Title

Bayesian Source Reconstruction and Non-Parametric Statistical Thresholding for Electromagnetic Data

Author

Owen, Julia Parsons

Publication Date

2011

Bayesian Source Reconstruction and Non-Parametric Statistical Thresholding for Electromagnetic Data

by

Julia Parsons Owen

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Bioengineering
in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

AND

[redacted]

Abstract

Bayesian Source Reconstruction and Non-Parametric Statistical Thresholding for
Electromagnetic Data

by

Julia Parsons Owen

In the last few decades there have been major advances in the technology of function brain imaging, allowing for insight into the functions of the human brain previously elusive to neuroscientists. These advances have been primarily on the hardware end and developing effective software to interpret the data collected by neuroimaging machines is a current challenge to the use of the technology. Magnetoencephalography (MEG), in particular, requires solving an ill-posed inverse problem in order to uncover the brain areas active during a task. While the solution to this inverse problem is not unique, there are many methods to estimate its solution and this is a field of active research. In Chapter Two of this thesis, we derive an algorithm that solves the inverse problem for MEG, and the related imaging method, electroencephalography (EEG). Our method improves upon existing algorithms in that it incorporates noise suppression into the estimation procedure and is theoretically and empirically robust to correlated sources. In Chapter Three, we show the results from extensive testing of our algorithm using simu-

lated and real M/EEG data and we show our algorithm's results in comparison to the benchmark algorithms. Chapter Four explores variants of the algorithm, including its application to data sets without pre-stimulus data. In Chapter Five, we present methods to statistically threshold the inverse solution results using nonparametric statistics. Finally, in Chapter Six, we provide some concluding remarks and ideas for future research directions. As a whole, the work presented in this thesis improves the interpretation and analysis of M/EEG data.

Acknowledgments

First, I would like to acknowledge and thank my advisor, Dr. Srikantan Nagarajan. In the early days, we spent many hours at the whiteboard working through the math that appears in Chapter Two of this thesis. While a busy professor, Sri always took time to explain things in great detail. I have learned a lot from him and I appreciate the amount of freedom he has allowed for me to explore research topics. I also greatly appreciate going to many international conferences on his dollar. He is always encouraging of new and exciting academic experiences.

Second, I would like to thank the other members of my thesis committee, Dr. Pratik Mukherjee and Dr. Stanley Klein. They both gave me much support and encouragement, buoying me through the final push to get my thesis done. I appreciate their thorough edits and insightful comments and questions.

Third, I must recognize the entity known as the BIL Lab, or the BI Lab if you are a stickler for redundancy. As a whole, you are a amazing group of friends and colleagues. I would like the thank Anne and Susanne for always showing me how to do things with a smile, even if it was the second, third, or fourth time. I would like to thank Dr. John Houde for teaching me much about MEG and data analysis in my first 6 months in the lab. Former members, Sarang and Johanna deserve my appreciation for the countless questions I have asked about NutMEG. Leighton has been a good friend and advisor through the murky waters of academia. Dr. Kensuke Sekihara, I consider you a lab member. Thank you for being an all-around great person to work with. I am always amazed at your good attitude, it makes research all the more enjoyable when you are around. Alex, you really know how to make a girl laugh - and occasionally cringe. Thank you for your friendship and support; you made the office a fun place to be. Dr. David Wipf, the people you (used) to hang out with miss you. I know you consider yourself some sort of failure for never successfully teaching me anything to do with graphical models. Last, but not least, Naomi I don't know if I could have done it without your support. Thank you for listening to me during moments of triumph and defeat. I hope I can be there, as reliably, for you in the future.

To the ladies of the nurseteria, thank you for keeping me fueled all these years.

I would like to take the opportunity to thank dear friends who have helped me along the way: Sam (Hamantha), Mike (Mikey Muffin), Jeff, Miki, and Janine. Janet, you are an amazing friend and you helped me more than you know. Lindsay and Maggie, I cannot t imagine better life-long friends.

I wouldn't be writing this without the support of my mother, Kim Borsavage.

