the multi-start approach, they stabilised the solutions with respect to the noise.

The solution that was obtained when the best constrained solution was restricted to lie within the mesh was found at the following location, $(-6.2, 0.9, 0.5)$ [cm], i.e. within a millimetre from the best unconstrained solution.

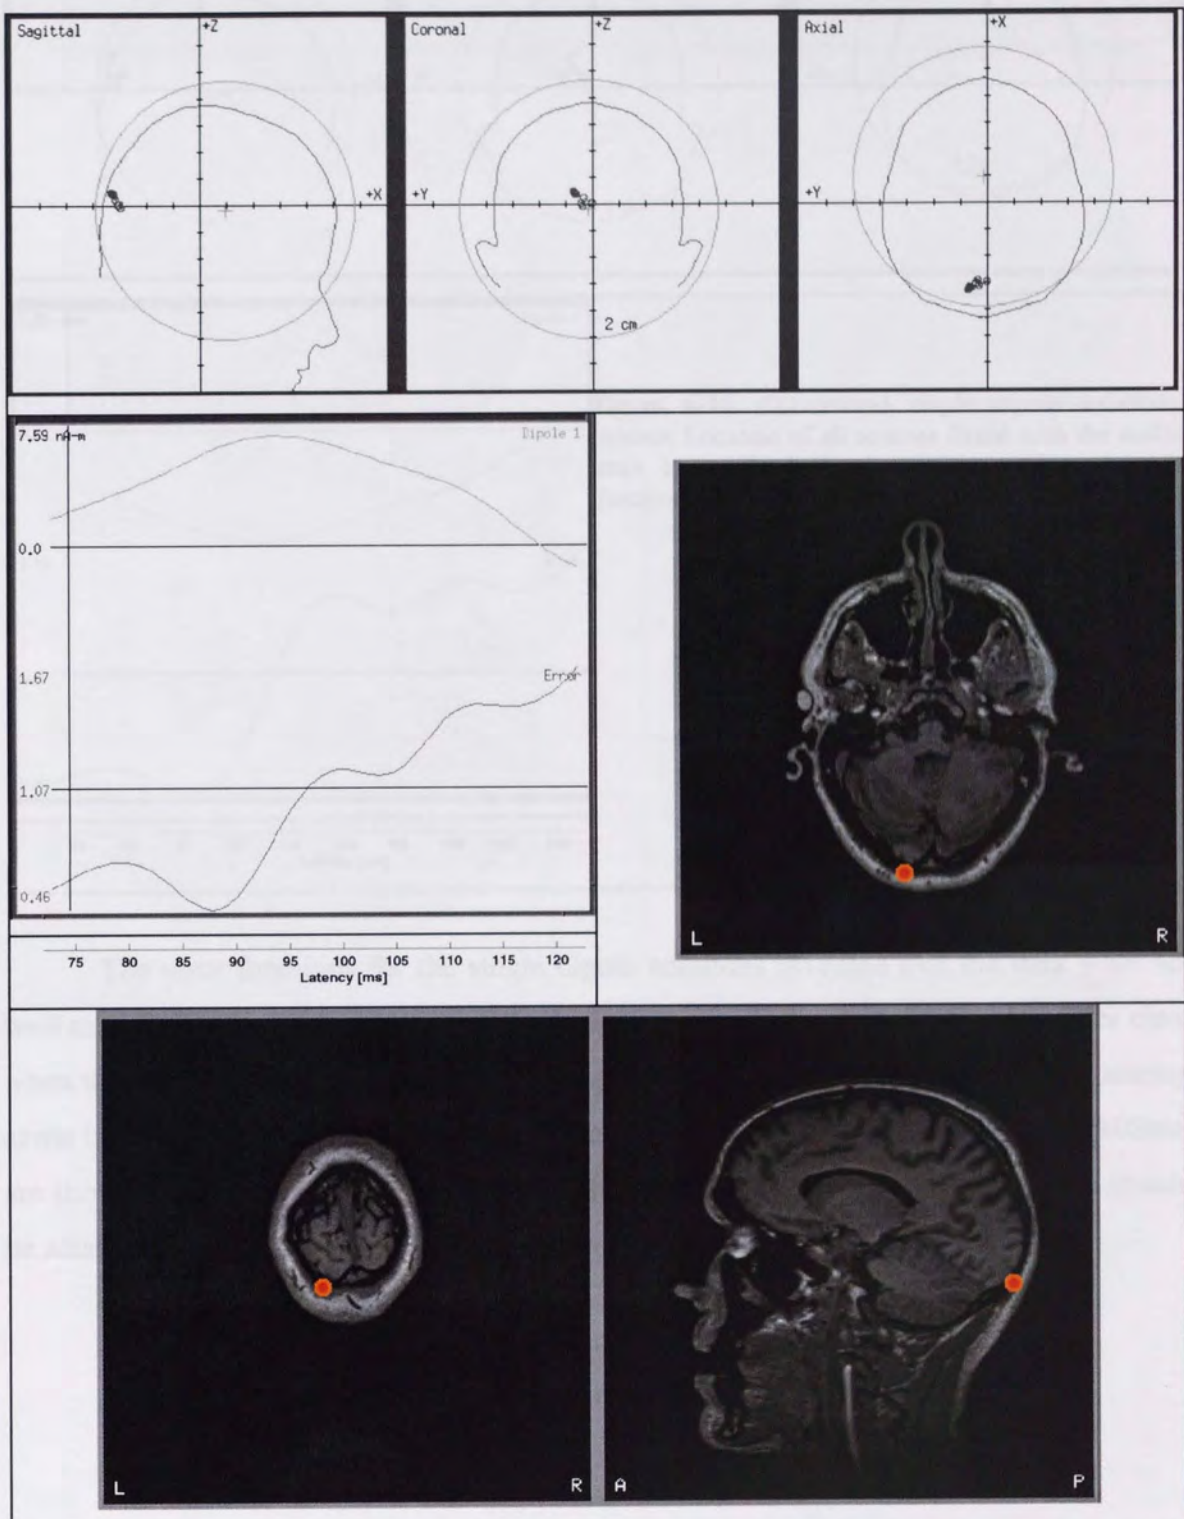


figure 6-15: Unconstrained single dipole solutions. Top row: Location of all sources found with the multi-start approach. Middle row, left panel: Activation curve and error function for the best solution. Note the tight cluster formed by the solutions. Middle row, right panel: Best solution overlaid on the subject's MRI (axial slice). Bottom row: Best solution overlaid on a coronal (left panel) and sagittal slice (right panel).

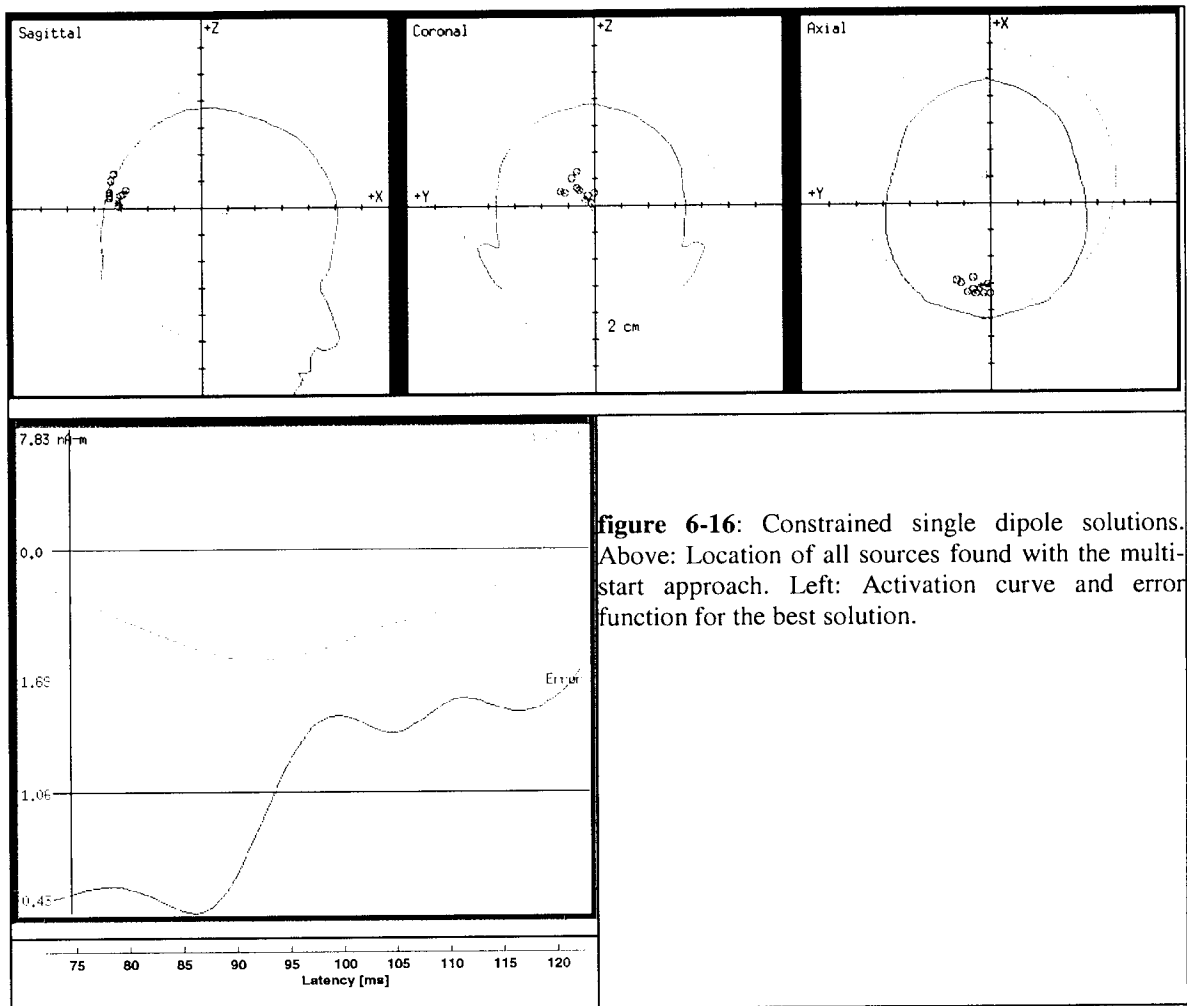


figure 6-16: Constrained single dipole solutions. Above: Location of all sources found with the multi-start approach. Left: Activation curve and error function for the best solution.

