

High Precision Anatomy for Magnetoencephalography

Luzia Rosa Tröbinger

UCL

*This thesis is submitted to University College London
for the degree of Doctor of Philosophy (PhD).*

Declaration

I, Luzia Rosa Tröbinger, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Acknowledgments

I would like to thank the J Jacob Astor Charitable Trust for sponsoring the work presented here. I would also like to thank the Brain Research Trust for the sponsorship and support it provided.

I would also like to thank the Wellcome Trust, as the work presented in this thesis was carried out at the Wellcome Trust Centre for Neuroimaging, which is supported by a strategic award from the Wellcome Trust.

Furthermore, I would like to thank Holly Rossiter for her support in proof reading this thesis, as well as David Lopez, Antoine Lutti, and my secondary supervisor Sven Bestmann.

Finally, I would like to thank my primary supervisor Gareth Barnes for his support, guidance and advice.

I would also like to note that the work presented in this thesis has led to a number of publications so far:

High Precision anatomy for MEG (Troebinger et al., 2014b)

Discriminating cortical laminae with MEG (Troebinger et al., 2014a)

Does function fit structure? A ground truth for non-invasive neuroimaging

(Stevenson et al., 2014)

Abstract

Magnetoencephalography (MEG) is a non-invasive brain imaging method with high temporal resolution but relatively poor spatial resolution as compared to some other non-invasive techniques. This thesis examines how the spatial resolution of MEG can be improved using new recording paradigms in which the subject's head position is fixed and known in advance. This is accomplished by using subject-specific head casts made using a combination of structural MRI and 3D printing technology. The resulting high-precision spatial models allow one to make inference at spatial scales of the order of cortical laminae. This thesis outlines the design of the head casts and examines the potential theoretical and empirical advances they offer. Specifically I outline simulation and then empirical investigations showing it is possible to non-invasively distinguish between electrophysiological signals in different layers of the cortex.

Table of Contents

Acknowledgments	3
Abstract	4
Table of Figures.....	8
1. Introduction.....	15
MEG Instrumentation	15
Physiological basis of the MEG signal.....	18
Structure of the neocortex	22
<i>Horizontal Organization of the Neocortex – Cortical Laminae</i>	22
MEG Scanning technique	29
Source localization error in MEG/EEG	33
Source Analysis	34
Forward model	34
<i>Spherical Models</i>	35
<i>Realistically shaped models</i>	36
Source reconstruction	40
Multiple Sparse Priors (MSP)	42
Conclusions.....	47
2. High Precision Anatomy for MEG.....	49
Abstract	49
Introduction	50
Methods and Materials.....	53
Head cast	53
MRI scanning	55
Optical scanning	56
Printing.....	57
Coregistration.....	57
<i>Native MRI to dewar-helmet</i>	58
<i>Dewar-helmet to MEG dewar</i>	58
<i>MEG dewar to MEG space</i>	59
<i>Bootstrapping for accuracy</i>	59
<i>Testing for reproducibility</i>	59
Empirical validation	60
<i>Subject task</i>	60

<i>Source reconstruction</i>	61
Results.....	64
Absolute coregistration error.....	64
Relative coregistration error.....	66
Cortical coregistration estimates	69
Relative cortical location	69
Simulation.....	71
Source reconstruction	76
Absolute cortical location	76
Discussion.....	78
3. Discrimination of cortical layers using MEG	84
Introduction	84
Methods and Materials.....	88
MRI acquisition	88
Freesurfer surface extraction.....	88
Multiple Sparse Priors (MSP)	89
Simulations.....	95
Experimental evaluation.....	99
Results.....	100
Simulations.....	100
Discussion.....	110
4. Distinguishing cortical laminae – Empirical Data	119
Introduction	119
Methods and Materials.....	125
MRI Acquisition	125
FreeSurfer Surface Extraction	126
Task	127
Source reconstruction and Multiple Sparse Priors.....	127
<i>Multiple Sparse Priors (MSP) and patch size variation</i>	128
Results.....	134
<i>Comparing outcomes based on Free Energy</i>	134
Discussion.....	148
5. Conclusions.....	154
Chapter 2.....	155

Chapter 3.....	156
Chapter 4.....	157
Future directions.....	159
References.....	162

Table of Figures

Figure 1: A. MEG system, including the chair in which patients/volunteers are seated during scanning. B. Gradiometers are used to reduce sensitivity to noise. They consist of two coils located 4-5 cm apart. While external noise results in similar field strengths at both coil locations, neural activity produces a stronger field at the coil location which is closest to the scalp (Source: <http://uuhsc.utah.edu/uumsi/ourmegsystem.html>). 17

Figure 2: The neurophysiological basis of the MEG signal. The figure illustrates the current produced by a large layer V pyramidal cell. Neural cells drive ionic currents, which arise from the difference in electrical potentials between basal and apical dendrites. These currents are referred to as primary currents. They can be modelled as equivalent current dipoles (ECD – black arrow). The dark plain lines represent secondary or volume currents. Both primary and secondary currents generate magnetic fields. The dotted lines illustrate the magnetic field lines induced by the primary current. (Source: <http://www.canada-meg-consortium.org/EN/MegBaillet2>).. 19

Figure 3: The mass effect of currents due to neural cells adds up at a larger spatial scale. This activity can also be represented as an ECD, illustrated by the red arrow in the figure above. The primary current source induces secondary currents, which are shown using purple lines. In addition, magnetic fields are induced (green lines). These can be measured using MEG (Source: <http://www.canada-meg-consortium.org/EN/MegBaillet2>). 20

Figure 4: Different types of glial cells visualized using silver staining (Source: Purves D, Augustine GJ, Fitzpatrick D, et al., (Editors)). 23

Figure 5: Different types of pyramidal cells. The figure on the left illustrates the structure of a pyramidal neuron from layers II/III, the figure on the right shows a pyramidal neuron from layer V (Source: Spruston, 2008; full length of layer II/III pyramidal neuron: unknown; full length of layer V pyramidal neuron: 1,180 μm). 26

Figure 6: Laminar organization of the cortex and layer specific inputs/outputs. Input from other cortical areas is received in layers I, II, IV and V, input from the brainstem by layers I-V. Layer IV also receives input from thalamic cells. In terms of output, layers I-III project to other cortical areas/the opposite hemisphere (layer III), layer V features outputs to subcortical structures, and layer VI connects to the thalamus (Source: DiFiore's Atlas of Histology with Functional Correlations, Victor P. Eroschenko). 27

Figure 7: Horizontal structure of the cerebral cortex. The figure illustrates the different cell types in each of the six layers visualized using silver staining (Source: <http://www.ncbi.nlm.nih.gov/pubmed/18270515>). 28

Figure 8: Different types of head models. The leftmost figure shows the simplest type of head model, which involves approximating the head as a sphere. The panel in the middle shows a single shell head model, in which just the inner skull boundary is modelled. Finally, the rightmost figure illustrates a three shell model, consisting of shells representing the skull, scalp, and inner skull boundary. 37

Figure 9: MEG Forward model. The rightmost panel illustrates the location of the pial and white matter surfaces in a structural MRI. Two cortical models are then derived from these surfaces, which form the basis of the forward model. The leftmost panel illustrates the cortical mesh (blue), the inner skull boundary (grey mesh), and the sensor positions (green dots). Note that while this model uses a single shell, further shells representing other tissue types can be included. 39

Figure 10: The left hand side illustrates the use of patches in MSP. The black dots correspond to MSP patch centre locations. On the right hand side, we show an example of source reconstruction using MSP. Following the search and optimisation process, only a subset of patches are 'left on', illustrating the estimated source locations. 45

Figure 11: Head cast/coregistration. Panel A shows the optical scan of the dewar helmet, which, in combination with an optical scan of the subject's head (Panel C), serves as the basis for creating the template of the head cast (Panel D). For coregistration purposes, MRI visible markers are placed on the outside of the dewar helmet, such that their position in an MRI structural image can be related to its counterpart in an optical scan (Panel E). To obtain a transform which allows us to relate MEG sensor space and MRI space, we use the 4 measurement coil positions in relation to the MEG system dewar in a manufacturer supplied image (Panel B). 54

Figure 12: Investigating sensitivity to cortical model used. Panel A shows the comparison between individual and canonical surface models in terms of log model evidence for each of 32 patch(/prior) sets used to fit the averaged evoked response from 9 separate sessions. The model evidence difference is generally positive, indicating that the individual cortical surface model is favoured in the majority of cases. Using a random effects analysis, Panel B shows the frequency with which the individual white surface is favoured over the canonical model within Panel A session 9. Panel C shows the exceedance probability (the area under the curve in Panel B where frequency $r > 0.5$) plotted over all sessions. Again, the evidence is overwhelmingly in favour of the individual cortical surface. Lastly, in Panel D, we investigated the effect of adding additional coregistration error and performing the same analyses. As coregistration error increases, the advantage of using a model based on the individual subject's cortical anatomy diminishes to the point where it is no longer favoured over the canonical model. 63

Figure 13: Evaluation of repeatability of coregistration results by means of bootstrapping. The bars show the standard deviation the three fiducial locations in x, y and z directions based on bootstrapped estimates of the transformation matrices. Panel A shows the likely error in fitting the dewar-helmet shape (A) to the MEG dewar (B). Panel B shows the fitting of the MRI fiducial locations to the dewar-helmet fiducials (E). Panel C shows an estimate of the error due to fitting the MEG sensor space (as measured using fiducial coils) to MEG dewar fiducial points (B, yellow arrows). Finally, all transforms were bootstrapped simultaneously (Panel D) to get an idea of the total expected absolute coregistration error— which we would expect to be of the order of 1 mm per fiducial. 65

Figure 14: Within (A) and between (B) session cast/head movement estimates for both subjects. Panel A. Within session estimates of variability in cast (left most bars) and head (right most bars) position. Note both cast-movement estimates are below 0.3 mm indicating that the external surface of the casts fit securely within the MEG dewar. Also note also that for subject 1 head movement is limited to within 0.3 mm but for subject 2, whose cast fit less well internally (most likely due to the optical scanning stage, see discussion), this figure is around 0.6 mm. Panel B shows between session cast/head movement estimates for both subjects. The between session errors in cast location were comparable with the within session errors for subject 1, but for subject 2 showed and increased to around 0.6 mm in the z (vertical) direction. In terms of head position, we were able to reposition subject 1 to within 0.3 mm between sessions. Whereas subject 2 was again more variable but could still be repositioned to within 1.4 mm, the error again predominant in the z direction indicating some freedom of vertical movement within the cast. 68

Figure 15: Repeatability of between session sensor level evoked responses. Panels A and B show the MEG channel with highest variance (MLT35) and its contralateral counterpart (MRT35) respectively. Solid black lines show the nine averaged evoked responses from each session over a six month period. Data were baseline corrected from -0.2 to 0 s. The shaded regions show the mean within session 95% confidence intervals on this mean. Note that the between session variability in the evoked response is of the same order as the within session variability. This between session variability includes physiological changes in the evoked response as well as noise due to changes in relative head position over scans. In panel C, current density estimates from the grand average evoked response are projected back out to the sensors using different forward models. Only the time period used for the source inversion is shown (hence the zeros outside $0-0.3$ s). The forward models differ in their coregistration error: the shaded blue area shows the sensor level standard deviation due to 1 mm error in coregistration over sessions; whilst coregistration errors of 5 mm (shaded red) give rise to considerably larger noise levels. Panel D shows sensor noise as a function of coregistration error. In the absence of any other noise, increasing the coregistration error from 1 to 5 mm incurs a 5 fold increase in RMS noise levels. The dotted line shows the within session standard error (the variability on the mean of any one session expected by chance) whilst the circles show the standard deviation of the evoked response (at peak latency) over sessions. Note that there is almost no difference between these two lines, which based on the simulation, should begin to diverge for coregistration errors of greater than 1 mm. 70

Figure 16: The repeatability of the MSP source level estimates. Panel A shows posterior probability maps of the current density (J) on the inflated cortical surface at 105 ms post stimulus at each of the 9 scanning sessions from within six month period. The colour scales show probability that J is non zero. Note the similarity over scanning sessions and the similarity to the MSP source level estimate based on the grand average (Panel B). Panel C shows the time-series estimates from the peak location in source space. Dotted black lines show source level estimates over the 9 scanning sessions and the average of these is shown as a blue solid line; the average 95% confidence estimates on any one session are shown in the blue shaded area. The estimate of current density on the grand mean sensor level data (red circles) and its reduced 95% confidence limits (red shaded) are also shown. 73

Figure 17: Field maps at time 0.105s post stimulus onset over the 9 scanning sessions. 74

Figure 18: Correlation matrix of field maps show in figure 19. The lowest correlation between any two field maps is 0.97. 75

Figure 19: Consistency across scanning sessions. The figure illustrates the consistency of MSP results over all 9 scanning sessions by showing the number of data sets with a posterior probability of >0.9 at each vertex. 75

Figure 20: Outline of the simulation process. We use FreeSurfer to extract the pial and white matter surfaces for a single male subject. The pial and the white matter surfaces define the superficial and deep generative surface models. Activity is simulated by randomly selecting source locations from a list of 21401 possible patch centres. Data are generated using two different patch sizes; corresponding to either $FWHM = 5$ mm or $FWHM = 10$ mm (top left inset panel). Sources are simulated on both surfaces, followed by coregistration of each dataset to both cortical models. To add coregistration error, a random perturbation is added to the fiducial locations at this stage, either taking the form of rotation or pure translation of zero mean and specified standard deviation (0, 1, 2, 5, 10, 20 mm/degrees). In both cases, coregistration error takes the form of a lateral shift. Finally, we use the MSP algorithm to perform source reconstruction. This

yields a free energy value as an outcome, which approximates the log model evidence, allowing us to compare between reconstructions based on the ‘correct’/‘incorrect’ surface models, as well as ‘correct’/‘incorrect’ patch sizes. 94

Figure 21: An example to show the link between the free energy metric used in this paper and traditional cross-validation approaches. The basic approach we used is illustrated in Panel A. Here, we simulate 3 sources on the deep surface, and reconstruct these data onto two candidate cortical models (superficial and deep). In this case the free energies (or log model evidences) for the two models differ by 4.3. This suggests that the pial candidate model is more likely with a probability of $1/(1 + \exp(-4.3)) = 0.9866$. An alternative method to judge between models would be to use cross-validation to see which model predicts new data more accurately. Based on the original set up we simulated a further 101 data sets, using the same source locations on the superficial surface (in other words these data had exactly the same underlying signal but different noise realisations). We now use the two candidate models to generate data and compare these predictions with new test data. Panel B shows the signal for a single MEG channel, for a single test dataset (blue solid) and predictions from the two candidate reconstructions (deep model green crosses, superficial model red circles). Note that the error between the superficial candidate model and the test model (based on superficial) is smaller than that of the deep. Panel C shows the ratio of these errors (over all channels and time points) for the two candidate models over 101 test datasets. The red line is at unity, points above the line show smaller error for the candidate deep surface model, points below indicate that the superficial model provides a better prediction of the test data. The incorrect model (deep surface model) is favoured in only two cases. This means that the deep model is more likely with a probability of $2/101 = 0.9802$, in accord with the analytically derived posterior probability based on free energy (0.9866). 98

Figure 22: A: The effect of lateral coregistration error (shift). The bars shows the average (over triplet simulations) free energy difference between true and incorrect surface models (evidence in favour of the true model is positive) for source triplets simulated on the superficial (left) and deep (right) surface models. Different coloured bars show different levels of coregistration error in mm. Both plots indicate that the ability to discriminate between models representing different cortical surfaces is destroyed once coregistration error exceeds 2 mm (free energy differences < 3). B: The effect of rotational coregistration error. The bars shows the average (over triplet simulations) free energy difference between true and incorrect surface models (evidence in favour of the true model is positive) for source triplets simulated on the superficial (left) and deep (right) surface models. Both for simulations based on the superficial as well as the deep surface model, the cut-off for being able to distinguish between the true/incorrect surface models (with 95% certainty or free energy > 3.0) lies around the 2-degree-mark. 102

Figure 23: The effect of varying SNR. The ability to distinguish between surface models improves with SNR. At 1 dB SNR, it is not possible to make a significant distinction between cortical models (free energy < 3). However, at a higher, yet still moderate SNR of 11 dB, there is strong evidence in favour of the correct surface model. 104

Figure 24: Investigating the effect of patch extent. Data were simulated on the superficial (left) or deep (right) surfaces models using either a patch extent of FWHM = 10 mm (upper panels), or FWHM = 5 mm (lower panels). Three sources were modelled, at a per trial SNR of 0 dB. In addition, rotational coregistration errors of 0, 2 and 5° were simulated. Panel A in the top row shows relative free energy differences for simulations based on the superficial surface and using a patch extent of FWHM = 10 mm. Looking at the leftmost set of bars (corresponding to reconstructions using

FWHM 10 mm), we observe strong evidence in favour of the correct model (positive values) for reconstructions at true patch extent, both for 0 and 1, but not 5° of coregistration error. This pattern is destroyed when underestimating patch extent (FWHM = 5 mm), as illustrated by the rightmost set of bars where no clear difference between surfaces emerges. However, looking at Panel B, which shows the comparison based on the same patch size (10 mm), but using the deep surface model, it is clear that even if we underestimate patch extent (5 mm), the strong evidence in favour of the true model is preserved. This suggests that when we underestimate patch extent, we are introducing a bias towards deeper surface models. The bottom row shows relative free energy differences for simulations based on smaller patch extent (FWHM = 5 mm). Panels C and D correspond to the true surface models being superficial and deep, respectively. Here, in the case of superficial-surface-based data, when overestimating patch size, the evidence in favour of the superficial surface model is preserved (Panel C). On the other hand, as shown in Panel D, in the case of deep-surface-based simulations, overestimating patch size decreases the evidence in favour of the deep surface model. This indicates that by overestimating patch extent, we are introducing a bias towards superficial surface models. 107

Figure 25: Investigation of layer and patch size models for an auditory evoked response paradigm in one subject over eight recording sessions. Panel A shows the difference in log model evidence between deep and superficial layer models for each of the eight runs for 5 (blue bars) and 10 mm (red bars) patch sizes. Positive values mean that the deep surface model is more likely; this is the case in six of the eight runs for both patch sizes. Panel B shows the probability of the two layer models based on all eight runs and both patch sizes. Note that the deeper cortical surface is more likely with a posterior probability of 0.9834. Panel C shows the probability of the different patch size models this time marginalising over layer models, the 5 mm patch size being more likely with a posterior probability of 0.8238. Panels D and E show the average posterior probability (probability that the current at any point is non zero) map for $t = 0.92$ s post stimulus over the eight sessions (combining both patch sizes) for deep and superficial layer models respectively..... 109

Figure 26: Time-frequency plot of the observed response when participants executed simple abductions of the right hand index finger cued by an auditory pip. Note the clear increase in gamma amplitude and the attenuation in beta amplitude around the time of movement onset, as well as the distinct beta rebound on movement cessation (Source: Muthukumaraswamy, 2010)..... 123

Figure 27: The strategy used for source reconstruction. Individual cortical surfaces were extracted from anatomical MRIs using Freesurfer, yielding a deep and superficial cortical model. Coregistration was performed using both models, followed by source reconstruction, yielding a free energy value as an outcome. We then repeated the process for different patch extents. 132

Figure 28: Results for subject OG, illustrated in terms of Free Energy differences. Each group of bars shows results for a particular patch size, with each individual bar corresponding to a single ten-minute-recording session. Positive values indicate a preference for the deep layer, negative values a preference for the superficial layer model. Panels A and B show results for Beta activity in time windows 1 (-500 to 500 ms) and 2 (500 to 1500 ms), respectively, Panel C shows results for gamma activity. In all cases, results were consistent with our hypothesis..... 136

Figure 29: Illustrating effect of varying patch size. Each bar represents the sum of Free Energy Differences across all runs at each patch size. 138

- Figure 30:** Outcomes in terms of Free Energy differences, subject ME. Positive values indicate a preference for the deep, negative values a preference for the superficial surface model. Again, results were consistent with our pre-formed hypothesis. 140
- Figure 31:** ME, illustrating the effect of varying patch extent. Bars represent the sum across Free Energy differences for all available runs at each patch size..... 141
- Figure 32:** Subject ES, outcomes illustrated in terms of Free Energy differences..... 143
- Figure 33:** ES, illustrating the effect of patch size. The pattern seems to indicate that as patch size increases, the absolute value of Free Energy differences increases, introducing a surface-specific bias. 144
- Figure 34:** GB, outcomes illustrated in terms of Free Energy differences. This was the only subject whose results did not agree with our a-priori hypothesis. However, we used a 'first-generation' type head cast for GB, with a slightly different coregistration scheme, which might be a possible reason for the striking differences. 146
- Figure 35:** GB, illustrating the effect of patch size. A clear, consistent pattern is visible - as patch extent increases, so does the absolute value of Free Energy Differences. . 147

1. Introduction

Magnetoencephalography (MEG) is a neuroimaging modality which measures the changes in magnetic fields induced by underlying neural activity. As a result, MEG benefits from excellent temporal resolution. Although spatial resolution has so far been limited, this is due to practical and empirical reasons rather than physiological constraints. This thesis sets out to address these practical and empirical concerns and aims to provide a perspective on the spatial resolution ultimately achievable with MEG.

Here, we will provide a brief introduction to MEG, including an overview of instrumentation and analysis methods. We will also discuss the neurophysiological basis of the MEG signal and the laminar organization of the neocortex.

MEG Instrumentation

MEG provides a direct measure of neural activity, resulting in excellent temporal resolution (in the millisecond-range). This sets it apart from other imaging modalities such as PET or fMRI, which infer neural activity from a hemodynamic response and have a temporal resolution in the range of seconds, much slower than the actual timescale of neural activity.

MEG measures the changes in magnetic fields induced by underlying neural activity, whose field strengths typically are in the fT-range (femto-Tesla). The measurement of such subtle field strengths requires extremely sensitive flux detectors. In the late 1960s, such detectors became available through the

invention of Superconducting Quantum Interference Devices (SQUIDs) (Josephson, 1962, Zimmermann, 1965, Hari and Salmelin, 2012).

Paired with magnetometers (pick-up coils), these detectors form the basic building blocks of MEG technology. A typical MEG system contains between 100-300 such sensors, which are housed inside a structure known as the dewar. Since superconducting technology requires extremely low temperatures (-269° C) to work, the dewar contains liquid helium for cooling purposes. A typical MEG system is illustrated in Figure 1.

In addition to highly sensitive measurement devices, the second important requirement for MEG is noise attenuation. Given that the magnetic fields we are interested in are much weaker than the earth's magnetic field, any external noise is likely to obscure them completely. For this reason, MEG systems are usually housed inside magnetically shielded rooms. This helps to eliminate noise from external sources such as moving vehicles/elevators, electronic devices or the power line. However, even factors such as mechanical vibration of the scanning system's dewar, or physiological aspects such as a subject's heartbeat can represent significant sources of noise. For this reason, gradiometers are used to reduce sensitivity to such signals. In their most basic form, they consist of two oppositely wound coils located about 4-5 cm apart, which measure the difference in magnetic flux between these two points in space. External noise is generated by sources which are approximately the same distance away from both coils, meaning that the strength of the resulting field will be the same at both coil locations. Thus, such signals would result in a zero net effect. However, the same is not true for magnetic fields generated by neural activity, since one of the coils is located closer to the scalp than the

other. Therefore, magnetic fields which are due to neural activity cause a signal to be registered at the sensor level, while the effect due to magnetic fields arising from external noise cancels out (Figure 1, Panel B).

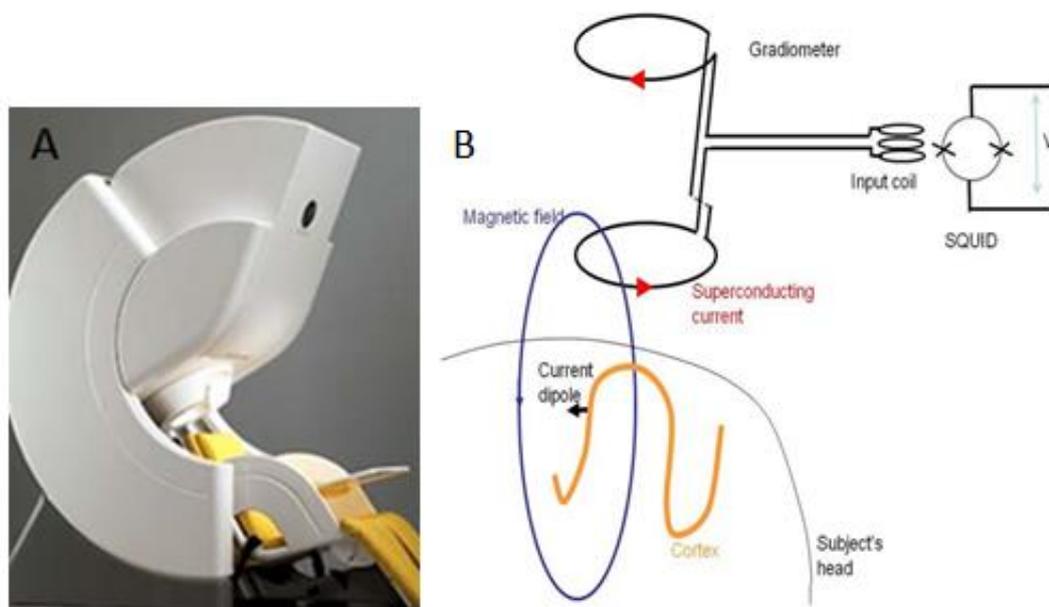


Figure 1: A. MEG system, including the chair in which patients/volunteers are seated during scanning. B. Gradiometers are used to reduce sensitivity to noise. They consist of two coils located 4-5 cm apart. While external noise results in similar field strengths at both coil locations, neural activity produces a stronger field at the coil location which is closer to the scalp (Source: <http://uuhsc.utah.edu/uumsi/ourmegsystem.html>).

Physiological basis of the MEG signal

Individual neurons in the brain have electrochemical properties which result in the flow of electrically charged ions through a cell. This slow ionic current flow generates electromagnetic fields, which form the basis of the MEG signal. However, the field generated by a single active neuron is not strong enough to produce a measurable magnetic field outside the head – it takes the combined effect of large populations of neurons ($\sim 10^4$) which exhibit simultaneous, synchronous activity to generate such a field (Murakami and Okada, 2006, Hari and Salmelin, 2012). Even so, the resulting magnetic fields are in the fT range, which means they are extremely subtle – about a billionth of the strength of the earth's magnetic field.

In addition, the cell geometry has to be such that it supports a large net current, while its orientation results in current flow that is tangential with respect to the scalp. Neurons with dendrites along the longitudinal axis possess what is commonly referred to as an 'open field geometry'. In essence this means that their electrical fields are not contained within the cell volume, resulting in a large net current. In contrast, cells with closed field configuration feature radially symmetric dendrites, resulting in weak or negligibly small net currents.

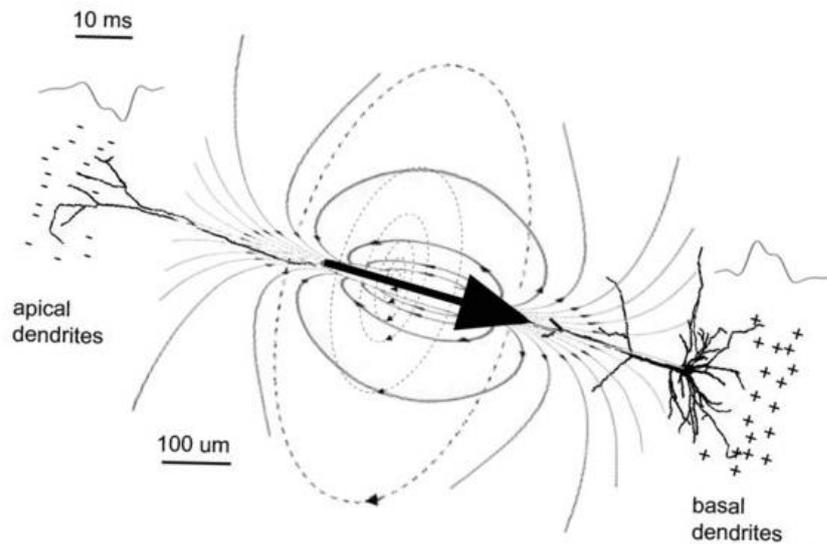


Figure 2: The neurophysiological basis of the MEG signal. The figure illustrates the current produced by a large layer V pyramidal cell. Neural cells drive ionic currents, which arise from the difference in electrical potentials between basal and apical dendrites. These currents are referred to as primary currents. They can be modelled as equivalent current dipoles (ECD – black arrow). The dark plain lines represent secondary or volume currents. Both primary and secondary currents generate magnetic fields. The dotted lines illustrate the magnetic field lines induced by the primary current. (Source: <http://www.canada-meg-consortium.org/EN/MegBaillet2>.)

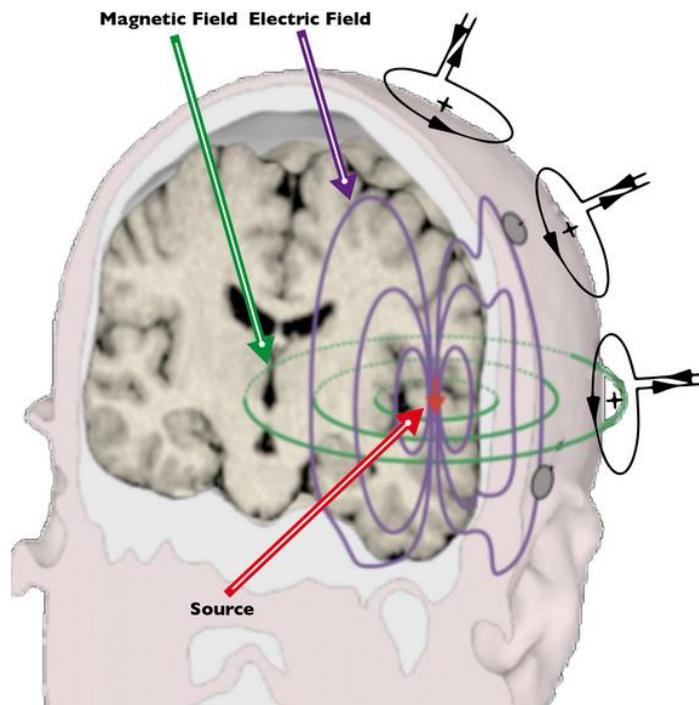


Figure 3: The mass effect of currents due to neural cells adds up at a larger spatial scale. This activity can also be represented as an ECD, illustrated by the red arrow in the figure above. The primary current source induces secondary currents, which are shown using purple lines. In addition, magnetic fields are induced (green lines). These can be measured using MEG (Source: <http://www.canada-meg-consortium.org/EN/MegBaillet2>).

The cortical surface is highly convoluted, consisting of a large number of gyri and sulci. The direction of a cell's current flow is perpendicular to the cortical surface, meaning that the net current will be either radial or tangential with respect to the scalp. In theory, a radial source will not produce a measurable magnetic field outside a spherical volume conductor. The human head is often modelled as a sphere, which would mean that MEG is 'blind' to radial sources. However, since the head is not really spherically symmetric, radial sources will not be completely 'silent', but will produce significantly weaker field strengths as compared to tangentially oriented sources (Ahlfors et al., 2010).

The third important factor which influences the strength of the resulting sensor level signal is source depth. The law of Biot and Savart states that the strength of a magnetic field decays as the square of distance. Hence, sources that are located further away from the sensors will produce significantly weaker fields – for this reason, there is much controversy about the detectability of sub-cortical sources with MEG (Ahlfors et al., 2010, Quraan et al., 2011).

In fact, source depth is often thought to exert a greater influence on detectability than source orientation. Hillebrand and Barnes (Hillebrand and Barnes, 2002) conducted an investigation of the sensitivity profile of MEG to sources of different depth and orientation. They found that source depth had a much greater influence on the ability to detect a given source than orientation. While there are thin strips of poor resolvability at the crests of gyri (~2mm wide), they are surrounded by elements with marginal tangential components, which benefit from their close proximity to the sensors. Thus, the superficial location of such sources results in relatively high detectability, making source depth the main limiting factor of MEG sensitivity.

In summary, a source producing a strong MEG signal must be formed by a large cortical population of neurons exhibiting synchronous, simultaneous activity. In order to produce a large net effect, the geometry of each cell must be such that it supports a large net current (open-field geometry). Therefore, the MEG signal ultimately not only depends on individual cells, but also on their local distribution across the cortex.

Structure of the neocortex

The neocortex is a highly convoluted sheet of tissue, which is between 2-4 mm thick and has a surface area of approximately 2600 cm². It is only found in mammals, and is organized into horizontal and vertical structures: In the horizontal sense, it consists of up to 6 layers, which can be distinguished based on their cellular architecture (Rockel et al., 1980). In addition, cells residing in different layers are linked through synaptic connections, forming vertical structures. In their simplest form, these vertical structures are grouped together in what is known as minicolumns – chains of neurons extending vertically from layers II-VI (Mountcastle, 1997). When large numbers of such minicolumns are bound together through short-range horizontal connections, they form cortical columns.

Horizontal Organization of the Neocortex – Cortical Laminae

Horizontally, the neocortex can be divided into up to six layers. Each layer is formed by distinctive neuronal populations, which feature cells of different densities, sizes and shapes. Layer IV is known as the granular layer, the layers above it (I – III) are called supra-granular, the layers beneath it (V-IV) infra-granular layers. The basic structure of the neocortex is illustrated in Figure 6.

Layer I, the uppermost cortical layer, is also known as the molecular layer, and is located directly underneath the pia mater. It is relatively cell-sparse, featuring mostly glial cells, some non-pyramidal neurons and dendrites of neurons belonging to other layers (Prieto et al., 1994). Since around 90-95% of Layer I cells stain for glutamate decarboxylase or GABA, it is likely that they may be GABA-ergic i.e. inhibitory. Additionally, the fact that most layer I neurons feature smooth dendrites with only very few spines suggests that they may be similar to interneurons and non-pyramidal cells in layers II-VI (Hestrin and Armstrong, 1996). Layer I receives input from thalamic cells, as well as from other cortical layers.