While I occasionally curse you for encouraging me to pursue engineering, I realize that you helped me to find a field that is exciting and challenging. You have always been there for me when I needed you and somehow, you pull the weight of two parents. Carter and Jack, my younger brothers deserve credit for being sweet, caring brothers. I would also like to acknowledge my paternal grandfather, John J. Owen (Pepe), another engineer. When I was a child, he gave me first hands-on taste of engineering in his workshop. Despite the gender inequity of the field, I always felt at home knowing that I inherited some (small) fraction of his brilliance. Julie and Sam Rea, my aunt and uncle, *merci beaucoup pour tout*.

Arthur! My sweet dog, I know you have little idea that something monumental has happened in my life, but you did help me to achieve it. Advice to anyone suffering from graduate school doldrums, get a dog! It is the best decision I ever made.

Shawn, thank you for all your love, kindness, support, advice, and sheer resilience to my ups and downs. I know that I could not have made it to the phinish line without you cheering me on. You have done so much for me; know that I will always be grateful.

I always imagined that the acknowledgment section was the most fun and sentimental part of the thesis to write. Now I know this to be true.

Contents

Acknowledgments **v**

List of Figures **xi**

1 Introduction **1**
- 1.1 Overview 1
- 1.2 Magnetoencephalography 3
- 1.3 Electroencephalography 4
- 1.4 Sensor-Space Analysis 5
- 1.5 Forward Problem: Lead Field 6
- 1.6 Inverse Problem: Source Localization 7
  - 1.6.1 Parametric Dipole Fitting 8
  - 1.6.2 Spatial Scanning Algorithms 9
  - 1.6.3 Tomographic Algorithms 9
  - 1.6.4 Bayesian Algorithms 10
  - 1.6.5 Sparse Algorithms 11
  - 1.6.6 Denoising 12
- 1.7 Difficult M/EEG Data Sets 16
- 1.8 Statistical Thresholding 19
  - 1.8.1 Correcting for Multiple Comparisons 20
  - 1.8.2 Nonparametric Statistics 21

2 The Champagne Algorithm: Theory **24**
- 2.1 Overview 24
- 2.2 The Champagne Algorithm 25
  - 2.2.1 Estimating the Hyperparameters (Γ) 28
  - 2.2.2 Learning the *Noise-plus-Interference* Factor (Σ_ϵ) 36
- 2.3 Discussion 38

3 The Champagne Algorithm: Performance Evaluation 39
3.1 Overview 39
3.2 Benchmark Source Localization Algorithms 41
3.3 Assessing Localization Accuracy 46
3.4 Generating Simulated Data 49
3.5 Performance on Difficult Source Configurations 53
3.5.1 Discriminating Two Sources 54
3.5.2 Detecting Multiple Sources: Vector Lead Field 56
3.5.3 Detecting Multiple Sources: High SNIR 62
3.5.4 Detecting Multiple Sources: No Noise 62
3.5.5 Detecting Multiple Sources: Inter-dipole Correlations and Source Time Courses 63
3.5.6 Detecting Multiple Sources: Scalar Lead Field 64
3.5.7 Rotating versus Fixed Dipole Model 68
3.5.8 Effect of Lead Field Errors 70
3.5.9 Deep Sources 75
3.5.10 Clusters 76
3.5.11 EEG Simulations: Vector Lead Field 80
3.5.12 EEG Simulations: Scalar Lead Field 81
3.5.13 Functional Connectivity 85
3.6 Performance on Real Data 97
3.6.1 Auditory Evoked Field 97
3.6.2 Audio-Visual Task 98
3.6.3 Face-Processing Task 106
3.6.4 Localizing Spiking Activity in an Epileptic Patient 114
3.6.5 Somatosensory Evoked Field 119
3.7 Discussion 123

4 The Champagne Algorithm: Sensitivity to Noise Estimation 133
4.1 Overview 133
4.2 Denoising the Post-Stimulus Data 136
4.2.1 Simulated Data 138
4.2.2 Real Data 141
4.3 Estimating the Noise Covariance without Pre-Stimulus Data 143
4.3.1 Simulated Data 144
4.3.2 Real Data 147
4.4 Discussion 153