The error functions for the single dipole solutions revealed that the data were not well modelled at the start and end of the latency range. The reason for this becomes clear when the measured and the estimated fieldmaps are examined. The field maps for a latency at the beginning of the latency range (83ms) and from the end of the latency range (108ms) are shown in figure 6-17. It is obvious from the difference maps that an extra dipole should be added to the source model in order to improve the fit to the data at these latencies.




figure 6-17: Recorded magnetic field, magnetic field induced by the (unconstrained) dipole model and the difference between the recorded and calculated magnetic fields. Left panel: field maps at 83ms. Right panel: field maps at 108ms. Note the dipolar patterns in the difference maps, indicating that an extra dipole should be added to the source model at these latencies.

Two Dipoles

The locations of the dipoles reconstructed with the unconstrained dipole fit algorithm are given in figure 6-18. The solutions obtained with the multi-start approach are first shown. It can be seen that two clusters were formed in the occipital pole. The one cluster was formed at a more anterior-inferior location than the other cluster. The activation curves for the two dipoles that form the best solution (reduced chi-square of 1.05) are of opposite polarity, but not completely self cancelling due differences in location, amplitude and orientation. The locations of the sources were given by: dipole 1: (-6.7, 2.2, 1.4) and dipole 2: (-7.4, 2.5, 1.0) [cm]. The peaks in the activation curves for the 2 sources were broader than for the single dipole solutions, hence the 2 sources accounted for a longer period of data. Note however that, although the mean value had come down compared with the single dipole solutions, the error function still indicated that the source model was incorrect at the start and end of the latency range.

The Monte Carlo simulations, resulting in a 95% confidence volume of 0.2 and 1.5mm³ for dipole 1 and dipole 2, respectively, confirmed the stability of the best solution.

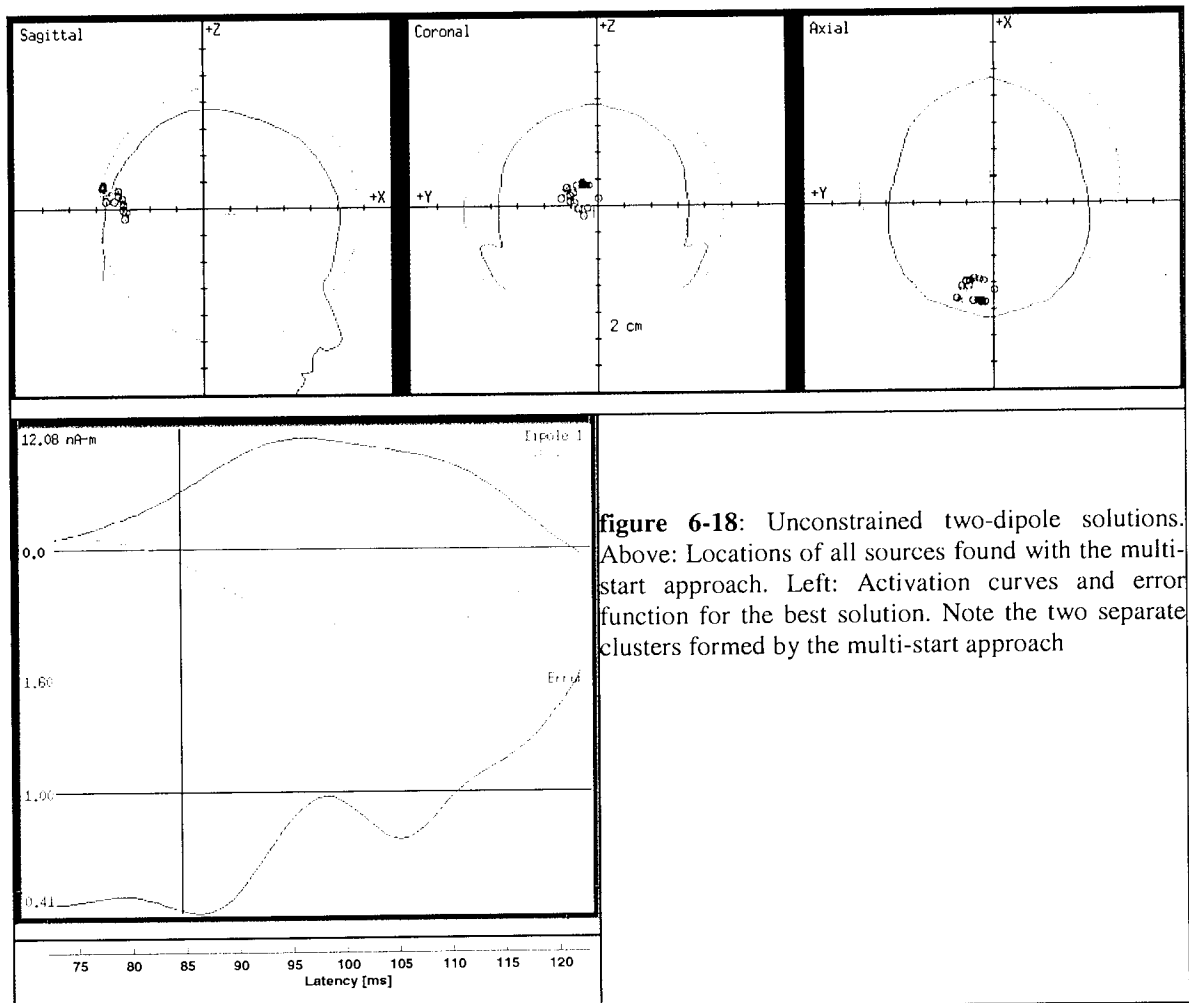


figure 6-18: Unconstrained two-dipole solutions. Above: Locations of all sources found with the multi-start approach. Left: Activation curves and error function for the best solution. Note the two separate clusters formed by the multi-start approach

In figure 6-19 the field maps for the best two-dipole solution are shown for the same latencies as in figure 6-17. The difference maps did not reveal a dipolar pattern. This suggested that the two-dipole model was an improvement on the single dipole model (see figure 6-17), at least for the data at these latencies.



figure 6-19: Recorded magnetic field, magnetic field induced by the (unconstrained) two-dipole model and the difference between the recorded and calculated magnetic fields. Left panel: field maps at 83ms. Right panel: field maps at 108ms. Note that the difference maps did not contain a distinct dipolar pattern (compare with figure 6-17).

The sources reconstructed with the constrained two-dipole fit algorithm are displayed in figure 6-20. This time there was only one cluster formed by the multi-start approach. An increase of the influence of noise in the data was revealed by larger 95% confidence intervals (7.3 and 399mm³ for dipole 1 and 2, respectively). However, the separation between the two sources in the best solution was clear. The first dipole was located at (-6.2, 1.9, 0.0) [cm] and the second at (-7.0, 0.8, 1.3) [cm]. The mean value of the reduced chi-square was 1.10. The error function decreased in value towards the end of the latency range.

The locations of the two dipoles that were found when the best constrained solution was confined to lie within the mesh were given by: (-5.8, 1.8, 0.4) and (-6.5, 0.8, 1.2) [cm]. The value of the reduced chi-square was 1.18 for this solution. Note that these dipoles were located more anterior than the dipoles in the best solution for the constrained fit.

The locations of these two dipoles are overlaid onto their respective MRI slices in figure 6-21.

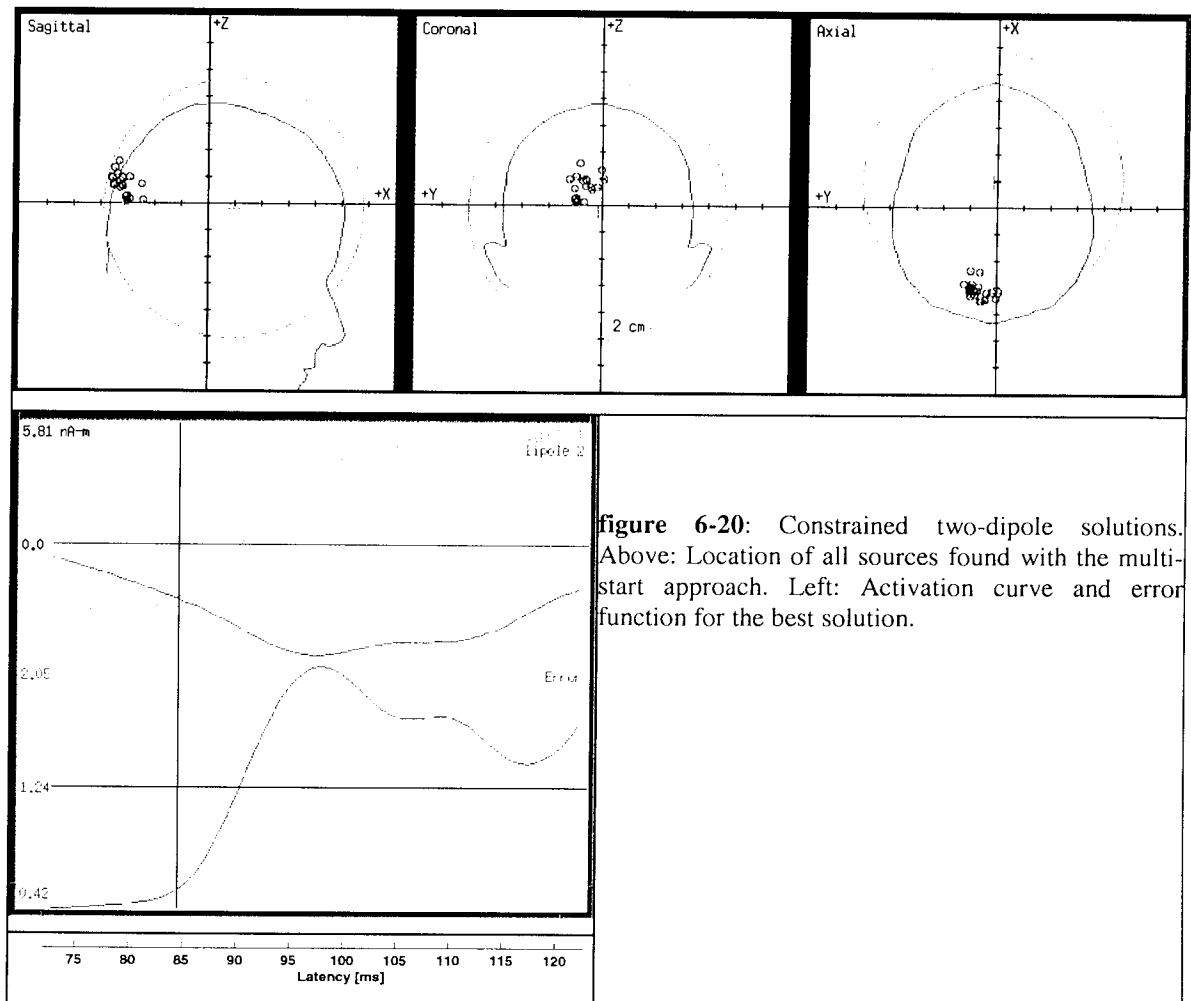


figure 6-20: Constrained two-dipole solutions. Above: Location of all sources found with the multi-start approach. Left: Activation curve and error function for the best solution.