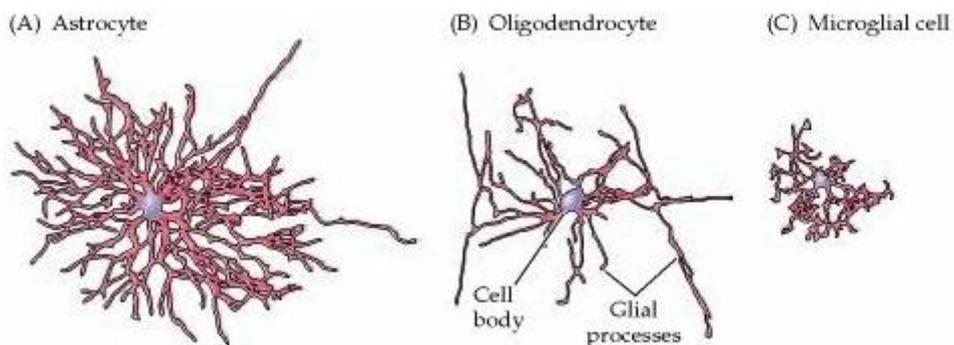


Figure 4: Different types of glial cells visualized using silver staining (Source: Purves D, Augustine GJ, Fitzpatrick D, et al., (Editors)).

Layer II, the external granular layer, consists mainly of different types of glial cells, small to medium sized pyramidal cells and spiny stellate cells. The pyramidal cells, due to their close proximity to the pia, do not form typical apical dendrites. As a result, their shape is more ovoid or round rather than triangular. Since this alters the layer's appearance, it is classed as 'granular', although it is

dominated by pyramidal cells (Peters et al., 1985). The molecular and neurochemical features of layer II are very similar to those of layer III, which has led to them often being combined and referred to as layer II/III. However, connectivity differs between the two layers: Layer II cells receive input from layers IV and V, as well as intralaminar input (Shepherd and Svoboda, 2005). Layer II pyramidal cells also feature denser, more extensive axonal arbours which target layers II and V.

Layer III is dominated by a morphologically homogeneous population of medium-sized pyramidal neurons. For this reason, it is also known as the external pyramidal layer. These pyramidal cells share almost identical somatodendritic morphological features – hence they are thought to represent the ‘purest’ specimens of pyramidal cells, together with those in layer V (Peters et al., 1985). The apical dendrites of layer III pyramidal neurons reach into layer I, where they form well-defined tufts, while the basal dendrites terminate in layers III and IV (Schroder and Luhmann, 1997, Lubke et al., 2000). Layer III also contains non-pyramidal cells, which are thought to be GABAergic interneurons (Helmstaedter et al., 2009).

Connectivity in layer III is dominated by inputs from other layers, as well as local intralaminar connections (Yoshimura and Callaway, 2005, Feldmeyer et al., 2006). Layer V is the main target structure for output from layer III.

Layer IV - the internal granular layer – is generally relatively thin, although its thickness varies across cortical regions. In some areas, it has a rudimentary presence, while being completely absent in others (e.g. the primary motor cortex). If a region contains no layer IV, it is usually referred to as being

'agranular'. The internal granular layer contains mostly small granule cells, some pyramidal cells and different types of glial cells. Specifically, this layer is populated by medium sized neurons with symmetrical or asymmetrical small dendritic trees. Some dendrites feature a large number of spines (spiny stellate cells), while others feature a smaller number or no spines at all (smooth/aspiny stellate cells) (Simons et al., 1984). The spine free and spine-sparse neurons are generally believed to be GABA-ergic interneurons. Layer IV also features pyramidal neurons, namely symmetric/asymmetric star pyramidal cells as well as 'classical' pyramidal neurons (Staiger et al., 2004). However, the number of pyramidal cells present in layer IV is highly region- (and species-) specific (Jones et al., 1975).

The internal granular layer receives input from two main sources; the thalamus and local neurons residing within the same region. It has also been suggested that, depending on cell type, layer IV neurons receive additional intracolumnar input from the supra- and infragranular layers.

Layer V, the internal pyramidal layer, is mainly populated by large pyramidal cells, whose axons extend to subcortical structures, such as the basal ganglia. It is sometimes 'split' into layer Va and layer Vb. Layer Va is populated by medium-sized pyramidal cells that feature slender dendritic trees with tufts terminating in layer I. The morphology of layer Vb is dominated by medium-sized and large pyramidal cells, with a much higher packing density than those in layer Va. Layer Va contains the largest pyramidal cells, particularly in the primary motor area, where they are known as Betz cells. These giant pyramidal neurons feature axons that extend through the internal capsule, the brain stem

and the cortico-spinal tract, forming the main pathway for voluntary motor control.

In addition, GABA-ergic interneurons are present at the layer IV/Va border (Ren et al., 1992). Local connections are dominated by intra-, as well as trans-columnar input.

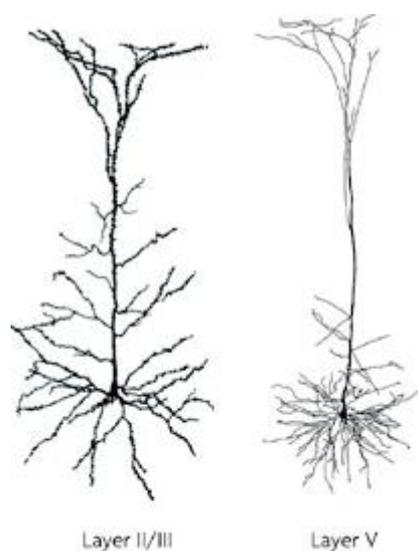


Figure 5: Different types of pyramidal cells. The figure on the left illustrates the structure of a pyramidal neuron from layers II/III, the figure on the right shows a pyramidal neuron from layer V (Source: Spruston, 2008; full length of layer II/III pyramidal neuron: unknown; full length of layer V pyramidal neuron: 1,180 μm).

Finally, layer VI is known as the multiform layer and is adjacent to the white matter. It is characterised by a population of small spindle-like and multiform pyramidal neurons, as well as some large pyramidal neurons. It also contains fusiform cells, granule cells, stellate cells and cells of Martinotti.

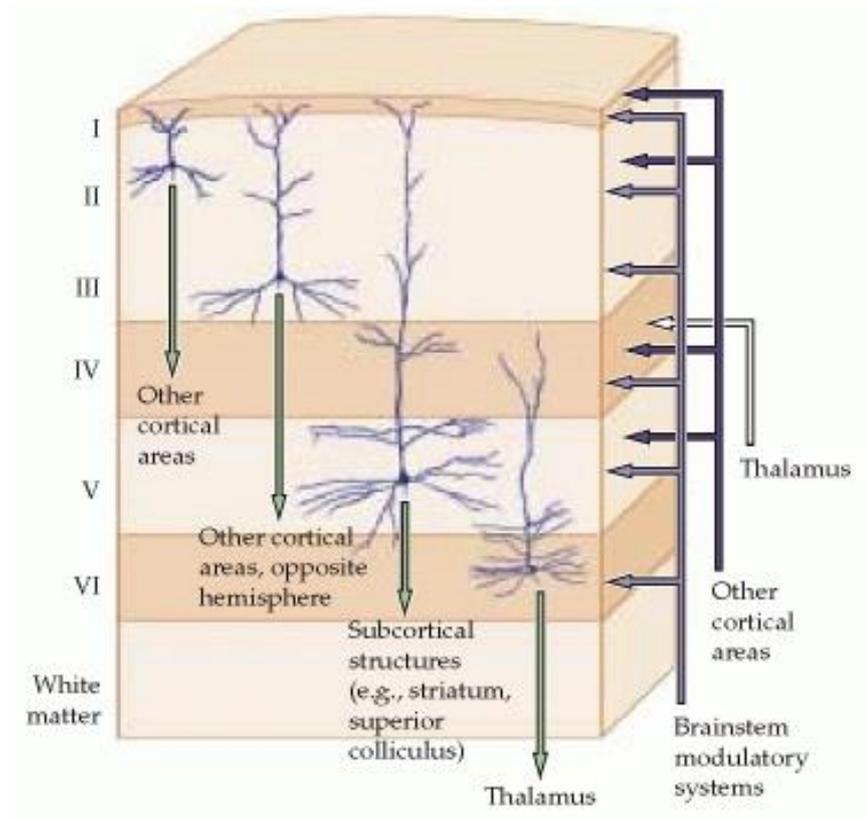


Figure 6: Laminar organization of the cortex and layer specific inputs/outputs. Input from other cortical areas is received in layers I, II, IV and V, input from the brainstem by layers I-V. Layer IV also receives input from thalamic cells. In terms of output, layers I-III project to other cortical areas/the opposite hemisphere (layer III), layer V features outputs to subcortical structures, and layer VI connects to the thalamus (Source: DiFiore's Atlas of Histology with Functional Correlations, Victor P. Eroschenko).

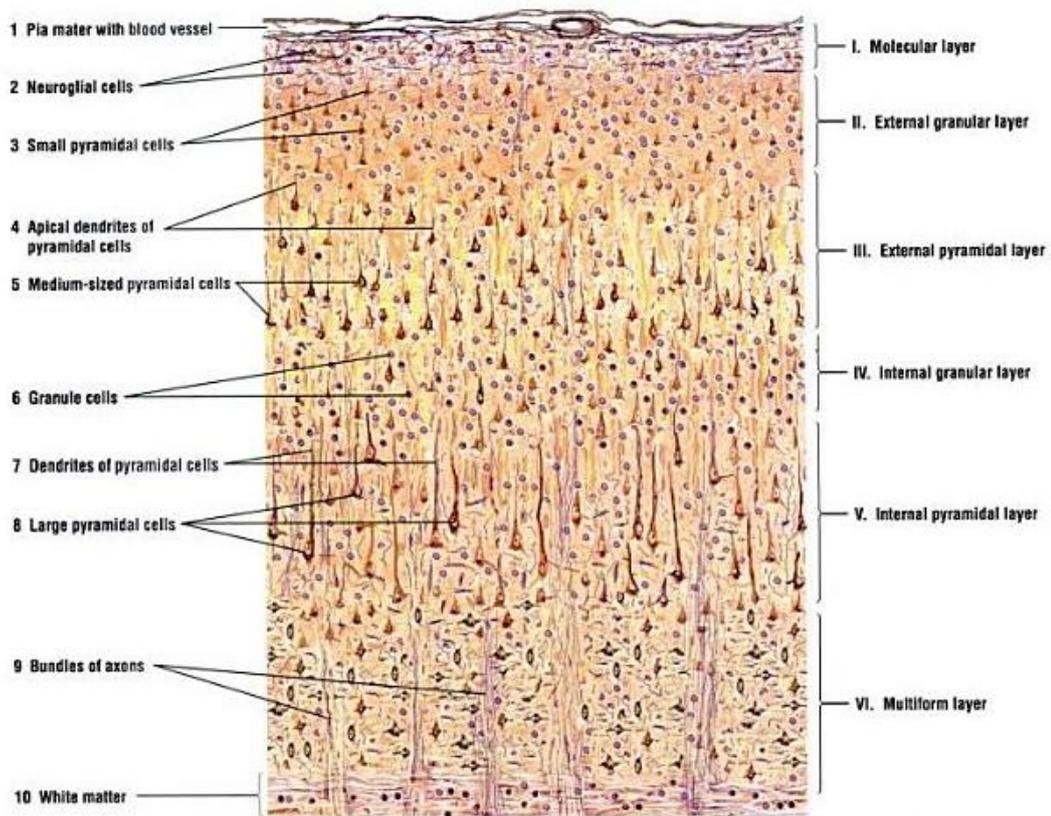


Figure 7: Horizontal structure of the cerebral cortex. The figure illustrates the different cell types in each of the six layers visualized using silver staining (Source: <http://www.ncbi.nlm.nih.gov/pubmed/18270515>).

In the previous section, we discussed the physiological basis of the MEG signal, and noted that it takes the synchronous, simultaneous activity of large populations of neurons to produce a signal that is measurable outside the head. In addition, we noted that such a signal most likely arises from cells whose geometry supports a large net current (open-field geometry), as well as being oriented tangentially with respect to the scalp.

The dense populations of pyramidal neurons in layers II/III and layer V satisfy these conditions, and they are generally thought to be the main contributors to the MEG signal (Hamalainen, 1992, Murakami and Okada, 2006, Ronnqvist et al., 2013). However, assuming that such populations of pyramidal neurons are the only contributors to the MEG signal is a rather simplistic view. From the invasive literature, we know that cell types other than pyramidal neurons are able to produce net currents of considerable strength (Murakami and Okada, 2006). In addition, several studies have investigated the sensitivity of MEG to subcortical sources, and have concluded that while source localization accuracy may suffer, it is by no means impossible to detect deep sources with MEG (Papadelis et al., 2009, Attal and Schwartz, 2013, Balderston et al., 2013).

Whilst this is very promising and warrants further investigation, the limitations on signal to noise ratio (SNR) and spatial resolution have so far limited our ability to investigate these issues more closely using MEG.

MEG Scanning technique

MEG benefits from being completely non-invasive, since it does not require the direct placement of electrodes onto a subject's scalp. Rather, participants are asked to sit in the scanning system's chair in an upright position with the head

protruding into the dewar. The height of this chair is then adjusted to position the subject's scalp in such a way that it is as close as possible to the sensors. However, often, no further restriction of head movement is applied, meaning a subject's head will remain free to move inside the dewar.

While this makes MEG comfortable for participants, it leads to problems for the experimenter. The first of these problems is that of within-session head movement. This type of movement leads to blurring of the sensor level data, decreasing the SNR. While participants are usually instructed to remain still while scanning is in progress, some head movement will typically still occur. Recent work by Stolk et al (2013) showed that head movement tends to increase as scanning time increases. For this reason, a single continuous session of MEG scanning should last no longer than 10 min.

In addition to instructing subjects to remain still and using visual cues or a reference system to aid in this process, most MEG systems are able to provide continuous head localization during scanning. There are a number of offline software modelling tools which allow for this information to be incorporated at the source reconstruction stage, in order to compensate for head movement related errors. Broadly speaking, such techniques are either based on interpolating the field topologies at the sensor level resulting from the 'true' head position, or compensating for head movement errors by using the lead field calculations that form part of the source reconstruction process and provide a model of the sensor-level topologies caused by the magnetic fields of neural sources, given certain properties such as the conductivity profile of the head (Numminen et al., 1995, Knosche, 2002, Taulu and Simola, 2006, Medvedovsky et al., 2007). However, these methods are not widely used, owing to their

complexity on the one hand, and their limitations on the other - in order to incorporate such head movement compensation schemes, a number of assumptions (e.g. about the amount of regularisation) must be made. Hence, the lack of head restriction makes it difficult to keep track of head movements on the one hand, and to incorporate this information at the analysis stage on the other. However, even given the availability of such data, their use is limited if we do not have precise information about the initial head position.

Thus, the second head movement related problem we are faced with in MEG is that of 'between session' head movement. If we cannot obtain precise information about the initial head location, this makes it impossible to reposition subjects between sessions. The ability to do this would allow experimenters to accumulate data from several sessions of scanning, resulting in higher SNR.

Repeatable head repositioning would allow high SNR recordings; however for source reconstruction we need to know how underlying cortical anatomy relates to the sensor positions. This process is referred to as coregistration. As the name suggests, the goal is to relate anatomical and functional data, which are obtained in two different reference systems: while anatomical information is based on structural MRI scans, the functional data are acquired with MEG. The most common coregistration strategy relies on identifying and marking a small number of fiducial points in a way that makes them visible in both reference systems. Usually, the nasion (between the eyebrows) and the left/right pre-auricular points (in front of the left/right ear) are used for this purpose. These points are identified by the experimenter, and markers are placed at his/her discretion - in MEG, so called fiducial coils are used (Gross et al., 2013). This results in considerable variations between subjects, experimenters, and

sessions, and is liable to introduce error at the coregistration stage. The scheme is also hampered by the fact that only three points are available for matching, which degrades the performance of tracking algorithms. However, since the head consists largely of smooth, round surfaces, it is difficult to increase the number of fiducial points.

Several efforts have been made in the past to improve coregistration methods, including the use of bite bars (Adjamian et al., 2004) and sophisticated surface matching techniques (Whalen et al., 2008). However, all of these methods give rise to coregistration errors of 4-5 mm, at best.

This coregistration error ultimately limits the precision with which we can estimate source locations at the source reconstruction stage. For instance, Hillebrand and Barnes showed that even if highly sophisticated cortical models are used in the source reconstruction process, this offers no benefit for coregistration errors of 2 mm or more (Hillebrand and Barnes, 2011). Similarly, Henson et al (2009) compared source reconstruction outcomes based on subject-specific and template cortical models. They found no evidence suggesting that cortical models derived from individual structural MRIs are superior to those based on a canonical mesh. Although this may seem surprising, one possible interpretation is that if we do not have precise knowledge of how underlying cortical anatomy relates to the sensor positions, sophisticated cortical models are of little benefit. In other words, if we do not know where the cortex is, no real advantage is derived from knowing exactly what it looks like.

As a result, MEG recorded data typically suffer from relatively low SNR and spatial resolution compared to other imaging modalities such as fMRI, even if highly sophisticated cortical models are used at the source reconstruction stage.

Source localization error in MEG/EEG

While both MEG and EEG are able to deliver excellent temporal resolution, both of these imaging modalities are thought to provide only limited spatial resolution, albeit for different reasons.

While MEG measures the changes in magnetic fields which are induced by the underlying neural activity and is therefore less sensitive to errors introduced by simplifying assumptions regarding the conductivity of the skull, scalp and brain, the same is not true for EEG. For instance, Cuffin et al (Cuffin, 1993) studied the effects of local variations in skull and scalp thickness on localization accuracy using a computational model. They found that such anatomical variations led to source localization errors of around 1 cm in EEG, whilst the effect on MEG recorded data was much smaller.

Leahy et al (Leahy et al., 1998) studied dipole localization accuracy in both EEG and MEG using a human skull phantom. They concluded that average localization errors over 32 dipoles for EEG were 7-8 mm, while they were around the 3 mm mark for MEG. The authors speculated that this difference was mainly due to forward modelling errors, which affect EEG to a greater extent than MEG.

Liu et al (Liu et al., 2002) used Monte Carlo simulations to remove the impact of forward modelling errors in order to compare source localization accuracy in EEG, MEG and for combined EEG/MEG data sets. The authors used crosstalk and point spread metrics to quantify either the amount of activity incorrectly localized onto a specified location from other locations, or the mis-localization of activity from a specified location to other locations in the brain, respectively. Their findings indicate that for the same number of sensors averaged over many source locations and orientations, source localization in EEG outperformed MEG, while combining both modalities yielded the best results in terms of localization accuracy for the same total number of sensors. The authors also noted that the use of a priori fMRI constraints reduces both crosstalk and point spread.

Source Analysis

At the source reconstruction stage, we want to estimate the number and location of neural sources which generate a given data set. However, before we can do so, it is necessary to develop a model that links the neural activity with the signals measured at the sensor level. This is referred to as the forward model.

Forward model

The forward model takes into account two main components: primary/microscopic currents due to neural sources, and secondary/macroscopic currents, which are determined by the conductivity profile of the head. The forward model gives rise to a gain matrix or a set of lead fields - which represent the field distribution caused by a current dipole of fixed

location and orientation across all sensors (Brette and Goodman, 2012, Destexhe and Bedard, 2012).

In MEG, the primary currents are the main generators of the externally measured magnetic fields, whereas the EEG signal is due to the scalp potential difference caused by the secondary currents.

The secondary currents depend on the conductivity profile of the head, which consists of different tissue types that can be modelled as 'shells' (Dannhauer et al., 2011). Typically, MEG head models take into account up to three shells, representing the skull, scalp, and the inner skull boundary. Several head models are available, which are based on different assumptions about head geometry and the conductivity of the different tissue types. As such, the available models can be split into spherical and realistically shaped models.

Spherical Models

These are the simplest head models available. As the name suggests, the spherical approach involves assuming that the head is spherically symmetric. The most basic spherical model uses the 'global sphere' approach. Here, a single sphere is fitted to the overall skull geometry. This produces a model that fits the geometry reasonably well around the central sulcus, but with poor fit elsewhere. In the locally optimal sphere approach, a sphere is fitted to the local skull geometry around a particular region of interest, and only sensors around this region are used. Finally, in the local spheres approach, separate spheres are fitted for each sensor, each following the local skull geometry. Using this approach results in improved fit as compared to the global sphere approach,

whilst allowing experimenters to take into account more than one region of interest.

Global or locally optimal sphere approaches are often referred to as 'single sphere' and multi-sphere models, respectively, and have been used in MEG source modelling for some time (Gallen et al., 1995). Although much more sophisticated, realistically shaped models are available now, this approach is still popular.

Realistically shaped models

Here, the conductivity profile is based on extracting conductivity boundaries from anatomical MRI scans. Again, there are several different models to choose from, depending on which tissue types are taken into account.

The simplest of the realistically shaped models is based on modelling only the shape of the inner skull boundary, while the volume inside the skull is assumed to be homogeneous. The skull and scalp are omitted (Hamalainen and Sarvas, 1987, 1989). This is motivated by the assumption that the conductivity of the skull and scalp is much lower than that of the brain. Thus, the effect of weak currents within these tissue types on the MEG signal is assumed to be negligible (Akhtari et al., 2002).

The most sophisticated model currently available as part of popular analysis toolboxes is the three shell model. It involves modelling the inner skull boundary, the skull, and the scalp, and uses the Boundary Element Method (BEM) to carry out numerical computations.

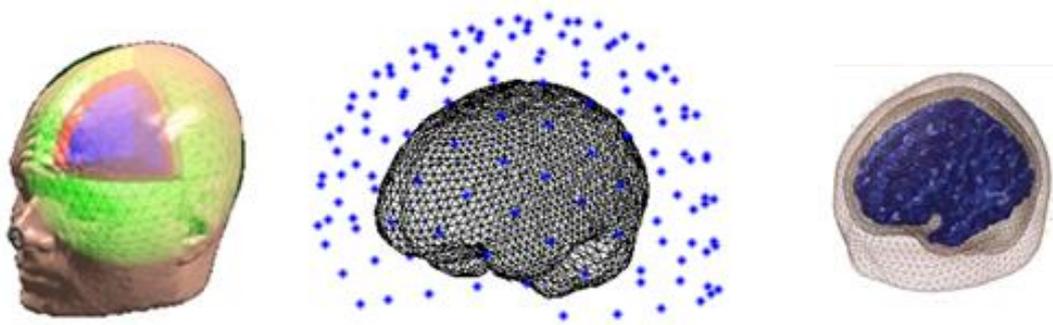


Figure 8: *Different types of head models. The leftmost figure shows the simplest type of head model, which involves approximating the head as a sphere. The panel in the middle shows a single shell head model, in which just the inner skull boundary is modelled. Finally, the rightmost figure illustrates a three shell model, consisting of shells representing the skull, scalp, and inner skull boundary.*

While realistically shaped models fit the skull geometry considerably better than their spherical counterparts, achieving a precise fit is still one of the biggest challenges in this case. The reason for this is the segmentation of the skull, which is based on T1 weighted MR-images. These images have excellent contrast for the brain, but rather poor contrast for the skull. As a result, segmentation tools struggle to produce consistent, accurate results when extracting the skull boundary, potentially affecting the fit of the resulting model (Dogdas et al., 2005, Lanfer et al., 2012, Perdue and Diamond, 2014).

Although there is a popular consensus that for the purpose of MEG analysis, spherical or single shell models are sufficient, this view has recently been contested in a number of papers. Stenroos et al (2014) carried out a thorough comparison of different forward models, including both spherical and realistically shaped variants. They concluded that the 3 shell BEM models far outperformed

spherical and single shell based approaches, whilst not significantly increasing computational cost.

Vorwerk et al. (Vorwerk et al., 2014) went a step further and developed a forward model which features six shells: skull compacta, skull spongiosa, cerebro-spinal fluid, grey matter, white matter, and white matter anisotropy. Investigating changes in MEG signal topography when including all or a subset of these shells, they found that while factors such as the distinction between skull compacta and spongiosa or white matter anisotropy had very little effect on the resulting signal topography, the inclusion of cerebro-spinal fluid as well as the gray/white matter distinction produced outcomes that differed significantly in terms of signal topography and magnitude.

However, whether a significant benefit may be gained from more sophisticated forward models ultimately depends on the quality of the available data (Henson et al., 2007, Mattout et al., 2007, Hillebrand and Barnes, 2011).

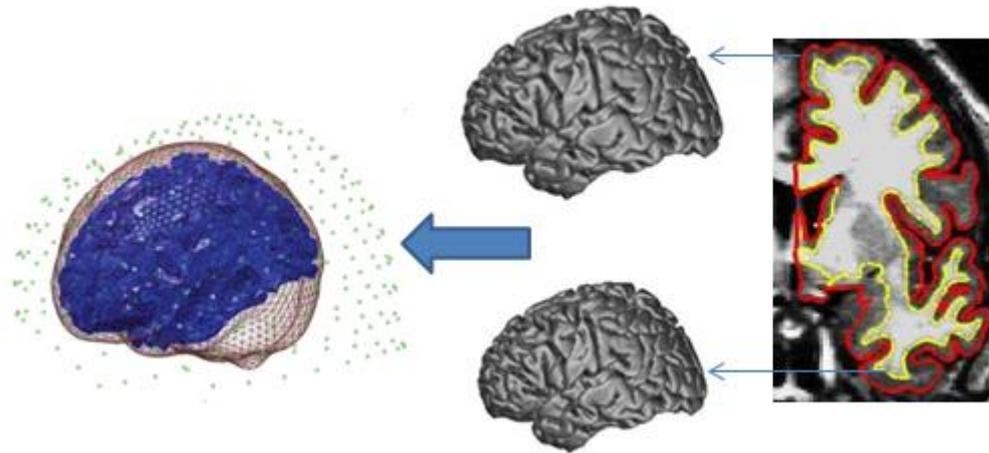


Figure 9: MEG Forward model. The rightmost panel illustrates the location of the pial and white matter surfaces in a structural MRI. Two cortical models are then derived from these surfaces, which form the basis of the forward model. The leftmost panel illustrates the cortical mesh (blue), the inner skull boundary (grey mesh), and the sensor positions (green dots). Note that while this model uses a single shell, further shells representing other tissue types can be included.

Source reconstruction

The previous section dealt with the selection of suitable forward models, which is the first step in the source reconstruction process. While the forward problem consists of establishing the relationship between active sources and the signal they produce at the sensor level, we are faced with an inverse problem at the source reconstruction stage: we need to estimate the locations of underlying sources based on this relationship and the recorded data.

As discussed at the start of this chapter, raw MEG data represent the activity of a discrete number of active sources in the brain. Each of these sources is formed by the synchronous, simultaneous activity of a large group of neurons (~10⁴), which can be modelled as an equivalent current dipole (ECD) (Lopez et al., 2013).

There are two basic methods of estimating underlying source locations. The first is to assume that the observed activity is caused by a small number of sources, which are fitted to the data using a non-linear search through the brain. This is referred to as the dipole fitting approach and, while being a powerful method, it suffers from the emergence of local extrema as the number of sources increases (Supek and Aine, 1993, Wendel et al., 2009).

In the second approach, the entire search space – the brain volume or cortical surface - is filled with a large number of dipoles, whose amplitudes are then estimated such that their combined activity best explains the recorded data. This is known as the distributed approach to source reconstruction, and benefits from the fact that the model is linear with respect to the neuronal currents. The main problem here is that the number of potential sources far outstrips the

number of MEG sensors, which means that there is no unique solution, since there are typically many source configurations which could explain a given data set. Hence, the MEG inverse problem is ill-posed, and we need to introduce additional constraints in order to obtain a unique solution.

Because of the linear mapping between sources and sensors and its independence of the number and characteristics of the activated regions, much effort has gone into developing strategies based on the distributed approach to source reconstruction which are more computationally feasible and robust. As a result, common MEG analysis toolboxes offer a variety of algorithms, which employ different assumptions about the search space and the nature of the underlying source activity.

One way of tackling source reconstruction using the distributed approach is by applying a Bayesian framework. In essence, source reconstruction means estimating source locations based on a given dataset and a model that links source activity to the signal detected at the sensors. In this context, applying a Bayesian framework basically means assuming that the source activity is a zero mean Gaussian process. In essence, within this framework, the source activity (J) is calculated based on the source co-variance matrix (Q), the sensor noise co-variance matrix (Q_ϵ), the recorded data set (Y) and the lead fields (L) which link the recorded activity to the underlying sources. Applying this assumption, the source reconstruction problem is reduced to finding a good estimate of the sensor- and source-level co-variance matrices, since Y and L are known (Mattout et al., 2006, Wipf and Nagarajan, 2009):

$$\hat{J} = QL^T(Q_\epsilon + LQL^T)^{-1}Y \quad (1.1)$$

The sensor noise co-variance matrix (QE) basically contains information about the noise affecting each sensor. In the absence of any information about sensor-specific noise variance, it is often assumed to be uniform across all sensors.

There are different approaches to estimating the source co-variance matrix, which are based on different assumptions about the underlying activity. These assumptions form the basis of many currently available source reconstruction algorithms. For instance, the simplest approach would be to assume that all sources are equally likely to have generated the data, and that their activity is not correlated (Hamalainen and Ilmoniemi, 1994). This is known as the Minimum Norm approach. To take into account the possibility of correlated sources, one could assume that sources vary smoothly over space. This assumption forms the basis of the LORETA algorithm.

However, the assumption that all dipoles are active at the same time can result in artefacts being interpreted as source activity – that is, the algorithm will try to explain noise in the source space. One way to avoid this is by using the Beamformer algorithm, the main assumption being that no two distant cortical areas will generate coherent local field potentials given long timescales (Robinson and Vrba, 2004).

Multiple Sparse Priors (MSP)

While most ‘classical’ approaches are based on a single co-variance matrix, the Multiple Sparse Priors (MSP) algorithm assumes that the prior source co-variance is based on several components. Each of these components consists

of a co-variance matrix (\mathbf{Q}) and a weight – a so-called hyperparameter (h), a relationship which is expressed in equation (1.1). If each component represents a potentially activated region of the cortex (\mathbf{C}), then the corresponding hyperparameter will carry information about the likelihood of this region having been involved in generating the observed data.

$$\mathbf{Q} = \sum_{i=1}^{N_q} h_i \mathbf{C}_i \quad (1.2)$$

An important advantage in this context is the fact that the use of such components, each of which has its own covariance matrix, allows the researcher to incorporate different assumptions about the underlying activity. In this way, the MSP framework allows for generalisation of other, single-covariance-matrix-based approaches.

The number of components used depends on anatomical and physiological assumptions. In theory, the most inclusive model would mean having a component for each dipole – around 8,000. However, given that there are certain general beliefs about information flow and processing within the brain and general neuroanatomical considerations, this is an over-inclusive approach. Typically, the number of components used is of the same order as the number of sensors – several hundred rather than in the thousands.

Since it is generally assumed that the contribution of sub-cortical structures to the MEG signal is negligible due to their neural anatomy and depth, the components are defined as regions or ‘patches’ on the cortical surface. The size and number of patches can be shaped according to prior information about source activity and physiology. Most currently used implementations are based

on a fixed number of patches. In the SPM software package, 512 patches are used, which cover the entire cortical surface and are based on a subset of the overall number of dipoles (8196). The size of these patches can be modified if there is prior information about the location or the size of the activated region (Lopez et al., 2012b).

Once the components or patches and their corresponding co-variance matrices have been selected, the next step is to determine which of the components are likely to have caused the recorded dataset. As mentioned above, this is expressed through the hyperparameters, which provide the weighting for the individual components.

This process aims to arrive at the optimal source model by assessing whether a component (and its associated hyperparameter) satisfies a specific cost or objective function. In current implementations of MSP (i.e. in SPM), Free Energy is used as the objective function. This value represents the evidence for a given model by assessing the accuracy of its fit to the data, while penalizing complexity.

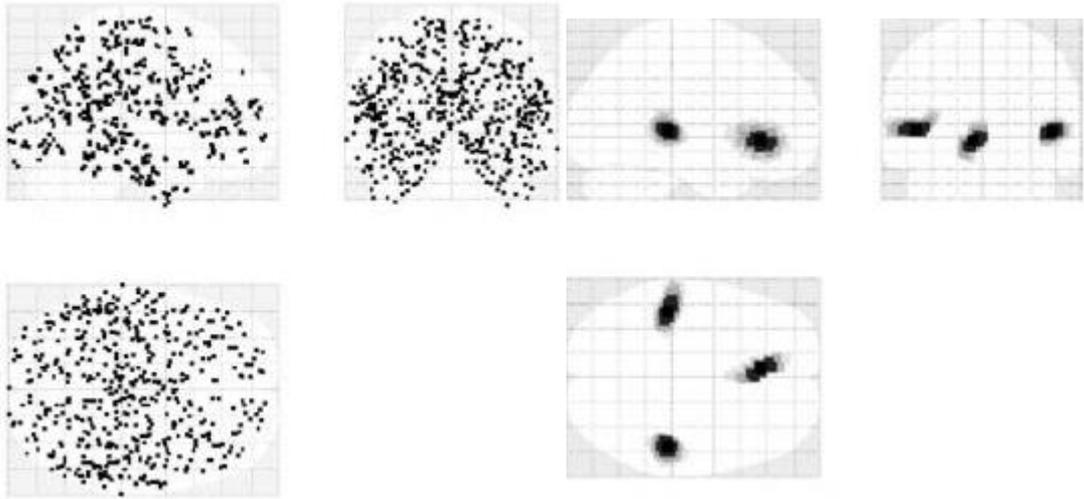


Figure 10: The left hand side illustrates the use of patches in MSP. The black dots correspond to MSP patch centre locations. On the right hand side, we show an example of source reconstruction using MSP. Following the search and optimisation process, only a subset of patches are 'left on', illustrating the estimated source locations.