5 Statistical Thresholding 156
5.1 Overview 156
5.2 Source Localization with Unaveraged Data 161
5.3 Statistical Thresholding 162
5.4 Performance on Simulated Data 168

5.5 Performance on Real Data 171
5.5.1 Audio-Visual Task 171
5.5.2 Somatosensory Evoked Field 176
5.5.3 Auditory Evoked Field 181
5.5.4 Face-Processing Task 185
5.5.5 Alternative to Maximal Statistic 189
5.6 Discussion 191
5.7 Appendix 199

6 Conclusions 201
6.1 Overview 201
6.2 Future Directions 202
6.2.1 Comparisons and Extensions of Champagne Algorithm . . 202
6.2.2 Applications to MEG Data Sets 204

List of Figures

2.1 Convergence rate: EM implementation compared to upper-bounding hyperplane implementation. 32
2.2 Demonstration of the pruning of the hyperparameters over iterations by Champagne. 36

3.1 Example of a simulated data at 10dB and 0dB. 51
3.2 Simulations with two dipoles seeded with varying inter-dipole distances. 57
3.3 Results from simulations testing localization on multiple dipoles (3 to 30) with a volumetric lead field. 59
3.4 A single example of the localization results with the vector lead field for 5 dipoles at SNIR=10dB and SNIR = 0dB. 60
3.5 A single example of the localization results with the vector lead field for 10 dipoles at SNIR=10dB and SNIR = 0dB. 61
3.6 Comparison of localization results at SNIR levels: 0dB, 10dB, and 20dB. 62
3.7 Simulation at SNIR=100dB with Gaussian time courses and Gaussian noise model. 63
3.8 Simulations at 10dB with different inter-dipole correlations and types of source time courses. 65
3.9 Results from simulations to test localization on multiple dipoles (2 to 100) with a scalar leadfield. 66
3.10 A single example of the localization results with the scalar lead field for 20 dipoles at SNIR=10dB. 67
3.11 Exploration of lead field errors with 10, 15, and 20 dipoles. 71
3.12 Localization results for 10, 15, and 20 dipoles with varying orientation errors introduced to the scalar lead field. 72
3.13 Simulations with un-shifted and shifted voxel grid at SNIR=10dB. . . . 73
3.14 Localization results in the presence of deep sources at SNIR=10dB and SNIR=0dB. 74
3.15 Performance on 5, 10 and 15 clusters of dipoles using a vector lead field. 77

3.16 A single example of the localization results for 10 clusters at SNIR = 10dB with the vector lead field. 78
3.17 Results for EEG simulations with a vector EEG lead field. 83
3.18 Results for EEG simulations with a scalar EEG lead field. 84
3.19 General connectivity schema for simulations 89
3.20 Source localization results for functional connectivity simulation. . . . 91
3.21 Connectivity results for functional connectivity simulation. 92
3.22 Plot of mixing matrix obtained from voxel time-course reconstructions using Champagne. 93
3.23 Networks with 2, 3, or 5 nodes. 94
3.24 Results from multiple functional connectivity simulations with 2, 3 and 5 nodes at 10dB and 0dB. 96
3.25 Auditory evoked field (AEF) results for 7 subjects. 99
3.26 Champagne localization results for the audio-visual task. 103
3.27 SL localization results for the audio-visual task. 104
3.28 MCE localization results for the audio-visual task. 105
3.29 MVAB localization results for the audio-visual task. 106
3.30 Champagne localization results for the face-processing task. 109
3.31 Champagne localization results for the face-processing task, ventral view of brain. 110
3.32 Champagne localization results for the contrast condition for face-processing task. 110
3.33 MVAB localization results for the face-processing task. 111
3.34 SL localization results for the face-processing task. 111
3.35 MCE localization results for the face-processing task. 112
3.36 Champagne's localization results on EEG face-processing data. 115
3.37 Benchmark algorithms' localization results on EEG face-processing data.116
3.38 Localization results for Champagne on a single spike in a epileptic patient. 117
3.39 Localization results for MCE on a single spike in a epileptic patient. . . 118
3.40 Localization results for MVAB on a single spike in a epileptic patient. . 118
3.41 Localization results for SL on a single spike in a epileptic patient. . . . 118
3.42 Localization results for Champagne, MVAB, SL/dSPM, and MCE for the SEF data. 120
3.43 Somatosensory processing for stimulation of the right index finger (RD2) and the right upper lip shown with unaveraged data. 122