Three Dipoles

Unstable solutions were obtained when a three-dipole model was used with the unconstrained dipole fit algorithm. This model was not used for additional source reconstructions.




figure 6-21: Solution for the constrained dipole fit algorithm (constrained to lie within the mesh) overlaid on the subject's MRI. Top row, from left to right: dipole 1 in an axial, coronal and sagittal slice, respectively. Bottom row, from left to right: dipole 2 in an axial, coronal and sagittal slice, respectively.

6.4 Discussion

In the study described in this chapter a whole head magnetometer was used to record the magnetic fields evoked by a visual stimulus. The advantage of using a whole head system was demonstrated by the surface maps of the figures of merit computed for the experimental set-up. The *Bias in Dipole Location* was negligible for most sources in the source space, hence most sources were shown to be able to be reconstructed without a large bias error. It was also found that the *Source Identifiability* was highest for gyral sources, due to their close proximity to the sensors.

Combining the results of the different localisation techniques, it is possible to find an answer to the question: how many sources produced the recorded MEG over the latency range from 71 to 122ms?

- Minimum Norm solution: A large occipital source was found (possibly made up out of 2 sources) and a parieto-occipital source.
- Moving Dipole model: The 'trace' formed by the moving dipole consisted of 2 stable clusters, suggesting two sources in the occipital cortex. The initial activity was found at a slightly more anterior-inferior position, closer to the midline than the second source.
- FOCUSS: The clusters found when all solutions were plotted suggested a single source in an anterior-superior region and an extended source in the occipital pole. The anterior-inferior source was found on the edge of the mesh, hence the location of this source should be interpreted with care. When the best solutions were plotted, it was revealed that the extended source consisted of two closely located sources.
- Dipole Fits: The solutions found with the constrained and unconstrained dipole fit algorithms were very similar, both in locations and activations of the sources. The two-dipole models resulted in lower mean chi-square values than the single dipole models. Additionally, the difference maps obtained for the single dipole model revealed the existence of a second source. The two sources were located in the occipital pole, at approximately the same locations as found with FOCUSS

Thus, it can be concluded that initially a source at an anterior-inferior location in the occipital pole was activated, followed 7ms later by the activation of a second, more posterior-superior source, shifted laterally. This sequence of activation and shift in location are in agreement with the findings by George and co-workers (George *et al*, 1989). A third

source is activated at a slightly later stage. The exact location of this anterior-superior source could not be determined, although activity in this area, related to similar visual stimuli, has been reported (Aine *et al*, 1996). It is possible that the source is related to the peaks in the evoked response found at later latencies and that the determination of its location can be improved by extending the latency range.

The reported inter-subject variability of shape and location of the primary visual cortex is large (Stensaas *et al.*, 1974). Yet, the locations and activation sequences of the occipital sources seem physiologically reasonable, both in location and activation sequence (dipole 1 and 2 corresponding to regions of cortex in the visual cortical areas V1 and V2, respectively (Kuffler *et al.*, 1984, Van Essen and Drury, 1997). Additionally, it is encouraging that sources at similar locations (see figure 6-22) and with similar differences in activation between the sources were found in experiments with the same subject, using a different MEG system (Barnes *et al*, 1998).



Aston University

Content has been removed for copyright reasons

figure 6-22: Sources reconstructed for the experiment described in Barnes *et al* (1998). Note that the locations of these sources correspond closely to the sources reconstructed in the study reported in this chapter (compare with figure 6-12), despite the use of a different (19 channel CES) MEG system. Figure taken from Barnes *et al*, 1998.

The error functions demonstrated that the data in the latency range of interest were well explained by the model up to about 110ms, where the error function began to increase. This suggests that a more complex source model is needed for the later period, possibly because of the onset of activity in other brain areas, related to the peaks seen in the evoked responses at later latencies.

The analysis with the FOCUSS algorithm could be expanded in several ways:

- Target weighting could be applied, using the locations found so far to tailor the initial weighting matrix.
- A mesh of the whole brain could be used, so that activity in more anterior regions could be faithfully reconstructed.
- Using the covariance matrix of the noise as part of the initial weighting matrix could reduce the influence of noise.

7 Discussion and Recommendations

In this work the incorporation of *a priori* information into the inverse problem for Magnetoencephalography (MEG) was evaluated.

It was found that the formulation of the FOCUSS algorithm (Gorodnitsky *et al*, 1995) enables the incorporation of *a priori* information in a simple and insightful way (chapter 2). Additionally, it does not rely on knowledge about the number of active sources.

The use of constraints formed by a realistically shaped representation of the cortical surface (chapter 3) relies on accurate co-registration of the MEG and MRI co-ordinate systems. The surface matching approach was found to provide this accuracy. The mean test-retest errors were smaller than 1.6mm. The maximum errors (2.1mm) were found for target points towards the back of the head (chapter 3).

Figures of merit designed for linearly distributed inverse problems gave insight in the characteristics of a measurement set-up and the usefulness of *a priori* information (chapter 4). It could be demonstrated that measurement with a small sensor array suffers from bias towards the sensors (chapter 4). This bias was reduced with a whole head system (chapter 6). Moreover, it was demonstrated that activity in most sulci and gyri is detectable. The ability to reconstruct activity in a certain area depends on a combination of the distance from the sensors, the orientation of the source and the norm of the lead fields (chapter 4).

The effect of *a priori* information was found to be largest on the *Bias in Dipole Location* measure, with only a small influence on the figures of merit that were derived from the resolution kernels instead of the resolution columns (chapter 4). The *Bias in Dipole Location* measure was shown to be a good indicator of the achievable accuracy of source reconstructions with the FOCUSS algorithm, enabling an objective, data independent, evaluation of the potential pay-off of any *a priori* information used (chapter

4). It was shown that target area weighting dramatically improved the accuracy of solutions for a specific area, without a decrease in performance in other areas (chapter 4).

The evaluation of FOCUSS, unconstrained and constrained dipole fit algorithms with a simulated cortical fold revealed that adding constraints improved the accuracy of reconstructed sources when noisy data were used (chapter 5). This improvement disappeared for constrained dipole fitting when the SNR of the data increased or when the surface normals were perturbed (chapter 5). With FOCUSS the advantage disappeared for data with high SNR and when the gridspacing was increased (chapter 5).

The accuracy of source reconstructions with FOCUSS decreased with increasing gridspacing (chapter 5). The influence was more pronounced for areas with large surface curvature (chapter 5). Due to the larger surface curvature a higher spatial sampling is needed in order to incorporate the information of the rapidly changing lead fields in the inverse problem. The performance of the constrained dipole fit algorithm was generally little influenced by an increase in gridspacing (chapter 5). An increase in error due to increasing gridspacing was only found for low SNR data and for an area with large surface curvature (chapter 5).

The error in the source reconstructions due to perturbations of the position of the mesh was in the order of 1mm for perturbations as large as 10mm, both for FOCUSS and constrained dipole fitting (chapter 5). This error was therefore negligible compared to other sources of error.

The performance of FOCUSS did not degrade due to perturbations of the surface normals (chapter 5). Interestingly, a small improvement was found, which can be explained by an improvement of the resolution kernels. Perturbations of the surface normals introduced errors in the source reconstructions as large as 5.6mm when the constrained dipole fit algorithm was used (chapter 5).

Finally, the application of the techniques developed in this thesis was demonstrated (chapter 6). The primary visual evoked response to a Gaussian windowed sinusoidal stimulus patch was modelled as two closely spaced sources at the occipital pole, activated in sequence (chapter 6). The locations and activation sequences were not only physiologically reasonable, complying with activation of areas in visual cortical areas V1 and V2, but also agreed with studies using a different MEG system.

Further Work

This thesis raised some interesting questions:

- Other types of *a priori* information could be obtained from the frequency domain (chapter 2).
- The accuracy of the solutions obtained with FOCUSS was dependent on the choice of the initial weighting matrix (chapter 4). The performance could be improved by computing the impulse response for many sources in the source space and optimising the weights to give the best overall performance, an approach ideally suited for a neural network approach.
- The angular distributions for gyral and sulcal sources across a number of cortical surfaces should be determined and correlated with the reconstruction accuracy for the two classes of sources (chapter 4).
- The increased accuracy found with target area weighting (chapter 4) suggests how MEG source localisation could benefit from information provided by other modalities, like fMRI or PET. The location of active sources in an fMRI/PET experiment could be incorporated by increasing the weights for these areas. FMRI/PET could then provide additional spatial information for the source localisation from MEG/EEG data, whereas fMRI/PET can benefit from the temporal information provided by MEG/EEG.