In essence, this optimisation process consists of testing different combinations of components to find the model which maximizes the Free Energy. Computational optimisation schemes are used for this purpose. In the SPM framework, the two basic strategies used are Automatic Relevance Detection (ARD) and the Greedy Search (GS) algorithm. Both of these schemes have the same objective, but accomplish it in different ways. While ARD arrives at the optimal solution by starting out with all possible components and pruning away those that do not contribute to the maximization of the Free Energy value, the Greedy Search splits and prunes different mixtures of components, selecting the most active ones at each step.

Using both schemes makes the resulting solutions more robust, as it takes into account the possibility of either of the search strategies getting stuck in local extrema.

Although MSP is one of the most sophisticated algorithms available and allows for taking into account the possibility of correlated sources as well as prior information, there are some limitations, one of these being high sensitivity to coregistration noise (Lopez et al., 2012a, Lopez et al., 2012b). The use of a fixed number of patches (which is smaller than the number of dipoles) in the absence of prior information about source locations implies a risk of omitting the 'true' patch from the search. However, this can be avoided by testing different patch combinations and combining the results using Bayesian Model Averaging (Lopez et al., 2013). In addition, this process can be carried out iteratively in such a way that on any given iteration, the results of the previous one will be taken into account. Another limitation inherent in the use of patches for source reconstruction is their size (or smoothness along the cortical surface). There is

currently no consensus on optimal patch size, and little is known about how this parameter affects source reconstruction outcomes.

Conclusions

In summary, MEG benefits from excellent temporal resolution, whilst being completely non-invasive and comfortable for subjects. However, the main limitation of this neuroimaging modality lies in the comparatively low SNR and spatial resolution of the resulting data. These limitations are largely caused by error introduced at the coregistration stage, as well as through within-session head movement. From this it follows that if we can develop techniques to reduce these sources of error, this has the potential to dramatically improve spatial resolution.

2. High Precision Anatomy for MEG

Abstract

Precise MEG estimates of neuronal current flow are undermined by uncertain knowledge of the head location with respect to the MEG sensors. This is either due to head movements within the scanning session or systematic errors in coregistration to anatomy. Here we show how such errors can be minimized using subject-specific head-casts produced using 3D printing technology. The casts fit the scalp of the subject internally and the inside of the MEG dewar externally, reducing within session and between session head movements. Systematic errors in matching to MRI coordinate system are also reduced through the use of MRI-visible fiducial markers placed on the same cast. Bootstrap estimates of absolute coregistration error were of the order of 1 mm. Estimates of relative coregistration error were < 1.5 mm between sessions. We corroborated these scalp based estimates by looking at the MEG data recorded over a 6 month period. We found that the between session sensor variability of the subject's evoked response was of the order of the within session noise, showing no appreciable noise due to between-session movement. Simulations suggest that the between-session sensor level amplitude SNR improved by a factor of 5 over conventional strategies. We show that at this level of coregistration accuracy there is strong evidence for anatomical models based on the individual rather than canonical anatomy; but that this advantage disappears for errors of greater than 5 mm. This work paves the way for source reconstruction methods which can exploit very high SNR signals and accurate

anatomical models; and also significantly increases the sensitivity of longitudinal studies with MEG.

Introduction

The first chapter discussed the many advantages Magnetoencephalography (MEG) offers compared to other scanning techniques. One of these advantages is the relatively high speed of measurement and subject comfort. The negative side of this pleasant scanning experience is that a large amount of coregistration error (uncertainty in head position) is introduced. This affects the forward modelling stage — the step used to relate current flow on the cortical surface to the signal measured by the sensors. Although this error is typically of the order of 5–10 mm it has a significant and detrimental effect on the estimate of electrical activity (Hillebrand and Barnes, 2011, Lopez et al., 2012b, Lopez et al., 2013). For example, investigating the effect of source extent estimation in MEG using beamformers, (Hillebrand and Barnes, 2011) concluded that when the cortical surface is known accurately, cortical surface models give accurate predictions of spatial extent of activation. However, if coregistration error exceeds 2 mm, this is not the case.

Typically, coregistration is performed using just 3 fiducial coils, which are placed on anatomical landmarks by the researcher (Gross et al., 2013). This introduces error, because the fiducial coil placement depends on the individual researcher's interpretation of these landmarks in each subject (both in the anatomical scan and in the flesh). One of the main drawbacks of fiducial-based coregistration is the small number of points available for matching. However, since the head does not offer a lot of sufficiently distinct landmarks, it is difficult

to increase this number. There have been several efforts in the past to improve coregistration. These include bite bars (Adjamian et al., 2004) with a small number of unequivocally defined points in both coordinate frames, or surface matching techniques, to match the many points on the subject's scalp surface (as measured by a digitisation device) to that of the structural MRI (Whalen et al., 2008). Typically however both approaches give rise to similar errors of the order of 4–5 mm for different reasons— the bite-bar because all of the fiducial points are at the front of the head (and small errors are magnified); the surface matching because the head is round and smooth giving the optimisation a rather shallow cost function and it in turn relies on keeping the subject's head still for the digitisation process. Furthermore, the accuracy of the surface point measuring device used further limits the resolution of such techniques (for example although systems like Polhemus are very accurate this can be undermined by small movements in the reference sensor which often needs to be fixed to the head).

The other source of error, which affects the data quality as well as the modelling is within session head movement (again of the order of 5–10 mm). This can be corrected for in software (de Munck et al., 2001, Uutela et al., 2001, Wilson, 2004, Taulu and Simola, 2006), but entails certain assumptions and an inevitable dimension reduction.

Uutela et al (Uutela et al., 2001) presented a method for measuring the head position during data acquisition, and then making use of this information to remove the effects of head movement on neuromagnetic source estimation and magnetic field alignment. By using minimum-norm estimate alignment and

forward calculation correction, the authors showed that the proposed method works with millimetre precision, whilst also being computationally efficient.

De Munck et al (de Munck et al., 2001) presented a scheme which involves the use of an algorithm to localize a set of simultaneously activated coils using MEG detectors. This information can then be used to monitor head movements and to reduce coregistration error by eliminating the need for using a 3D digitizer and mapping MEG directly onto other modalities. Using a computational model, the authors concluded that the localization accuracy was of the order of several millimetres.

If one could improve between session repositioning of the subject then it would also be possible to build up very high signal to noise ratio sensor level data by scanning the same individual on a number of occasions. Recently, Stolk et al. (Stolk et al., 2013) introduced a method using continuous head localization during scanning, with an option of incorporating this information in the offline data analysis. They exploited this head localization method to reposition subjects at onset of scanning between sessions, reducing intersession head distances up to 2 mm. Interestingly, they found a significant source of error in the form of a slow but steady drift away from the initial head position during the recording sessions.

It is often head movement considerations that constrain the amount of time devoted to an MEG scanning session. Ultimately, these sources of error, compounded with the uncertainties in the inverse problem lead to the rhetoric that MEG has poor spatial resolution.

In this chapter, we introduce a new technique based on subject-specific head casts made using a combination of optical scans/MRI scans and 3D printing to constrain head movement. They limit between and within session movement and provide a fiducial reference frame in the form of MRI visible markers. The paper begins with a description of the construction of the head-cast and the procedures required to co-register the data. We go on to make empirical estimates of reproducibility of head position. We then make an estimate of brain movement based on repeatability of evoked response measures over several sessions. Finally, we show that the anatomical information is accurate to such a degree that we are able to distinguish between cortical surfaces using the MEG data alone.

Methods and Materials

Head cast

Figure 11 shows the major components in the process of construction of the head cast. We used a manufacturer supplied dewar-helmet (A) with the same internal dimensions as the MEG dewar (typically used to check that a subject would fit within the MEG scanner). A manufacturer supplied surface image of our MEG dewar (B) was used in order to link the helmet (A) to the four reference points on the exterior of the MEG system (yellow arrows). The subject-specific head-cast was designed to fill the space between the subject's scalp (C) and the inside of the dewar-helmet (A). During the anatomical MRI scan, the subject wore both the head-cast and the dewar-helmet. The outer surface of the dewar helmet contained 12 MR visible markers that were also optically digitized (E).

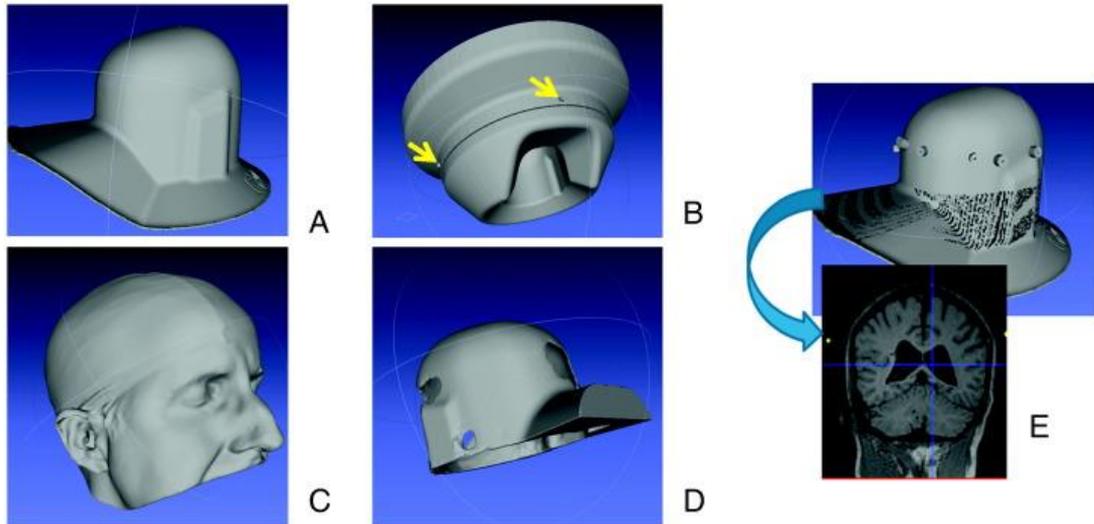


Figure 11: Head cast/coregistration. Panel A shows the optical scan of the dewar helmet, which, in combination with an optical scan of the subject's head (Panel C), serves as the basis for creating the template of the head cast (Panel D). For coregistration purposes, MRI visible markers are placed on the outside of the dewar helmet, such that their position in an MRI structural image can be related to its counterpart in an optical scan (Panel E). To obtain a transform which allows us to relate MEG sensor space and MRI space, we use the 4 measurement coil positions in relation to the MEG system dewar in a manufacturer supplied image (Panel B).

MRI scanning

MRI data was acquired using a Siemens Tim Trio 3 T system (Erlangen, Germany). The subject lay in the supine position wearing the head-cast fit into the dewar-helmet. Two types of cylindrical markers with different diameters (4 mm/8 mm; depth 10 mm, with wall thickness 1 mm of these measurements for side, top and bottom) were made using the 3D printer and filled with a copper sulphate water solution (concentration: 1 g per litre). The solution was mixed with agar to give it a gel-like consistency, minimizing leaks and making the filling of the cylinders easier. The cylinders were glued to the outside of the dewar-helmet in an approximately circular arrangement (Figure 11E). The body-transmit coil was located inside the bore of the scanner for detection of the MRI signal. The MRI data was acquired using a 3D FLASH sequence for optimal scanning efficiency (Frahm et al., 1986). The following acquisition parameters were used: field-of view: (256,256,208) mm along the (phase (A–P), read (H–F), partition (R–L)) directions, image resolution: 1 mm³. The repetition time TR was set to 23.7 ms and the excitation flip angle was set to 20° to yield good T1-weighted contrast, standard in most anatomical applications (Helms et al., 2008). 8 echoes were acquired following each excitation and averaged offline to produce one anatomical image with optimal signal-to-noise. Potential sources of error from the MR acquisition should be carefully considered. A high readout bandwidth was used to preserve brain morphology and no significant geometric distortions are expected in the images. Padding was used to minimize subject motion but some residual effects might remain present in the MRI images. These effects might be further reduced by use of navigator echo techniques. The total acquisition time was 21 min 07 s.

Optical scanning

We obtained an optical scan of the subject's head/face anatomy (Figure 11C), using an Artec 3D handheld optical scanner (http://www.artec3d.com/3d_scanners/artec-mht). This type of scanner has a resolution of up to 0.5 mm, and captures images at a rate of 15 frames per second. It has a 3D point accuracy of up to 0.1 mm. The handheld scanner is moved around the subject's scalp at a constant, moderate speed, whilst the subject's head is kept still using a bite-bar (Adjajian et al., 2004). It should be noted that the scanning device we used generally fails when attempting to capture images of dark, shiny objects. Also, some areas of the head and face are difficult to capture, such as ears and hair. For these reasons the subjects wore swimming caps to cover their hair. Since the ears were not particularly important for our proposed design (and also difficult to scan), we only ensured that enough space was left to accommodate them in the finished head cast. Note that swimming caps were only used at this stage, purely to avoid problems with the optical scanner, these caps had a thickness of 0.27 mm and therefore served mainly to constrain the hair without adding appreciable thickness to the scalp estimate.

In order to fit this image to the MEG system dewar surface, a similar scan was obtained of the inside of the MEG dewar-helmet. Following this, we also scanned the outside of the dewar-helmet with MR visible markers attached (panel E). We used the Magics (<http://software.materialise.com/magics>) CAD package to process the images.

Printing

The images of the internal dewar surface (Figure 11A) and the scalp (Figure 11C) were aligned by eye to place the subject's head in a typical recording position. These two surfaces were converted into solids and the Boolean difference between the two structures was used to make the head cast (Figure 11D). However, a few practical issues had to be considered: to ensure as tight a fit as possible, the head cast had to have some 'anchor' points, consisting of distinct anatomical marks — such as the inion bone or the bridge of the nose. We therefore decided to let the cast come down past the eyes, covering the bridge of the nose. In order to allow the subject to put on and take off the helmet, the image was split vertically down the centre, incorporating a simple locking mechanism. This image was passed on to the 3D printer. For our prototype, we used an ZPrinter 350, which has a resolution of 350 × 400 dpi (dots per inch).

Coregistration

The following process consists of 3 distinct registration stages. We wish to estimate the transformation from MRI to dewar-helmet space; dewar-helmet space to MEG dewar space and finally MEG dewar space to MEG measurement space. Given all these stages we can then produce a single rigid body transformation between MRI anatomy and MEG sensor space.

In each case we used the ICP (Iterative Closest Point) algorithm to match point clouds (<http://www.mathworks.co.uk/matlabcentral/fileexchange/24301-finite-iterative-closest-point>).

ICP algorithms are used to minimize the distance between two point clouds. While the reference point cloud is kept fixed, the source cloud is transformed in order to achieve the best fit. In addition to the two point clouds, the algorithm can be supplied with an initial 'guess', or starting point, for the transformation, as well as specific criteria for stopping iterations. For each iteration, the closest reference cloud point for each point in the source cloud is identified. The algorithm then estimates the combination of rotation and translation to best align these points using a means squared error cost function. This transformation is then applied, and the process is repeated until some specified stopping criteria are met.

Native MRI to dewar-helmet

The anatomical MRI scan gives an image showing the position of the cortex relative to the registration markers (MRI space) (Figure 11E). These registration markers were extracted from the MRI scan through a simple thresholding operation. A similar process was performed to extract the locations of the registration markers in the optical image. This resulted in two point clouds which were aligned using the ICP algorithm outlined above.

Dewar-helmet to MEG dewar

The next step is to transform from the (optical) dewar-helmet space to the MEG dewar space (the external surface of the MEG measurement system). Again, we need to find a set of appropriate reference points. Since the dewar-helmet is an exact replica of the inside of the MEG dewar, we can use the subset of points of the manufacturer supplied data (Figure 11 B) that describes the same shape. We then follow a similar procedure as outlined above, bringing the

images into approximate alignment, then following this up with ICP alignment (Figure 11B).

MEG dewar to MEG space

In the last stage, we need to acquire a transformation matrix between MEG dewar space and MEG sensor space. The relationship between these two spaces will depend on the locations of the sensors within the measurement dewar. In order to do this we matched the four coil calibration points based on the MEG dewar specification (arrows in Figure 11B) to those measured using the MEG sensors to locate fiducial coils attached to these reference locations (we also accounted for coil thickness in these estimates). We made 5 measurements of each coil position giving rise to four point clouds to fit to four single points.

Bootstrapping for accuracy

In order to make an estimate of the robustness of these various fitting procedures we bootstrapped (using 100 bootstrap iterations) each fitting stage independently and then all fitting stages at once. Each of these 100 iterations gave rise to a different transformation matrix that we used to estimate a new fiducial location (on the subject's scalp). This distribution of points gave a measure of the expected variance in fiducial location due to measurement and fitting error. In a first step, bootstrapping was applied to a single transformation stage in turn, whilst all others were held constant. Next, bootstrapping was applied to all transforms simultaneously, to get an estimate of the likely absolute coregistration error.

Testing for reproducibility

In order to test for within and between run reproducibility in head position and head-cast position we attached one fiducial coil to the cast and one firmly to the subject's head (just above the right ear). Over a ten minute period, we made eight recordings of the coil positions whilst the subject was seated within the dewar. Over the subsequent forty minute period we then made eight further measurements, of these coil positions but removing the subject from the dewar and the head-cast in between runs.

Empirical validation

We have now performed the same experiment on subject 1 nine times over a six month period. We wanted to use this functional data, based on the location of the cortex, to corroborate our estimates of coregistration error based on the scalp surface.

Subject task

We used a single, male, right-handed subject for this study, who was instructed to perform a simple finger movement task adapted from Muthukumaraswamy (Muthukumaraswamy, 2011), after we had obtained approval from UCL's research ethics committee (application number 3419/001). The task involved performing simple abductions of the right hand index finger performed to an auditory cue. The cue consisted of a simple auditory tone (1000 Hz), played via a piezo electric device connected via plastic tubing to ear-inserts, followed by an inter-stimulus interval of 3.4–4.5 s. This gave approximately 130 epochs of data per ten minute recording session. EMG traces of the first dorsal interosseus (FDI) were used to track finger movements (although we did not make use of this information directly for the purpose of this paper). Each

session of scanning was split into 4 sections of 10 min each. During the first ten minutes, the subject was told to perform simple abductions of the right hand index finger on hearing the pip. Following these 10 min of task performance was a rest period of equal length, during which the subject was simply told to remain as relaxed as possible. This was followed by another 10 min of task performance and 10 min of rest (the rest data were not used in this study). We repeated this same recording five times over a six month period. In the first recording run we had a triggering problem and only one of the two task runs was recorded giving a total of nine task runs.

Source reconstruction

We used averaged evoked responses from 0 to 300 ms (0–80 Hz) time-locked to the auditory cue, baseline corrected based on –200 to 0 ms. The only artefact removal was the removal of trials containing large jumps (due to loss of lock in the feedback electronics) which could be clearly seen by eye (as the subject was unable to open their eyes we had no need to correct for eyeblinks). These data were projected into 100 (for the majority of this work) or 250 (for Figure 12) orthogonal spatial (lead field) modes and 16 temporal modes. We used the greedy search option of the MSP algorithm (Friston et al., 2008) implemented as outlined in (Lopez et al., 2012b). The MSP algorithm requires a set of covariance matrix priors corresponding to cortical patches to be defined a-priori. We used a pseudo-random (the same sequences were used when comparing different surface models) selection of 512 mesh vertices to define the patch centres and produced 32 such solutions (each based on different patch sets) per dataset. There were no symmetric priors used. The MSP algorithm returns a Free energy value which approximates the model evidence

for the final generative model. This generative model includes the cortical surface used and the lead fields and therefore model evidence can be used to select between different models of the geometry (Henson et al., 2009, Lopez et al., 2012a). We reconstructed the same MEG data onto two cortical surfaces: the subject's cortical surface or the canonical mesh (Mattout et al., 2007). The canonical mesh is simply an extracted mesh from a template brain in standard space warped into the subject space. The algorithm returns posterior estimates for the mean and variance of current density at each cortical location. These posterior estimates can be converted into either posterior probability maps (probability that current density is non-zero) or plots showing the mean and 95% confidence intervals ($\pm 1.66^*$ posterior standard deviation).

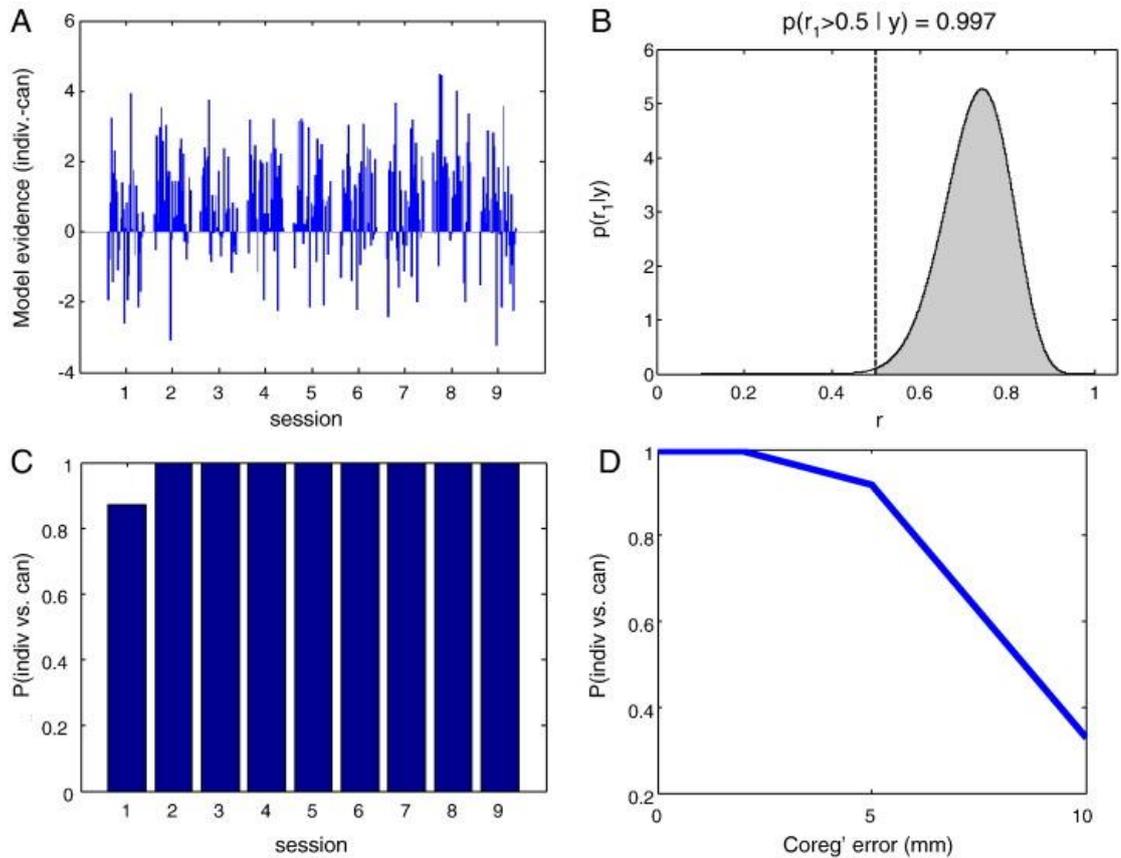


Figure 12: Investigating sensitivity to cortical model used. Panel A shows the comparison between individual and canonical surface models in terms of log model evidence for each of 32 patch(/prior) sets used to fit the averaged evoked response from 9 separate sessions. The model evidence difference is generally positive, indicating that the individual cortical surface model is favoured in the majority of cases. Using a random effects analysis, Panel B shows the frequency with which the individual white surface is favoured over the canonical model within Panel A session 9. Panel C shows the exceedance probability (the area under the curve in Panel B where frequency $r > 0.5$) plotted over all sessions. Again, the evidence is overwhelmingly in favour of the individual cortical surface. Lastly, in Panel D, we investigated the effect of adding additional coregistration error and performing the same analyses. As coregistration error increases, the advantage of using a model based on the individual subject's cortical anatomy diminishes to the point where it is no longer favoured over the canonical model.

Results

Absolute coregistration error

In this first stage we want to get an estimate of the absolute coregistration error due to the fitting process. We bootstrapped the calculation of the different transformation matrices and looked at the variability in the estimates of fiducial locations to get an estimate of absolute coregistration error. This variability was estimated per transform (Figure 13 A-C) and for all transforms simultaneously (Figure 13D). The bars indicate standard deviations for three fiducial locations in the x (blue), y (green) and z (red) direction. We found that the transformation between the MEG dewar space and the MEG sensor space (panel C) potentially represents the biggest source of error — most likely because it was based on the smallest number of unique points. However in all cases, and in particular for the case when all stages were bootstrapped simultaneously (panel D), the expected errors in any one dimension are still within the sub-millimetre range.

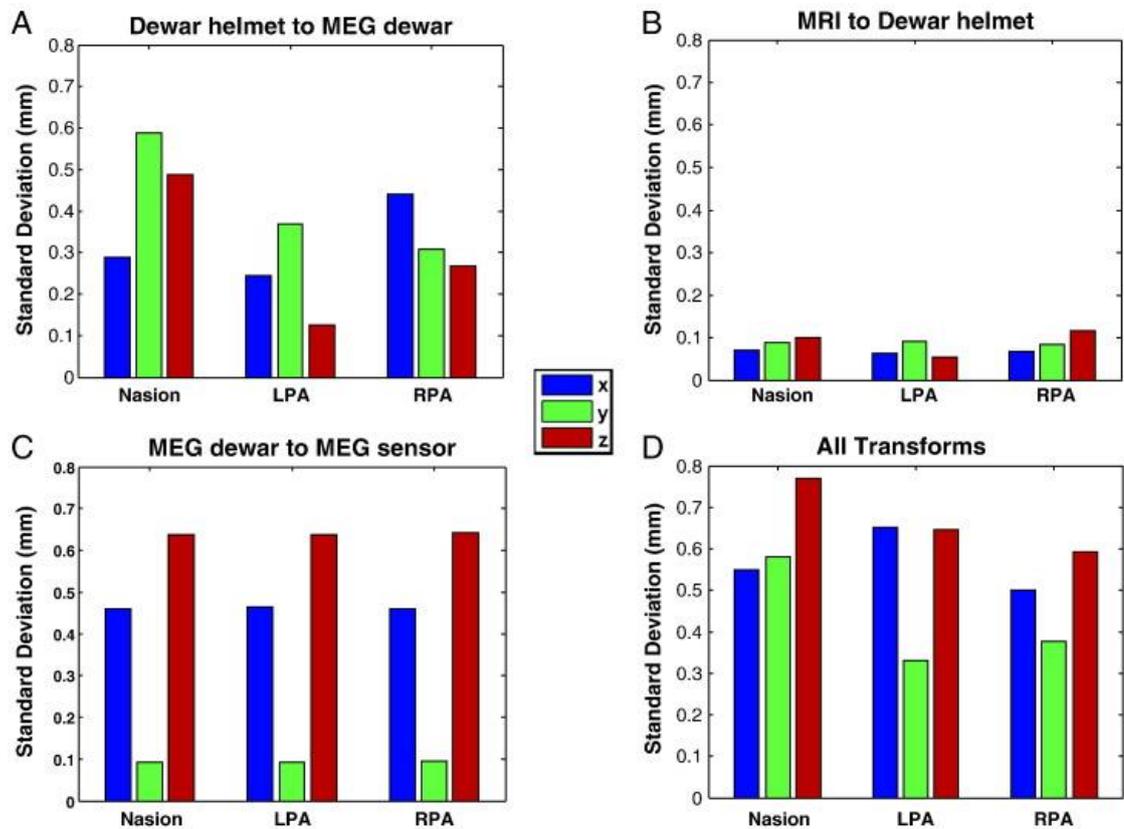


Figure 13: Evaluation of repeatability of coregistration results by means of bootstrapping. The bars show the standard deviation the three fiducial locations in x, y and z directions based on bootstrapped estimates of the transformation matrices. Panel A shows the likely error in fitting the dewar-helmet shape (A) to the MEG dewar (B). Panel B shows the fitting of the MRI fiducial locations to the dewar-helmet fiducials (E). Panel C shows an estimate of the error due to fitting the MEG sensor space (as measured using fiducial coils) to MEG dewar fiducial points (B, yellow arrows). Finally, all transforms were bootstrapped simultaneously (Panel D) to get an idea of the total expected absolute coregistration error— which we would expect to be of the order of 1 mm per fiducial.

Note in Figure 13C (MEG dewar to sensor) there is less error in the y than x and z directions. This transformation is the matching of 4 fiducial markers on the dewar surface with 4 fiducial coil locations (estimated by the MEG system). As we were working with a set of 3 fiducial coils we measured left, right and front for half the runs; and left, right and back for the other half. This meant that we have twice as many points for left and right fiducials than for front and back which most likely gives rise to the increase in precision in the y (left–right) direction. The average distance of an MEG estimated fiducial location to its corresponding point on the manufacturer supplied dewar surface was 0.23 mm.

Relative coregistration error

We used two subjects to test for within and between run reproducibility in head position and head-cast position. In each case, we attached one fiducial coil to the cast (close to the left pre-auricular) and one firmly to the subject's head (just above the right ear). Over a ten minute period, we made eight recordings of the coil positions whilst the subject was seated within the dewar. Over the subsequent forty minute period we then made eight further measurements, of these coil positions but removing the subject from the dewar and the head-cast in between runs.

Figure 14 shows the standard deviation of coil position as measured within (panel A) and between sessions (panel B). The leftmost sets of bars in each figure indicate variability in head cast position with respect to the dewar for both subjects. The head-cast location varies by less than 0.3 mm within runs; showing that the cast fit rigidly within the dewar and gives us an upper bound on the error (due to reproducibility) of the MEG system estimate of the fiducial coil locations. For the coil attached to the subject's head the expected errors are

larger and give an idea of the fit of the cast to the subject's head internally. The first subject shows very little within run movement the expected error was below 0.2 mm predominantly in the vertical (z) direction. For the second subject the errors were larger yet still sub-millimetre (expected (14D) Euclidean displacement of 0.7 mm from the mean). Figure 14B illustrates variability in head and head cast position between sessions. Again the left-most panels shows repositioning errors on the two head-casts inside the dewar; below 0.6 mm but dominant in the vertical (z) direction (i.e. there is little rotation, but some error introduced when the cast is not pushed firmly up inside the dewar). Similarly, for both subjects, the variability of head cast position with respect to the MEG dewar (14B, right-most bars) is most prominent in the z direction, reflecting the fact that the subject (and cast) has some vertical freedom of movement, but little room to rotate. Note that the variability is larger for our second subject (1.4 mm in the z direction). Although there is some difference in cast-repositioning error (compare two left-most groups of bars) it would seem that the majority of the increase in error for the second subject is due to a relatively poor fit of the cast internally. Most likely this is due to errors in the optical scanning (see Discussion). Within and between session movement estimates are very similar for subject number one — moreover, they seem to follow the head cast. In other words, subject one's head is moving little within the cast and most of the error is due to replacing the cast within the dewar. Expressing these data as multivariate Gaussian probability distributions, this means that the expected coregistration errors (as a single Euclidean distance) are 0.29 and 1.3 mm for subjects 1 and 2 respectively with corresponding 95% confidence limits of 0.6 and 2.7 mm.

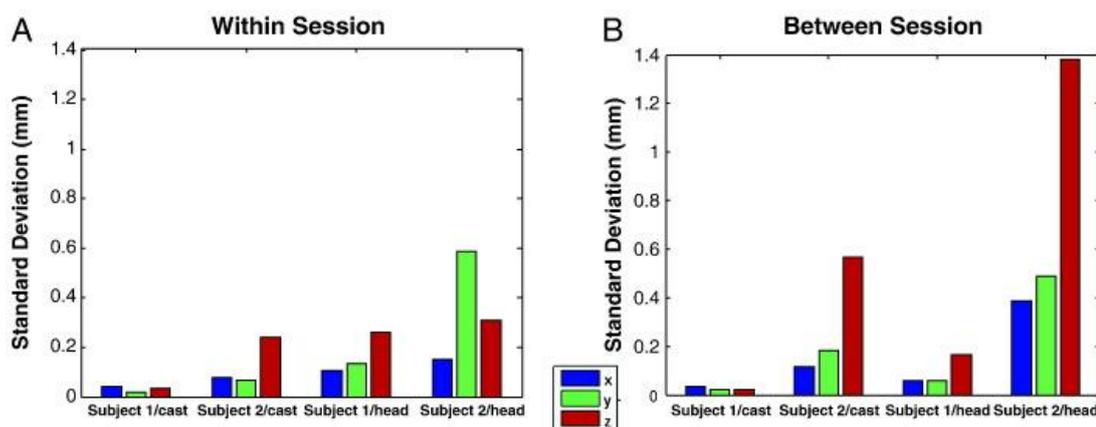


Figure 14: Within (A) and between (B) session cast/head movement estimates for both subjects. Panel A. Within session estimates of variability in cast (left most bars) and head (right most bars) position. Note both cast-movement estimates are below 0.3 mm indicating that the external surface of the casts fit securely within the MEG dewar. Also note also that for subject 1 head movement is limited to within 0.3 mm but for subject 2, whose cast fit less well internally (most likely due to the optical scanning stage, see discussion), this figure is around 0.6 mm. Panel B shows between session cast/head movement estimates for both subjects. The between session errors in cast location were comparable with the within session errors for subject 1, but for subject 2 showed and increased to around 0.6 mm in the z (vertical) direction. In terms of head position, we were able to reposition subject 1 to within 0.3 mm between sessions. Whereas subject 2 was again more variable but could still be repositioned to within 1.4 mm, the error again predominant in the z direction indicating some freedom of vertical movement within the cast.

Cortical coregistration estimates

Our previous measurements were based on relative movements of the subject's scalp within and between sessions. Ultimately, however, it is the error in the location of the cortex with respect to the MEG scanner that we wish to estimate. In this section we use functional data, recorded over multiple sessions, to make estimates of the relative and absolute cortical location.

Relative cortical location

Changes in the location of the head over scanning runs will give rise to increased sensor level variability. Figure 15 A,B show the averaged evoked responses from two (the channel with maximum variance, and its contralateral partner) sensor channels over the six month scanning period. The solid black lines show the averaged evoked responses from the 9 recording sessions. The solid blue line shows the grand mean (over all 9 sessions), the shaded area shows 95% confidence intervals on this mean based on the average within (not between) session variance. Note that the variability over sessions is of the same order as the within session variance; suggesting that little additional coregistration noise has been added to the data. In order to quantify the amount of coregistration noise one would expect for different levels of coregistration error we performed a series of simulations.