4.1 Eigenspectrum of simulated data. 139
4.2 Cleaned versus Full Data 140
4.3 Champagne's results on the cleaned signal covariance for the SEF data. 141
4.4 Champagne's results on the cleaned signal covariance for the AV data. 142
4.5 Champagne's results on the cleaned signal covariance for the AEF data. 142
4.6 *Isotropic Noise* results 145

4.7 *Non-Isotropic, Diagonal Noise* results 146
4.8 *Non-Isotropic, Non-Diagonal Noise* 147
4.9 Results for Champagne run on the SEF data without access to the pre-stimulus period to learn the noise covariance. 149
4.10 Results for Champagne run on the AV data without access to the pre-stimulus period to learn the noise covariance, part 1. 150
4.11 Results for Champagne run on the AV data without access to the pre-stimulus period to learn the noise covariance, part 2. 151
4.12 Results for Champagne run on the AEF data without access to the pre-stimulus period to learn the noise covariance. 152

5.1 Cartoon of thresholding procedure. 165
5.2 Methods for thresholding 166
5.3 Simulated data example (0dB) with surrogates generated with resampling pre-stimulus data and computing post-stimulus power, Method 3. 169
5.4 Simulated data example (-5dB) with surrogates generated with resampling pre-stimulus data and computing post-stimulus power, Method 3. 170
5.5 AV data with surrogates generated with resampling pre-stimulus data and computing post-stimulus power around auditory activation, Method 3. 175
5.6 P-values computed from the maximal statistic distribution for the window around the auditory activation for the AV data. 176
5.7 AV data with surrogates generated with resampling pre-stimulus data and computing post-stimulus power around visual activation, Method 3.177
5.8 P-values computed from the maximal statistic distribution for the window around the visual activation for the AV data. 178
5.9 AV results for MVAB. 178
5.10 Somatosensory (SEF) data with surrogates generated with resampling pre-stimulus data and computing post-stimulus power, Method 3. . . . 180
5.11 P-values for the SEF data computed from the maximal statistic distribution. 181
5.12 Auditory evoked field (AEF) data with surrogates generated with resampling pre-stimulus data and computing post-stimulus power, Method 3. 183
5.13 P-values for the AEF data computed from the maximal statistic distribution. 184
5.14 AEF results for MVAB. 184
5.15 Face processing data with surrogates generated with resampling pre-stimulus data and computing post-stimulus power, Method 3. 187
5.16 P-values for the face data computed from the maximal statistic distribution. 188

5.17 MVAB results on face processing data set. 188
5.18 AV results comparing maximal statistic to less stringent alternative. . . 189
5.19 SEF results comparing maximal statistic to less stringent alternative. . 190
5.20 AV results comparing maximal statistic to less stringent alternative. . . 191
5.21 Histograms of the post-stimulus power. 198
5.22 Audio-visual data with surrogates generated with pre- and post-stimulus switching and calculating t-values around auditory activation, Method 1. 200
5.23 AV data with surrogates generated with pre- and post-stimulus switching and calculating post-stimulus power around auditory activation, Method 2. 200
5.24 Plot of standard deviation of the pre-stimulus period for the AV data set.200

Chapter 1

Introduction

1.1 Overview

The human brain is a relatively undiscovered scientific territory. This is due, in part, to its sheer complexity. There are approximately 100 billion neurons (Kandel et al., 2000) in the human brain; these neurons govern our physiological processes, our perceptions of the world, our emotions, our memories, some might say, our *selves*. How, exactly, these neurons perform all these functions is a mystery that neuroscientists (and philosophers) have been chipping away at for centuries. The other part of the story is that until recently, we could not observe neuronal activity within a live, functioning human brain. Modern technological advances in functional brain imaging have given researchers their first view into the inner workings of the brain and opened up possibilities to understand how individual neurons act in concert to produce the human experience.