References

- Achim A, Richer F and Saint-Hilaire J-M (1991) Methodological considerations for the evaluation of spatio-temporal source models. *Electroen. Clin. Neuro.* **79**: 227-240.
- Achim A, Richer F and Saint-Hilaire JM (1988) Methods for Separating Temporally Overlapping Sources of Neuroelectric Data. *Brain Topogr.* **1(1)**: 22-28.
- Ahlfors S and Ilmoniemi RJ (1989) Magnetometer position indicator for multichannel MEG. In: *Advances in Biomagnetism* (Eds Williamson SJ, Hoke M, Stroink G and Kotani M). Plenum Press, New York.
- Aine C, George J, Medvick P, Supek S, Flynn E and Bodis-Wollner I (1989) Identification of multiple sources in transient visual evoked neuromagnetic responses. In: *Advances in Biomagnetism* (Eds Williamson SJ, Hoke M, Stroink G and Kotani M). Plenum Press, New York.
- Aine CJ (1995) A Conceptual Overview and Critique of Functional Neuroimaging Techniques in Humans: I. MRI/fMRI and PET. *Crit. Rev. Neurobiol.* **9(2&3)**: 229-309.
- Aine CJ, Supek S, George JS, Ranken D, Lewine J, Sanders J, Best E, Tice W, Flynn ER and Wood CC (1996) Retinotopic Organization of Human Visual Cortex: Departures from the Classical Model. *Cereb. Cortex* **6(3)**: 354-361.
- Anderson SJ, Holliday IE, Singh KD and Harding GFA (1996) Localization and functional analysis of human cortical area V5 using magneto-encephalography. *P. Roy. Soc. Lond. B* **263**: 423-431.
- Arun KS, Huang TS and Blostein SD (1987) Least-Squares Fitting of Two 3-D Point Sets. *IEEE T. Pattern Anal.* **9(5)**: 698-700.
- Ary JP, Klein SA and Fender DH (1981) Location of Sources of Evoked Scalp Potentials: Corrections for Skull and Scalp Thicknesses. *IEEE T. Biomed. Eng.* **28(6)**: 447-452.
- Backus G and Gilbert F (1968) The Resolving Power of Gross Earth Data. *Geophys. J. R. Astr. Soc.* **16**: 169-205.
- Bafra AK, Ersoy OK and Hansen DJ (1994) Image processing and computer graphics algorithms for surface rendering from MRI data. Purdue University, USA. Technical Report TR-EE 94-39.
- Baillet S and Garnero L (1997) A Bayesian Approach to Introducing Anatomic-Functional Priors in the EEG/MEG Inverse Problem. *IEEE T. Biomed. Eng.* **44(5)**: 374-385.
- Ball T, Schreiber A, Feige B, Wagner M, Lücking CH and Kristeva-Feige R (1999) The role of higher-order motor areas in voluntary movement as revealed by high-resolution EEG and fMRI. *NeuroImage* **10(6)**: 682-694.
- Bamidis PD, Hellstrand E, Lidholm H, Abraham-Fuchs K and Ioannides AA (1995) MFT in complex partial epilepsy: spatio-temporal estimates of interictal activity. *NeuroReport* **7**: 17-23.
- Bamidis PD and Ioannides AA (1996) Combination of point and surface matching techniques for accurate registration of MEG and MRI. Tenth International Conference on Biomagnetism - Biomag96, Santa Fe, New Mexico, USA.
- Barnes GR, Holliday IE and Harding GFA (1998) Human cortical magnification in the central two degrees of the visual field. *Invest. Ophtal. and Visual Sci.* **39(4)**: S181 abstract no.B784.

- Baule G and McFee R (1965) Theory of Magnetic Detection of the Heart's Electrical Activity. *J. Appl. Phys.* **36(6)**: 2066-2073.
- Bedford JL (1995) Neuro-electromagnetic imaging of the human visual cortex. PhD Thesis, The University of Aston in Birmingham.
- Biemond J, Lagendijk RL and Mersereau RM (1990) Iterative methods for image deblurring. *Proc. IEEE* **78(5)**: 856-883.
- Böcker KBE, Brunia CHM and Cluitmans PJM (1994) A spatio-temporal dipole model of the readiness potential in humans. I. Finger movement. *Electroen. Clin. Neuro.* **91**: 275-285.
- Böcker KBE, Brunia CHM and Cluitmans PJM (1994) A spatio-temporal dipole model of the readiness potential in humans. II. Foot movement. *Electroen. Clin. Neuro.* **91**: 286-294.
- Boissonnat J-D (1988) Shape Reconstruction from Planar Cross Sections. *Comput. Vision Graph* **44**: 1-29.
- Bomans M, Höhne K-H, Tiede U and Riemer M (1990) 3-D Segmentation of MR Images of the Head for 3-D Display. *IEEE T. Med. Imaging* **9(2)**: 177-183.
- Borgefors G (1986) Distance Transformations in Digital Images. *Comput. Vision Graph* **34**: 344-371.
- Bowyer A (1981) Computing Dirichlet tessellations. *The Computer Journal* **24(2)**: 162-166.
- Brauer H, Kosch O, Tenner U, Weichmann H and Arlt A (1996) A Modified Linear Estimation Approach for Solving Biomagnetic Inverse Problems. *IEEE T. Magn.* **32(3)**: 1298-1301.
- Brinkmann BH, O'Brien TJ, Dresner MA, Lagerlund TD, Sharbrough FW and Robb RA (1998) Scalp-Recorded EEG Localization in MRI Volume Data. *Brain Topogr.* **10(4)**: 245-253.
- Brooks DH, Ahmad GF, MacLeod RS and Maratos GM (1999) Inverse Electrocardiography by Simultaneous Imposition of Multiple Constraints. *IEEE T. Biomed. Eng.* **46(1)**: 3-18.
- Clarke CJS and Janday BS (1989) The solution of the biomagnetic inverse problem by maximum statistical entropy. *Inverse Probl.* **5**: 483-500.
- Clarke LP, Velthuisen RP, Camacho MA, Heine JJ, Vaidyanathan M, Hall LO, Thatcher RW and Silbiger ML (1995) Review of MRI segmentation: Methods and applications. *Magn. Reson. Imaging* **13(3)**: 343-368.
- Cohen D (1968) Magnetoencephalography: Evidence of Magnetic Fields Produced by Alpha-Rhythm Currents. *Science* **161**: 784-786.
- Cohen D (1972) Magnetoencephalography: Detection of the Brain's Electrical Activity with a Superconducting Magnetometer. *Science* **175**: 664-666.
- CTF Systems Inc. Personal Communications.
- CTF Systems Inc. (1999). Users manual.
- Cuffin BN and Cohen D (1977) Magnetic Fields of a Dipole in Special Volume Conductor Shapes. *IEEE T. Biomed. Eng.* **24(4)**: 372-381.
- Dale AM, Fischl B and Sereno MI (1999) Cortical Surface-Based Analysis I: Segmentation and Surface Reconstruction. *NeuroImage* **9**: 179-194.
- Dale AM and Sereno MI (1993) Improved Localization of Cortical Activity by Combining EEG and MEG with MRI Cortical Surface Reconstruction: A Linear Approach. *J. Cognitive Neurosci.* **5(2)**: 162-176.
- de Munck JC (1990) The estimation of time varying dipoles on the basis of evoked potentials. *Electroen. Clin. Neuro.* **77**: 156-160.

- de Munck JC (1992) A Linear Discretization of the Volume Conductor Boundary Integral Equation Using Analytically Integrated Elements. *IEEE T. Biomed. Eng.* **39(9)**: 986-990.
- de Munck JC, van Dijk BW and Spekreijse H (1988) Mathematical Dipoles are Adequate to Describe Realistic Generators of Human Brain Activity. *IEEE T. Biomed. Eng.* **35(11)**: 960-966.
- Dogandzic A and Nehorai A (2000) Estimating Evoked Dipole Responses in Unknown Spatially Correlated Noise with EEG/MEG Arrays. *IEEE T. Signal Proces.* **48(1)**: 13-25.
- Duffy FH (1982) Topographic display of evoked potentials: clinical applications of brain electrical activity mapping (BEAM). *Ann. N.Y. Acad. Sci.* **388**: 183-196.
- Duffy FH, Bartels PH and Burchfiel JL (1981) Significance probability mapping: an aid in the topographic analysis of brain electrical activity. *Electroen. Clin. Neuro.* **51**: 455-462.
- Ferguson AS, Zhang X and Stroink G (1994) A Complete Linear Discretization for Calculating the Magnetic Field Using the Boundary Element Method. *IEEE T. Biomed. Eng.* **41(5)**: 455-459.
- Fernández DC, Grave de Peralta Menéndez R and González Andino SL (1995) Some Limitations of Spatio Temporal Source Models. *Brain Topogr.* **7(3)**: 233-243.
- Fitzpatrick JM, West JB and Maurer Jr. CR (1998) Predicting Error in Rigid-Body Point-Based Registration. *IEEE T. Med. Imaging* **17(5)**: 694-702.
- Fuchs M, Wagner M, Wischmann H-A, Köhler T, Theißen A, Drenckhahn R and Buchner H (1998) Improving source reconstructions by combining bioelectric and biomagnetic data. *Electroen. Clin. Neuro.* **107**: 93-111.
- Fuchs M, Wagner M, Wischmann H-A, Ottenberg K and Dössel O (1994) Possibilities of functional brain imaging using a combination of MEG and MRT. In: *Oscillatory Event-Related Brain Dynamics* (Eds Pantev C, Elbert T and Lütkenhöner B). Plenum Press, New York.
- Fuchs M, Wischmann H-A, Wagner M and Krüger J (1995) Coordinate System Matching for Neuromagnetic and Morphological Reconstruction Overlay. *IEEE T. Biomed. Eng.* **42(4)**: 416-420.
- Geiger B (1993) Three-dimensional modeling of human organs and its application to diagnosis and surgical planning. PhD Thesis, .
- Geiger B (1996). NUAGES. Version 4.1. The French National Institute For Research In Computer Science And Control (INRIA).
- Gençer NG and Williamson SJ (1998) Differential Characterization of Neural Sources with the Bimodal Truncated SVD Pseudo-Inverse for EEG and MEG Measurements. *IEEE T. Biomed. Eng.* **45(7)**: 827-838.
- Gençer NG, Williamson SJ, Gueziec A and Hummel R (1996) Optimal reference electrode selection for electric source imaging. *Electroen. Clin. Neuro.* **99**: 163-173.
- George JS, Aine CJ, Medvick PA and Flynn ER (1989) Spatial/temporal resolution of multiple sources: paths of activation in human visual cortex. In: *Advances in Biomagnetism* (Eds Williamson SJ, Hoke M, Stroink G and Kotani M). Plenum Press, New York.
- George JS, Jackson PS, Ranken DM and Flynn ER (1989) Three-dimensional volumetric reconstruction for neuromagnetic source localization. In: *Advances in Biomagnetism* (Eds Williamson SJ, Hoke M, Stroink G and Kotani M). Plenum Press, New York.
- Gerson J, Cardenas VA and Fein G (1994) Equivalent dipole parameter estimation using simulated annealing. *Electroen. Clin. Neuro.* **92**: 161-168.
- Geselowitz DB (1970) On the Magnetic Field Generated Outside an Inhomogeneous Volume Conductor by Internal Current Sources. *IEEE T. Magn.* **6(2)**: 346-247.