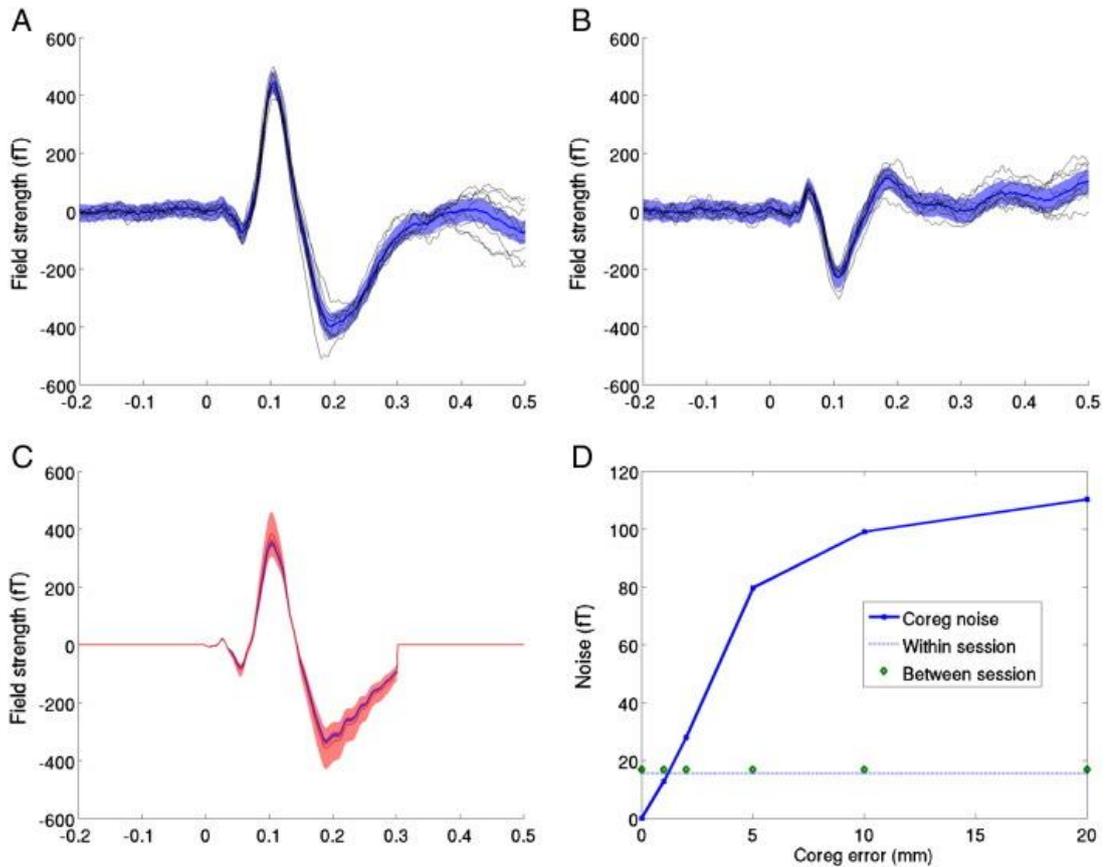


Figure 15: Repeatability of between session sensor level evoked responses. Panels A and B show the MEG channel with highest variance (MLT35) and its contralateral counterpart (MRT35) respectively. Solid black lines show the nine averaged evoked responses from each session over a six month period. Data were baseline corrected from -0.2 to 0 s. The shaded regions show the mean within session 95% confidence intervals on this mean. Note that the between session variability in the evoked response is of the same order as the within session variability. This between session variability includes physiological changes in the evoked response as well as noise due to changes in relative head position over scans. In panel C, current density estimates from the grand average evoked response are projected back out to the sensors using different forward models. Only the time period used for the source inversion is shown (hence the zeros outside $0-0.3$ s). The forward models differ in their coregistration error: the shaded blue area shows the sensor level standard deviation due to 1 mm error in coregistration over sessions; whilst coregistration errors of 5 mm (shaded red) give rise to considerably larger noise levels. Panel D shows sensor noise as a function of coregistration error. In the absence of any other noise, increasing the coregistration error from 1 to 5 mm incurs a 5 fold increase in RMS noise levels. The dotted line shows the within session standard error (the variability on the mean of any one session expected by chance) whilst the circles show the standard deviation of the evoked response (at peak latency) over sessions. Note that there is almost no difference between these two lines, which based on the simulation, should begin to diverge for coregistration errors of greater than 1 mm.

Simulation

In order to estimate the amount of sensor noise one would expect due purely to coregistration error we projected a realistic current density (based on the inversion of the grand average, see next section and Figure 16B) back through a forward model based on different fiducial locations. We took 16 random fiducial locations drawn from a Gaussian distribution of standard deviation equal to the coregistration error (0,1,2,5,10 and 20 mm) and computed the estimated sensor level signal. Figure 15C shows the effect of coregistration errors on sensor noise levels (only the data from inverted time period 0–300 ms has been re-estimated). Solid lines show the mean and the shaded regions show plus and minus one standard deviation. For 1 mm coregistration error (blue) the standard deviation of the signal is of the order of 15fT (blue shaded); for 5 mm coregistration error (red) it is around 80fT. Figure 15D shows the relationship between coregistration error (solid blue) and the expected sensor level standard deviation. Intuitively, the more coregistration error, the more between session signal variance. By decreasing coregistration error from 5 mm to 1 mm one can increase the sensitivity to between session effects by a factor of 5. This simulated curve (solid blue) gives us an estimate of the amount of sensor level noise one would expect due to additional between session coregistration error. For example, for between session repositioning errors of the order of 1 cm we would expect between session variability to increase by 100 fT over within session variability (assuming this is negligible). We can now work backwards to get an estimate for coregistration error based on the difference in between and within session noise levels. The two measured values corresponding to mean within session (signal) variability between session (signal) variability are shown

as dots and circles respectively. Importantly, the multiple scanning sessions have hardly added to the within session noise level suggesting that minimal coregistration error has been introduced (as observed also in Figure 15A,B). Assuming (in the worst case) that any difference in within and between session variability can be attributed solely to head movement (rather than other physiological factors), these functional estimates based on the cortex corroborate with the scalp estimates and put the between session coregistration error at the sub-millimetre level.

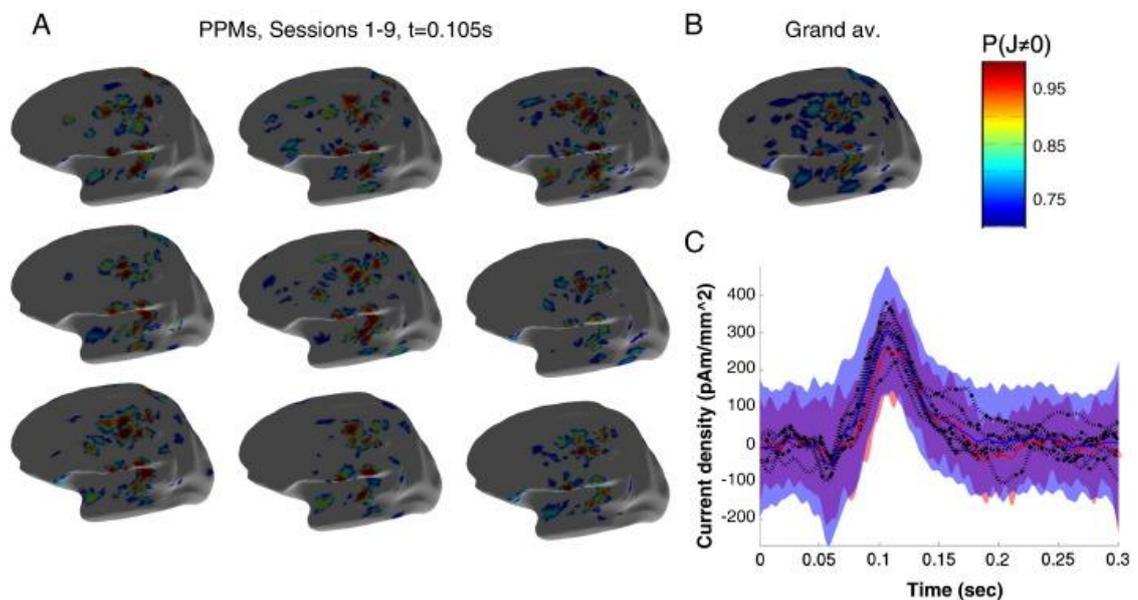


Figure 16: The repeatability of the MSP source level estimates. Panel A shows posterior probability maps of the current density (J) on the inflated cortical surface at 105 ms post stimulus at each of the 9 scanning sessions from within six month period. The colour scales show probability that J is non zero. Note the similarity over scanning sessions and the similarity to the MSP source level estimate based on the grand average (Panel B). Panel C shows the time-series estimates from the peak location in source space. Dotted black lines show source level estimates over the 9 scanning sessions and the average of these is shown as a blue solid line; the average 95% confidence estimates on any one session are shown in the blue shaded area. The estimate of current density on the grand mean sensor level data (red circles) and its reduced 95% confidence limits (red shaded) are also shown.

Sessions 1-9, $t=0.105s$

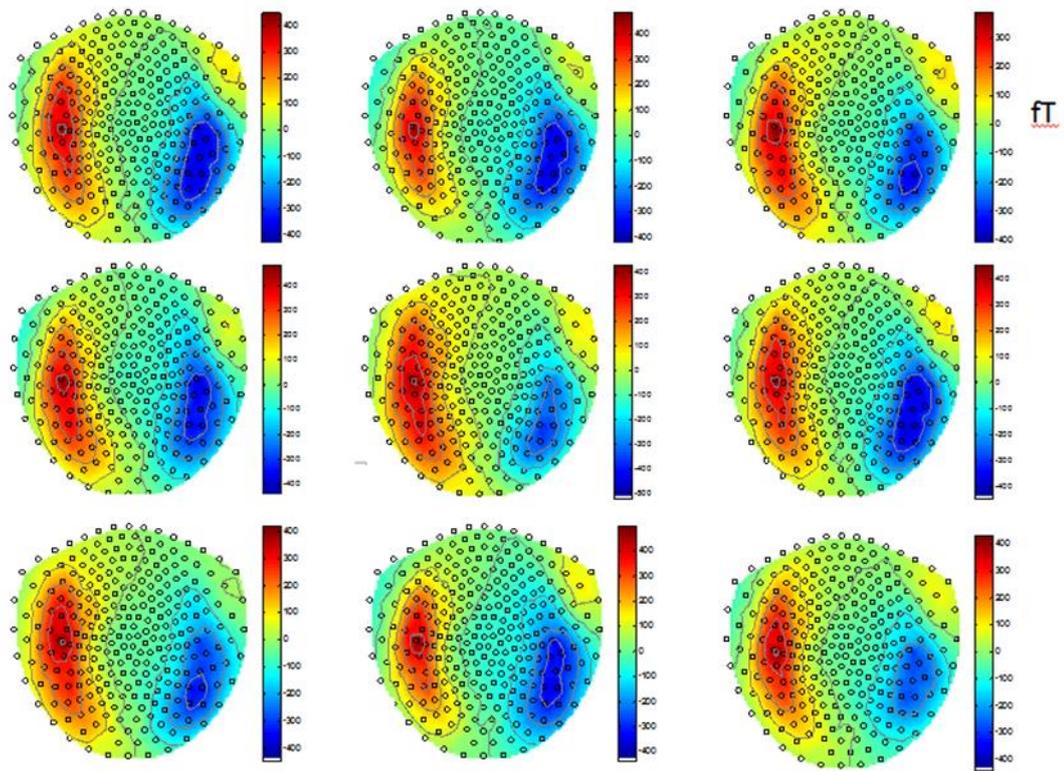


Figure 17: Field maps at time 0.105s post stimulus onset over the 9 scanning sessions.

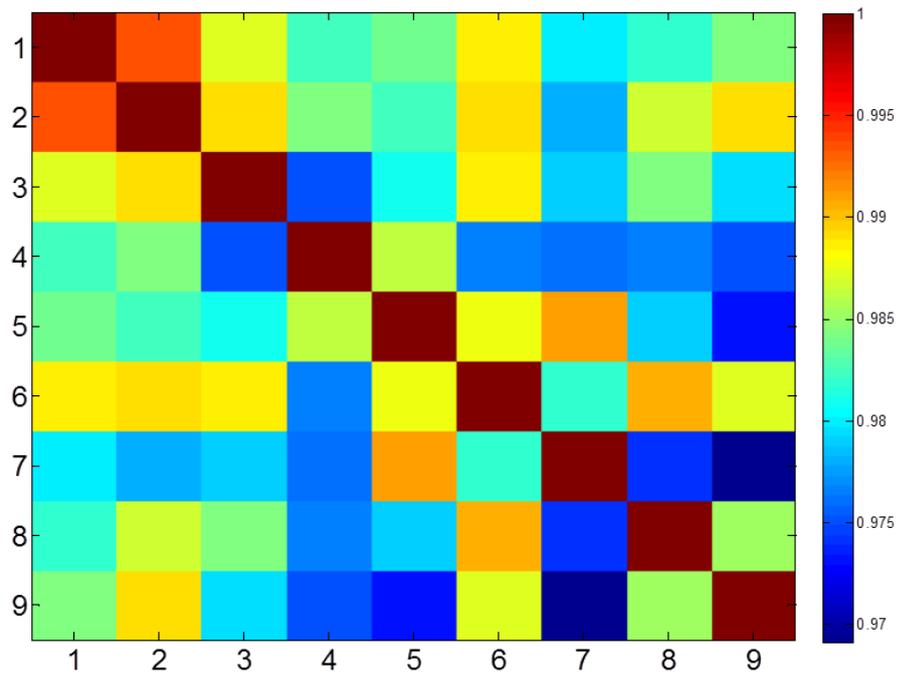


Figure 18: Correlation matrix of field maps show in figure 19. The lowest correlation between any two field maps is 0.97.

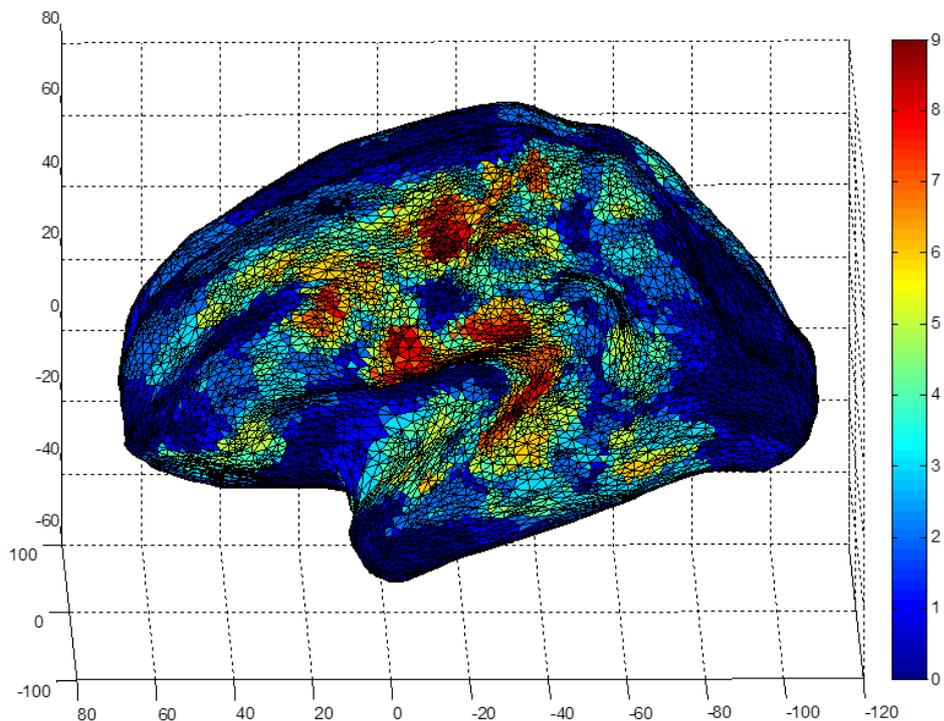


Figure 19: Consistency across scanning sessions. The figure illustrates the consistency of MSP results over all 9 scanning sessions by showing the number of data sets with a posterior probability of >0.9 at each vertex.

Source reconstruction

Figure 16A shows the source level posterior probability maps (PPM), for the MSP reconstructions based on a 0–300 ms window (baseline corrected based on – 200 to 0 ms), at the time of the peak evoked response over sessions on the inflated cortex. Note the consistency of these images in which no fiducial coils were used. Also apparent is the patchy nature of the reconstruction due to the discrete set of MSP source priors used. Figure 16B shows the corresponding PPM of the MSP inversion of grand average sensor data, again consistent with all the individual runs. Figure 16C shows the current density over time estimates at the peak of this grand average PPM (at MNI coordinates: $x = 34.3623$ $y = -31.6975$ $z = 62.0449$ mm). The dotted black lines show individual current estimates from each of the 9 sessions; the mean of these current estimates is shown as a solid blue line and the 95% confidence interval (expected in any one session) on this mean is shown as the blue shaded area. The MSP estimate based on the grand (sensor level) mean (red circles) closely follows this estimate, which would not be the case had any non-linearities been introduced by substantial head movement. The confidence intervals on this grand mean MSP estimate (red shaded) are decreased with respect to the signal session (blue shaded) estimates; but not by as one would expect (a factor of 3 or the square root of the number of scanning sessions). Presumably this difference is due to non-linear behaviour both at the level of the brain, and at the non-linear MSP optimisation stage.

Absolute cortical location

Finally, in order to get some idea of whether we now have useful and precise anatomical information registered to the MEG sensors, we tried replacing the

subject's cortical surface with someone else's. The use of a canonical mesh (Mattout et al., 2007) involves the use of a generic cortical mesh in standard space which is then warped into subject space. The method is attractive as it removes the need to extract per-subject cortical meshes, and has been shown to have little impact on the overall source reconstruction (Henson et al., 2009). We therefore reconstructed the MEG averaged evoked response data for each session onto both the subject's individual mesh and the canonical mesh. One of the advantages of using the Multiple Sparse Priors (MSP) framework is that it returns the negative variational free energy, an approximation to the log model evidence. Free energy will increase with improved model fit to the data, but will decrease as the model becomes more complex (i.e. the simplest models that explain most of the data are favoured). The relative likelihood of one model over another can be assessed by comparing their Free Energies. A Free Energy difference between two models of 3 or more equates to one model being 20 times more likely than the other. The model in this case includes both the forward model (cortical sheet, head location) as well as priors (MSP patches) used to explain the data (Henson et al., 2009). The difference in log model evidence between the two cortical surfaces for each of 32 reconstructions (using different patch sets) for each of 9 sessions is shown in panel A of Figure 12. It is clear that the evidence is in favour of the individual mesh (predominantly positive). An alternative way to summarise this information is to ask if we were to draw a simulation from the set at random how often would it be in favour of the individual mesh (Stephan et al., 2009). This can be more robust than typical fixed effects analysis as it is relatively immune to outliers (a very large free energy difference in just one simulation). Panel B shows the

frequency (r) at which we would chose the individual over the canonical mesh if we were to chose a patch set (from run 9 in Panel A) at random. Summing the area under the curve above 0.5 gives the probability that the individual mesh is the better model (an exceedance probability). These exceedance probabilities are then plotted per session in Figure 12C. Note again that there is strong and consistent evidence for the individual mesh over a six month scanning period. We were interested to see how if we could reconcile this finding with previous work by Henson et al. (Henson et al., 2009) who found no evidence for an improvement using individual meshes. We repeated the same analysis, yet this time we added different amounts of coregistration noise to our head model for each session (Figure 12D). Note here that as the coregistration increases beyond 5 mm (typical coregistration error) there is little to distinguish between the two meshes.

Discussion

In this chapter, we have presented a new and conceptually simple approach to reduce between and within session coregistration errors in MEG. The head casts provide a robust fiducial frame reducing absolute and relative errors for coregistration to anatomy to within 1–2 millimetre levels. By constraining head movement within and between sessions they also allow us to build up very high SNR data sets. The between session repeatability improves our longitudinal sensor level sensitivity to signal amplitude change by a factor of 5. This equates to scanning 25 times fewer sessions (at 1 mm coregistration error) and retaining the same sensitivity (as at 5 mm coregistration error).

Both coil based and functional estimates put the relative coregistration error at around 1 mm. The absolute coregistration error (i.e. systematic errors between MEG cortical and true cortical location) is more difficult to assess. Here we used bootstrapping to test the robustness of our estimates but this cannot account for some physiological issues: for example, the movement of the cortex within the skull between upright MEG and supine MRI scanning. With this in mind we tested whether we were able to distinguish between the subject's individual anatomy and a warped version of canonical anatomy within that same individual's skull. We were able to do this and furthermore show that adding back coregistration error destroyed this ability to discriminate. Our results explain the findings of Henson et al. (Henson et al., 2009) who found that using a cortical model based on the individual subject's anatomical MRI does not necessary lead to better source reconstruction outcomes. However we were surprised that we had to add over 5 mm error (to the existing error) before the two models became equally likely. We can say therefore that our absolute coregistration error lies somewhere in between the bootstrap estimate (1 mm) and this upper bound.

It would seem that the biggest source of error in the head cast production process is the optical scanning— which has difficulty dealing with hair (we used a swim cap) and also difficulty in piecing together a number of smooth round surfaces. We are currently revising our procedure to use scalp surface information directly from the anatomical MRI. This could also help to reduce the time it takes to produce a single head cast. At present, having to obtain an optical scan of the subject's head increases overall production times. This could be greatly reduced by extracting surfaces directly from an anatomical MRI scan.

Also, in this first iteration no cut outs were left for the eyes, the aim being to constrain the head as much as possible. This restricted us to using only non-visual paradigms, and may also be intimidating to some subjects who are prone to claustrophobia. However, the design can easily be made to accommodate openings or periscopes for the eyes, which is a consideration for future incarnations of these head casts.

At present, the design uses nylon-based material, resulting in a rigid head casts that could contribute to subject discomfort and claustrophobia. In addition, the rigidity makes fitting more difficult, both in terms of positioning subjects within the scanning system, as well as fitting the head cast to a subject's head anatomy. The rigidity of the design could present health and safety concerns regarding fitting and positioning of subjects. In addition, the idea of being 'locked' into a rigid structure could contribute to subject discomfort. Pressure exerted by the head cast itself while in the system could mean having to restrict scanning times for single sessions of scanning. Thus, we need to explore further options regarding the basic head cast design, with a view of increasing subject comfort and addressing health and safety concerns, while still preserving the benefits resulting from precise fit and improved coregistration.

The coregistration scheme presented here still suffers from some drawbacks, such as having to extract point clouds for coregistration markers from MRIs, and being vulnerable to errors introduced at each individual transformation stage. The optical scanning stage in particular is liable to introducing variable amounts of error, as it depends strongly on image acquisition and processing. However, we hope that for future incarnations of the head casts, these errors will be reduced by including fixed coregistration markers on the head casts

themselves. Issues relating to optical scanning could be removed by using MRI extracted surfaces as the basis of the head cast design. Importantly, the whole coregistration process presented here can be automated by using custom-written scripts that only require information on individual cortical anatomy (from a structural MRI) as an input.

Given the improvements using the head cast technique, it will be interesting to look into further factors which may present sources of error. For instance, error could be introduced by the fact that subjects are in the supine position for MRI acquisition, but seated upright in the MEG scanning system. The location of the cortical surface may well differ between these two positions, and more work is needed to investigate this. Furthermore, the accuracy of the anatomical MRI scan may also present a source of error. In our case, the choice of MRI sequence to be used was limited by the fact that we had to accommodate the dewar-helmet in the MRI scanner, which is too large to fit most head coils. Whilst we found the resulting image to be sufficiently accurate for our purposes, further work is needed to assess the influence and possible contribution of error of anatomical MRI resolution.

Besides the improved longitudinal sensitivity and increased precision on the underlying volume conductor model parameters, we believe a major application of this work will be to build up very high SNR datasets to help define the limits of MEG source localisation and spatial resolution. Given accurate anatomical information, and the ability to record over long periods of time, we can reduce errors introduced by poor co-registration to underlying anatomy and excessive head motion within and between sessions. Since MEG spatial resolution is not bounded by physiology, there is no theoretical physiological limit on the

potential spatial resolution achievable with MEG, and practical constraints will ultimately depend on the improvements which can be gained from reducing within and between session head movement and better coregistration techniques. In future work we will describe the use of these head-casts in order to estimate which cortical laminae give rise to specific signal characteristics.

We believe that the above characteristics will render the use of the head cast technique particularly valuable in the context of studies which depend on high quality data for a limited number of subjects, e.g. longitudinal drug studies. The method also has application in any clinical setting in which precise localisation of function or pathology is required; for example, as a tool for pre-surgical mapping in epilepsy patients. While the cost of producing a single head-cast may seem high at around £500 per subject, this is roughly of the same order as the cost of a single session of scanning. We therefore feel that the potential benefits outweigh the increased cost in studies which depend on high-quality longitudinal data. In addition, the rapid advances in 3D printing technology and its increasing popularity could drive down costs in the long term. At present, 3D printing is already extremely accessible, with a wide range of professional providers and an increase in printing kits targeted at the individual user. While at present, the latter lack the precision and capacity required for making head casts, the increasing popularity and sophistication of such devices could make this possible in the future.

The only real constraint we presently foresee is that the head-casts can induce feelings of claustrophobia and we are working on new less intimidating designs.

3. Discrimination of cortical layers using MEG

Abstract

In the previous chapter, we have introduced the use of subject-specific head casts to improve SNR and spatial resolution in MEG. Here, we aim to establish the extent of the improvements gained by the use of this technique in terms of its impact on source reconstruction outcomes.

Typically MEG source reconstruction is used to estimate the distribution of current flow on a single anatomically derived cortical surface model. In this study we use two such models representing superficial and deep cortical laminae. We establish how well we can discriminate between these two different cortical layer models based on the same MEG data in the presence of different levels of coregistration noise, Signal-to-Noise Ratio (SNR) and cortical patch size. We demonstrate that it is possible to make a distinction between superficial and deep cortical laminae for levels of coregistration noise of less than 2 mm translation and 2° rotation at SNR > 11 dB. We also show that an incorrect estimate of cortical patch size will tend to bias layer estimates. We then use a 3D printed head-cast (Troebinger et al., 2014) to achieve comparable levels of coregistration noise, in an auditory evoked response paradigm, and show that it is possible to discriminate between these cortical layer models in real data.

Introduction

Magnetoencephalography (MEG) non-invasively measures the changes in magnetic fields outside the head which are caused by neuronal current flow. As

MEG is a direct measure of this current flow, the achievable spatial resolution is not bounded by physiology or anatomy (for example, there are no physical limitations due to local vasculature or capillary bed structure, as in PET or fMRI). Rather, the bounds on MEG spatial resolution come from the models used to explain the MEG data. The higher the signal to noise ratio (SNR) – due to better experimental design, more sensors or longer recording sessions – the more complex the model we can propose. Consequently, modelling error becomes the main limiting factor to MEG spatial resolution. In practice, a great deal of this modelling error is due to limited knowledge of cortical anatomy with respect to sensor positions.

Head movement during MEG recordings is usually not fully constrained and coregistration to anatomy relies upon matching a small number of fiducial locations or matching two smooth round surfaces. This means that the accuracy of the information we have regarding true head position – an important factor our source and forward models rely on – is often limited. For example, using conventional coregistration techniques (i.e. fiducial markers placed on anatomical landmarks), coregistration error is typically assumed to be of the order of 5 mm (Singh et al., 1997, Stolk et al., 2013).

We recently introduced a technique which leads to dramatic reductions of coregistration error by restricting head movement through the use of 3D printed, subject-specific head casts (Troebinger et al., 2014b). This gives access to very high signal to noise ratio (SNR) data sets in which the mapping onto underlying anatomy can be achieved with high precision. Furthermore, our work showed how the MEG data could be used to discriminate between an anatomical model based on the subject's cortex and an anatomical model based on another

individual's cortex warped into the same space. In the present paper we sought to take this a step further and attempt to discriminate between cortical models representing the deep and superficial cortical laminae of the same individual.

In order to be able to detect changes in electromagnetic fields induced by neural activity outside the scalp, two conditions have to be met by the neuronal generators of that signal. First, that the architecture of the neuronal cell is such that it supports and gives rise to a large net current flow, and second, that neighbouring cells drive their intracellular currents with a high degree of group synchronization, so that activity adds up and produces a signal detectable at some distance. These criteria are satisfied by the large populations of pyramidal cells in layers II/III and V of the neo-cortex. Traditionally, they are assumed to form the largest contribution to the MEG signal detected at the sensors (Hamalainen, 1992, Murakami et al., 2002). However, possible contributions from other cell types have been discussed in recent years. For instance, Murakami and Okada (Murakami and Okada, 2006) found that, whilst layer V and II/III pyramidal cells were capable of producing very strong (0.29–0.90 pA) electrical dipoles (Q), both spiny and aspiny stellate cells, which make up the majority of layer IV neurons, produce a dipolar current of sizeable magnitude (0.27 pA and 0.06 pA, respectively). Other distinctions between cell types can be observed in their local connectivity. Layer V pyramidal neurons receive the greatest lateral trans-columnar input (Schubert et al., 2007), whereas in LII/III, the probability of a lateral connection significantly decays over a spatial scale of $\sim 150 \mu\text{m}$ (Holmgren et al., 2003). In addition to this, layer V pyramidal neurons are generally longer with thicker dendrites than those in

layer III and will therefore have a greater dipole moment for the same current flow per unit volume (Murakami and Okada, 2006, Jones et al., 2007).

A number of elegant biophysical time-series models have been proposed and implemented (Jones et al., 2007, Bastos et al., 2012) to make inference on the contributions of these different cell types. For example, Jones et al. (Jones et al., 2007) used a biophysically realistic computational model of the primary sensory cortex in order to make laminar specific predictions of the generators of the first 175 ms of the sensory evoked response.

Our aim here is to establish the conditions necessary to spatially distinguish between activity arising in different layers of the cortex non-invasively using MEG. We use purely spatial information and assume that we are dealing with neuronal populations that are oriented normal to the cortical surface. We simulate data on one of two possible surface models corresponding to the most superficial (pial) and deepest (white/grey boundary) cortical surfaces. We then attempt to reconstruct these data onto both of these surface models and examine the evidence in favour of the correct model. We then examine this discrimination under different levels of coregistration noise and signal to noise ratio. As we know that cells in different cortical layers also have distinct lateral connectivity domains (Schubert et al., 2007) we were interested in how our prior expectations of the extent of a cortical source (Hillebrand and Barnes, 2002, Jerbi et al., 2004, Cottureau et al., 2007) would influence our ability to distinguish between cortical layers. Having established the constraints in simulation we then go on to show an example of an auditory evoked response, recorded using a head cast, in which significant differences between layer models are evident. This is single subject data for which we had no strong prior

hypothesis; we use it principally as a vehicle to describe some of the additional methodological steps that will have to be dealt with in real data.

Methods and Materials

MRI acquisition

MRI data was acquired using a Siemens Tim Trio 3T system (Erlangen, Germany). The subject lay in a supine position. The body-transmit coil was located inside the bore of the scanner for detection of the MRI signal. The MRI data was acquired using a 3D FLASH sequence for optimal scanning efficiency (Frahm et al., 1986). The following acquisition parameters were used: field of view: (256, 256, 208) mm along the (phase (A–P), read (H–F), partition (R–L)) directions, and image resolution: 1 mm³. The repetition time TR was set to 23.7 ms and the excitation flip angle was set to 20° to yield good T1-weighted contrast, standard in most anatomical applications (Helms et al., 2008). 8 echoes were acquired following each excitation and averaged offline to produce one anatomical image with optimal signal-to-noise. A high readout bandwidth was used to preserve brain morphology and no significant geometric distortions are expected in the images. Padding was used to minimize subject motion but some residual effects might remain present in the MRI images. The total acquisition time was 21 min 07 s.

Freesurfer surface extraction

FreeSurfer (Fischl, 2012) was used to perform extraction of cortical surfaces from the anatomical MRI of the individual subject. FreeSurfer breaks this task down into several steps. First of all, intensity correction is performed to deal with any intensity variations due to magnetic field inhomogeneity. This results in a

normalized intensity image. Extracerebral voxels are then removed using a 'skull stripping' approach. Segmentation was performed based on the geometric structure of the grey–white interface, and cutting planes are computed to separate the cortical hemispheres and disconnect the subcortical structures. A single filled volume is produced for each of the hemispheres, each of which is covered with a triangular tessellation and deformed such that an accurate, smooth representation is formed of both the white/grey matter boundary and the pial surface. For more detailed information on the algorithmic procedures, see Dale, Fischl, and Sereno (Dale et al., 1999).

The above process yields surface extractions for the pial surface (the most superficial layer of the cortex adjacent to the cerebro-spinal fluid (CSF)), and the white/grey matter boundary (the deepest cortical layer). Each of these surfaces is represented by a mesh comprising 21,401 vertices. For the remainder of this paper, we will refer to these two surfaces as deep (white/grey interface) and superficial (grey-CSF interface).

Multiple Sparse Priors (MSP)

We used the greedy search option of the MSP algorithm (Friston et al., 2008) implemented as outlined in Lopez et al (Lopez et al., 2012b). The MSP algorithm requires a set of covariance matrix priors corresponding to cortical patches to be defined a priori. Each of these covariance priors corresponds to a single smooth cortical patch of activity (we did not use bilateral patches as in the original MSP formulation) and is therefore determined by the location of an impulse response on the cortical surface and a local smoothness operator determining the full-width half maximum (FWHM) or spatial extent of the source. In this case we had $N = 48$ such priors for the simulation studies and 32

randomly chosen sets of $N = 512$ patches/priors for the real data (see later). We also varied the smoothness of the cortical distribution to see if under- or overestimating this parameter would tend to bias us towards deep or superficial layers.