There are several functional brain imaging technologies available to neuroscientists, of which, functional magnetic resonance imaging (fMRI), electroencephalography (EEG), and magnetoencephalography (MEG) are the most widely used in human neuroscience. fMRI measures the local ratio of oxyhemoglobin to deoxyhemoglobin; an increase in this ratio, i.e. an increase in oxyhemoglobin, is an indication of neuronal activity. fMRI does not directly measure brain activity, rather it measures a secondary, metabolic effect. As such it is limited by perfusion of the blood and can be dictated by the blood supply to a particular region. The time resolution of fMRI is on the order of a second, as the perfusion of the blood to the firing neurons is not instantaneous and the images are acquired in a systematic, and somewhat time-intensive fashion. While a powerful and widely-available imagining technique, fMRI is not always the best imaging method to use. The temporal resolution of fMRI is long compared to the time for axonal conductance, which occurs in the range of tens of milliseconds. If examining the precise timing of neuronal activity, MEG or EEG are powerful imaging tools as they directly and noninvasively measure the magnetic fields or electric potentials generated by neuronal currents with a fine temporal resolution (sub-millisecond timescale).

1.2 Magnetoencephalography

Magnetoencephalography measures the magnetic fields generated by neuronal activity with superconducting quantum interference devices (SQUIDs). There are typically 275 SQUIDs sensors positioned in a hemisphere around the head. These sensors detect changes in the magnetic field (with either magnetometers or gradiometers) on the order of 10 femto Tesla; these fields are 1 billionth the strength of the Earth's magnetic field.(Lu and Kaufman, 2003). Magnetic fields arise from the pyramidal cells in the cortex, which are arranged in a parallel fashion along the cortical sheath. When there is brain activity in a certain region, tens of thousands of neurons fire synchronously, producing the changes in the magnetic field detectable by the sensors. As such, in order to be detected by the SQUIDs, neuronal activity in a small patch of cortex (roughly 1-8mm^3) must be active, producing magnetic fields approximately 10-100fT outside the head.

MEG is a relatively expensive technology; commercial systems typically cost more than a million USD. Also, MEG can only be used on the brain and in some rare instances, on the fetal heart, whereas a MRI can be used on any part of human body. While the sampling rate is high, around 1200Hz, many trials (around 50-100) must be collected in order to have an appreciable signal to noise ratio. While in the MEG, the patient cannot move as muscle artifacts can greatly distort the measurements. Additional sources of noise include eye blinks, metal fillings in the teeth, and any medical devices implanted in the subject. MEG is a powerful imaging technique under ideal circumstances, but its effectiveness is fleeting in the face of noise.

1.3 Electroencephalography

Electroencephalography is a related imaging method; it measures voltage potentials (also originating from neuronal activity) with electrodes placed on the scalp, as opposed to measuring magnetic fields, as in MEG. The magnetic fields pass undisturbed through the biological tissue (brain, scalp, and skin), but the electrical fields must conduct through these materials. The effect of this physical property of electric fields is that the potentials are smeared as they emanate from their point of origin to the sensors on the scalp. This smearing of voltage potential across the different conductivity layers makes EEG data harder to interpret (as it is difficult to know exactly where a signal originates). As such, MEG has been reported to have a higher spatial precision in its ability to estimate source location (Leahy et al., 1998). EEG must have a reference electrode since the electric potential can only be measured relative to a reference, whereas MEG does not need to be referenced (it is internally referenced). The placement of the reference electrode is not obvious and can greatly affect the measurements. While the spread of the electric field and the need to reference makes EEG less favorable to use, EEG has some advantages over MEG. The technology for EEG is much simpler, the cost of a commercial MEG system is roughly 50 times that of an EEG system, EEG is portable, and the subject can move his/her head as needed once the electrodes are placed. EEG is used widely for clinical diagnosis and treatment of neurological diseases and conditions, but its sensitivity falls short of MEG's ability to localize brain function.