- Geva A, Pratt H and Zeevi Y (1993) Wavelet decomposition of multichannel evoked potentials. *Electroen. Clin. Neuro.* **87**: S25.
- Geva AB, Pratt H and Zeevi YY (1995) Spatio-temporal multiple source localization by wavelet-type decomposition of evoked potentials. *Electroen. Clin. Neuro.* **96**: 278-286.
- Gevins A, Le J, Brickett P, Reutten B and Desmond J (1991) Seeing Through the Skull: Advanced EEGs Use MRIs to Accurately Measure Cortical Activity from the Scalp. *Brain Topogr.* **4(2)**: 125-131.
- Gevins A, Le J, Mertin NK, Brickett P, Desmond J and Reutter B (1994) High resolution EEG: 124-channel recording, spatial deblurring and MRI integration methods. *Electroen. Clin. Neuro.* **90**: 337-358.
- Ghahremani D, Makeig S, Jung T-P, Bell AJ and Sejnowski TJ (1996) Independent Component Analysis of Simulated EEG using a Three-shell Spherical Head Model. San Diego, La Jolla CA, University of California. Institute for Neural Computation Technical Report 96-01.
- Gorodnitsky I (1998) Personal Communications.
- Gorodnitsky I, George JS, Schlitt HA and Lewis PS (1992) A Weighted Iterative Algorithm for Neuromagnetic Imaging. Satellite Symposium on Neuroscience and Technology, 14th Annual International Conference, Lyon, France.
- Gorodnitsky IF (1994) Truncated Total Least Squares Regularization Algorithm for Underdetermined Problems. Proc. of 7th SP Workshop on Statistical Signal and Array Processing.
- Gorodnitsky IF, George JS and Rao BD (1995) Neuromagnetic source imaging with FOCUSS: a recursive weighted minimum norm algorithm. *Electroen. Clin. Neuro.* **95**: 231-251.
- Gorodnitsky IF and Rao BD (1992) A New Iterative Weighted Norm Minimization Algorithm and its Applications. Proc. of 6th SP Workshop on Statistical Signal and Array Processing.
- Gorodnitsky IF and Rao BD (1994) Analysis of Error Produced by Truncated SVD and Tikhonov Regularization Methods. 28th Asilomar Conference on Signals, Systems and Computers.
- Gorodnitsky IF and Rao BD (1997) Sparse Signal Reconstruction from Limited Data using FOCUSS: A Re-weighted Minimum Norm Algorithm. *IEEE T. Signal Proces.* **45(3)**: 600-616.
- Grave de Peralta Menendez R and Gonzalez Andino SL (1999) Backus and Gilbert Method for Vector Fields. *Hum. Brain Mapp.* **7**: 161-165.
- Grave de Peralta Menendez R, Hauk O, Gonzales Andino S, Vogt H and Michel C (1997) Linear Inverse Solutions With Optimal Resolution Kernels Applied to Electromagnetic Tomography. *Hum. Brain Mapp.* **5**: 454-467.
- Grave de Peralta Menendez R, Lütkenhöner B and Gonzales Andino SL (1996) Figures of Merit to Compare Distributed Linear Inverse Solutions. *Brain Topogr.* **9(2)**: 117-124.
- Grave de Peralta-Menendez R and Gonzales-Andino SL (1998) A Critical Analysis of Linear Inverse Solutions to the Neuroelectromagnetic Inverse Problem. *IEEE T. Biomed. Eng.* **45(4)**: 440-448.
- Gray CM, König P, Engel AK and Singer W (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* **338**: 334-337.
- Greenblatt RE (1993) Probabilistic reconstruction of multiple sources in the bioelectromagnetic inverse problem. *Inverse Probl.* **9**: 271-284.
- Greitz T, Bohm C, Hellstrand E, Hämäläinen MS, Hari R, Ilmoniemi LJ, Seitz R and Lounasmaa OV (1994) Use of a computerized brain atlas in magnetoencephalographic activation studies. *NeuroReport* **5**: 449-452.

- Gull SF and Daniell GJ (1978) Image reconstruction from incomplete and noisy data. *Nature* **272**: 686-690.
- Hämäläinen M (1987) Interpretation of neuromagnetic measurements: modeling and statistical considerations. PhD Thesis, Helsinki University of Technology.
- Hämäläinen M (1991) Anatomical correlates for magnetoencephalography: integration with magnetic resonance images. *Clin. Phys. Physiol. Meas.* **12(Suppl. A)**: 29-32.
- Hämäläinen M, Hari R, Ilmoniemi RJ, Knuutila J and Lounasmaa OV (1993) Magnetoencephalography - theory, instrumentation, and applications to noninvasive studies of the working human brain. *Rev. Mod. Phys.* **65(2)**: 413-497.
- Hämäläinen M and Sarvas J (1989) Realistic Conductivity Geometry Model of the Human Head for Interpolation of Neuromagnetic Data. *IEEE T. Biomed. Eng.* **36(2)**: 165-171.
- Hämäläinen MS and Ilmoniemi RJ (1984) Interpreting measured magnetic fields of the brain: estimates of current distributions. Helsinki, Helsinki University of Technology. TKK-F-A559.
- Haneishi H, Ohyama N, Sekihara K and Honda T (1994) Multiple Current Dipole Estimation Using Simulated Annealing. *IEEE T. Biomed. Eng.* **41(11)**: 1004-1009.
- Hara J, Musha T and Shankle WR (1999) Approximating Dipoles from Human EEG Activity: The Effect of Dipole Source Configuration on Dipolarity Using Single Dipole Models. *IEEE T. Biomed. Eng.* **46(2)**: 125-129.
- Hari R (1991) A neurophysiologist's view on biomagnetic source localization. In: Biomagnetic Localisation and 3D Modelling (eds J Nenonen, H-M Rajala and T Katila). Helsinki, University of Technology. TKK-F-A689.
- Hari R and Ilmoniemi RJ (1986) Cerebral Magnetic Fields. *Crit. Rev. Biomed. Eng.* **14(2)**: 93-126.
- Harner RN (1990) Singular Value Decomposition - A General Linear Model for Analysis of Multivariate Structure in the Electroencephalogram. *Brain Topogr.* **3(1)**: 43-47.
- Harner RN and Riggio S (1989) Application of Singular Value Decomposition to Topographic Analysis of Flash-Evoked Potentials. *Brain Topogr.* **2(1/2)**: 91-98.
- Heinonen T, Dastidar P, Frey H and Eskola H (1999) Applications of MR Image Segmentation. *Electroen. Clin. Neuro.* **1(1)**:
- Hildreth EC (1983) The Detection of Intensity Changes by Computer and Biological Vision Systems. *Comput. Vision Graph* **22**: 1-27.
- Hjorth B (1989) Eigenvectors and Eigenfunctions in Spatiotemporal EEG Analysis. *Brain Topogr.* **2(1/2)**: 57-61.
- Höhne KH and Hanson WA (1992) Interactive 3D Segmentation of MRI and CT Volumes using Morphological Operations. *J. Comput. Assist. Tomo.* **16(2)**: 285-294.
- Huang J-C, Nicholson C and Okada YC (1990) Distortion of magnetic evoked fields and surface potentials by conductivity differences at boundaries in brain tissue. *Biophys. J.* **57**: 1155-1166.
- Huang M, Aine CJ, Supek S, Best E, Ranken D and Flynn ER (1998) Multi-start downhill simplex method for spatio-temporal source localization in magnetoencephalography. *Electroen. Clin. Neuro.* **108**: 32-44.
- Huppertz H-J, Otte M, Grimm C, Kristeva-Feige R, Mergner T and Lücking C (1998) Estimation of the accuracy of a surface matching technique for registration of EEG and MRI data. *Electroen. Clin. Neuro.* **106**: 409-415.
- Hurley PD (1999). Individual Bust - Capturing the moment in 3-D, Individual Bust Ltd., Birmingham, UK.
- Ioannides AA, Bolton JPR and Clarke CJS (1990) Continuous probabilistic solutions to the biomagnetic inverse problem. *Inverse Probl.* **6**: 523-542.