MSP returns a Free Energy value which approximates the model evidence for the final generative model. Since this generative model includes the cortical surface model used as well as the lead fields, it can be used to compare between different models (Henson et al., 2009, Lopez et al., 2012b). In this paper we show mean log model evidence differences over simulations, i.e. the difference in log model evidence which one would expect on any single realisation (a much finer distinction being possible for the whole group of simulations).

In brief, the MEG data can be related to the neural activity that generates it using the linear model:

$$\mathbf{Y} = \mathbf{K} \mathbf{J} + \boldsymbol{\varepsilon} \quad (3.1)$$

where $\mathbf{Y} \in \mathfrak{R}^{N_c \times N_t}$ is the sensor data, where $N_c = 274$ is the number of sensors (normally 275 but one channel turned off) and N_t is the number of time samples; $\mathbf{K} \in \mathfrak{R}^{N_c \times N_d}$ is the lead field matrix that maps the N_d source locations to the N_c channels; $\mathbf{J} \in \mathfrak{R}^{N_d \times N_t}$ is the current distribution at each source location; and $\boldsymbol{\varepsilon}$ is zero mean Gaussian noise. We used a single shell (Nolte, 2003) based on the inner surface of the skull to define the forward model .

Under Gaussian assumptions, for known source covariance, \mathbf{Q} , the current distribution $\hat{\mathbf{J}}$ can be estimated directly:

$$\mathbf{J} = \mathbf{Q}\mathbf{K}^T(\Sigma_\epsilon + \mathbf{K}\mathbf{Q}\mathbf{K}^T)^{-1}\mathbf{Y} \quad (3.2)$$

Where T denotes a matrix transpose. Here we assume that sensor noise $\Sigma_\epsilon = h_0 I_{N_c}$ is independent and uniformly distributed, with I_{N_c} an $(N_c \times N_c)$ identity matrix and h_0 a hyperparameter effectively controlling the regularisation. Different M/EEG algorithms entail different choices of the prior source covariance \mathbf{Q} (Friston et al., 2008b; Wipf et al., 2010). For the minimum norm (MNM) solution, \mathbf{Q} is simply an $(N_d \times N_d)$ identity matrix; for the Multiple Sparse Prior (MSP) solution, \mathbf{Q} comprises an optimised mixture of a library of N_q source covariance components $\mathbf{C} = \{\mathbf{C}_1, \dots, \mathbf{C}_{N_q}\}$:

$$\mathbf{Q} = \sum_{i=1}^{N_q} h_i \mathbf{C}_i \quad (3.3)$$

Each source covariance component \mathbf{C}_i is generated as a bell-shaped smoothed region with a maximum over its centre (Harrison et al., 2007). This shape is formed with a Green's function over a graph Laplacian. The Green's function \mathbf{Q}_G of $(N_d \times N_d)$ is defined as:

$$\mathbf{Q}_G = e^{\sigma \mathbf{G}_L} \quad (3.4)$$

With σ a parameter that defines the size of the bell, and \mathbf{G}_L of $(N_d \times N_d)$ the graph Laplacian generated with the vertices and faces provided by the cortical surface:

$$\mathbf{G}_L = \begin{cases} -\sum_{k=1}^{N_d} A_{ik}, & \text{for } i = j, \text{ with } A_i \text{ the } i\text{-th row of } A \\ A_{ij} & , \text{ for } i \neq j \end{cases}$$

Where A is an adjacency matrix of $(N_d \times N_d)$, with $A_{ij} = 1$ if there is face connectivity between vertices i and j (maximum six neighbours for each vertex). For a more detailed discussion of these patches and their implementation in SPM see (Belardinelli et al., 2012, Lopez et al., 2014).

The algorithm then uses a non-linear search to optimise the hyperparameters using the variational free energy as a cost function (Friston et al., 2008a). Briefly, the negative variational free energy is a trade-off between the accuracy of the model in explaining the data, and the complexity of achieving that accuracy (Penny et al., 2010).

Here, optimization is performed at two levels: within layer and across layers. Initially, each \mathbf{C}_i corresponds to single local smooth patch of cortex. We used two possible levels of smoothing corresponding to FWHM=5 and 10mm over the cortex. In this case, where there are many hyperparameters, the optimization is achieved using a Greedy Search algorithm (Friston et al., 2008a). This optimisation is carried out independently (on the same data) for each cortical layer model and returns lower bound on the free energy or an approximation to the evidence for that surface model. The selection of the centres of these patches is a problem in practice, as there is a trade off between having a large number of patches entailing a large computational and optimization burden; and a small number of patches meaning that the cortex will be subsampled. For simulations we used the same prior and simulated source locations so this was not an issue. For the empirical data however we took 32 sets of 512 randomly selected vertex indices to comprise the prior set (Lopez et al., 2012a) and used the prior with the highest model evidence.

In the empirical section, we wanted to pool the source estimates from across two independent layer models (each with N_d vertices) into a single two layer model (with $2N_d$ vertices). In this case each \mathbf{C}_i became $(2N_d \times 2N_d)$ diagonal matrix with either the upper or lower N_d elements set to the source covariance estimate from one of the two (individually MSP optimized) source covariance matrices \mathbf{Q} from equation (3.3). After optimization this gives a new $(2N_d \times 2N_d)$ source covariance estimate \mathbf{Q}_{both} which can be substituted into equation (3.2) to produce a cortical current estimate distributed over two cortical layers.

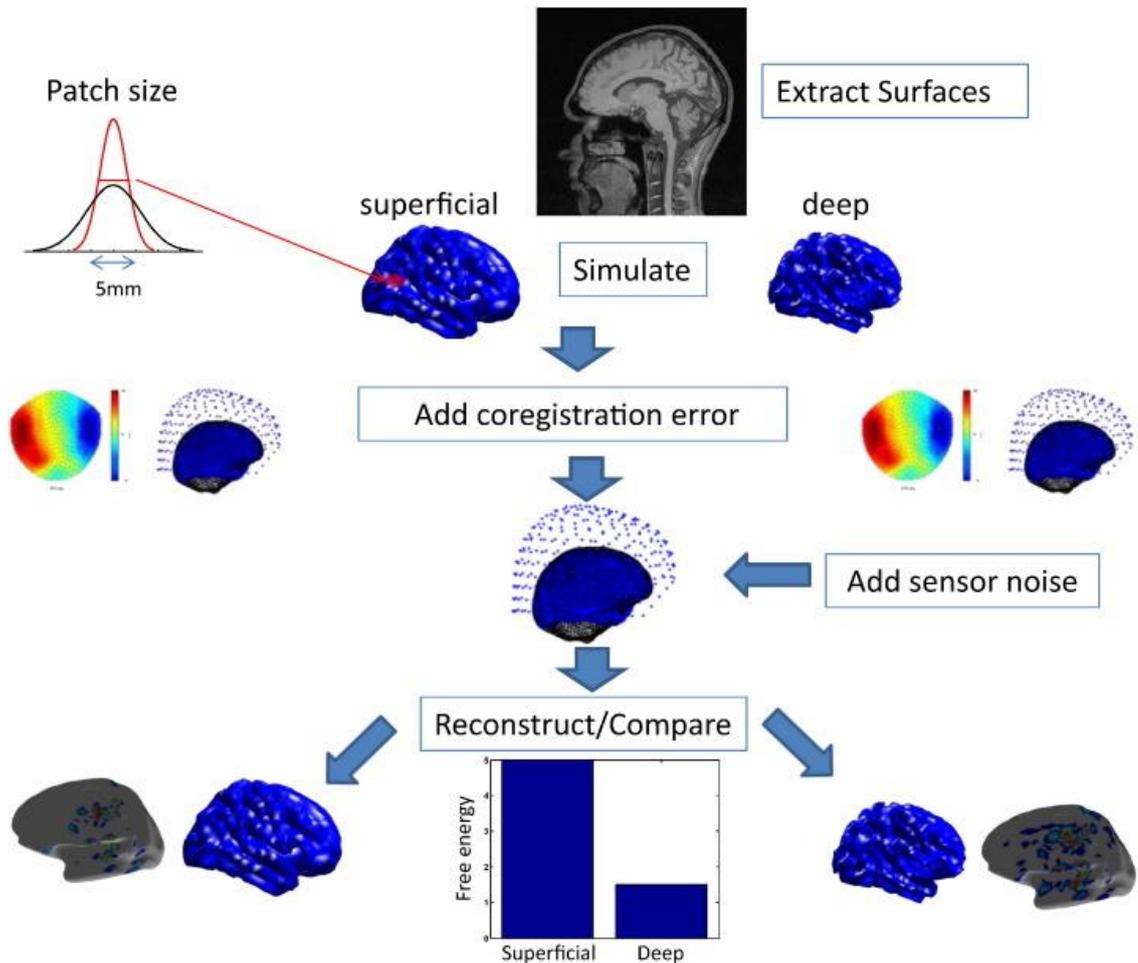


Figure 20: Outline of the simulation process. We use FreeSurfer to extract the pial and white matter surfaces for a single male subject. The pial and the white matter surfaces define the superficial and deep generative surface models. Activity is simulated by randomly selecting source locations from a list of 21401 possible patch centres. Data are generated using two different patch sizes; corresponding to either $FWHM = 5\text{ mm}$ or $FWHM = 10\text{ mm}$ (top left inset panel). Sources are simulated on both surfaces, followed by coregistration of each dataset to both cortical models. To add coregistration error, a random perturbation is added to the fiducial locations at this stage, either taking the form of rotation or pure translation of zero mean and specified standard deviation (0, 1, 2, 5, 10, 20 mm/degrees). In both cases, coregistration error takes the form of a lateral shift. Finally, we use the MSP algorithm to perform source reconstruction. This yields a free energy value as an outcome, which approximates the log model evidence, allowing us to compare between reconstructions based on the ‘correct’/‘incorrect’ surface models, as well as ‘correct’/‘incorrect’ patch sizes.

Simulations

The basic experimental procedure is outlined in Figure 20. On any single iteration, the first step is to generate a synthetic data set. All simulations were based on a single dataset acquired from a real experimental recording using a CTF 275 channel Omega system (Troebinger et al., 2014b). The sampling rate was 600 Hz and the dataset consisted of 145 trials each with 1201 samples. We used a single shell forward model (Nolte, 2003).

At each iteration, we drew three vertex indices from a total of 21,401 indices to specify simulated source centres. These same indices were used to index cortical locations on both superficial and deep laminae (in Freesurfer the vertices remain ordered such that the vertex n on the pial surface is adjacent to vertex n on the white matter surface). We simulated sinusoidal activity profiles of 20, 22, and 24 Hz for the three sources over a time window from -100 to 300 ms. This meant that for each iteration, we had two datasets: one corresponding to data generated by sources on either superficial or deep cortical layers, sitting at the same approximate location (i.e. differing by the cortical thickness).

We repeated this process for 16 iterations giving 32 data sets post-simulation (16 source triplets simulated on the superficial surface, and 16 triplets on the deep surface). Gaussian white Noise was then added to the simulated data to give per-trial amplitude SNR of 0 (default), -10 or -20 dB. Coregistration noise was added by either translating or rotating the fiducial locations by a random amount drawn from a distribution of zero mean and specified standard deviation (20, 10, 5, 2, 1, and 0 mm and 20, 10, 5, 2, 1, and 0° respectively).

We then reconstructed these data (with both sensor and coregistration noise) onto both surface models using the MSP algorithm (above). For each surface model we recorded the negative free energy.

In order to give some insight into the free energy metric of model fit used in this paper we show the parallels with cross-validation accuracy in Figure 21. Cross validation involves the partition of the data to be fit into training and test portions. The idea is to fit models to the training data and then to judge between models based on the test data. Models which are too complex will overfit the training data (i.e. fitting the noise) and therefore perform poorly on the test data. Similarly models which are too simple will also not fit the test data. In other words there is the same accuracy-complexity trade off as with free energy. Much like arguments for parametric and non-parametric statistics, the free energy approximation will be more powerful (as it uses all the data) when the underlying assumptions are met, whereas cross validation will be not quite as sensitive, more time consuming, yet robust. Here (Figure 21A) we generate 102 sets of data based on a triplet of sources sitting on the superficial surface model. All these datasets have the same underlying signal but differ in their noise content. We then fit just the first (training) dataset using superficial and deep candidate surface models. Each of these fits returns a free energy value and the difference between them gives us an analytic estimate of the posterior probability of one model over the other. In this case the log free energy difference between models is 4.3, meaning that the superficial candidate is the most likely with posterior probability of 0.9866 ($1/(1 + \exp(-4.3))$). We can now pursue an alternative route to evaluate the two candidate models by comparing the prediction of the data made by each of the models with the remaining 101

test datasets. In Figure 21B one channel from one test data set (blue solid) is plotted alongside predictions from the superficial (red circles) and deep (green crosses) candidate models (which were fit to the training data). Note that the superficial model makes the more accurate prediction of the new test data. Now we can compare the errors (over all channels and time points) for each test dataset from the two models. Figure 21C shows the ratio of these errors for the 101 test datasets. In 99 of these cases the error for the superficial model is smaller than that of the deep model (points below unity) — i.e., a cross validation accuracy (or non-parametric posterior model probability) of 0.9802.

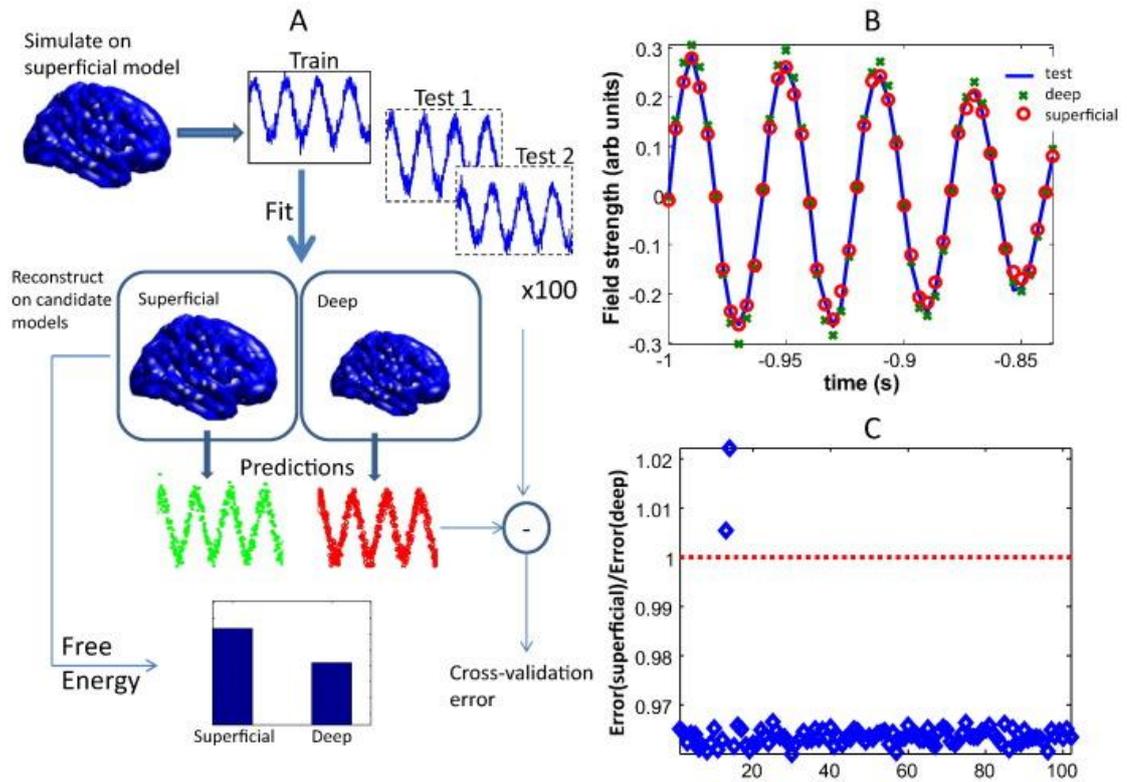


Figure 21: An example to show the link between the free energy metric used in this paper and traditional cross-validation approaches. The basic approach we used is illustrated in Panel A. Here, we simulate 3 sources on the deep surface, and reconstruct these data onto two candidate cortical models (superficial and deep). In this case the free energies (or log model evidences) for the two models differ by 4.3. This suggests that the pial candidate model is more likely with a probability of $1/(1 + \exp(-4.3)) = 0.9866$. An alternative method to judge between models would be to use cross-validation to see which model predicts new data more accurately. Based on the original set up we simulated a further 101 data sets, using the same source locations on the superficial surface (in other words these data had exactly the same underlying signal but different noise realisations). We now use the two candidate models to generate data and compare these predictions with new test data. Panel B shows the signal for a single MEG channel, for a single test dataset (blue solid) and predictions from the two candidate reconstructions (deep model green crosses, superficial model red circles). Note that the error between the superficial candidate model and the test model (based on superficial) is smaller than that of the deep. Panel C shows the ratio of these errors (over all channels and time points) for the two candidate models over 101 test datasets. The red line is at unity, points above the line show smaller error for the candidate deep surface model, points below indicate that the superficial model provides a better prediction of the test data. The incorrect model (deep surface model) is favoured in only two cases. This means that the deep model is more likely with a probability of $2/101 = 0.9802$, in accord with the analytically derived posterior probability based on free energy (0.9866).

Experimental evaluation

Subject task

We collected data from a single, male, right-handed subject using a head cast. Full ethical approval was obtained from the UCL research ethics committee (application number 3419/001). We used a simple finger movement task adapted from Muthukumaraswamy (Muthukumaraswamy, 2011), which involved abductions of the right hand index finger performed to an auditory cue. The cue consisted of a simple auditory tone (1000 Hz), played via a piezoelectric device connected via plastic tubing to ear-inserts, followed by an inter-stimulus interval of 3.4–4.5 s. This gave approximately 145 epochs of data per ten minute recording session. EMG traces of the first dorsal interosseus (FDI) were used to track finger movements (although we did not make use of this information directly for the purpose of this paper). Each session of scanning was split into 4 ten minute sections, during which the subject performed the finger movement task described above. Two such recording sessions were performed on separate days, giving 8 runs of task performance in total.

We used a new version of the head cast technique described in Troebinger et al. (Troebinger et al., 2014b). Here, rather than building a solid nylon cast we used a 3D printer (Zprinter 510) to construct a positive image of the subject's scalp surface including fiducial markers. We then used this positive image and a replica internal dewar surface to construct a polyurethane foam head-cast. Because the fiducial markers were printed onto the subject's positive head surface used to make the casts, these new head casts included designated locations for placement of fiducial coils. The standard deviation of the fiducial coil locations over the eight scanning runs was 1.1, 0.4, 0.9 mm (Nasion, nas),

0.5, 0.1, 0.3 mm (left pre-auricular point, lpa) and 0.4, 0.2, 0.6 mm (right pre-auricular point, rpa), in x, y and z, respectively.

Results

Simulations

Coregistration error

In practice, coregistration error will typically occur as a result of a combined lateral shift and rotation of the subject's model anatomy with respect to the true anatomy. Here, we simulated both of these types of error separately.

Figure 22A shows the mean Free Energy difference between the reconstruction onto the true surface (e.g., simulated on the superficial layer and reconstructed on the superficial layer) and the reconstruction on the incorrect (e.g., simulated on superficial layer and reconstructed on the deep layer) surface for different amounts of coregistration error (lateral shift). The left and right sections of the plot correspond to the true models being superficial and deep surfaces respectively; evidence in favour of the true model is positive. Firstly, note that the picture from both sets of simulations is very similar (there is little asymmetry). The smaller the coregistration error the greater the model evidence difference between surface models. Mean free energy differences of greater than 3 mean that the true surface model is (on average) twenty times more likely than the incorrect one. If we take this as our significance level then note that discrimination between surfaces is only possible when coregistration error is less than 2 mm. At typical coregistration errors (of sigma = 5–10 mm) there is only marginal evidence for the true model whereas at coregistration errors of standard deviation of 20 mm, there is slightly more evidence in favour of the incorrect model. It is also interesting to note the steep rise in mean Free Energy

difference moving from values of 5 mm to 1 mm of coregistration error, indicating the large amount of spatial information carried by the MEG data. It also suggests that considerable improvement lies just beyond typical (5 mm) coregistration error.

In Figure 22B we show the effect of random rotations (rather than translations) of the MEG coordinate frame. An orientation error of 2° (with 0 translation error) would correspond, for example, to 1.4 mm lateral shift in the nasion coil with corresponding 1.4 mm shifts in the opposite direction of the left and right pre-auriculars (assuming an 8 cm head radius). For both simulations based on the superficial and deep surfaces, we can only confidently ($p < 0.05$) discriminate between surfaces for rotation errors of 2° or less. Although we see very little difference in the relative change of free energy for the two different surface models it does seem that the data simulated on the deeper cortical layers are more sensitive to this parameter, presumably as this deeper surface is more tightly convoluted and will produce magnetic field signatures of higher spatial frequency content.

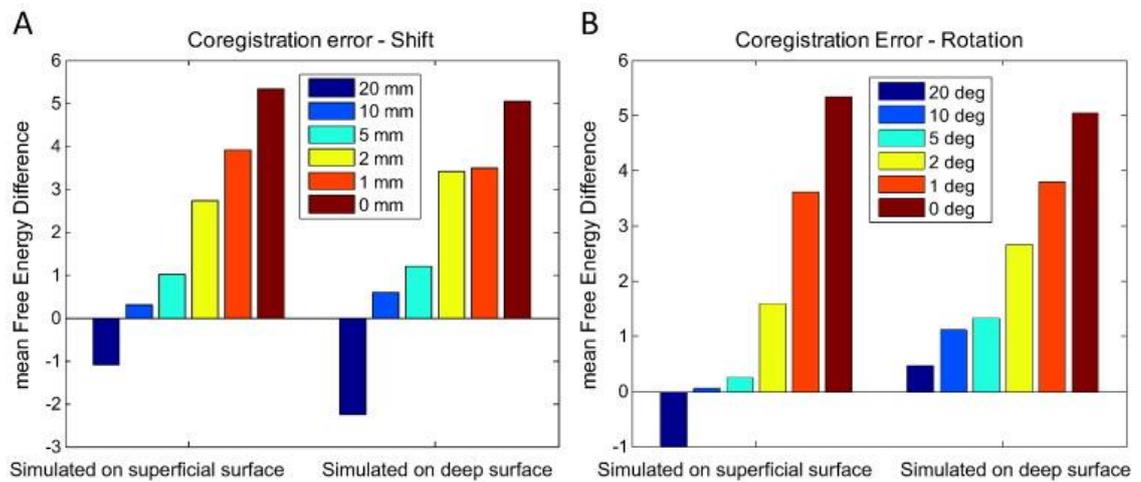


Figure 22: **A:** The effect of lateral coregistration error (shift). The bars shows the average (over triplet simulations) free energy difference between true and incorrect surface models (evidence in favour of the true model is positive) for source triplets simulated on the superficial (left) and deep (right) surface models. Different coloured bars show different levels of coregistration error in mm. Both plots indicate that the ability to discriminate between models representing different cortical surfaces is destroyed once coregistration error exceeds 2 mm (free energy differences < 3). **B:** The effect of rotational coregistration error. The bars shows the average (over triplet simulations) free energy difference between true and incorrect surface models (evidence in favour of the true model is positive) for source triplets simulated on the superficial (left) and deep (right) surface models. Both for simulations based on the superficial as well as the deep surface model, the cut-off for being able to distinguish between the true/incorrect surface models (with 95% certainty or free energy > 3.0) lies around the 2-degree-mark.

Signal-to-Noise Ratio (SNR)

Typically, MEG recorded data exhibit relatively low SNR (~ 0 dB or equivalent signal and noise r.m.s. amplitude in any single trial of data), limiting the spatial resolution at the source reconstruction stage. We wanted to test how critical a factor sensor level SNR was in the selection of the true cortical model. We simulated data at -20 , -10 and 0 dB per trial amplitude SNR with a typical number of 145 trials. We looked at averaged data and this boosted our SNR by 21 dB to 1, 11 and 21 dB. Figure 20 shows the average model evidence difference for reconstruction of 3 sources with zero coregistration noise. It is clear that even at 11 dB average SNR (i.e. with per trial signal amplitude at -10 dB or 1/10th of the r.m.s. noise) it was still possible to make a clear distinction between cortical models. Whilst this may seem surprising at first glance, it is important to note that there was no coregistration noise in this case and as the simulated data were temporally smooth, the temporal dimension reduction in MSP would have also boosted the signal relative to the noise.

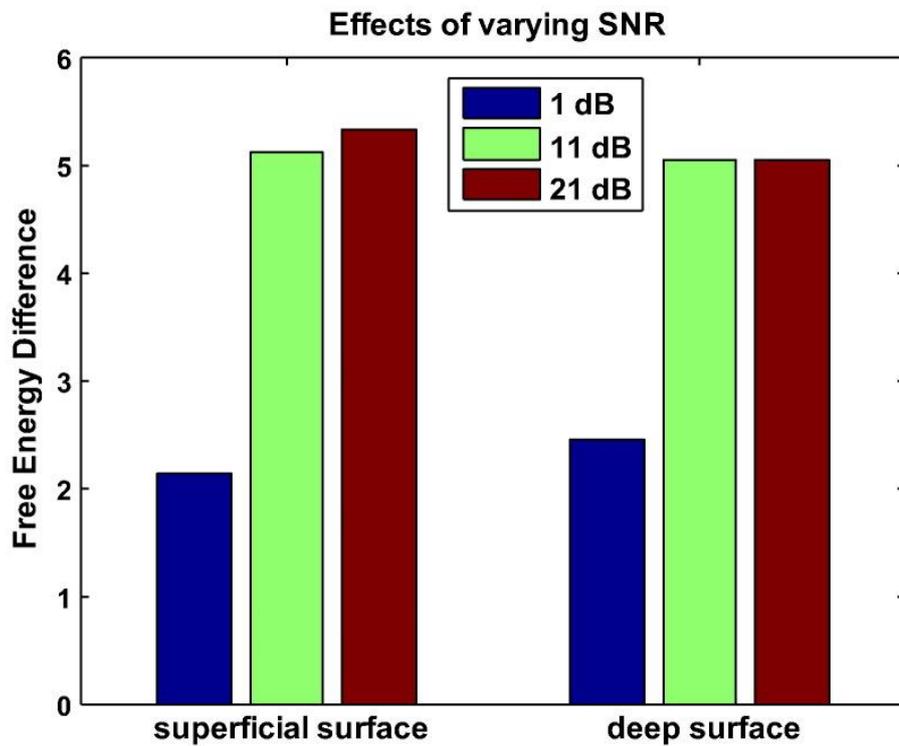


Figure 23: *The effect of varying SNR. The ability to distinguish between surface models improves with SNR. At 1 dB SNR, it is not possible to make a significant distinction between cortical models (free energy < 3). However, at a higher, yet still moderate SNR of 11 dB, there is strong evidence in favour of the correct surface model.*

The MSP algorithm depends on a prior library of candidate cortical patches formed from the smoothing of a cortical impulse response. The default smoothing translates to an FWHM (or effective patch diameter) of 10 mm. We simulated sources of both FWHM = 10 and FWHM = 5 mm and then attempted to identify the true cortical surface (i.e., the surface sources were simulated on) either assuming FWHM = 10 or 5 mm. In addition to varying patch size and cortical layer we randomly perturbed rotational coregistration by $\sigma = 1, 2$ and 5° (as in Figure 22B). All other parameters were set to default: 0 dB per trial SNR, 3 sources, 16 iterations etc.

The top row of Figure 24 shows the relative model evidence differences between true (A superficial, B deep) and incorrect cortical models for simulated sources of extent 10 mm when the MSP has the correct (FWHM = 10 mm) and an underestimate (FWHM = 5 mm) of patch extent. Looking to the right of Figure 24A we can see that if our extent prior is smaller than that simulated there is no strong evidence (free energy difference is less than 3.0) to support even superficial models with no coregistration error over deep ones. Conversely panel B right column, in which the data are simulated on the deeper surface, show very strong evidence for this surface even at coregistration errors of five degrees. Taking panels A and B together, underestimating the patch extent (FWHM = 5 rather than 10 mm) has biased the evidence in favour of the deeper cortical model. Now looking at panels C and D in which the true model is the 5 mm patch diameter, we see that there is similar discrimination between surfaces using the correct patch size priors (left most columns) as we observed for the 10 mm patches. In this case however, there is relatively strong evidence in favour of the superficial model when patch size is over-estimated (panel C, FWHM = 10 mm) but no strong evidence in favour of the deeper model even when it is the true model (panel D, FWHM-10 mm). In other words overestimation of patch extent biases the evidence towards superficial cortical layer models.

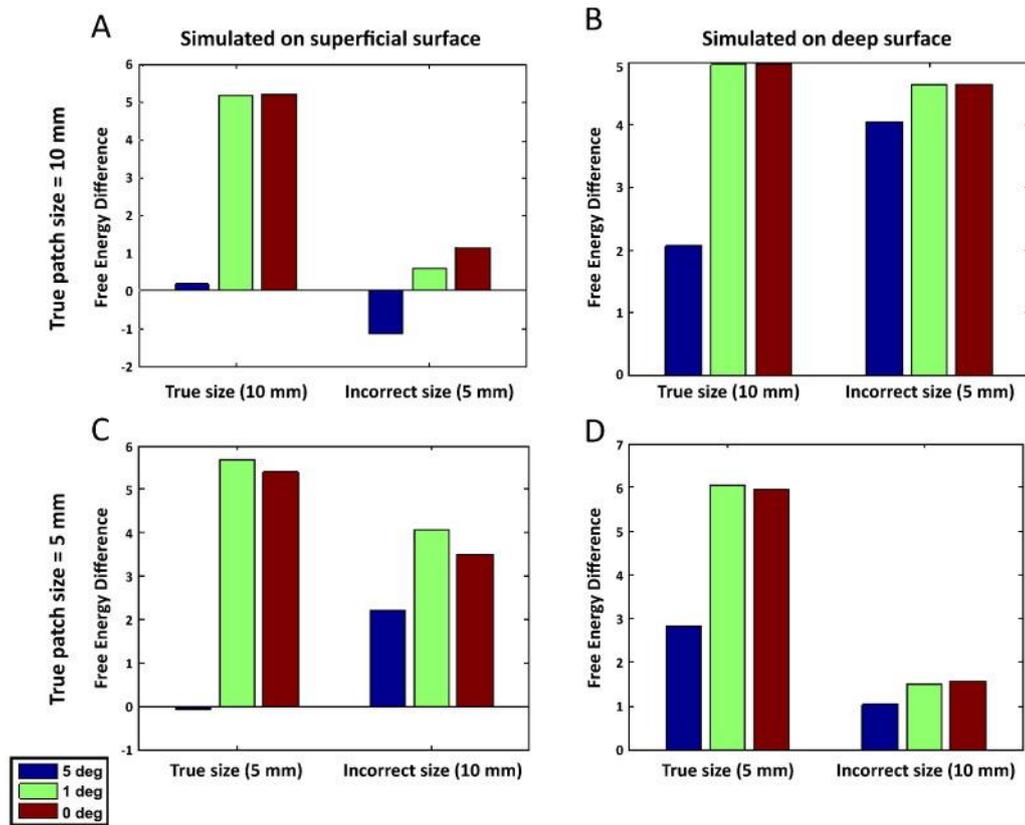


Figure 24: Investigating the effect of patch extent. Data were simulated on the superficial (left) or deep (right) surfaces models using either a patch extent of FWHM = 10 mm (upper panels), or FWHM = 5 mm (lower panels). Three sources were modelled, at a per trial SNR of 0 dB. In addition, rotational coregistration errors of 0, 2 and 5° were simulated. Panel A in the top row shows relative free energy differences for simulations based on the superficial surface and using a patch extent of FWHM = 10 mm. Looking at the leftmost set of bars (corresponding to reconstructions using FWHM 10 mm), we observe strong evidence in favour of the correct model (positive values) for reconstructions at true patch extent, both for 0 and 1, but not 5° of coregistration error. This pattern is destroyed when underestimating patch extent (FWHM = 5 mm), as illustrated by the rightmost set of bars where no clear difference between surfaces emerges. However, looking at Panel B, which shows the comparison based on the same patch size (10 mm), but using the deep surface model, it is clear that even if we underestimate patch extent (5 mm), the strong evidence in favour of the true model is preserved. This suggests that when we underestimate patch extent, we are introducing a bias towards deeper surface models. The bottom row shows relative free energy differences for simulations based on smaller patch extent (FWHM = 5 mm). Panels C and D correspond to the true surface models being superficial and deep, respectively. Here, in the case of superficial-surface-based data, when overestimating patch size, the evidence in favour of the superficial surface model is preserved (Panel C). On the other hand, as shown in Panel D, in the case of deep-surface-based simulations, overestimating patch size decreases the evidence in favour of the deep surface model. This indicates that by overestimating patch extent, we are introducing a bias towards superficial surface models.

Experimental data

The task we selected consisted of a button press to an auditory cue as described in Muthukumaraswamy (Muthukumaraswamy, 2011). We recorded 8 ten minute runs (of approximately 145 trials each) from a single male subject, spread over two sessions conducted on separate days. We used averaged evoked responses from 0–300 ms (0–80 Hz) time-locked to the auditory cue, baseline corrected based on – 200 to 0 ms. These data were projected into 274 orthogonal spatial (lead field) modes and 8 temporal modes. We used two different cortical layer models (deep and superficial), each with two possible patch extents (FWHM = 5 and 10 mm), to model the average auditory evoked response 0–300 ms post-stimulation in the 0–80 Hz band. For each of the 8 runs, we performed 32 MSP iterations with 512 random patch locations, using each of the two layer models, with two possible patch smoothness values. We took the best model (in terms of free energy) for each of these 4 parameter combinations. In Figure 25, we show the difference in free energy between the best deep and superficial models for 5 mm and 10 mm patch extent. Note that the deep layer model tends to improve upon the superficial and that, although this difference is influenced by patch size, the deep layer model wins in 6 of the 8 runs. Taking account of all these free energy differences in a fixed effects manner (Stephan et al., 2009, Penny et al., 2010) allows us to compute the posterior probability of the two surface models marginalising over patch smoothness (panel B), or of the two patch smoothness options marginalising over layers (panel C). We find that the most likely origin of these data over the group of recordings is the deeper cortical layer with a posterior probability of 0.9834. The patch comparison is more equivocal with the smaller patch size having slightly greater posterior probability at 0.8238.