1.4 Sensor-Space Analysis

Sensor-space analysis describes data analysis that only considers the sensor time courses and does not go as far as localizing neuronal activity within the brain volume, referred to as *source localization*. A surprising amount can be learned from only looking at the sensor time sources. When trials of data are aligned to the stimulus and averaged over many trials (50 to 100), characteristic peaks can be observed for certain paradigms. For instance, with an auditory evoked field stimulus (single frequency tones are presented binaurally), one can expect to see a peak in the averaged sensor data at around 100ms after the presentation of a single tone. The peak, called the M100 in MEG and the N100 in EEG, can then be used as a dependent variable in an experimental design. For instance, the amplitude and latency of the M100/N100 can be compared between conditions or between a control and patient population. While a viable data analysis technique, sensor-space analyses do not provide (strong) conclusions about the parts of the brain involved in processing information or controlling a behavior. In the case of the M100/N100, one can only say that the greatest peak is seen in sensors over or near auditory cortex. As such, performing source localization has become a cornerstone of data analysis with M/EEG data and an active field of research.

1.5 Forward Problem: Lead Field

The sensor measurements are taken near the surface of the head in both MEG and EEG and a model is used to determine the mapping between activity at brain locations and measurements at the sensors, called the *lead field*. Typically the lead field is calculated using an equivalent current dipole (ECD) model. As discussed above, the pyramidal cells in the cortex are arranged in a parallel fashion and when a group of pyramidal neurons fire synchronously, the sum of their activity can be approximated as a single current dipolar source. Assuming the head is roughly spherical, the magnetic field expected at each sensor can be calculated given a current dipole at every location in the brain from Maxwell's equations (Sarvas, 1987). The candidate locations throughout the brain are called *voxels* and the resolution of the voxel grid is at the discretion of the user, but typically ranges from 3-10mm.

The magnetic fields pass through the brain and skull unperturbed, meaning that the conductances of the biological materials are not parameters in the calculation of the lead field for MEG. The electric field, however, does not pass freely through the biological tissues and therefore, the lead field calculation for EEG is more complicated than the calculation for MEG. The EEG lead field depends greatly on the estimation of tissue conductivities to account for the current distortion across tissue boundary layers (from brain to skull to scalp). The conductivities are not consistent across subjects and there can be anisotropies in the conductances, further muddying the problem. Forward modeling is a field in and of itself and there are more sophisticated methods than the ECD model, such as boundary element models and finite element models available (Mosher et al., 1999). These methods, while potentially more precise, are more computationally expensive and require more accurate knowledge of the brain geometry and segmentation of the structural MRI. The ECD model, while theoretically simple, is a powerful method of relating underlying brain activity to sensor readings.

1.6 Inverse Problem: Source Localization

Solving the inverse problem in M/EEG involves transforming the recorded sensor data to brain activations at specific locations. This inverse problem is inherently ill-posed as the number of voxels (typically 3,000 to 10,000) is greater than the number of sensors (typically 275 sensors for MEG systems and 64, 128, or 256 electrodes for EEG systems). Finding a solution to the inverse problem is complicated by the presence of correlated sources, sensor noise, and interference (both from within and outside the brain). Obtaining accurate source locations and time course estimates is central to the validity of M/EEG as a functional brain imaging method. Currently, a wide variety of source localization algorithms exist for estimating source activity. These algorithms can be grouped into three types: parametric dipole fitting, spatial scanning, and tomographic.