- Ioannides AA, Bolton JPR and Hasson R (1989) Localised and distributed source solutions for the biomagnetic inverse problem II. In: *Advances in Biomagnetism* (Eds Williamson SJ, Hoke M, Stroink G and Kotani M). Plenum Press, New York.
- Ioannides AA, Liu MJ, Liu LC, Bamidis PD, Hellstrand E and Stephan KM (1995) Magnetic field tomography of cortical and deep processes: examples of "real-time mapping" of averaged and single trial MEG signals. *Int. J. Psychophysiol.* **20**: 161-175.
- Itti L, Chang L, Mangin J-F, Darcourt J and Ernst T (1997) Robust Multimodality Registration for Brain Mapping. *Hum. Brain Mapp.* **5**: 3-17.
- Jeffs B, Leahy R and Singh M (1987) An Evaluation of Methods for Neuromagnetic Image Reconstruction. *IEEE T. Biomed. Eng.* **34(9)**: 713-723.
- John ER, Ruchkin DS and Villegas J (1964) Experimental background: signal analysis and behavioral correlates of evoked potential configurations in cats. *Ann. N.Y. Acad. Sci.* **112**: 362-420.
- Johnston PR and Gulrajani RM (1997) A New Method for Regularization Parameter Determination in the Inverse Problem of Electrocardiography. *IEEE T. Biomed. Eng.* **44(1)**: 19-39.
- Kaufman L and Williamson SJ (1986) *The Neuromagnetic Field*. In: *Evoked Potentials*. Alan R Liss, Inc.,
- Kavanagh RN, Darcey TM, Lehmann D and Fender DH (1978) Evaluation of Methods for Three-Dimensional Localization of Electrical Sources in the Human Brain. *IEEE T. Biomed. Eng.* **25(5)**: 421-429.
- Khosla D, Singh M and Don M (1997) Spatio-Temporal EEG Source Localization Using Simulated Annealing. *IEEE T. Biomed. Eng.* **44(11)**: 1075-1091.
- Knösche T (1997) Solutions of the neuroelectromagnetic inverse problem. PhD Thesis, University of Twente.
- Koles ZJ and Soong AC (1998) EEG source localization: implementing the spatio-temporal decomposition approach. *Electroen. Clin. Neuro.* **107**: 343-352.
- Konyshev VA, Maragey RA, Kholodov YA, Verkhlutov VM and Gorbach AM (1989) Constructing a realistically shaped model of the human head. In: *Advances in Biomagnetism* (Eds Williamson SJ, Hoke M, Stroink G and Kotani M). Plenum Press, New York.
- Kozinska D, Tretiak OJ, Nissanov J and Ozturk C (1997) Multidimensional Alignment Using the Euclidean Distance Transform. *Graph. Model Im. Proc.* **59(6)**: 373-387.
- Kuffler SW, Nicholls JG and Martin AR (1984). From neuron to brain, a cellular approach to the function of the nervous system. Sunderland, Mass., Sinauer Associates.
- Kullmann WH (1991) Can volume conductor modelling improve biomagnetic distributed source imaging? In: *Biomagnetic Localisation and 3D Modelling* (eds J Nenonen, H-M Rajala and T Katila). Helsinki, University of Technology. TKK-F-A689.
- Lantz G, Michel CM, Pasqual-Marqui RD, Spinelli L, Seeck M, Seri S, Landis T and Rosen I (1997) Extracranial localization of intracranial interictal epileptiform activity using LORETA (low resolution electromagnetic tomography). *Electroen. Clin. Neuro.* **102**: 414-422.
- Le Goualher G, Procyk E, Collins DL, Venugopal R, Barillot C and Evans AC (1999) Automated Extraction and Variability Analysis of Sulcal Neuroanatomy. *IEEE T. Med. Imaging* **18(3)**: 206-217.
- Leahy RM, Mosher JC, Spencer ME, Huang MX and Lewine JD (1998) A study of dipole localization accuracy for MEG and EEG using a human skull phantom. *Electroen. Clin. Neuro.* **107**: 159-173.
- Lemoine D, Lussot E, Legeard D and Barillot C (1994) Multimodal Registration System for the Fusion of MRI, CT, MEG and 3D or Stereotactic Angiographic Data. *SPIE Med. Imaging* **2164**: 46-56.

- Liang Z, MacFall JR and Harrington DP (1994) Parameter Estimation and Tissue Segmentation from Multispectral MR Images. *IEEE T. Med. Imaging* **13(3)**: 441-449.
- Liu AK, Belliveau JW and Dale AM (1998) Spatiotemporal imaging of human brain activity using functional MRI constrained magnetoencephalography data: Monte Carlo simulations. *Proc. Natl. Acad. Sci. USA* **95**: 8945-8950.
- Lorensen WE and Cline HE (1987) Marching cubes: A high resolution 3D surface construction algorithm. *Computer Graphics* **21(4)**: 163-169.
- Lütkenhöner B and Grave de Peralta Menendez R (1997) The resolution-field concept. *Electroen. Clin. Neuro.* **102**: 326-334.
- Lütkenhöner B, Greenblatt R, Hämläinen M, Mosher J, Scherg M, Tesche C and Sosa PV (1996) Comparison between different approaches to the biomagnetic inverse problem - workshop report. Tenth International Conference on Biomagnetism - Biomag96, Santa Fe, New Mexico, USA.
- Lütkenhöner B, Menninghaus E, Steinsträter O, Wienbruch C, Gißsler HM and Elbert T (1995) Neuromagnetic Source Analysis Using Magnetic Resonance Images for the Construction of Source and Volume Conductor Model. *Brain Topogr.* **7(4)**: 291-299.
- Maier J, Dagnelie G, Spekrijse H and van Dijk BW (1987) Principal components analysis for source localization of VEPs in man. *Vision Res.* **27(2)**: 165-177.
- Makeig S, Bell AJ, Jung T-P and Sejnowski TJ (1996b) Independent Component Analysis of Electroencephalographic Data. In: *Advances in Neural Information Processing Systems 8* (Eds Touretzky D, Mozer M and Hasselmo M). MIT Press, Cambridge, MA.
- Makeig S, Jung T-P, Ghahremani D, Bell AJ and Sejnowski TJ (1997) Blind Separation of Event-related Brain Responses into Independent Components. *Proc. Natl. Acad. Sci. USA* **94**: 10979-10984.
- Makeig S, Jung T-P, Ghahremani D and Sejnowski TJ (1996a) Independent Component Analysis of Simulated ERP Data. San Diego, La Jolla CA, University of California. Institute for Neural Computation Technical Report 96-06.
- Malmivuo J and Plonsey R (1995). *Bioelectromagnetism*. Oxford, Oxford University Press.
- Malmivuo J, Suihko V and Eskola H (1997) Sensitivity Distributions of EEG and MEG Measurements. *IEEE T. Biomed. Eng.* **44(3)**: 196-208.
- Marquardt DW (1963) An algorithm for least-squares estimation of nonlinear parameters. *J. Soc. Indust. Appl. Math.* **11(2)**: 431-441.
- Matsuura K and Okabe Y (1995) Selective Minimum-Norm Solution of the Biomagnetic Inverse Problem. *IEEE T. Biomed. Eng.* **42(6)**: 608-615.
- Matsuura K and Okabe Y (1997) A Robust Reconstruction of Sparse Biomagnetic Sources. *IEEE T. Biomed. Eng.* **44(8)**: 720-726.
- Maurer Jr. CR, Maciunas RJ and Fitzpatrick JM (1998) Registration of Head CT Images to Physical Space Using a Weighted Combination of Points and Surfaces. *IEEE T. Med. Imaging* **17(5)**: 753-761.
- Medvick PA, Lewis PS, Aine C and Flynn ER (1989) Monte Carlo analysis of localization errors in magnetoencephalography. In: *Advances in Biomagnetism* (Eds Williamson SJ, Hoke M, Stroink G and Kotani M). Plenum Press, New York.
- Meijs JWH, ten Voorde BJ, Peters MJ, Stok CJ and Lopes da Silva FH (1988) The Influence of Various Head Models on EEGs and MEGs. In: Hans Huber Publishers (Eds Pfurtscheller G and Lopes da Silva FH)
- Michel CM, Grave de Peralta R, Lantz G, Gonzalez Andino S, Spinelli L, Blanke O, Landis T and Seeck M (1999) Spatiotemporal EEG Analysis and Distributed Source Estimation in Presurgical Epilepsy Evaluation. *J. Clin. Neurophysiol.* **16(3)**: 239-266.
- Miltner W, Braun C, Johnson Jr. R, Simpson GV and Ruchkin DS (1994) A test of brain electrical source analysis (BESA): a simulation study. *Electroen. Clin. Neuro.* **91**: 295-310.

- Mitzdorf U (1985) Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. *Physiol. Rev.* **65**: 37-100.
- Möcks J and Verleger R (1986) Principal component analysis of event-related potentials: a note on misallocation of variance. *Electroen. Clin. Neuro.* **65**: 393-398.
- Mondt JP (1989) On the effects on source localization of volume currents in neuroelectric and neuromagnetic signals. *Phys. Med. Biol.* **34(8)**: 1073-1088.
- Mosher JC and Leahy RM (1998) Recursive MUSIC: A framework for EEG and MEG source localization. *IEEE T. Biomed. Eng.* **45(11)**: 1342-1354.
- Mosher JC, Leahy RM and Lewis PS (1997) Matrix Kernels for the Forward Problem in EEG and MEG. Los Alamos, Los Alamos National Laboratory. LA-UR-97-3812.
- Mosher JC, Leahy RM and Lewis PS (1999) EEG and MEG: Forward Solutions for Inverse Methods. *IEEE T. Biomed. Eng.* **46(3)**: 245-259.
- Mosher JC, Lewis PS and Leahy RM (1992) Multiple Dipole Modeling and Localization from Spatio-Temporal MEG Data. *IEEE T. Biomed. Eng.* **39(6)**: 541-557.
- Nagano T, Ohno Y, Uesugi N, Ikeda H, Ishiyama A and Kasai N (1998) Multi-Source Localization by Genetic Algorithm using MEG. *IEEE T. Magn.* **34(5 Pt1)**: 2976-2979.
- Nelder JA and Mead R (1965) A simplex method for function minimization. *Comput. J.* **7**: 308-313.
- Nicholson C and Freeman JA (1975) Theory of Current Source-Density Analysis and Determination of Conductivity Tensor for Anuran Cerebellum. *J. Neurophysiol.* **38**: 356-368.
- Nunez PL (1986) The brain's magnetic field: some effects of multiple sources on localization methods. *Electroen. Clin. Neuro.* **63**: 75-82.
- Odesanya OS, Waggenspack Jr. WN and Thompson DE (1993) Construction of Biological Surface Models from Cross-Sections. *IEEE T. Biomed. Eng.* **40(4)**: 329-334.
- Okada Y (1982) Neurogenesis of evoked magnetic fields. In: Biomagnetism: An Interdisciplinary Approach (Eds Williamson S, Romani GL, Kaufman L and Modena I). Pergamon Press, New York.
- Okada YC (1989) Recent developments on the physiological basis of magnetoencephalography (MEG). In: Advances in Biomagnetism (Eds Williamson SJ, Hoke M, Stroink G and Kotani M). Plenum Press, New York.
- Okamoto Y, Teramachi Y and Musha T (1983) Limitation of the Inverse Problem in Body Surface Potential Mapping. *IEEE T. Biomed. Eng.* **30(11)**: 749-754.
- Oster HS and Rudy Y (1992) The Use of Temporal Information in the Regularization of the Inverse Problem of Electrocardiography. *IEEE T. Biomed. Eng.* **39(1)**: 65-75.
- Pal NP and Pal SK (1993) A review on image segmentation techniques. *Pattern Recognition* **26(9)**: 1277-1294.
- Pasqual-Marqui RD, Michel CM and Lehmann D (1994) Low Resolution Electromagnetic Tomography: A New Method for Localizing Electrical Activity in the Brain. *Int. J. Psychophysiol.* **18**: 49-65.
- Pelizzari CA, Chen GT, Spelbring DR, Weichselbaum RR and Chen C-T (1989) Accurate Three-Dimensional Registration of CT, PET, and/or MR Images of the Brain. *J. Comput. Assist. Tomo.* **13(1)**: 20-26.
- Penrose R (1955) A generalized inverse for matrices. *Math. Proc. Cambridge* **51**: 406-413.
- Phillips JW, Leahy RM and Mosher JC (1997) MEG-Based Imaging of Focal Neuronal Current Sources. *IEEE T. Med. Imaging* **16(3)**: 338-348.
- Plonsey R and Heppner DB (1967) Considerations of quasi-stationarity in electrophysiological systems. *Bull. Math. Bio.* **29**: 657-664.
- Polhemus (1992). 3Space, ISOTRAK, User's Manual.
- Powell MJD (1964) An efficient method for finding the minimum of a function of several variables without calculating derivatives. *Comput. J.* **7**: 155-162.