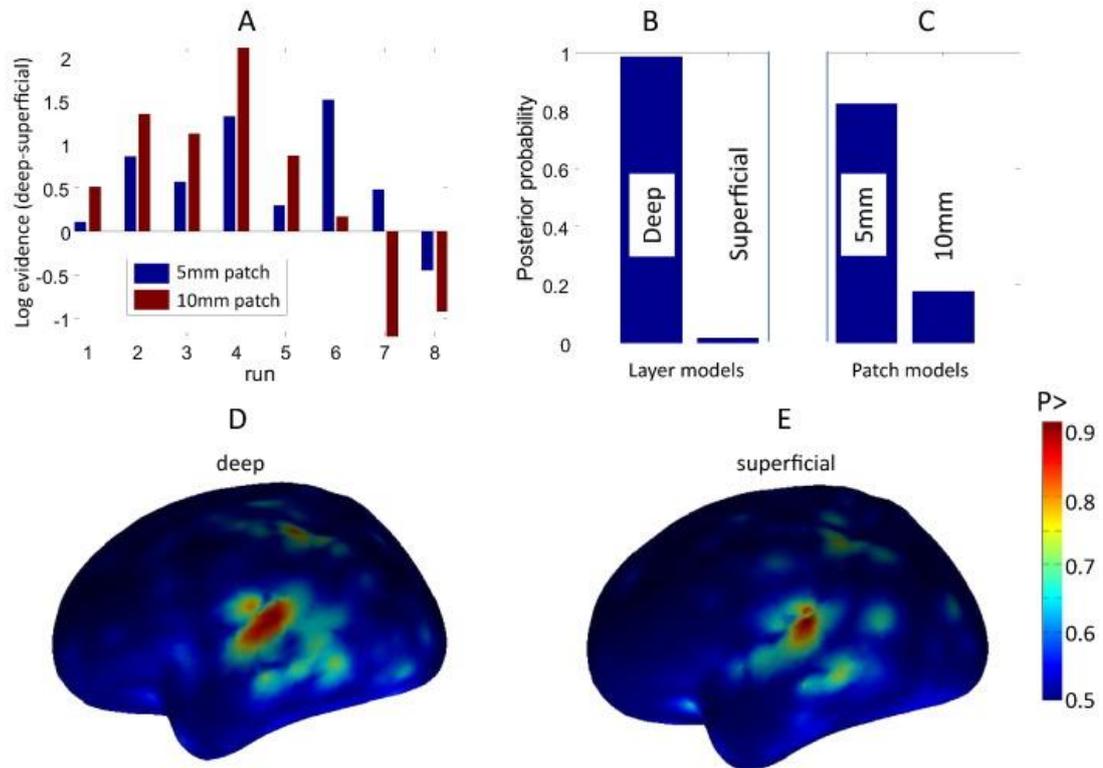


Figure 25: Investigation of layer and patch size models for an auditory evoked response paradigm in one subject over eight recording sessions. Panel A shows the difference in log model evidence between deep and superficial layer models for each of the eight runs for 5 (blue bars) and 10 mm (red bars) patch sizes. Positive values mean that the deep surface model is more likely; this is the case in six of the eight runs for both patch sizes. Panel B shows the probability of the two layer models based on all eight runs and both patch sizes. Note that the deeper cortical surface is more likely with a posterior probability of 0.9834. Panel C shows the probability of the different patch size models this time marginalising over layer models, the 5 mm patch size being more likely with a posterior probability of 0.8238. Panels D and E show the average posterior probability (probability that the current at any point is non zero) map for $t = 0.92$ s post stimulus over the eight sessions (combining both patch sizes) for deep and superficial layer models respectively.

For each run we then pooled the source covariance estimates for each of the four best source models (deep and FWHM 5 mm, superficial and FWHM 5 mm, deep and FWHM 10 mm, superficial and FWHM 10 mm) to create a single current distribution extending over both cortical layers. In panels D and B we show the average of the eight posterior probability maps (the probability of non-zero current flow) at 92 ms post stimulus (the peak in global field power over the eight runs). Note the different current distributions on the two cortical layers; note also that although at the group level the evidence in favour of the deep layer model was overwhelming; at the individual run level the marginal (< 3) differences in model evidence between the two layer models means that the posterior current density will straddle both surfaces.

Discussion

In this chapter, we have shown that for low values of coregistration error (below 2 mm and two deg) and moderate SNR (11 dB for average data), it is theoretically possible to distinguish superficial and deep cortical laminae using MEG. In addition to exploring this in a theoretical context using simulations, we have also performed a first demonstration that differences between layer models can be observed in real MEG data, recorded using a head cast.

Here, we have added coregistration error in the form of lateral shift, both as translational as well as rotational movement. It could be argued that in practice, this is likely to occur as a combined movement, and that there could be additional movement in the vertical direction. However, for the purpose of the work presented here, where we wanted to get an estimate of possible improvements, the simplifications outlined above were deemed permissible. The question remains whether vertical movement would have a qualitatively different effect on source reconstruction outcomes than lateral shifts, and this could be addressed in future work. It should also be possible to practically account for any small and systematic errors in coregistration by optimization cortical location based on Free Energy (Lopez et al., 2012a).

It is important to note that here, we used a new version of the head cast design, consisting of latex-based 'soft' casts. This not only allowed us to improve subject comfort, but also enabled us to improve on our previous coregistration scheme, by including designated placeholders for fiducial coils into the head cast design. Thus, we were able to enter fiducial coordinates directly into SPM, rather than having to resort to the complex custom-built scheme presented in our previous paper, improving ease-of-use and reducing the risk of introducing additional errors by having to match point clouds in different reference systems (Troebinger et al., 2014b).

The results presented in this paper provide evidence that, with only slight modifications to current scanning paradigms and procedures, non-invasive human laminar electrophysiology is within reach. In this study we have concentrated on the distinction between the upper and lower surfaces of the cortical sheet but it is not inconceivable that finer distinctions between individual cortical laminae will be possible given the SNR levels (built up over multiple sessions) and low coregistration errors that can now be achieved using individual head casts. These laminar models could be defined through interpolation between the surface models used here or based on new in-vivo MRI techniques such as in-vivo histology (Dick et al., 2012).

One important finding was that the distinction between laminae will be biased by assumptions made about source extent: in our simulations, we showed that an underestimation of the true patch extent will tend to bias estimates towards deeper layers, whereas an overestimation of patch extent will tend to bias layer estimates more superficially. Given the pre-existing differences in lateral connectivity between cell populations in different layers (Schubert et al., 2007) this will be an important factor to marginalise out of the inference on layer specificity in future work.

The simplest explanation of this bias due to patch size is that larger patches of extended cortex create field maps which are very similar to deeper (and smaller) patches. However, since patch extent and its influence on source reconstruction outcomes has not been widely studied to date, it will be necessary to repeat the work presented here in the context of a study using a greater number of subjects, and testing further values for patch extent to establish whether the bias we observed here is consistent across patch sizes and subjects. Following such an investigation, methods to prevent a bias – such as testing several patch extent values and combining the results using Bayesian Model Averaging – could be developed.

Also, in the same way that removal of coregistration error improves the forward model, one would expect improved layer distinction with more comprehensive volume conductor models. For example, recent work (Stenroos et al., 2014) showed that a three-shell boundary element model outperformed the single-shell model used here.

Here we used a head-cast both to remove head movements during scanning and to provide a precise mapping between the structural and functional data. If this precise mapping (between MEG fiducial coils and anatomy) was known for a subject, and the head-movements compensated for using techniques such as Signal Space Separation (SSS) (Taulu and Simola, 2006), then in principle the same layer distinctions should be possible. The only potential issue is that the distinction between laminae may depend on the number of orthogonal spatial modes in the data (which are inevitably diminished in software based approaches to head movement correction).

The finding that discrimination between layers is not particularly sensitive to SNR is encouraging. Importantly SNR here is not the signal change between conditions, or from one cortical source with respect to all others, but rather all cortical sources with respect to the system (plus environmental) noise. Here we assume Gaussian white noise for typical MEG recordings above 1 Hz; but other noise models (e.g. pink) might be more appropriate in situations where one is trying to link layer models to specific frequency bands (see later).

The empirical data raised a number of methodological issues. Classical statistics based on random field theory for multiple lamina models will become an issue here as the surfaces are very highly correlated (in MEG space), although not connected (in our anatomical model), hence random field theory (on a surface) will overestimate the intrinsic volume (or become more conservative) by at least a factor of two (for two surfaces). There could be a number of ways around this — the use of non-parametric approaches (Nichols and Holmes, 2002), the use of heuristic estimates of the effective number of independent elements (Barnes et al., 2011), or, alternatively, the warping of the set of surfaces into a common space defined by their lead fields. At the moment we are uncertain how to optimally combine data from multiple runs. Here we have opted to plot the average posterior probability map. This means pooling data from the different models (superficial, deep, 5 and 10 mm patch size) at a per-run level. Consequently the balance of current flow between the two layers is also determined by the relative model evidences for the different surface models on each individual run. One alternative approach would be to take the winning model from the group (the deep layer model in this case) and present all data from all runs on this single surface. Related to this issue, we currently

perform the MSP prior optimization for each layer model independently. This allows us to make categorical statements about the most likely layer model. One could imagine a third model that would support prior current distributions over both cortical layers (or even multiple layer models). This scenario, which would allow some cortical areas to consist of predominantly deeper layer sources and others more superficial, would certainly be more realistic (especially given that typically, we are interested in time windows of hundreds of milliseconds), although potentially more computationally demanding.

We had no strong hypotheses about the layer specificity of the auditory evoked response. In this single subject, the preference for the deep layer model over the superficial layer model is clear, but we note that the invasive literature here is equivocal. Whilst some invasive studies in other species report similar magnitude responses across layers (Steinschneider et al., 2008, Ogawa et al., 2011) other studies do show a deep layer bias (Sakata and Harris, 2009, Harris et al., 2011, Profant et al., 2013). Profant et al. (Profant et al., 2013) studied the multi-unit response properties of the auditory cortex in the rat. They found the weakest response in superficial layers, which they attribute to the lowest direct thalamic input (as well as possible influences from their recording procedure). Perhaps the strongest arguments come from the general physiology and anatomy. Investigating the laminar organization of the auditory cortex and its response to stimuli (tones, clicks) using single cell recordings, Sakata and Harris found that while layer II/III pyramidal cells exhibit selective responses in both spectral and temporal domains, layer V thick pyramidal cells do not share this selectivity, and exhibit a much more general response pattern. This specificity of superficial and generality of deeper neuronal population is linked

by these authors to their distinct lateral connectivity profiles (narrow and broad respectively). Indeed we think the most plausible explanation of the preference for the deep layer model we found here is that the increased trans-columnar connectivity in the deeper layers will result in a larger synchronous population and hence a larger dipole moment. Coupled with this, layer V pyramidal neurons are generally longer with thicker dendrites than those in layer III and will therefore have a greater dipole moment for the same current flow per unit volume (Murakami and Okada, 2006, Jones et al., 2007).

In these data we also find a preference for smaller patch sizes. This parameter needs further investigation. Reliable estimates of patch size would allow us to compare directly with anatomical connectivity estimates (which may also help distinguish between layers); provide a non-invasive estimate of cortical current density; and generally help inform the basis set for other algorithms which tend to be either based around point-like (e.g. beamformers) or intrinsically smooth (e.g. LORETA) solutions.

Both empirical and modelling works show that the MEG signal is a result of multiple cell populations in different layers. Identifying a paradigm to conclusively demonstrate this layer specificity empirically will be the next hurdle. See for example recent work (Ronnqvist et al., 2013) comparing human MEG measurements with laminar in-vitro recordings.

One promising avenue is a focus on beta and gamma oscillations, for which both empirical and theoretical works suggest a predominance in deep and superficial layers respectively (Bastos et al., 2012).

This is only the first step towards empirical validation and requires replication using empirical data from several subjects and different paradigms. We should also note that here we have used purely spatial models and made no use of the time-series information (Jones et al., 2007, Bastos et al., 2012). Future work might consider the comparison and eventual combination of such techniques. To conclude, here we have provided first evidence that non-invasive human laminar electrophysiology is practically achievable.

4. Distinguishing cortical laminae – Empirical Data

In chapters 2 and 3, we described the development of the head cast technique, and discussed how it could be used in MEG recordings to improve the SNR of the resulting data. We also showed that these improvements in SNR lead to sufficient spatial resolution to allow us to distinguish between cortical models based on superficial and deep surfaces, respectively. These findings confirm that errors resulting from poor coregistration and within session head movement are factors constraining MEG spatial resolution. The fact that we were able to distinguish between cortical models representing different cortical layers was particularly encouraging, as it suggests a spatial resolution in the millimetre-range. However, since our previous findings were based on data recorded from a single subject, we could not evaluate whether our findings were physiologically meaningful. Here, we aim to corroborate previous results in the context of real MEG data recorded from a small number of subjects (4) with multiple runs per subject.

Introduction

We have already discussed the two main limitations associated with MEG scanning protocols. On the one hand, within-session head movement leads to blurring of the sensor level data. On the other hand, fiducial-based coregistration schemes, which aim to assess the location of underlying cortical anatomy with respect to the sensors, typically introduce coregistration errors of 5 mm or more. This leads to considerable reductions in SNR and spatial resolution, and compromises our ability to localize the neuronal sources of the

observed activity at the source reconstruction stage (Hillebrand and Barnes, 2011, Lopez et al., 2012b).

To overcome these limitations, we introduced the use of subject-specific head casts (Troebinger et al., 2014b), which restrict within session head movement and enable precise repositioning of subjects between sessions of scanning. By integrating designated spaces for fiducial coil placement into the head cast design, coregistration error is significantly reduced (<2mm) (Troebinger et al., 2014b). This results in more precise information about how cortical anatomy relates to the sensor positions, and also allows us to accumulate data from several sessions of scanning to build up high-SNR datasets.

Testing its potential to improve spatial resolution using both simulations and MEG recorded data, we demonstrated that this approach allows us to distinguish between source reconstruction outcomes based on cortical models representing deep and superficial laminae. This is encouraging, because it indicates that coregistration- and within-session head movement errors rather than any physiological constraints limit MEG spatial resolution, and that the head cast technique successfully reduces both of these sources of error. Moreover, since we were able to distinguish between models representing different cortical layers, this indicates that the resulting spatial resolution is in the millimetre range. To put this into perspective, MEG spatial resolution (using 'standard' scanning protocols) is usually thought to be in the centimetre range.

Reconstructing data onto models based on different cortical layers offers another exciting possibility, as the differences we observed in source reconstruction outcomes might be indicative of an underlying physiological

mechanism. If this were the case, then MEG recorded data could be used to complement (or even substitute) invasive techniques when investigating cortical information flow and processing. Moreover, since the invasive literature on these topics is mainly based on animal studies, it would offer the possibility of collecting data from human subjects in-vivo, opening up many exciting new avenues for research.

The aim of the work described in this chapter is to assess whether our previous findings are physiologically meaningful. Collecting data from several subjects (4) allowed us to test whether the differences in source reconstruction outcomes between cortical layer models were present and consistent across subjects. To test the physiological significance of our findings, we aimed to select a task eliciting a well-characterized induced response, allowing us to develop a clear hypothesis about its origin in terms of cortical depth. To this end, we turned to the invasive literature on the one hand, and to current theories about the laminar origin of the MEG signal on the other.

Although the laminar profile of sensory induced responses and laminar attributes of information flow within the cortex have been the focus of numerous studies to date, the resulting literature is equivocal. In addition, the physiological basis of the MEG signal is still disputed. Pyramidal neurons have been identified as the main generators of the MEG signal based on their cellular architecture, which gives rise to a comparatively large net current. Furthermore, the activity of populations of pyramidal neurons exhibits a high degree of synchronicity, thereby producing a net effect that can be detected outside the head. Therefore, it is a generally accepted view that the pyramidal neurons in layers II/III and V are the main generators of the MEG signal (Hamalainen,

1992, Murakami and Okada, 2006). However, this interpretation is far from comprehensive, and to some extent based on the reputedly low SNR and spatial resolution of MEG recorded data. It should therefore be noted that other populations of neurons are also capable of producing dipoles of considerable strength, for instance the spiny and aspiny stellate cells in layer IV (Murakami and Okada, 2006). Hence, using the head cast technique to improve spatial resolution and SNR, we must be conscious of the fact that current hypotheses about the origin of the MEG signal may no longer be comprehensive enough.

In a bid to reduce this ambiguity, we focused on the motor cortex: It is 'agranular', meaning that layer IV is completely absent, removing the possible influence of spiny and aspiny stellate cells in this particular cortical region. In addition, the overwhelming majority of excitatory cells found in the motor cortex are pyramidal cells, with the largest specimens to be found in layer V.

Overall, the electrophysiological profile of motor responses has been widely studied and is therefore well-characterized. The motor cortex of humans shows a transient increase in gamma-band activity around the time of movement onset, while beta oscillations are present during movement preparation and are suppressed during movement execution, followed by a post-movement rebound (Igarashi et al., 2013, Watanabe et al., 2014).

For the purpose of the work presented here, we chose to use a simple motor task adapted from Muthukumaraswamy et al (2010) which consists of performing abductions of the right-hand index finger cued by an auditory pip. As illustrated in Figure 23, Muthukumaraswamy et al observed a clear increase in gamma amplitude around the time of movement onset, whilst beta frequency

oscillations exhibited a distinct pattern of amplitude attenuation around movement onset, with a subsequent rebound on movement cessation. These findings are consistent with previous results from the invasive (Pfurtscheller et al., 2003, Miller et al., 2007) and electrophysiological literature (Cheyne et al., 2008, Gaetz et al., 2010).

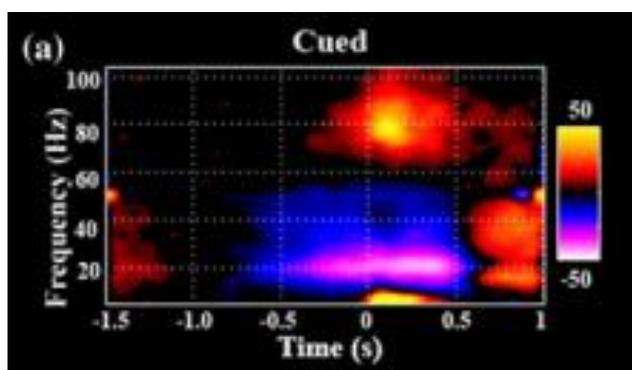


Figure 26: Time-frequency plot of the observed response when participants executed simple abductions of the right hand index finger cued by an auditory pip. Note the clear increase in gamma amplitude and the attenuation in beta amplitude around the time of movement onset, as well as the distinct beta rebound on movement cessation (Source: Muthukumaraswamy, 2010).

Oscillatory activity in the beta frequency band (15-30 Hz) is generally associated with complex behaviours in the sensorimotor domain, but is particularly prevalent in the primary motor cortex (M1). Yamawaki et al (2008) investigated the properties of M1 network oscillations in coronal slices taken from the rat brain, and reported that beta frequency activity was present in all five layers. The authors also observed that the power of beta activity increased from superficial to deep layers, and noted that this could be linked to greater numbers of pyramidal cells in deep as compared to superficial layers. In

general, they found that activity in layer V preceded activity in the more superficial layers.

Similar findings were reported by Lacey et al (2014), who used an in vitro pharmacological approach to obtain persistent oscillations in slices of the rat primary motor cortex. They not only showed that this network preferentially oscillates at beta frequency, but also concluded that these oscillations were generated in deep layers.

The available literature for gamma oscillations is more ambiguous. Using cell recordings to study populations of pyramidal neurons in layers II/III and V of the rat motor cortex, Tsubo et al (Tsubo et al., 2013) found that synaptic connections promote synchronized neuronal firing in layer II/III pyramidal neurons in the gamma band, but not in layer V, indicating that these oscillations are generated in superficial cortical layers. However, Lee and Jones (Lee and Jones, 2013) used a computational model to simulate features that would arise from the types of current dipole signals studied using MEG or EEG. While their findings do not directly contradict the theory that gamma oscillations originate in superficial layers, they found that gamma frequency rhythms were dominated by apical dendrites of layer V pyramidal cells, which obscured simultaneous independent gamma activity in layer II/III networks. Findings from the animal literature suggest that in the rat motor cortex, slow gamma oscillations are particularly strong in superficial networks, whilst fast gamma oscillations dominate in layer V (Igarashi et al., 2013).

Thus, if the differences in source reconstruction outcomes we previously observed are indeed physiologically meaningful, we would expect to find a clear

preference for the deep layer model in the case of beta oscillations, while for gamma activity, we tentatively assume a preference for the superficial layer model.

Methods and Materials

MRI Acquisition

MRI data was acquired using a Siemens Tim Trio 3T system (Erlangen, Germany). The subject lay in the supine position. The body-transmit coil was located inside the bore of the scanner for detection of the MRI signal. The MRI data was acquired using a 3D FLASH sequence for optimal scanning efficiency (Frahm et al., 1986). The following acquisition parameters were used: field-of-view: (256,256,208) mm along the (phase (A-P), read (H-F), partition (R-L)) directions, image resolution: 1mm³. The repetition time TR was set to 23.7ms and the excitation flip angle was set to 20° to yield good T1-weighted contrast, standard in most anatomical applications (Helms et al., 2008). 8 echoes were acquired following each excitation and averaged offline to produce one anatomical image with optimal signal-to-noise. Potential sources of error from the MR acquisition should be carefully considered. A high readout bandwidth was used to preserve brain morphology and no significant geometric distortions are expected in the images. Padding was used to minimize subject motion but some residual effects might remain present in the MRI images. These effects might be further reduced by use of navigator echo techniques. The total acquisition time was 21min 07s.

FreeSurfer Surface Extraction

FreeSurfer (Fischl, 2012) was used to extract cortical surfaces from the anatomical MRIs of individual subjects. FreeSurfer breaks this task down into several steps. First of all, intensity correction is performed to deal with any intensity variations due to magnetic field inhomogeneity. This results in a normalized intensity image. Extracerebral voxels are then removed using a 'skull stripping' approach. Segmentation was performed based on the geometric structure of the grey-white interface, and cutting planes are computed to separate the cortical hemispheres and disconnect the subcortical structures. A single filled volume is produced for each of the hemispheres, each of which is covered with a triangular tessellation and deformed such that an accurate, smooth representation is formed of both the white/grey matter boundary and the pial surface. For more detailed information on the algorithmic procedures, see Dale, Fischl and Sereno (Dale et al., 1999).

The above process yields surface extractions for the pial surface (the most superficial layer of the cortex adjacent to the cerebro-spinal fluid (CSF), and the white/grey matter boundary (the deepest cortical layer). Each of these surfaces is represented by a mesh comprising 21401 vertices. For the purpose of this paper, we will use these two surfaces to represent deep (white/grey interface) and superficial (grey-CSF interface) cortical models.

Task

As described above, we used a simple finger movement task, adapted from Muthukumaraswamy et al (Muthukumaraswamy, 2011). We used the same task for the work presented in chapter 3 (Troebinger et al., 2014a). However, while in chapter 3, we focused on the auditory evoked response, we were interested in motor-evoked activity in this case. In brief, the task consists of performing simple abductions of the right hand index finger, cued by a pip (1000 Hz), which is followed by an inter-stimulus interval of 3.5 – 4.5 s. EMG traces of the first dorsal interosseus (FDI) were used to track finger movements. We used 4 male, right-handed subjects for this study, with ages ranging from 19-46 (3 subjects 19-36). For three of our subjects, we used second-generation head casts ('soft' latex-based head casts), while we used a rigid, first-generation head cast for the remaining subject, since we already had a cast available for this subject from a previous study. Subjects attended 2 recording sessions, each of which lasted one hour and was split into several 10 min recordings. Due to variations in preparatory time, between 2 and 4 ten-minute recordings were acquired during each visit. As a result, the number of runs available for analysis for each of the 4 subjects ranges from 5 to 8 (ME and GB: 8 runs, ES: 6 runs, OG: 5 runs).

Source reconstruction and Multiple Sparse Priors

Source Reconstruction and basic procedure

Trials were identified based on the EMG trigger. In accordance with the preprocessing strategy used by Muthukumaraswamy et al (Muthukumaraswamy, 2011), increases in the rectified EMG signal that exceed a value of three times the standard deviation above the noise floor were used to mark movement onset . Data were epoched from 1500 ms before to 1500 ms

after EMG onset, and subsequently filtered using a low pass fifth order Butterworth filter with a cut off of 120 Hz. Data were baseline corrected, using the period of 1000 to 500 ms before movement onset. The only artefact removal was the removal of trials containing large jumps (due to loss of lock in the feedback electronics) which could be clearly seen by eye (as the subject was unable to open their eyes we had no need to correct for eyeblinks).

The resulting data were projected into 274 orthogonal spatial (lead field) modes and 8 temporal modes. We used the greedy search option of the MSP algorithm (Friston et al., 2008) implemented as outlined in Lopez et al (Lopez et al., 2012b). The MSP algorithm requires a set of covariance matrix priors corresponding to cortical patches to be defined a-priori. We used a pseudo-random (the same sequences were used when comparing different surface models) selection of 512 mesh vertices to define the patch centres and produced 32 such solutions (each based on different patch sets) per dataset. There were no symmetric priors used.

Activity was reconstructed using two different frequency bands. For beta activity, the specified frequency range was 15-30 Hz, for gamma it was between 60-90 Hz. In the case of beta activity, we used two different time windows. The first time window was chosen in the range between -500 to 500 ms around the stimulus onset, while the second was between 500 and 1500 ms after movement onset (post movement beta rebound). For activity in the gamma frequency band, we used a time window from -200 to 500 ms around the time of movement onset.

Multiple Sparse Priors (MSP) and patch size variation

Since we used the same reconstruction approach as in chapter 3, we refer the reader back to this chapter for more in-depth information on the basics of this algorithm.

MSP returns a Free Energy value which approximates the model evidence for the final generative model. Since this generative model includes the cortical surface model used as well as the lead fields, it can be utilized to compare between different models (Henson et al., 2009, Lopez et al., 2012b). Briefly, the negative variational free energy represents a trade-off between the accuracy of the model in explaining the data, and the complexity of achieving that accuracy (Penny et al., 2010).

Our goal here was two-fold: in addition to comparing cortical models based on deep/superficial laminae, we were also interested in testing the influence of patch size, a specific parameter of the MSP algorithm. The basic procedure we used for source reconstruction is outlined in Figure 24. In essence, we reconstructed data using two different cortical models, each based on Freesurfer extractions of the individual subject's pial/white matter surface, respectively. In each case, we used the MSP algorithm to perform source reconstruction, repeating this process for four different patch size values (FWHM ~ 3, 7, 10, 13 mm). Patch size (the putative extent of active cortex) is elaborated in equations 4 and 5 below.

Basically, the MSP algorithm requires a set of covariance matrix priors corresponding to cortical patches to be defined a priori. Each of these covariance priors corresponds to a single smooth cortical patch of activity (we did not use bilateral patches as in the original MSP formulation) and is therefore

determined by the location of an impulse response on the cortical surface and a local smoothness operator determining the full-width half maximum (FWHM) or spatial extent of the source. We varied the smoothness of the cortical distribution to see if under- or overestimating this parameter would tend to bias us towards deep or superficial layers.

In brief, the MEG data can be related to the neural activity that generates it using the linear model:

$$\mathbf{Y} = \mathbf{K}\mathbf{J} + \boldsymbol{\varepsilon} \quad (4.1)$$

where $\mathbf{Y} \in \mathfrak{R}^{N_c \times N_t}$ is the sensor data, where $N_c = 274$ is the number of sensors (normally 275 but one channel turned off) and N_t is the number of time samples; $\mathbf{K} \in \mathfrak{R}^{N_c \times N_d}$ is the lead field matrix that maps the N_d source locations to the N_c channels; $\mathbf{J} \in \mathfrak{R}^{N_d \times N_t}$ is the current distribution at each source location; and $\boldsymbol{\varepsilon}$ is zero mean Gaussian noise. We used a single shell (Nolte, 2003) based on the inner surface of the skull to define the forward model.

Under Gaussian assumptions, for known source covariance, \mathbf{Q} , the current distribution $\hat{\mathbf{J}}$ can be estimated directly:

$$\mathbf{J} = \mathbf{Q}\mathbf{K}^T (\boldsymbol{\Sigma}_\varepsilon + \mathbf{K}\mathbf{Q}\mathbf{K}^T)^{-1} \mathbf{Y} \quad (4.2)$$

Where T denotes a matrix transpose. Here we assume that sensor noise $\boldsymbol{\Sigma}_\varepsilon = h_0 \mathbf{I}_{N_c}$ is independent and uniformly distributed, with \mathbf{I}_{N_c} an $(N_c \times N_c)$ identity matrix and h_0 a hyperparameter effectively controlling the regularisation. Different M/EEG algorithms entail different choices of the prior source covariance \mathbf{Q} (Friston et al., 2008b; Wipf et al., 2010). For the minimum norm (MNM) solution, \mathbf{Q} is simply an $(N_d \times N_d)$ identity matrix; for the Multiple Sparse

Prior (MSP) solution, \mathbf{Q} comprises an optimised mixture of a library of N_q source covariance components $\mathbf{C} = \{C_1, \dots, C_{N_q}\}$:

$$\mathbf{Q} = \sum_{i=1}^{N_q} h_i \mathbf{C}_i \quad (4.3)$$

Each source covariance component C_i is generated as a bell-shaped smoothed region with a maximum over its centre (Harrison et al., 2007). This shape is formed with a Green's function over a graph Laplacian. The Green's function Q_G of $(N_d \times N_d)$ is defined as:

$$Q_G = e^{\sigma G_L} \quad (4.4)$$

With σ a parameter that defines the size of the bell, and G_L of $(N_d \times N_d)$ the graph Laplacian generated with the vertices and faces provided by the cortical surface:

$$G_L = \begin{cases} -\sum_{k=1}^{N_d} A_{ik}, & \text{for } i = j, \text{ with } A_i \text{ the } i\text{-th row of } A \\ A_{ij}, & \text{for } i \neq j \end{cases} \quad (4.5)$$

Where A is an adjacency matrix of $(N_d \times N_d)$, with $A_{ij} = 1$ if there is face connectivity between vertices i and j (maximum six neighbours for each vertex). For a more detailed discussion of these patches and their implementation in SPM see (Belardinelli et al., 2012, Lopez et al., 2014).

Therefore, by varying the value of σ in equation (4), we can effectively vary patch size. We already exploited this property in chapter 3, using just two values for σ , corresponding to a patch extent of 5 and 10 mm, respectively. Here, we extended this to using four different values (0.2, 0.4, 0.6, 0.8), allowing us to further investigate the impact of over- or underestimating patch size.

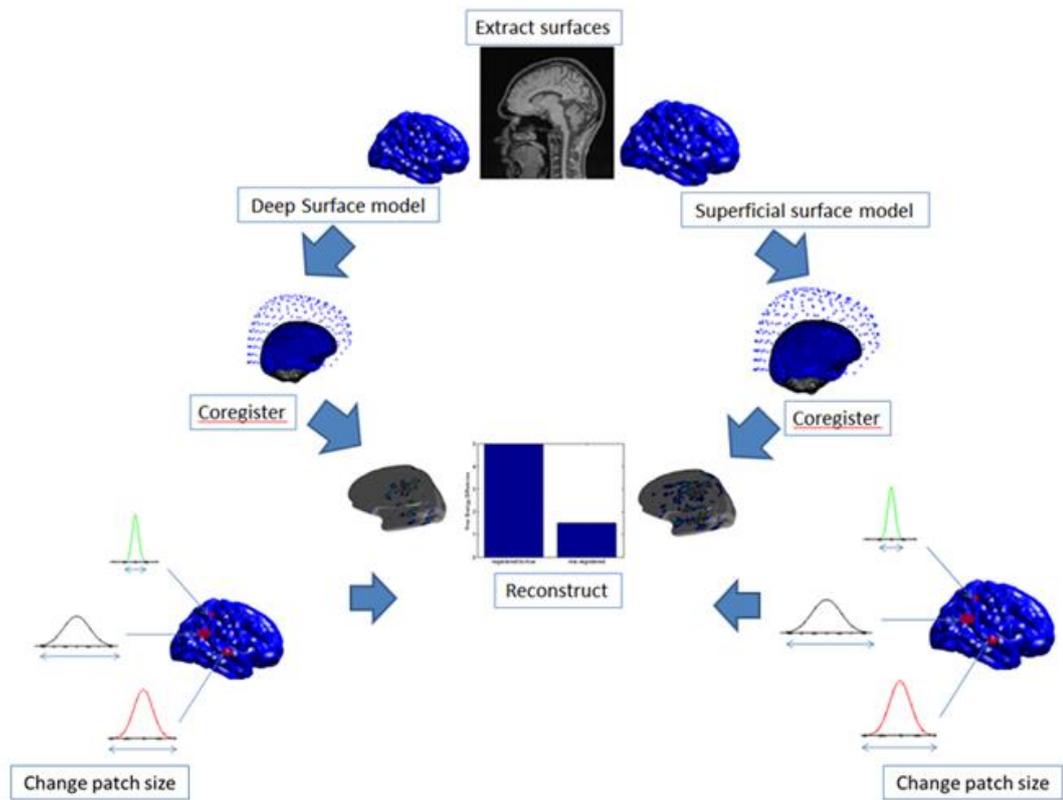


Figure 27: The strategy used for source reconstruction. Individual cortical surfaces were extracted from anatomical MRIs using Freesurfer, yielding a deep and superficial cortical model. Coregistration was performed using both models, followed by source reconstruction, yielding a free energy value as an outcome. We then repeated the process for different patch extents.

To link this investigation with our previous work, there were two main questions to consider. In the second chapter, we used both simulated and real MEG data to probe whether it is possible to distinguish between cortical models representing deep/superficial laminae, based on Free Energy following source reconstruction. We were able to answer this question in the affirmative, but our investigation was hampered by the fact that we only had MEG data from a single subject. Here, we address the same basic questions, but using real MEG data recorded from multiple subjects (4).