1.6.1 Parametric Dipole Fitting

Parametric dipole fitting, or the equivalent current dipole (ECD) method, is a common technique, where a small number of point dipole sources are assumed to generate the M/EEG data. The problem reduces to determining the parameters of the dipoles, such as the location, orientation, and amplitude, usually by non-linear optimization or search. This method was first introduced in 1983 (Tuomisto et al., 1983) for use with localizing auditory dipoles. Dipole fitting is robust when there is one underlying dipole, but searching methods scale exponentially with the number of dipoles. For two or more sources, the results can be highly dependent on initialization of the locations of the dipoles and on the total number of dipoles being fit and often, trying to fit three or more dipoles will fit noise rather than true source locations. Non-linear optimization methods have major issues with local minima and can be sensitive to the initialization of the number and location of dipoles (Mosher et al., 1992; Uutela et al., 1998). Dipole fitting also has trouble resolving dipoles that have anti-parallel orientations, are close together, or are deep within the brain (Lutkenhoner, 1998). In addition, the traditional ECD model does not account for extended sources, i.e. sources that are larger than one voxel; variants on the dipole fitting technique have been developed to account for the extended sources, both extended in 1-D, as in a line sources (Yetik et al., 2005) or extended in 2-D, as in a patch (Schmidt et al., 1999; Jerbi et al., 2002). It is also possible to estimate the noise and/or signal covariances from data (Mosher and Leahy, 1998; Jun et al., 2005, 2006b,a; Huizenga et al., 2002; de Munck et al., 2004; Bijma et al., 2005), which enables the ECD approach to better localize brain activity.

1.6.2 Spatial Scanning Algorithms

Spatial scanning techniques estimate the time course at every candidate location while suppressing the interference from activity at the other candidate source locations. Some examples of scanning techniques are minimum-variance adaptive beamforming (MVAB) and other variants of beamformers (Sekihara and Nagarajan, 2008), multiple signal classification (MUSIC)(Mosher and Leahy, 1998), synthetic aperture morphometry (SAM) (Vrba and Robinson, 2001), dynamic imaging of coherent sources (DICS) (Gross et al., 2001), and source activity using knowledge of event timing for independence from noise and interference (SAKETINI) (Zumer et al., 2007). Beamforming methods, such as MVAB, are widely used in the field of MEG. Although computationally easy to implement, MVAB is confounded by correlated sources and has the tendency to produce diffuse activity even when the underlying activity is focal.

1.6.3 Tomographic Algorithms

The third class of algorithms is the tomographic techniques, which model the activity at all candidate source locations simultaneously. There is a variety of tomographic methods, such as minimum-norm estimation (MNE)(Hämäläinen and Ilmoniemi, 1994), dynamic statistical parametric mapping (dSPM)(Dale et al., 2000), and standardized low resolution brain electromagnetic tomography (sLORETA) (Pascual-Marqui, 2002). Among these, MNE is the most widely-used class of inverse algorithms; the sLORETA algorithm is a variant of MNE. Empirically, SLORETA and dSPM, although differing in their theoretical properties, perform similarly with simulated and experimental data. Like MVAB, the inverse solution with sLORETA and dSPM are easy to compute. But, also like MVAB, their solutions are often overly diffuse and their ability to reconstruct multiple dipolar or regional source clusters has not been rigorously tested.

1.6.4 Bayesian Algorithms

Most of the source reconstruction algorithms from the three classes can be framed in a Bayesian schema. This perspective is useful as, at a high level, the prior distribution implicitly or explicitly imposed for each method can be used to differentiate and compare the various source localization methods. Algorithms such as MNE, minimum current estimate (MCE), sLORETA, dSPM, and MVAB (and other beamformers) assume a known, fixed prior. Alternatively, the parameters of the prior distribution (hyperparameters) can be learned from the data, referred to as *empirical Bayes*. These ideas are extensively discussed in (Wipf et al., 2009). Often empirical Bayesian algorithms are more robust in that they include a full statistical model; some examples include (Friston et al., 2008; Kiebel et al., 2008; Mattout et al., 2006; Nummenmaa et al., 2007; Phillips et al., 2005; Sahani and Nagarajan, 2004; Zumer et al., 2007; Schmidt et al., 1999; Jun et al., 2005).