- Press WH, Teukolsky SA, Vetterling WT and Flannery BP (1992). Numerical Recipes in C: The Art of Scientific Computing. Cambridge, University Press.
- Psaltikidou M, Chouvarda I and Anogianakis G (-) A fully-automatic algorithm for Magnetic Resonance Image segmentation. Thessaloniki, Aristotelian University of Thessaloniki, Dept. of Physiology, Faculty of Medicine. Internal Report
- Ramon C, Meyer MG, Nelson AC, Spelman FA and Lamping J (1993) Simulation Studies of Biomagnetic Computed Tomography. *IEEE T. Biomed. Eng.* **40(4)**: 317-322.
- Reitz JR, Milford FJ and Christy RW (1992). Foundations of Electromagnetic Theory. United States of America, Addison-Wesley Publishing Company, Inc.
- Robinson SE (1989) Theory and properties of lead field synthesis analysis. In: Advances in Biomagnetism (Eds Williamson S, Hoke M, Stroink G and Kotani M). Plenum Press, New York.
- Robinson SE (1997) Functional imaging of language cortex by MEG. The Japanese Biomagnetism Conference, Japan.
- Robinson SE (1999) Synthetic Aperture Magnetometry (SAM). Personal Communications.
- Robinson SE and Rose DF (1992) Current source image estimation by spatially filtered MEG. In: Biomagnetism: Clinical Aspects (Eds Hoke M, Ern  SN, Okada YC and Romani GL). Elsevier Science Publishers, Amsterdam.
- Robinson SE and Vrba J (1998) Functional Neuroimaging by Synthetic Aperture Magnetometry (SAM). 11th International Conference on Biomagnetism, Sendai, Japan.
- Romani GL (1989) Fundamentals on neuromagnetism. In: Advances in Biomagnetism (Eds Williamson SJ, Hoke M, Stroink G and Kotani M). Plenum Press, New York.
- Romani GL (1989) The use of SQUIDS in the study of biomagnetic fields. In: Superconducting Electronics (Eds Weinstock H and Nisenhoff M). Springer-Verlag Berlin, Heidelberg.
- Roth BJ and Wikswo JP (1986) Electrically silent magnetic fields. *Biophys. J.* **50**: 739-745.
- Sakuma K, Sekihara K and Hashimoto I (1999) Neural source estimation from a time-frequency component of somatic high-frequency magnetic oscillations to posterior tibial nerve stimulation. *Clin. Neurophysiol.* **110**: 1585-1588.
- Sarvas J (1987) Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem. *Phys. Med. Biol.* **32(1)**: 11-22.
- Scherg M and Berg P (1991) Use of Prior Knowledge in Brain Electromagnetic Source Analysis. *Brain Topogr.* **4(2)**: 143-150.
- Scherg M and von Cramon D (1985a) A new interpretation of the generators of BAEP waves I-V: results of a spatio-temporal dipole model. *Electroen. Clin. Neuro.* **62**: 290-299.
- Scherg M and von Cramon D (1985b) Two bilateral sources of the late AEP as identified by a spatio-temporal dipole model. *Electroen. Clin. Neuro.* **62**: 32-44.
- Scherg M and von Cramon D (1986) Evoked dipole source potentials of the human auditory cortex. *Electroen. Clin. Neuro.* **65(344-360)**:.:
- Schimmel H (1967) The (\pm) Reference: Accuracy of Estimated Mean Components in Average Response Studies. *Science* **157**: 92-94.
- Schmidt DM, George JS and Wood CC (1999) Bayesian Inference Applied to the Electromagnetic Inverse Problem. *Hum. Brain Mapp.* **7**: 195-212.
- Schmidt RO (1986) Multiple Emitter Location and Signal Parameter Estimation. *IEEE T. Antenn. Propag.* **34(3)**: 276-280.
- Schneider M (1974) Effect of Inhomogeneities on Surface Signals Coming from a Cerebral Current-Dipole Source. *IEEE T. Biomed. Eng.* 52-54.
- Schwartz D, Lemoine D, Poiseau E and Barillot C (1996) Registration of MEG/EEG Data with 3D MRI: Methodology and Precision Issues. *Brain Topogr.* **9(2)**: 101-116.

- Sekihara K, Haneishi H and Ohyama N (1992) Details of Simulated Annealing Algorithm to Estimate Parameters of Multiple Current Dipoles Using Biomagnetic Data. *IEEE T. Med. Imaging* **11(2)**: 293-299.
- Sekihara K, Nagarajan S, Poeppel D and Miyashita Y (1999) Time-Frequency MEG-MUSIC Algorithm. *IEEE T. Med. Imaging* **18(1)**: 92-97.
- Sekihara K, Poeppel D, Marantz A, Koizumi H and Miyashita Y (1999) MEG Spatio-Temporal Analysis Using a Covariance Matrix Calculated from Nonaveraged Multiple-Epoch Data. *IEEE T. Biomed. Eng.* **46(5)**: 515-521.
- Sekihara K and Scholz B (1995) Average-Intensity Reconstruction and Wiener Reconstruction of Bioelectric Current Distribution Based on Its Estimated Covariance Matrix. *IEEE T. Biomed. Eng.* **42(2)**: 149-157.
- Sekihara K, Scholz B, Bruder H and Graumann R (1994) Reconstructing Current Distributions from Biomagnetic Measurements Under Large External Noise Disturbances. *IEEE T. Med. Imaging* **13(1)**: 144-151.
- Shareef N, Wang DL and Yagel R (1999) Segmentation of Medical Images Using LEGION. *IEEE T. Med. Imaging* **18(1)**: 74-91.
- Shim YS and Cho ZH (1981) SVD Pseudoinversion Image Reconstruction. *IEEE T. Acoust. Speech* **29(4)**: 904-909.
- Simpson GV, Pflieger ME, Foxe JJ, Ahlfors SP, Vaughan Jr. HG, Hrabe J, Ilmoniemi RJ and Lantos G (1995) Dynamic Neuroimaging of Brain Function. *J. Clin. Neurophysiol.* **12(5)**: 432-449.
- Singer W and Gray CM (1995) Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience* **18**: 555-586.
- Singh KD and Harding GFA (1996) Monte-Carlo Analysis and Confidence Region Ellipsoids for Equivalent Current Dipole Solutions. Tenth International Conference on Biomagnetism - Biomag96, Santa Fe, New Mexico, USA.
- Singh KD, Holliday IE, Furlong PL and Harding GFA (1996) Comparison of Head and Bite-Bar based MEG/MRI Co-registration. Tenth International Conference on Biomagnetism - Biomag96, Santa Fe, New Mexico, USA.
- Singh KD, Holliday IE, Furlong PL and Harding GFA (1997) Evaluation of MRI - MEG/EEG co-registration strategies using Monte Carlo simulation. *Electroen. Clin. Neuro.* **102**: 81-85.
- Singh M, Doria D, Henderson VW, Huth GC and Beatty J (1984) Reconstruction of images from neuromagnetic fields. *IEEE T. Nucl. Sci.* **31(1)**: 585-589.
- Smith WE (1992) Estimation of the Spatio-Temporal Correlations of Biological Electrical Sources from Their Magnetic Fields. *IEEE T. Biomed. Eng.* **39(10)**: 997-1004.
- Smith WE and Barrett HH (1988) Linear estimation theory applied to the evaluation of *a priori* information and system optimization in coded-aperture imaging. *J. Opt. Soc. Am. A* **5(3)**: 315-330.
- Smith WE, Dallas WJ, Kullmann WH and Schlitt HA (1990) Linear estimation theory applied to the reconstruction of a 3-D vector current distribution. *Appl. Optics* **29(5)**: 658-667.
- Snyder AZ (1991) Dipole source localization in the study of EP generators: a critique. *Electroen. Clin. Neuro.* **80**: 321-325.
- Srebro R (1994) Continuous Current Source Inversion of Evoked Potential Fields in a Spherical Model Head. *IEEE T. Biomed. Eng.* **41(11)**: 997-1003.
- Srebro R (1996a) An iterative approach to the solution of the inverse problem. *Electroen. Clin. Neuro.* **98**: 349-362.
- Srebro R (1996b) Iterative Refinement of the Minimum Norm Solution of the Bioelectric Inverse Problem. *IEEE T. Biomed. Eng.* **43(5)**: 547-552.