Secondly, we previously highlighted the fact that MSP patch size has a significant impact on source reconstruction outcomes: in the second chapter, we showed that both over- as well as under-estimating patch size biases layer specific inference. However, in chapter 3, we only looked at two different patch extents, corresponding to FWHM = 5/10 mm, respectively. Hence, the second question we ask here concerns patch size. In essence, we wanted to be able to factor out the influence of patch size and also provide a quantitative statement on the most likely patch size given the MEG data.

The basic procedure we followed is outlined in

Figure 27. For each subject, we used two cortical models, based on FreeSurfer extractions of the pial/white matter surfaces for each individual. The pial surface model was used to represent superficial layers, while the white matter surface model was used to represent deep layers.

Each available dataset was coregistered to both models. We used the MSP algorithm to perform source reconstruction on the resulting datasets, which yielded a Free Energy value for each case. Crucially, although we based the 32

MSP iterations for each dataset/cortical model on a set of 512 randomly selected patches, the same set of patches was used for deep/superficial model constructions for a given iteration/cortical model/patch size/subject, allowing us to compare source reconstruction outcomes for different models.

Investigating patch size, we followed the same basic guidelines. Again, datasets were coregistered to both models in turn. The only difference here was that in addition to having two different cortical models, we repeated the source reconstruction process for each of the four smoothness values.

Results

Comparing outcomes based on Free Energy

Here, we compare results in terms of Free Energy outcomes. For each subject and each available dataset (each representing one 10-minute run), we selected the maximum Free Energy value (F_{max}) from the 32 available values (based on 32 MSP iterations). Subtracting the Free Energy values for deep and superficial surfaces then allows us to evaluate whether one cortical model outperformed the other: the larger (ie, the more positive) the Free Energy, the better the performance of the model in question. We repeat this process for the four different patch extents and for both frequency bands.

For each subject, we had a different number of runs available. Each subject attended two or more sessions of recording, each of which was split into separate ten-minute runs. Due to differing setup times and number of visits, the overall number of runs varied between subjects. Hence, the bar graphs below show a differing number of bars for each patch size, each bar corresponding to the Free Energy difference between the deep and superficial cortical models for

a single ten-minute run. A positive value indicates that the deep model was preferred, whilst a negative value means that the superficial model outperformed the deep one.

First of all, we look at Free Energy outcomes for each subject individually.

OG

Figure 28 illustrates the results for subject OG. Each of the four 'groups' of bars represents outcomes for a specific patch size, with the leftmost group representing results for the smallest patch extent used, increasing in size as we move along the x-axis. Each group consists of five bars, corresponding to the five ten-minute datasets available for this subject. The y axis shows the differences in Free Energy; a positive value indicates a preference for the deep layer model, whilst a negative value suggests a preference in favour of the superficial model.

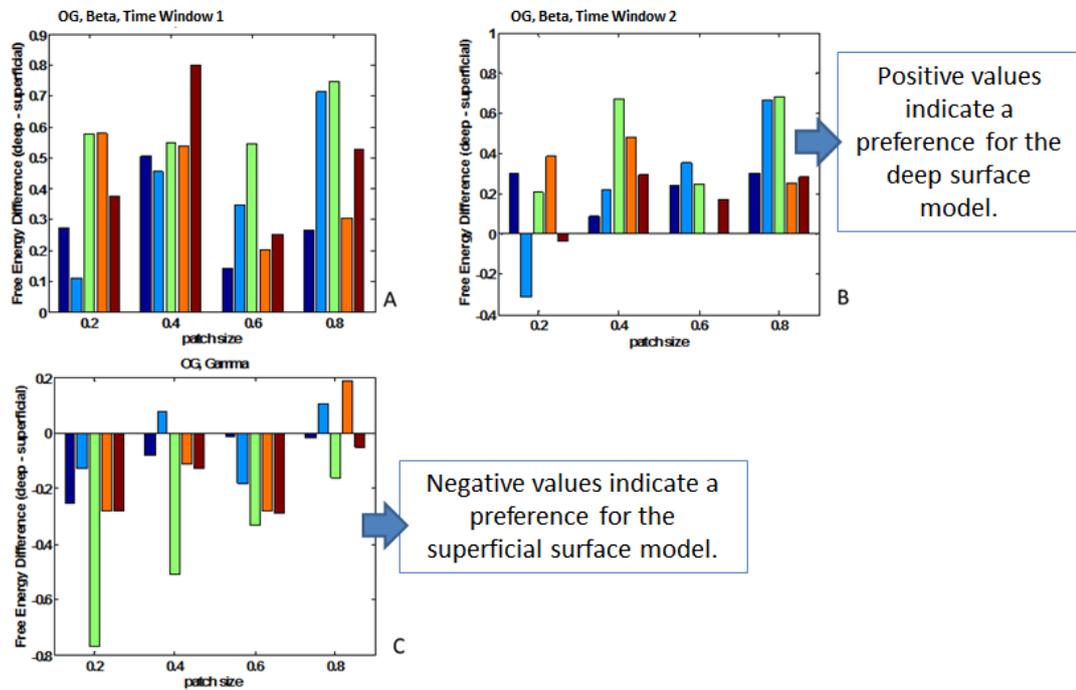


Figure 28: Results for subject OG, illustrated in terms of Free Energy differences. Each group of bars shows results for a particular patch size, with each individual bar corresponding to a single ten-minute-recording session. Positive values indicate a preference for the deep layer, negative values a preference for the superficial layer model. Panels A and B show results for Beta activity in time windows 1 (-500 to 500 ms) and 2 (500 to 1500 ms), respectively, Panel C shows results for gamma activity. In all cases, results were consistent with our hypothesis.

What is striking here is that the results seem largely in favour of our hypothesis: for beta activity in both time windows (Time window 1: -500 to 500 ms – see panel A; Time window 2: 500 to 1500 ms – see Panel B), the deep surface model outperformed the superficial model, regardless of patch extent. For gamma activity, the opposite was the case.

Looking at patch size, the results are most striking for gamma activity: as patch size increases, the initially overwhelmingly negative Free Energy differences decrease, eventually ‘flipping’ to positive values. As noted above, negative Free Energy values suggest a preference for the superficial model, whilst positive values mean that the deep model was favoured. Thus, the effect we observe due to increasing patch extent seems to indicate that by increasing patch size, we are also introducing a bias towards deep surfaces. Figure 29 shows the sum of Free Energy differences over all runs in each of the time window to give an overall impression of the effect of patch size.

For Beta activity, the effect of patch size is less obvious – however, in the second time window, Free Energy differences seem to increase with increasing patch size, again suggesting a bias towards deeper layers with increasing patch size. For beta activity in the first time window on the other hand, we did not observe a consistent tendency.

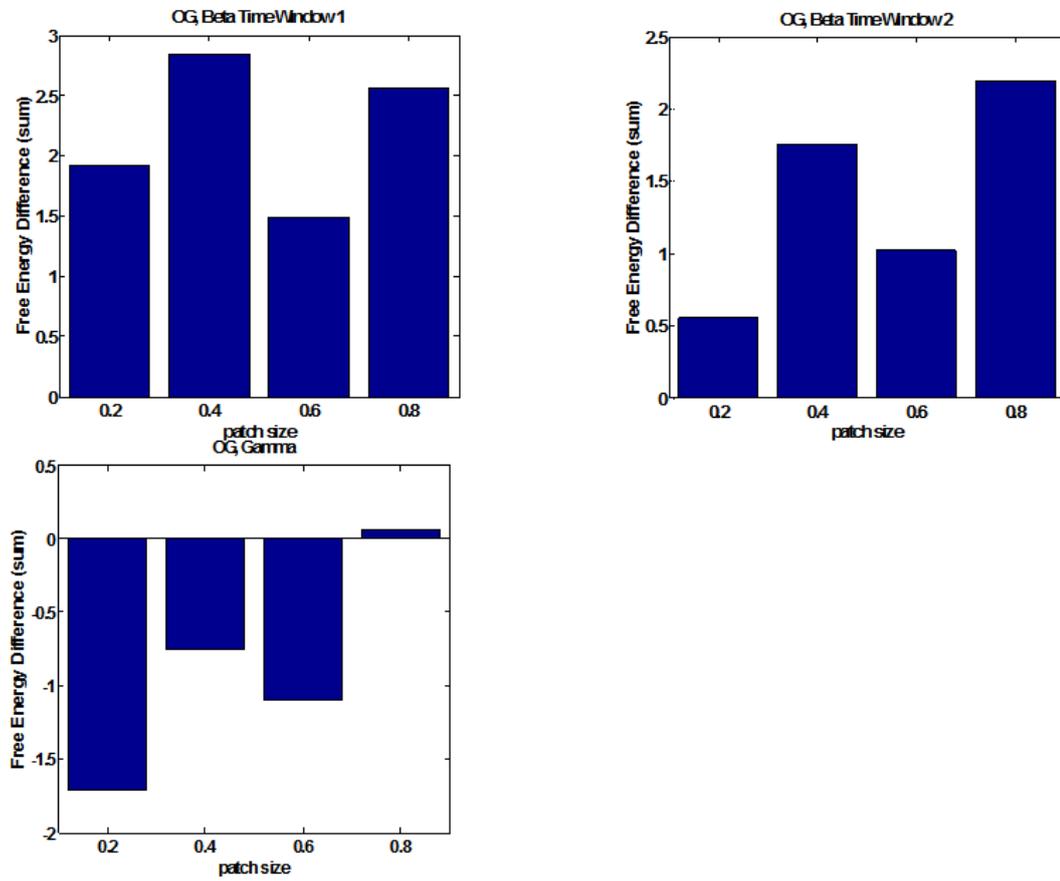


Figure 29: Illustrating effect of varying patch size. Each bar represents the sum of Free Energy Differences across all runs at each patch size.

ME

We had 8 ten-minute runs of data available for this subject. Looking at Free Energy differences, we once more see a preference for the deep layer model in the case of beta band activity in both time windows (Figure 30, Panels A and B), while the superficial model is favoured in the case of gamma activity. (Panel C).

Figure 31 illustrates the same results, but showing the information in a slightly different way: Here, we took the sum of Free Energy differences across all runs for each patch size in turn, giving a single bar for each size. While there seems to be a clear effect of patch size for beta activity in time window 2 – as patch size increases, Free Energy differences decrease – we did not observe a similarly clear and consistent trend for gamma activity and beta activity in the first time window.

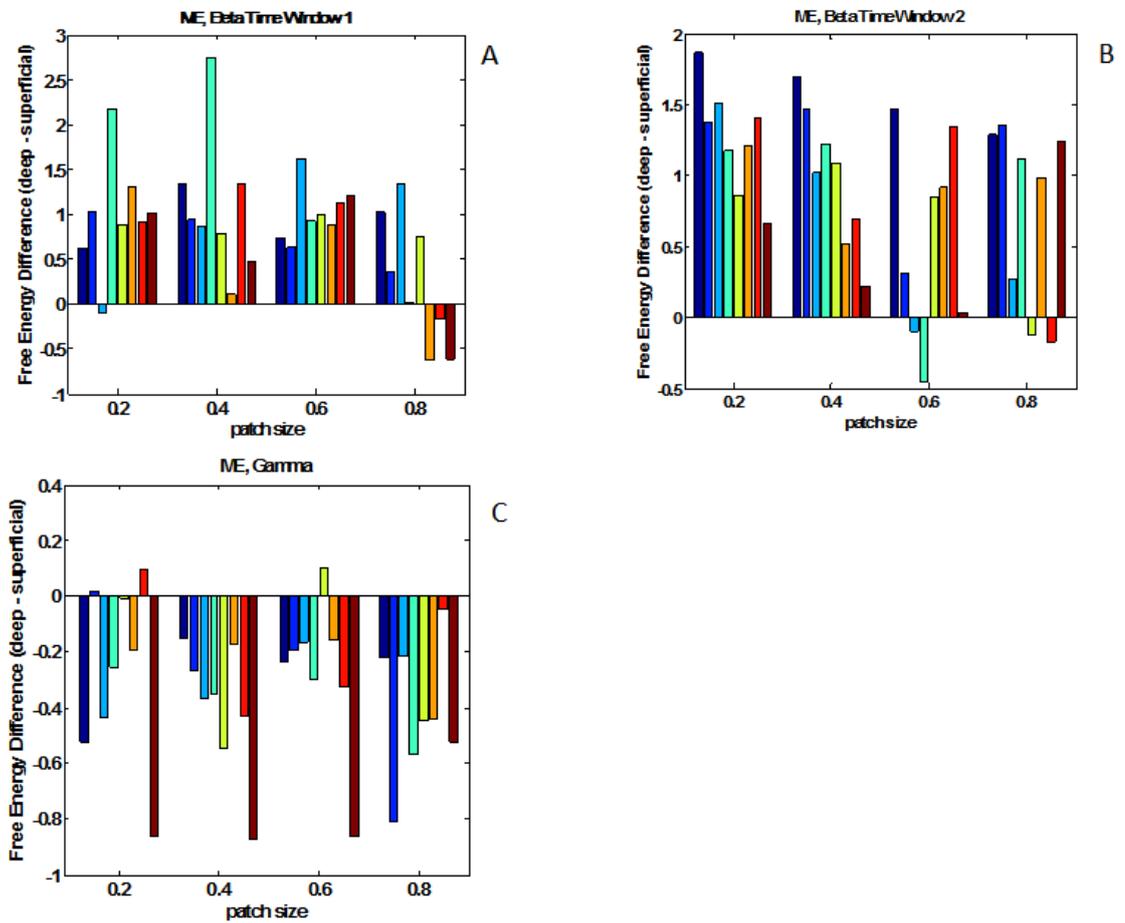


Figure 30: Outcomes in terms of Free Energy differences, subject ME. Positive values indicate a preference for the deep, negative values a preference for the superficial surface model. Again, results were consistent with our pre-formed hypothesis.

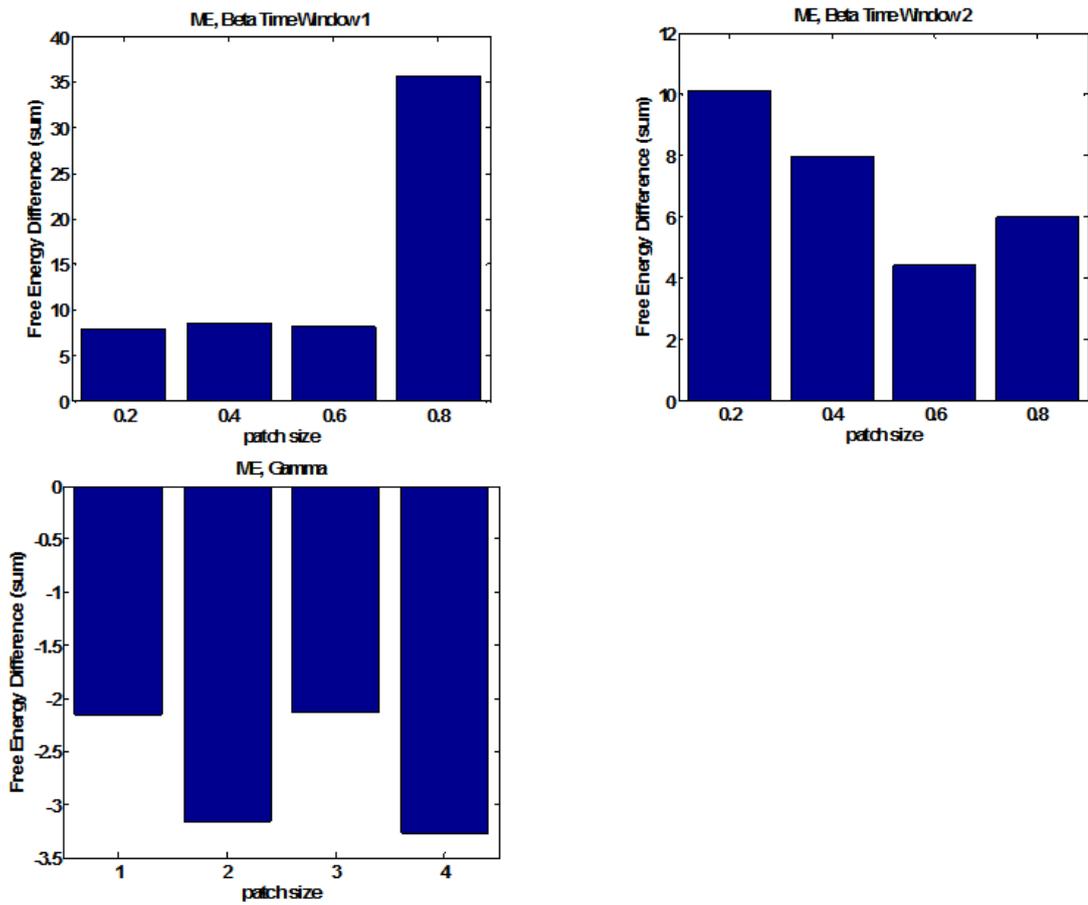


Figure 31: ME, illustrating the effect of varying patch extent. Bars represent the sum across Free Energy differences for all available runs at each patch size.

ES

There were 6 ten-minute runs available for this subject. Comparing Free Energy differences (Figure 32), we find the same preference for the deep layer model for beta band activity that we observed for the previous two subjects, as well as the same preference for the superficial layer model in the case of gamma frequency activity.

There also appears to be a clear effect of patch size here: As patch size increases, the absolute value of Free Energy differences increases (Figure 33).

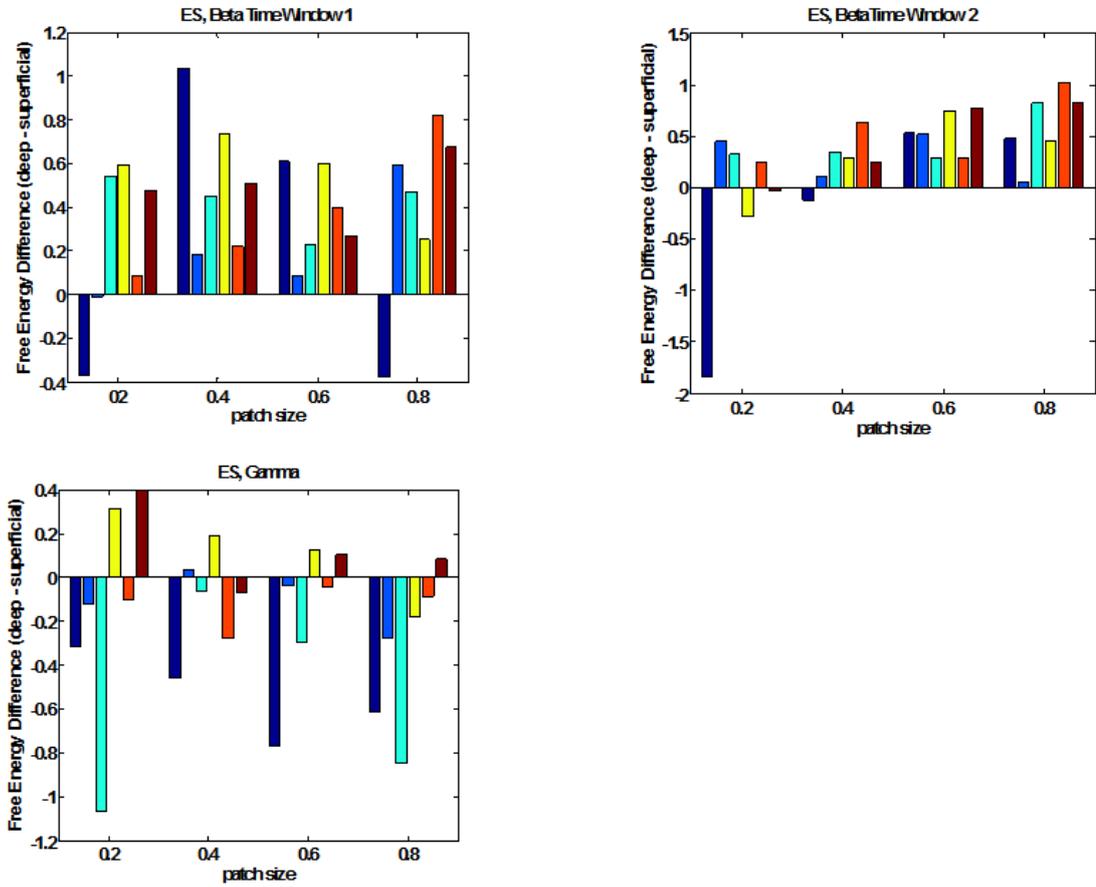


Figure 32: Subject ES, outcomes illustrated in terms of Free Energy differences.

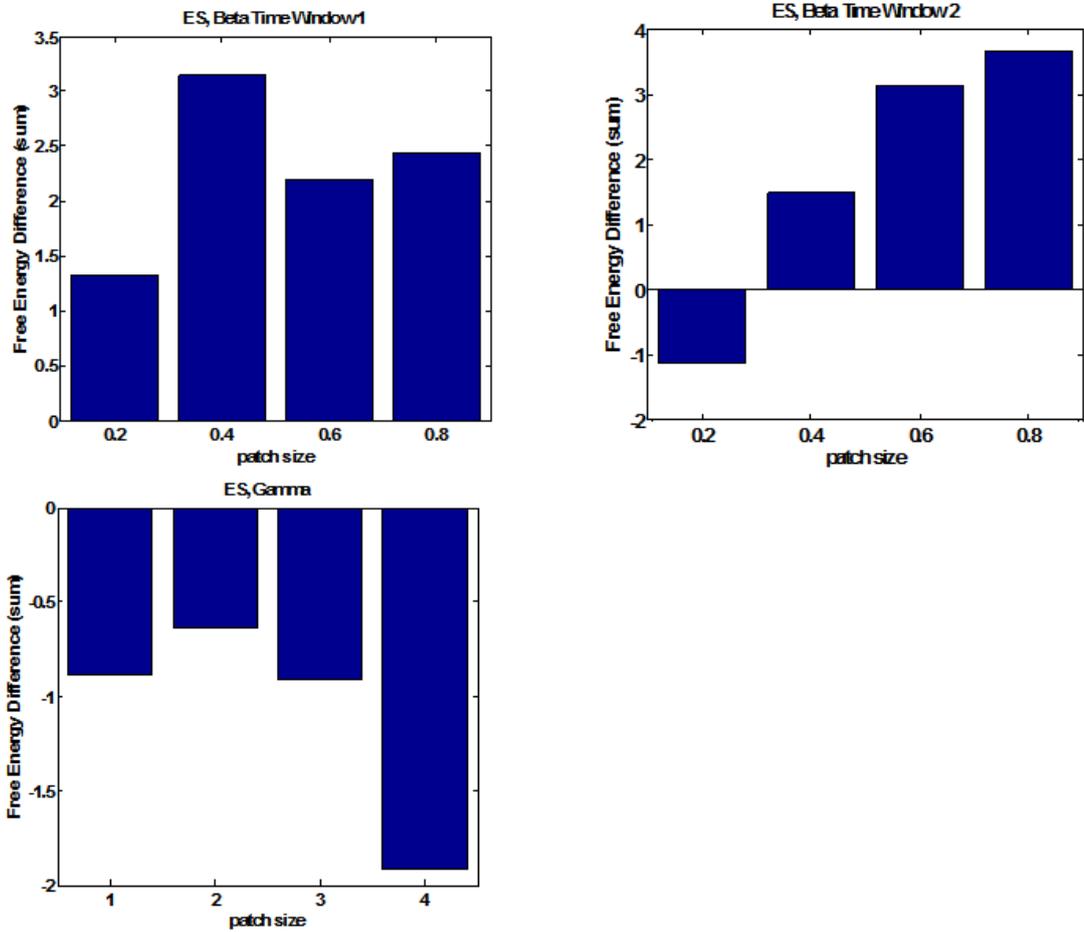


Figure 33: ES, illustrating the effect of patch size. The pattern seems to indicate that as patch size increases, the absolute value of Free Energy differences increases, introducing a surface-specific bias.

GB

There were 8 ten-minute runs available for this subject. However, contrary to the previous three subjects, we did not observe the predicted preference for the deep surface model for beta activity. Here, there seems to be a clear preference for the superficial surface model across all frequency bands and time windows (Figure 34).

As far as patch size is concerned, we again observe the effect that as patch size increases, so does the absolute value of Free Energy Differences (Figure 35).

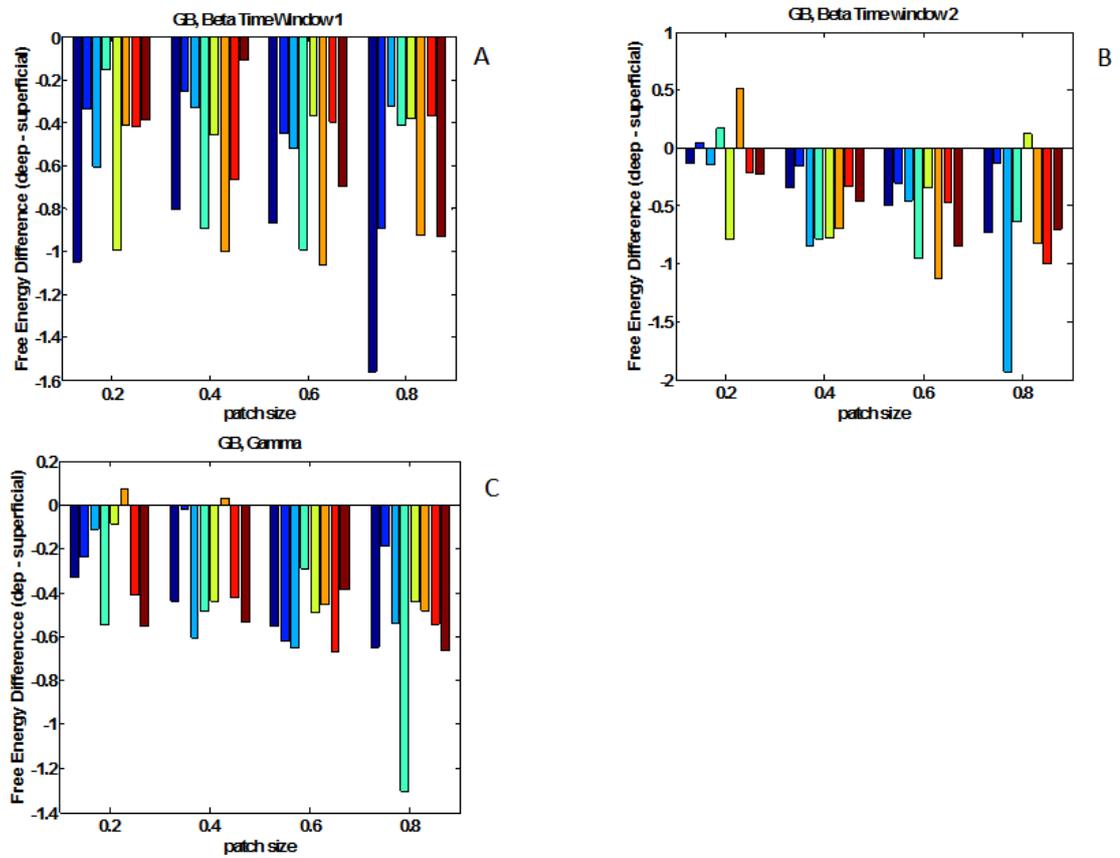


Figure 34: GB, outcomes illustrated in terms of Free Energy differences. This was the only subject whose results did not agree with our a-priori hypothesis. However, we used a 'first-generation' type head cast for GB, with a slightly different coregistration scheme, which might be a possible reason for the striking differences.

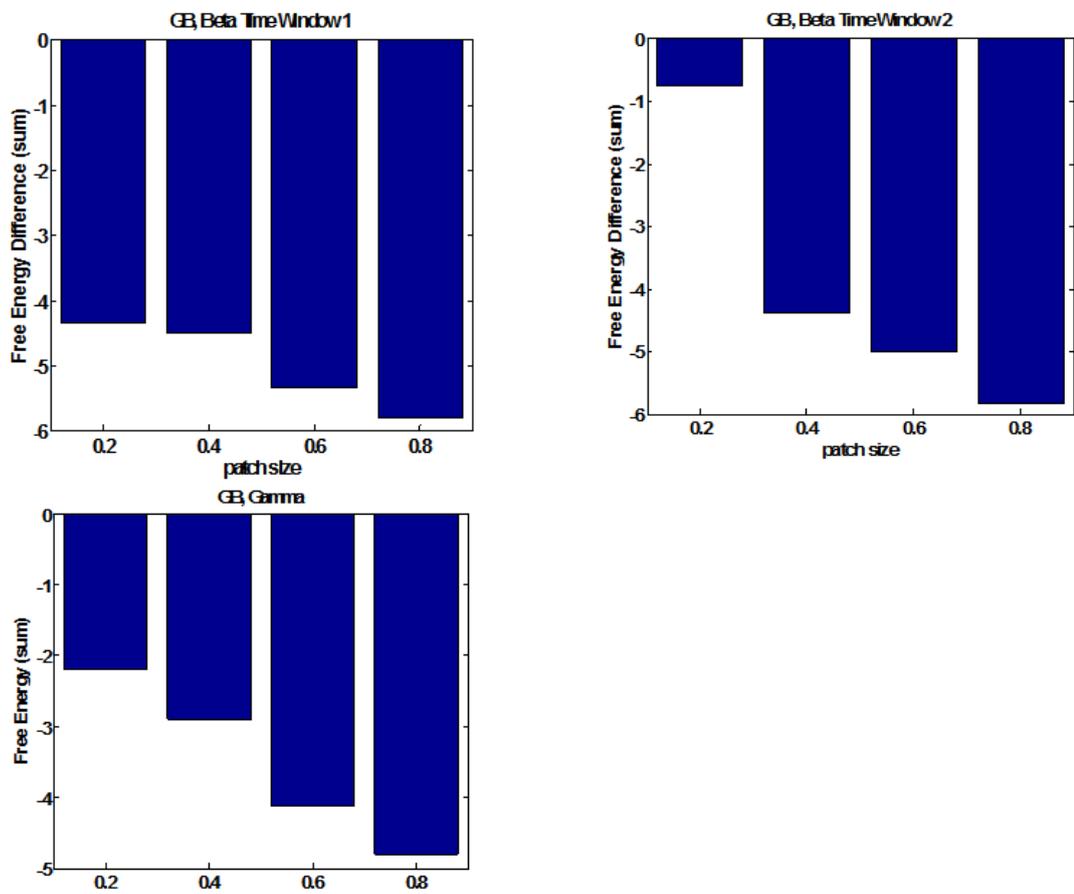


Figure 35: GB, illustrating the effect of patch size. A clear, consistent pattern is visible - as patch extent increases, so does the absolute value of Free Energy Differences.

Discussion

The aim of this chapter was to broaden the scope of previous work, and to establish the nature of the effects we observed in previous investigations, in particular the physiological significance – or lack thereof – of differences in source reconstruction outcomes based on cortical models which represent different cortical layers.

As described in the introduction to this paper, there exists a large body of literature – based on animal studies and cell recordings – which focuses on cortical information flow and the laminar origin of sensory evoked responses. To test the physiological significance of our results, we established an a-priori hypothesis based on the findings of these studies.

The fact that three of the four subjects we studied showed effects which were consistent with this a priori hypothesis – a preference for the deep surface model for beta band activity and a preference for the superficial model in the case of gamma activity – is extremely encouraging and suggests that the differences in source reconstruction outcomes are indeed due to physiological effects, rather than e.g. algorithmic shortcomings.

The fact that the effects we observed for one of our subjects were not consistent with this hypothesis is perhaps not entirely discouraging either. There are several reasons for this. On the one hand, although the available literature on information flow within the motor cortex is by no means sparse, the majority of these studies are based on invasive techniques, which do not share the same limitations as MEG.

The cellular makeup of the MEG signal and the composition of information flow within the cortex represent another unknown: We may be able to develop clear models based on findings from invasive studies, but since there is uncertainty regarding the contributions of different types of neurons to the MEG signal, we do not yet fully understand whether this introduces a bias towards a particular layer model. Lastly, we also have to take into account all aspects of the head cast technique - from MRI surface extraction to the resolution available from even the most state-of-the art 3D printers, and the possible influence of such factors on head cast fit and, as a consequence, coregistration error. This is an especially important point, since we essentially used two different types of head casts for this study: while we used latex-foam-based 'second generation' head casts for subjects OG, ES and ME, the head cast we used for subject GB was a 'first generation' type head cast – in fact, the very first one we made. While these first generation head casts were made from rigid material, the second generation head casts were based on a slightly different production technique. Although less rigid, the design allows for easier – and perhaps better – placement, offering greater control and less safety concerns. In addition, place holders for fiducial coils were integrated into the head casts. The locations were based on points chosen from the MRI-extracted image of each individual subject's head, allowing for highly accurate fiducial placement. The strategy we used for first generation head casts on the other hand was based on MRI visible markers and calculating transformation matrices to link the different reference systems. Since we have consistently found that coregistration error is crucial in determining spatial resolution, it would not be unreasonable to assume that the

differences in outcomes might be linked to the different types of head cast design/coregistration scheme.

Lastly, we have to consider another important point: Typically, when we talk about coregistration error, we focus on the scalp rather than the actual cortical anatomy. This is mainly due to the fact that traditionally, coregistration was accomplished by placing markers directly on the scalp. However, since the brain is suspended in fluid, its position within the skull is not fixed. This is particularly important in the context of MEG studies, since subjects are usually positioned sitting upright in the scanning system's chair. In contrast, MRI studies – including the acquisition of the type of structural MRIs which form the basis of cortical surface extraction for the purpose of head cast design and coregistration – usually require participants to assume a supine position. Thus, the position of the cortical surface is likely to differ between these two imaging modes, possibly introducing a significant source of error. In other words, by integrating fixed fiducials into the head cast design and relating them to MRI anatomy, we may have succeeded in getting a lot closer to knowing where the brain is than other fiducial-based coregistration schemes, but perhaps this has merely allowed us to highlight that we need get even closer.

We also have to consider aspects which are not directly connected to the head cast technique, such as individual variations in cortical anatomy and age-related differences. There was a considerable age difference between subject GB and the three remaining subjects – it is not unreasonable to assume that this age difference could have contributed to the striking differences in analysis outcomes. Age-related differences in cortical thickness have been widely studied (Resnick et al., 2003, Rettmann et al., 2006, Thambisetty et al., 2010) –

for instance, Thambisetty et al studied longitudinal changes in the cortical thickness of 66 older adults, and concluded that age-related decline in cortical thickness is widespread, with frontal and parietal regions showing greater rates of decline than temporal and occipital (Thambisetty et al., 2010). For the purpose of the data presented here, we did not specifically look at differences in cortical thickness between subjects – however, perhaps one way to further test the validity of the hypotheses presented here would be to restrict the age group subjects are chosen from.

Of course, it may also be the case that our models of the laminar origin of the different oscillatory responses we studied here are not yet sophisticated enough to really allow us to judge whether our findings are physiologically meaningful. In particular, Lee and Jones (Lee and Jones, 2013) used a computational model of a laminar neocortical network to investigate the underlying neural mechanisms of two distinct types of gamma generation. In their investigation, they focused on modelling the cellular level biophysics of macroscopic current dipole signals – i.e. the type of signals recorded in MEG studies. Interestingly, they found that the net current dipole signals were dominated by the gamma frequency activity in Layer V, obscuring any other activity in Layers II/III. They further concluded that signals from Layer II/III networks could only be observed if the network was of considerable size, as it would otherwise be obscured by simultaneous activity in deeper layers (in the same frequency band).

In the same paper, Lee et al mention the distinction between evoked and induced gamma activity. Tallon-Baudry et al (Tallon-Baudry et al., 1997, Tallon-Baudry, 2009) suggested that evoked gamma activity is generated by pyramidal neurons in layer V, while induced gamma activity stems from radially oriented

sources in Layers II/III – a further point to consider when refining experimental setups designed to study the laminar origin of gamma activity.

Nevertheless, the results we observed are extremely encouraging. The fact that three of four subjects showed the types of effects predicted by our hypothesis indicates that MEG is indeed capable of excellent spatial resolution, and that the use of the head cast technique allows us to make use of this property.

Another important point highlighted by the work presented here concerns the bias potentially introduced by under- or overestimating patch extent. Here, we aimed to build on previous work, in which we used a single subject and two different values for patch extent to establish whether this parameter has a significant effect on source reconstruction outcomes. Here, our results were ambiguous; while variations in patch extent had an effect on source reconstruction outcomes, its nature was less obvious. We therefore believe that this particular point needs further investigation, and that perhaps the development of a scheme which allows researchers to ‘test’ different patch extents in the same way that MSP allows for evaluating different patch locations may be necessary.

However, the work presented here also highlights the fact that we are only just getting started: many unknowns remain, from the physiological basis of the MEG signal, to algorithmic shortcomings of computational methods which were designed with low SNR data in mind.

5. Conclusions

In the introduction to this thesis, I provided an overview of Magnetoencephalography (MEG), including instrumentation, scanning protocol and analysis techniques, as well as discussing the physiological basis of the MEG signal. I concluded that while MEG benefits from being non-invasive, offering excellent temporal resolution and representing a direct measure of neural activity, its spatial resolution is limited by non-physiological factors associated with the scanning protocol: Coregistration error introduces uncertainty about the location of the cortex relative to the sensors, and within-session head movement leads to blurring of the sensor level data, resulting in low Signal-to-Noise Ratio (SNR) and poor spatial resolution.

These limitations provided the starting point for the work presented in this thesis. Our goal was two-fold: First of all, we aimed to develop a helmet-like device to fit the individual subject, which would allow us to immobilize the head during scanning, and which could also be exploited for coregistration purposes. Secondly, we wanted to establish the bounds on MEG spatial resolution given the availability of high SNR data. In other words, we wanted to test the hypothesis that in the absence of coregistration error and within session head movement, MEG is capable of delivering excellent spatial resolution. We further wanted to investigate whether given such conditions, spatial resolution could be sufficient to distinguish between sources on different cortical layers.

Here, I will present a final summary of our findings, and I will discuss their implications in terms of future directions.

Chapter 2

The second chapter introduced the concept of using head casts for MEG scanning. The head casts we developed are helmet-like devices, which are made to fit the individual subject's head internally, and the MEG system's dewar externally. Combining optical scanning and structural MRI scans of the individual subject, we used a combination of CAD and 3D printing techniques to produce the head casts. We showed that these casts enabled us to reduce within session head movement to sub-millimetre levels. In addition, we observed a reduction in coregistration error to levels of less than 1.5 mm.

We used the head casts to record data from the same subject over a 6-month-period, and showed that given the availability of high SNR data, MEG is much more sensitive in the spatial sense than previously assumed. We demonstrated this by testing whether source reconstruction outcomes differ when basing them on different cortical models. This was prompted by the generally accepted view that no clear benefit is derived from using cortical models based on the individual subject's anatomy rather than the canonical model, which is simply a 'template' warped into the subject's space (Mattout et al., 2007, Henson et al., 2009).

Using the same approach as Henson et al, we performed source reconstruction based on both the canonical and an individual cortical model. We observed clear differences in the outcomes, with the individual cortical model emerging as superior to the canonical one. However, using simulations to add different amounts of coregistration noise, we also showed that the ability to distinguish between these cortical models is destroyed for levels of coregistration error of 5

mm and above – values we would typically expect when using conventional coregistration and recording techniques.

This is an important finding since on the one hand, it shows that MEG is capable of much better spatial resolution than previously thought, whilst illustrating on the other hand that the choice of cortical model and analysis strategy depends on the quality of the available data. In other words, if we don't know where the brain is, there is little benefit in knowing exactly what it looks like.

Chapter 3

In the third chapter, we aimed to build on the work described above by probing the theoretical limits of MEG spatial resolution given the absence of coregistration error and the availability of high SNR data sets. Spatial resolution in MEG is not limited by physiological factors – thus, it should in theory be possible to distinguish between sources on different cortical layers. Since the neocortex, which consists of up to 6 layers, is between 2 and 4 mm thick, spatial resolution would have to be in the mm-range to successfully distinguish between them. First of all, we used simulations and synthetic data sets to test this hypothesis. Using two different cortical models, which were both based on the individual subject's cortical anatomy, we employed a similar approach as in the first chapter to compare source reconstruction outcomes. We found that for levels of coregistration error of less than 2mm translation/2 degrees rotation, we could successfully distinguish between these models. We further showed that under- or overestimating patch extent – a parameter associated with the

Multiple Sparse Priors (MSP) algorithm for MEG source reconstruction - biased layer estimates.

We then used data recorded from a single subject using an auditory paradigm to show that the ability to distinguish between cortical models representing deep and superficial laminae is preserved even in the context of real MEG data. While we observed a clear preference for the deep layer model, we concluded that the question of whether this is indicative of an underlying physiological mechanism remains open, and that further studies using multiple subjects are needed.

Chapter 4

In Chapter 4, we built on the work presented in the preceding chapters to corroborate previous findings in the context of data recorded from multiple subjects and using a simple motor paradigm. We chose a task with a well-characterized response in order to minimize uncertainty, and to establish a hypothesis about the laminar origin of this activity. We then used the same approach as in Chapter 3 – that is, carrying out the source reconstruction process using two different models, representing deep and superficial laminae in the individual subject. We were interested in several factors: First of all, whether we would be able to distinguish between these models in each of our subjects; secondly, whether we would see a preference for the same model in all of our subjects and lastly, whether this preference would conform to an a-priori hypothesis regarding the laminar origin of the observed activity. In addition, we also tested a range of MSP patch sizes to characterize the nature of its effect on source reconstruction outcomes.

We adapted a task from Muthukumaraswamy et al (Muthukumaraswamy, 2011), providing us with a well characterized response. We focused on activity in the beta- and gamma-bands, using two different time windows for beta oscillations to capture the effects of pre-movement attenuation and post-movement rebound. Based on the invasive literature, we expected a preference for the deep surface model for beta band activity (Yamawaki et al., 2008, Lacey et al., 2014). For gamma activity, the available evidence is slightly more equivocal (Tallon-Baudry, 2009, Lee and Jones, 2013), but overall suggests a preference for the superficial surface model (Tsubo et al., 2013).

Comparing outcomes in terms of Free Energy Differences, we found that three out of four subjects showed layer preferences which were in agreement with our hypothesis – namely, a preference for the deep surface model for beta band activity, and a preference for the superficial surface model in the case of gamma activity. The fact that the fourth subject did not exhibit the predicted layer preferences was not completely discouraging either since this was the only subject with a ‘first-generation’ head cast, meaning that the cast did not have integrated place holders for MEG fiducial coils. Perhaps this provides further evidence suggesting that coregistration error is crucial in determining the spatial resolution of MEG data.

Our results also suggested that patch size undeniably has an impact on source reconstruction outcomes. Further investigation will be needed to determine whether there is such a thing as an ‘optimal’ patch size, or whether this parameter has to be adjusted and fine tuned for each study and subject.

Future directions

The nature of the work presented in this thesis is highly exploratory. We have never before had access to high SNR data from MEG recordings and, as a consequence, we were dealing with many unknowns.

For instance, the available analysis tools – excellent as they may be – have been designed with low SNR data in mind, or have at least only been used only on such data so far. Therefore, it is possible that there are certain characteristics and flaws which have previously gone unnoticed, and which only emerge when we use them on high SNR data.

Here, we have specifically probed the impact of MSP patch size on source reconstruction outcomes. While our results showed that patch size variations influence source reconstruction outcomes, it has also generated more questions: Is there such a thing as an optimal patch size? Does this vary depending on the type of experimental paradigm or even the cortical anatomy of the individual subject – in other words, is the current ‘one patch fits all’-approach valid or not? The MSP algorithm in its current incarnation aims to establish the location of underlying sources by using dedicated search algorithms to find the patch or combination of patches which maximizes model evidence (Friston et al., 2008). Perhaps it is not only the number and location of patches, but also their spatial extent which needs optimizing.

Physiological aspects represent another big unknown in this equation. The basis of the MEG signal and the laminar nature of responses are not undisputed, making it difficult to present a convincing case for the physiological significance of our findings. How can we tell whether what we are seeing is

physiologically meaningful? We need a firm, a-priori hypothesis about the laminar pattern of information flow underlying the observed activity. However, the invasive literature is often ambiguous. For instance, in the case of auditory evoked activity of the type we studied in Chapter 3, several conflicting theories exist. Some studies report responses of similar magnitude across all layers (Ogawa and Inui, 2011), whilst others report a distinct deep-layer-bias (Sakata and Harris, 2009, Harris et al., 2011, Profant et al., 2013). Similarly, there are conflicting findings regarding the laminar pattern of gamma band activity evoked by motor tasks of the type we used for the work presented in chapter 3. Here, the physiological basis of the MEG signal may be of particular significance. Lee and Jones (Lee and Jones, 2013) used computational models to show that while there is evidence for gamma band activity in the superficial layers of the motor cortex (Layers II/III), this activity is likely to be 'obscured' if there is simultaneous activity from layer V pyramidal neurons. In addition, Tallon-Baudry et al (Tallon-Baudry, 2009) noted that induced and evoked gamma responses may have different laminar origins, with induced activity originating in the superficial layers, while evoked activity is generated in layer V. On the other hand, Lee and Jones reported synchronized firing in the gamma band in Layer II/III pyramidal neurons of the rat motor cortex, but not in the deeper layers.

This, however, is only one part of the problem. When we attempt to draw conclusions about the laminar origin of the observed activity, we need to consider that we are 'filtering' it through the lens of MEG – in other words, we need to also consider the populations of neurons which contribute to the MEG signal, and how their contributions are weighted. Again, we are facing the problem of an outdated view of MEG spatial resolution: Current theories about

the origin of the MEG signal are based on assuming low spatial sensitivity and low SNR.

Nevertheless, the work presented in this thesis is extremely encouraging. Particularly the fact that we observed consistent layer preferences in a majority of our subjects, and that these layer preferences were in agreement with the hypothesis we formed a priori, suggests that in spite of the many unknowns, we are on the right track.

But what we have achieved here is valuable, even if we disregard the question of whether or not we will be able to use MEG to answer questions about lamina-specific information flow and processing: The head casts are evolving into even more sophisticated designs, allowing their use in a wider range of studies, and have been gaining in popularity across several MEG laboratories. The early successes that are documented here show that MEG is capable of much greater spatial resolution than previously thought – something that will hopefully help introduce this technique to a much wider audience, opening up the scope of projects and experiments which MEG can be applied to.

References

- Adjamian P, Barnes GR, Hillebrand A, Holliday IE, Singh KD, Furlong PL, Harrington E, Barclay CW, Route PJ (2004) Co-registration of magnetoencephalography with magnetic resonance imaging using bite-bar-based fiducials and surface-matching. *Clin Neurophysiol* 115:691-698.
- Ahlfors SP, Han J, Belliveau JW, Hamalainen MS (2010) Sensitivity of MEG and EEG to source orientation. *Brain Topogr* 23:227-232.
- Akhtari M, Bryant HC, Mamelak AN, Flynn ER, Heller L, Shih JJ, Mandelkern M, Matlachov A, Ranken DM, Best ED, DiMauro MA, Lee RR, Sutherling WW (2002) Conductivities of three-layer live human skull. *Brain Topogr* 14:151-167.
- Attal Y, Schwartz D (2013) Assessment of subcortical source localization using deep brain activity imaging model with minimum norm operators: a MEG study. *PLoS One* 8:e59856.
- Balderston NL, Schultz DH, Baillet S, Helmstetter FJ (2013) How to detect amygdala activity with magnetoencephalography using source imaging. *J Vis Exp*.
- Barnes GR, Litvak V, Brookes MJ, Friston KJ (2011) Controlling false positive rates in mass-multivariate tests for electromagnetic responses. *Neuroimage* 56:1072-1081.
- Bastos AM, Usrey WM, Adams RA, Mangun GR, Fries P, Friston KJ (2012) Canonical microcircuits for predictive coding. *Neuron* 76:695-711.
- Belardinelli P, Ortiz E, Barnes G, Noppeney U, Preissl H (2012) Source reconstruction accuracy of MEG and EEG Bayesian inversion approaches. *PLoS One* 7:e51985.
- Brette R, Goodman DF (2012) Simulating spiking neural networks on GPU. *Network* 23:167-182.
- Cheyne D, Bells S, Ferrari P, Gaetz W, Bostan AC (2008) Self-paced movements induce high-frequency gamma oscillations in primary motor cortex. *Neuroimage* 42:332-342.
- Cottereau B, Jerbi K, Baillet S (2007) Multiresolution imaging of MEG cortical sources using an explicit piecewise model. *Neuroimage* 38:439-451.
- Cuffin BN (1993) Effects of local variations in skull and scalp thickness on EEG's and MEG's. *IEEE Trans Biomed Eng* 40:42-48.
- Dale AM, Fischl B, Sereno MI (1999) Cortical surface-based analysis. I. Segmentation and surface reconstruction. *Neuroimage* 9:179-194.
- Dannhauer M, Lanfer B, Wolters CH, Knosche TR (2011) Modeling of the human skull in EEG source analysis. *Hum Brain Mapp* 32:1383-1399.
- de Munck JC, Verbunt JP, Van't Ent D, Van Dijk BW (2001) The use of an MEG device as 3D digitizer and motion monitoring system. *Phys Med Biol* 46:2041-2052.
- Destexhe A, Bedard C (2012) Do neurons generate monopolar current sources? *J Neurophysiol* 108:953-955.
- Dick F, Tierney AT, Lutti A, Josephs O, Sereno MI, Weiskopf N (2012) In vivo functional and myeloarchitectonic mapping of human primary auditory areas. *J Neurosci* 32:16095-16105.
- Dogdas B, Shattuck DW, Leahy RM (2005) Segmentation of skull and scalp in 3-D human MRI using mathematical morphology. *Hum Brain Mapp* 26:273-285.
- Feldmeyer D, Lubke J, Sakmann B (2006) Efficacy and connectivity of intracolumnar pairs of layer 2/3 pyramidal cells in the barrel cortex of juvenile rats. *J Physiol* 575:583-602.
- Fischl B (2012) FreeSurfer. *Neuroimage* 62:774-781.
- Frahm J, Haase A, Matthaei D (1986) Rapid three-dimensional MR imaging using the FLASH technique. *J Comput Assist Tomogr* 10:363-368.
- Friston K, Harrison L, Daunizeau J, Kiebel S, Phillips C, Trujillo-Barreto N, Henson R, Flandin G, Mattout J (2008) Multiple sparse priors for the M/EEG inverse problem. *Neuroimage* 39:1104-1120.

- Gaetz W, Macdonald M, Cheyne D, Snead OC (2010) Neuromagnetic imaging of movement-related cortical oscillations in children and adults: age predicts post-movement beta rebound. *Neuroimage* 51:792-807.
- Gallen CC, Hirschkoff EC, Buchanan DS (1995) Magnetoencephalography and magnetic source imaging. Capabilities and limitations. *Neuroimaging Clin N Am* 5:227-249.
- Gross J, Baillet S, Barnes GR, Henson RN, Hillebrand A, Jensen O, Jerbi K, Litvak V, Maess B, Oostenveld R, Parkkonen L, Taylor JR, van Wassenhove V, Wibral M, Schoffelen JM (2013) Good practice for conducting and reporting MEG research. *Neuroimage* 65:349-363.
- Hamalainen MS (1992) Magnetoencephalography: a tool for functional brain imaging. *Brain Topogr* 5:95-102.
- Hamalainen MS, Ilmoniemi RJ (1994) Interpreting magnetic fields of the brain: minimum norm estimates. *Med Biol Eng Comput* 32:35-42.
- Hamalainen MS, Sarvas J (1987) Feasibility of the homogeneous head model in the interpretation of neuromagnetic fields. *Phys Med Biol* 32:91-97.
- Hamalainen MS, Sarvas J (1989) Realistic conductivity geometry model of the human head for interpretation of neuromagnetic data. *IEEE Trans Biomed Eng* 36:165-171.
- Hari R, Salmelin R (2012) Magnetoencephalography: From SQUIDs to neuroscience. *Neuroimage 20th anniversary special edition. Neuroimage* 61:386-396.
- Harris KD, Bartho P, Chadderton P, Curto C, de la Rocha J, Hollender L, Itskov V, Luczak A, Marguet SL, Renart A, Sakata S (2011) How do neurons work together? Lessons from auditory cortex. *Hear Res* 271:37-53.
- Helms G, Dathe H, Dechent P (2008) Quantitative FLASH MRI at 3T using a rational approximation of the Ernst equation. *Magn Reson Med* 59:667-672.
- Helmstaedter M, Sakmann B, Feldmeyer D (2009) Neuronal correlates of local, lateral, and translaminar inhibition with reference to cortical columns. *Cereb Cortex* 19:926-937.
- Henson RN, Mattout J, Phillips C, Friston KJ (2009) Selecting forward models for MEG source-reconstruction using model-evidence. *Neuroimage* 46:168-176.
- Henson RN, Mattout J, Singh KD, Barnes GR, Hillebrand A, Friston K (2007) Population-level inferences for distributed MEG source localization under multiple constraints: application to face-evoked fields. *Neuroimage* 38:422-438.
- Hestrin S, Armstrong WE (1996) Morphology and physiology of cortical neurons in layer I. *J Neurosci* 16:5290-5300.
- Hillebrand A, Barnes GR (2002) A quantitative assessment of the sensitivity of whole-head MEG to activity in the adult human cortex. *Neuroimage* 16:638-650.
- Hillebrand A, Barnes GR (2011) Practical constraints on estimation of source extent with MEG beamformers. *Neuroimage* 54:2732-2740.
- Holmgren C, Harkany T, Svennenfors B, Zilberter Y (2003) Pyramidal cell communication within local networks in layer 2/3 of rat neocortex. *J Physiol* 551:139-153.
- Igarashi J, Isomura Y, Arai K, Harukuni R, Fukai T (2013) A theta-gamma oscillation code for neuronal coordination during motor behavior. *J Neurosci* 33:18515-18530.
- Jerbi K, Baillet S, Mosher JC, Nolte G, Garnero L, Leahy RM (2004) Localization of realistic cortical activity in MEG using current multipoles. *Neuroimage* 22:779-793.
- Jones EG, Burton H, Porter R (1975) Commissural and cortico-cortical "columns" in the somatic sensory cortex of primates. *Science* 190:572-574.
- Jones SR, Pritchett DL, Stufflebeam SM, Hamalainen M, Moore CI (2007) Neural correlates of tactile detection: a combined magnetoencephalography and biophysically based computational modeling study. *J Neurosci* 27:10751-10764.
- Josephson BD (1962) Possible new effects in superconductive tunneling. *Physics Letters* 1:251-253.

- Knosche TR (2002) Transformation of whole-head MEG recordings between different sensor positions. *Biomed Tech (Berl)* 47:59-62.
- Lacey MG, Gooding-Williams G, Prokic EJ, Yamawaki N, Hall SD, Stanford IM, Woodhall GL (2014) Spike firing and IPSPs in layer V pyramidal neurons during beta oscillations in rat primary motor cortex (M1) in vitro. *PLoS One* 9:e85109.
- Lanfer B, Scherg M, Dannhauer M, Knosche TR, Burger M, Wolters CH (2012) Influences of skull segmentation inaccuracies on EEG source analysis. *Neuroimage* 62:418-431.
- Leahy RM, Mosher JC, Spencer ME, Huang MX, Lewine JD (1998) A study of dipole localization accuracy for MEG and EEG using a human skull phantom. *Electroencephalogr Clin Neurophysiol* 107:159-173.
- Lee S, Jones SR (2013) Distinguishing mechanisms of gamma frequency oscillations in human current source signals using a computational model of a laminar neocortical network. *Front Hum Neurosci* 7:869.
- Liu AK, Dale AM, Belliveau JW (2002) Monte Carlo simulation studies of EEG and MEG localization accuracy. *Hum Brain Mapp* 16:47-62.
- Lopez JD, Espinosa JJ, Barnes GR (2012a) Random location of multiple sparse priors for solving the MEG/EEG inverse problem. *Conf Proc IEEE Eng Med Biol Soc* 2012:1534-1537.
- Lopez JD, Litvak V, Espinosa JJ, Friston K, Barnes GR (2014) Algorithmic procedures for Bayesian MEG/EEG source reconstruction in SPM. *Neuroimage* 84:476-487.
- Lopez JD, Penny WD, Espinosa JJ, Barnes GR (2012b) A general Bayesian treatment for MEG source reconstruction incorporating lead field uncertainty. *Neuroimage* 60:1194-1204.
- Lopez JD, Troebinger L, Penny W, Espinosa JJ, Barnes GR (2013) Cortical surface reconstruction based on MEG data and spherical harmonics. *Conf Proc IEEE Eng Med Biol Soc* 2013:6449-6452.
- Lubke J, Egger V, Sakmann B, Feldmeyer D (2000) Columnar organization of dendrites and axons of single and synaptically coupled excitatory spiny neurons in layer 4 of the rat barrel cortex. *J Neurosci* 20:5300-5311.
- Mattout J, Henson RN, Friston KJ (2007) Canonical source reconstruction for MEG. *Comput Intell Neurosci* 67613.
- Mattout J, Phillips C, Penny WD, Rugg MD, Friston KJ (2006) MEG source localization under multiple constraints: an extended Bayesian framework. *Neuroimage* 30:753-767.
- Medvedovsky M, Taulu S, Bikmullina R, Paetau R (2007) Artifact and head movement compensation in MEG. *Neurol Neurophysiol Neurosci* 4.
- Miller KJ, Leuthardt EC, Schalk G, Rao RP, Anderson NR, Moran DW, Miller JW, Ojemann JG (2007) Spectral changes in cortical surface potentials during motor movement. *J Neurosci* 27:2424-2432.
- Mountcastle VB (1997) The columnar organization of the neocortex. *Brain* 120 (Pt 4):701-722.
- Murakami S, Okada Y (2006) Contributions of principal neocortical neurons to magnetoencephalography and electroencephalography signals. *J Physiol* 575:925-936.
- Murakami S, Zhang T, Hirose A, Okada YC (2002) Physiological origins of evoked magnetic fields and extracellular field potentials produced by guinea-pig CA3 hippocampal slices. *J Physiol* 544:237-251.
- Muthukumaraswamy SD (2010) Functional properties of human primary motor cortex gamma oscillations. *J Neurophysiol* 104:2873-2885.
- Muthukumaraswamy SD (2011) Temporal dynamics of primary motor cortex gamma oscillation amplitude and piper corticomuscular coherence changes during motor control. *Exp Brain Res* 212:623-633.
- Nichols TE, Holmes AP (2002) Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Hum Brain Mapp* 15:1-25.

- Nolte G (2003) The magnetic lead field theorem in the quasi-static approximation and its use for magnetoencephalography forward calculation in realistic volume conductors. *Phys Med Biol* 48:3637-3652.
- Numminen J, Ahlfors S, Ilmoniemi R, Montonen J, Nenonen J (1995) Transformation of multichannel magnetocardiographic signals to standard grid form. *IEEE Trans Biomed Eng* 42:72-78.
- Ogawa K, Inui T (2011) Neural representation of observed actions in the parietal and premotor cortex. *Neuroimage* 56:728-735.
- Ogawa T, Riera J, Goto T, Sumiyoshi A, Nonaka H, Jerbi K, Bertrand O, Kawashima R (2011) Large-scale heterogeneous representation of sound attributes in rat primary auditory cortex: from unit activity to population dynamics. *J Neurosci* 31:14639-14653.
- Papadelis C, Poghosyan V, Fenwick PB, Ioannides AA (2009) MEG's ability to localise accurately weak transient neural sources. *Clin Neurophysiol* 120:1958-1970.
- Penny WD, Stephan KE, Daunizeau J, Rosa MJ, Friston KJ, Schofield TM, Leff AP (2010) Comparing families of dynamic causal models. *PLoS Comput Biol* 6:e1000709.
- Perdue KL, Diamond SG (2014) T1 magnetic resonance imaging head segmentation for diffuse optical tomography and electroencephalography. *J Biomed Opt* 19:026011.
- Peters A, Kara DA, Harriman KM (1985) The neuronal composition of area 17 of rat visual cortex. III. Numerical considerations. *J Comp Neurol* 238:263-274.
- Pfurtscheller G, Graimann B, Huggins JE, Levine SP, Schuh LA (2003) Spatiotemporal patterns of beta desynchronization and gamma synchronization in corticographic data during self-paced movement. *Clin Neurophysiol* 114:1226-1236.
- Prieto JJ, Peterson BA, Winer JA (1994) Morphology and spatial distribution of GABAergic neurons in cat primary auditory cortex (A1). *J Comp Neurol* 344:349-382.
- Profant O, Burianova J, Syka J (2013) The response properties of neurons in different fields of the auditory cortex in the rat. *Hear Res* 296:51-59.
- Quraan MA, Moses SN, Hung Y, Mills T, Taylor MJ (2011) Detection and localization of hippocampal activity using beamformers with MEG: a detailed investigation using simulations and empirical data. *Hum Brain Mapp* 32:812-827.
- Ren JQ, Aika Y, Heizmann CW, Kosaka T (1992) Quantitative analysis of neurons and glial cells in the rat somatosensory cortex, with special reference to GABAergic neurons and parvalbumin-containing neurons. *Exp Brain Res* 92:1-14.
- Resnick SM, Pham DL, Kraut MA, Zonderman AB, Davatzikos C (2003) Longitudinal magnetic resonance imaging studies of older adults: a shrinking brain. *J Neurosci* 23:3295-3301.
- Rettmann ME, Kraut MA, Prince JL, Resnick SM (2006) Cross-sectional and longitudinal analyses of anatomical sulcal changes associated with aging. *Cereb Cortex* 16:1584-1594.
- Robinson SE, Vrba J (2004) Cleaning fetal MEG using a beamformer search for the optimal forward model. *Neurolog Clin Neurophysiol* 2004:73.
- Rockel AJ, Hiorns RW, Powell TP (1980) The basic uniformity in structure of the neocortex. *Brain* 103:221-244.
- Ronnqvist KC, McAllister CJ, Woodhall GL, Stanford IM, Hall SD (2013) A multimodal perspective on the composition of cortical oscillations. *Front Hum Neurosci* 7:132.
- Sakata S, Harris KD (2009) Laminar structure of spontaneous and sensory-evoked population activity in auditory cortex. *Neuron* 64:404-418.
- Schroder R, Luhmann HJ (1997) Morphology, electrophysiology and pathophysiology of supragranular neurons in rat primary somatosensory cortex. *Eur J Neurosci* 9:163-176.
- Schubert D, Kotter R, Staiger JF (2007) Mapping functional connectivity in barrel-related columns reveals layer- and cell type-specific microcircuits. *Brain Struct Funct* 212:107-119.

- Shepherd GM, Svoboda K (2005) Laminar and columnar organization of ascending excitatory projections to layer 2/3 pyramidal neurons in rat barrel cortex. *J Neurosci* 25:5670-5679.
- Simons DJ, Durham D, Woolsey TA (1984) Functional organization of mouse and rat Sml barrel cortex following vibrissal damage on different postnatal days. *Somatosens Res* 1:207-245.
- Singh KD, Holliday IE, Furlong PL, Harding GF (1997) Evaluation of MRI-MEG/EEG co-registration strategies using Monte Carlo simulation. *Electroencephalogr Clin Neurophysiol* 102:81-85.
- Staiger JF, Flagmeyer I, Schubert D, Zilles K, Kotter R, Luhmann HJ (2004) Functional diversity of layer IV spiny neurons in rat somatosensory cortex: quantitative morphology of electrophysiologically characterized and biocytin labeled cells. *Cereb Cortex* 14:690-701.
- Steinschneider M, Fishman YI, Arezzo JC (2008) Spectrotemporal analysis of evoked and induced electroencephalographic responses in primary auditory cortex (A1) of the awake monkey. *Cereb Cortex* 18:610-625.
- Stenroos M, Hunold A, Haueisen J (2014) Comparison of three-shell and simplified volume conductor models in magnetoencephalography. *Neuroimage* 94:337-348.
- Stephan KE, Penny WD, Daunizeau J, Moran RJ, Friston KJ (2009) Bayesian model selection for group studies. *Neuroimage* 46:1004-1017.
- Stevenson C, Brookes M, Lopez JD, Troebinger L, Mattout J, Penny W, Morris P, Hillebrand A, Henson R, Barnes G (2014) Does function fit structure? A ground truth for non-invasive neuroimaging. *Neuroimage* 94:89-95.
- Stolk A, Todorovic A, Schoffelen JM, Oostenveld R (2013) Online and offline tools for head movement compensation in MEG. *Neuroimage* 68:39-48.
- Supek S, Aine CJ (1993) Simulation studies of multiple dipole neuromagnetic source localization: model order and limits of source resolution. *IEEE Trans Biomed Eng* 40:529-540.
- Tallon-Baudry C (2009) The roles of gamma-band oscillatory synchrony in human visual cognition. *Front Biosci (Landmark Ed)* 14:321-332.
- Tallon-Baudry C, Bertrand O, Delpuech C, Permier J (1997) Oscillatory gamma-band (30-70 Hz) activity induced by a visual search task in humans. *J Neurosci* 17:722-734.
- Taulu S, Simola J (2006) Spatiotemporal signal space separation method for rejecting nearby interference in MEG measurements. *Phys Med Biol* 51:1759-1768.
- Thambisetty M, Wan J, Carass A, An Y, Prince JL, Resnick SM (2010) Longitudinal changes in cortical thickness associated with normal aging. *Neuroimage* 52:1215-1223.
- Troebinger L, Lopez JD, Lutti A, Bestmann S, Barnes G (2014a) Discrimination of cortical laminae using MEG. *Neuroimage* 102 Pt 2:885-893.
- Troebinger L, Lopez JD, Lutti A, Bradbury D, Bestmann S, Barnes G (2014b) High precision anatomy for MEG. *Neuroimage* 86:583-591.
- Tsubo Y, Isomura Y, Fukai T (2013) Neural dynamics and information representation in microcircuits of motor cortex. *Front Neural Circuits* 7:85.
- Uutela K, Taulu S, Hamalainen M (2001) Detecting and correcting for head movements in neuromagnetic measurements. *Neuroimage* 14:1424-1431.
- Vorwerk J, Cho JH, Rampp S, Hamer H, Knosche TR, Wolters CH (2014) A guideline for head volume conductor modeling in EEG and MEG. *Neuroimage* 100:590-607.
- Watanabe H, Takahashi K, Nishimura Y, Isa T (2014) Phase and magnitude spatiotemporal dynamics of beta oscillation in electrocorticography (ECoG) in the monkey motor cortex at the onset of 3D reaching movements. *Conf Proc IEEE Eng Med Biol Soc* 2014:5196-5199.

- Wendel K, Vaisanen O, Malmivuo J, Gencer NG, Vanrumste B, Durka P, Magjarevic R, Supek S, Pascu ML, Fontenelle H, Grave de Peralta Menendez R (2009) EEG/MEG source imaging: methods, challenges, and open issues. *Comput Intell Neurosci* 656092.
- Whalen C, Maclin EL, Fabiani M, Gratton G (2008) Validation of a method for coregistering scalp recording locations with 3D structural MR images. *Hum Brain Mapp* 29:1288-1301.
- Wilson HS (2004) Continuous head-localization and data correction in a whole-cortex MEG sensor. *Neurol Clin Neurophysiol* 2004:56.
- Wipf D, Nagarajan S (2009) A unified Bayesian framework for MEG/EEG source imaging. *Neuroimage* 44:947-966.
- Yamawaki N, Stanford IM, Hall SD, Woodhall GL (2008) Pharmacologically induced and stimulus evoked rhythmic neuronal oscillatory activity in the primary motor cortex in vitro. *Neuroscience* 151:386-395.
- Yoshimura Y, Callaway EM (2005) Fine-scale specificity of cortical networks depends on inhibitory cell type and connectivity. *Nat Neurosci* 8:1552-1559.
- Zimmermann AHSaJE (1965) Quantum transitions and loss in multiply connected superconductors. *Physical Review Letters* 15.