- Srebro R and Oguz RM (1997) Estimating cortical activity from VEPS with the shrinking ellipsoid inverse. *Electroen. Clin. Neuro.* **102**: 343-355.
- Stensaas SS, Eddington DK and Dobelle WH (1974) The topography and variability of the primary visual cortex in man. *J. Neurosurg.* **40**: 747-755.
- Stoica P and Nehoria A (1989) MUSIC, Maximum Likelihood, and Cramer-Rao Bound. *IEEE T. Acoust. Speech* **37(5)**: 720-741.
- Suetens P, Bellon E, Vandermeulen D, Smet M, Marchal G, Nuyts J and Mortelmans L (1993) Image segmentation: methods and applications in diagnostic radiology and nuclear medicine. *Eur. J. Radiol.* **17**: 14-21.
- Supek S and Aine CJ (1993) Simulation Studies of Multiple Dipole Neuromagnetic Source Localization: Model Order and Limits of Source Resolution. *IEEE T. Biomed. Eng.* **40(6)**: 529-540.
- Supek S and Aine CJ (1997) Spatio-Temporal Modeling of Neuromagnetic Data: I. Multi-Source Location Versus Time-Course Estimation Accuracy. *Hum. Brain Mapp.* **5**: 139-153.
- Swinney KR and Wikswo Jr. JP (1980) A calculation of the magnetic field of a nerve action potential. *Biophys. J.* **32**: 719-732.
- Taxt T and Lundervold A (1994) Multispectral Analysis of the Brain Using Magnetic Resonance Imaging. *IEEE T. Med. Imaging* **13(3)**: 470-481.
- Taylor JG, Ioannides AA and Müller-Gärtner H-W (1999) Mathematical Analysis of Lead Field Expansions. *IEEE T. Med. Imaging* **18(2)**: 151-163.
- Teo PC, Sapiro G and Wandell BA (1997) Anatomically Consistent Segmentation of the Human Cortex for Functional MRI Visualization. Hewlett-Packard Labs. Technical Report HPL-97-03.
- The Mathworks - Matlab, Version 5.3.
- Tilg B and Wach P (1995) An iterative approach on magnetic source imaging within the human cortex - a simulation study. *Int. J. BioMed. Comput.* **40**: 51-57.
- Tilg B, Wach P and Rucker W (1995) A Simulation Study on Magnetic Source Imaging with a Realistic Model of the Entire Human Heart. *Biomed. Tech.* **40**: 168-174.
- Tilg B, Wach P, Rucker W and Kynor D (1995) Biomagnetic functional localisation: iterative approach to estimation of electrical sources within the human heart from the magnetocardiogram. *Med. Biol. Eng. Comput.* **33**: 238-240.
- Treece GM, Prager RW and Gee AH (1998) Regularised marching tetrahedra: improved iso-surface extraction. Cambridge, Cambridge University Engineering Department. Report CUED/F-INFENG/TR 333.
- Tu YX, Wernsdörfer A, Honda S and Tomita Y (1996) Estimation of Conduction Velocity Distribution by Regularized-Least-Squares Method. *IEEE T. Biomed. Eng.* **44(11)**: 1102-1106.
- Uutela K, Hämäläinen M and Salmelin R (1998) Global Optimization in the Localization of Neuromagnetic Sources. *IEEE T. Biomed. Eng.* **45(6)**: 716-723.
- van den Broek SP (1997) Volume Conduction Effects in EEG and MEG. PhD Thesis, University of Twente.
- van den Elsen PA and Viergever MA (1991) Marker guided registration of electromagnetic dipole data with tomographic images. *Lect. Notes Comput. Sc.* **511**: 142-153.
- Van Essen DC and Drury HA (1997) Structural and Functional Analysis of Human Cerebral Cortex Using a Surface-Based Atlas. *J. Neurosci.* **17(18)**: 7079-7102.
- Van Veen BD and Buckley KM (1988) Beamforming: A Versatile Approach to Spatial Filtering. *IEEE ASSP Magazine* **5**: 4-24.
- Vrba J (1997) Baseline optimization for noise cancellation systems. IEEE-EMBS, Chicago.
- Vrba J (1998) Optimization of SQUID Detectors for MEG Application. The Japanese Bio-Magnetics Society, Tokyo, Japan.

- Vrba J, Angus V, Betts K, Burbank MB, Cheung T, Fife AA, Haid G, Kubik PR, Lee S, Ludwig W, McCubbin J, McKay J, McKenzie D, Robinson SE, Smith M, Spear P, Taylor B, Tillotson M, Cheyne D and Weinberg H (1996) 143 Channel Whole-Cortex MEG System. 10th International Conference on Biomagnetism, Santa Fe, New Mexico.
- Vrba J and McKay J (1997) Character and Acquisition of Multichannel Biomagnetic Data. ISEC, Berlin, Germany.
- Vrba J and Robinson SE (1998) Detection Probability Curves for Evaluating Localization Algorithms and Comparing Array Types. 11th International Conference on Biomagnetism, Sendai, Japan.
- Wang B, Toro C, Zeffiro TA and Hallett M (1994) Head Surface Digitization and Registration: A Method for Mapping Positions on the Head onto Magnetic Resonance Images. *Brain Topogr.* **6(3)**: 185-192.
- Wang J-Z (1993) Minimum-Norm Least-Squares Estimation: Magnetic Source Images for a Spherical Model Head. *IEEE T. Biomed. Eng.* **40(4)**: 387-396.
- Wang J-Z (1994) MNLS Inverse Discriminates Between Neuronal Activity on Opposite Walls of a Simulated Sulcus of the Brain. *IEEE T. Biomed. Eng.* **41(5)**: 470-479.
- Wang J-Z, Williamson SJ and Kaufman L (1992) Magnetic Source Images Determined by a Lead-Field Analysis: The Unique Minimum-Norm Least-Squares Estimation. *IEEE T. Biomed. Eng.* **39(7)**: 665-675.
- Wieringa HJ (1993) MEG, EEG and the integration with Magnetic Resonance Imaging. PhD Thesis, The University of Twente, The Netherlands.
- Wieringa HJ and Peters MJ (1993) Processing MRI Data for Electromagnetic Source Imaging. *Med. Biol. Eng. Comput.* **31(6)**: 600-606.
- Wikswa JP (1989) Biomagnetic sources and their models. In: *Advances in Biomagnetism* (Eds Williamson SJ, Hoke M, Stroink G and Kotani M). Plenum Press, New York.
- Williamson SJ and Kaufman L (1987) Analysis of Neuromagnetic Signals. In: *Handbook of Electrencephalography and Clinical Neurophysiology* (Eds Gevins AS and Rémond A). Elsevier, Amsterdam.
- Williamson SJ, LLü Z-L, Karron D and Kaufman L (1991) Advantages and Limitations of Magnetic Source Imaging. *Brain Topogr.* **4(2)**: 169-180.
- Wolters CH, Beckmann RF, Rienäcker A and Buchner H (1999) Comparing Regularized and Non-Regularized Nonlinear Dipole Fit Methods: A Study in a Simulated Sulcus Structure. *Brain Topogr.* **12(1)**: 3-18.
- Wood CC, Cohen D, Cuffin BN, Yarita M and Allison T (1985) Electrical Sources in Human Somatosensory Cortex: Identification by Combined Magnetic and Potential Recordings. *Science* **227**: 1051-1053.
- Zanow F (1997) Realistically shaped models of the head and their application to EEG and MEG. PhD Thesis, University of Twente.
- Zanow F and Peters MJ (1995) Individually shaped volume conductor models of the head in EEG source localisation. *Med. Biol. Eng. Comput.* **33**: 582-588.
- Zhou H and van Oosterom A (1992) Computation of the Potential Distribution in a Four-Layer Anisotropic Concentric Spherical Volume Conductor. *IEEE T. Biomed. Eng.* **39(2)**: 154-158.
- Ziólkowski M and Brauer H (1996) Methods of Mesh Generation for Biomagnetic Problems. *IEEE T. Magn.* **32(3)**: 1345-1348.

Appendices

Appendix A

MNLS solution in terms of SVD

In order to derive a formulation for the MNLS solution in terms of SVD, we start of with the matrix formulation of the forward problem.

$$\mathbf{B} = \mathbf{L}\mathbf{Q} + \mathbf{N}, \quad (\text{A-1})$$

with \mathbf{N} the vector containing the noise in the data.

The MNLS solution is given by:

$$\hat{\mathbf{Q}} = \mathbf{L}^+ \mathbf{B} = \mathbf{L}^+ \mathbf{L} \mathbf{Q} + \mathbf{L}^+ \mathbf{N}, \quad (\text{A-2})$$

Using the orthonormality of the eigenvectors obtained with the SVD, the truncated SVD of \mathbf{L} and \mathbf{L}^+ becomes:

$$\mathbf{L} = \mathbf{U} \Delta \mathbf{V}^T = \sum_{i=1}^t \lambda_i^{1/2} \mathbf{u}_i \mathbf{v}_i^T, \quad (\text{A-3})$$

$$\mathbf{L}^+ = \mathbf{V} \Delta^+ \mathbf{U}^T = \sum_{i=1}^t \lambda_i^{-1/2} \mathbf{v}_i \mathbf{u}_i^T, \quad (\text{A-4})$$

with t the truncation index.

Substituting these two equations into equation A-2 gives:

$$\hat{\mathbf{Q}} = \sum_{i=1}^l (\lambda_i^{-1/2} \mathbf{v}_i \mathbf{u}_i^T \lambda_i^{1/2} \mathbf{u}_i \mathbf{v}_i^T \mathbf{Q}) + \sum_{i=1}^l \lambda_i^{-1/2} \mathbf{v}_i \mathbf{u}_i^T \mathbf{N} = \sum_{i=1}^l (\mathbf{Q}^T \mathbf{v}_i \mathbf{v}_i + \lambda_i^{-1/2} \mathbf{N}^T \mathbf{u}_i \mathbf{v}_i), \quad (\text{A-5})$$

With the definitions $\mathbf{q}_i = \mathbf{Q}^T \mathbf{v}_i$ and $\mathbf{n}_i = \mathbf{N}^T \mathbf{u}_i$ this becomes:

$$\hat{\mathbf{Q}} = \sum_{i=1}^l (\mathbf{q}_i \mathbf{v}_i + \lambda_i^{-1/2} \mathbf{n}_i \mathbf{v}_i), \quad (\text{A-6})$$

The error function that is minimised by the truncated MNLS solution is given by:

$$\Delta = \mathbf{Q} - \hat{\mathbf{Q}} = \sum_{i=1}^r \mathbf{q}_i \mathbf{v}_i - \left\{ \sum_{i=1}^l (\mathbf{q}_i \mathbf{v}_i + \lambda_i^{-1/2} \mathbf{n}_i \mathbf{v}_i) \right\} = \sum_{i=(r+1)}^l \mathbf{q}_i \mathbf{v}_i - \sum_{i=1}^l \lambda_i^{-1/2} \mathbf{n}_i \mathbf{v}_i, \quad (\text{A-7})$$

with r the rank of \mathbf{L} .