Noise estimation is an important part of the source localization procedure. Electromagnetic sensor data are susceptible to many sources of noise, both internal and external. There is internal electronic noise at the individual sensor-level, which is independent across sensors. Many methods for estimating this noise have been proposed, ranging in the sophistication of their approaches (Sekihara et al., 1997; Waldorp et al., 2001; Sekihara et al., 1994, 2008); formulating the generative model in a Bayesian framework is often a convenient and effective means to estimate this noise (Jun et al., 2005; Zumer et al., 2007; De Munck et al., 2002; Jun et al., 2006b). There is also noise that arises from external source and is seen across some or all of the sensors. This type of noise is more difficult to model as it can be difficult to separate from the true brain signals measured by the sensors, which can be quite similar. Some methods tailored to remove the external noise sources are discussed in the *Denoising* section below.

1.6.5 Sparse Algorithms

Some tomographic methods promote sparseness in the solution, i.e. the majority of the candidate locations do not have significant activity. Empirical evidence shows that a sparse source model can improve the accuracy of the localization in a noisy environment. Examples of sparse methods include the focal underdetermined system solver (FOCUSS) (Gorodnitsky and Rao, 1997) and minimum-current estimate (MCE) (Uutela et al., 1999; Wipf et al., 2009). Sparsity can be obtained in different ways. One way to achieve a sparse solution to the inverse problem is to assume a sparse prior for the sources. A sparse distribution used is the Laplace probability distribution and can be used to gain sparsity in the inverse solution (Gerven et al., 2009). Another way to promote sparsity in the source activity is to use a ℓ_1-norm instead of a ℓ_2-norm in the minimized cost function (Uutela et al., 1999; Wipf et al., 2009; Haufe et al., 2008); one such algorithm is MCE. MCE solves the inverse problem by finding the solution that minimizes the ℓ_1-norm of the sources at every time point. While the solution obtained by MCE is inherently sparse, it is not empirically robust to correlated sources or able to localize a large number of sources in a noisy environment. The common shortcomings of MCE and many other sparse tomographic algorithms is that they are not explicitly designed to handle complex, correlated source configurations in the presence of background interference (e.g., spontaneous brain activity, sensor noise, etc.).

1.6.6 Denoising

The sensor data collected in M/EEG is noisy, meaning that they do not only arise from stimulus-evoked activity (the activity of interest in the experiment), but also from other sources termed *interference sources* and from noise and variability in the sensors. As discussed above, the internal or sensor noise is fairly easy to model as it is independent across sensors. The external noise or interference sources are more difficult to suppress as they have varied origin including: spontaneous brain activity (not evoked by the stimulus), biological sources, such as eye blinks, muscle artifacts from the head and neck, and heartbeats, and non-biological sources, such as powerlines, elevators, metal dental fillings, and implanted medical devices in the subject. These interference sources often overlap with the stimulus-evoked activity and make localizing the source of this activity more difficult. As such, being able to limit the effect of interference and sensor noise on source localization is an essential step to improving source localization.

The most basic approach to removing interference from the data is to collect a large number of trials and then average the data across the trials (after aligning to the stimulus). The underlying assumption is that the interference sources do not have a set phase to the stimulus and will be averaged out if sufficient trials are averaged. This approach requires longer data acquisition times and does not allow for analysis of high frequency activity (high frequency oscillations can be averaged out as they are not necessarily phase-locked to the stimulus). In addition, some sources of noise are phase-locked to the stimulus such as, movement artifacts, so the averaging approach does not produce clean sensor data in these cases.

Data-driven approaches, such as principal component analysis (PCA), independent component analysis (ICA), maximum-likelihood estimation (MLE), and Wiener filtering, have also been applied to data to suppress the interference sources (Ossadtchi et al., 2004; Ungan P, 1976; Nagarajan et al., 2006; Baryshnikov et al., 2004). These approaches require some subjective choices; in the case of PCA a truncation threshold must be chosen and with ICA, the relevant number of components must be selected. There are principled ways to this for both PCA and ICA. In addition, they do not exploit the experimental design. The majority of M/EEG experiments have a stimulus, which prompts a sensation, perception, or behavior. Between each trial, or stimulus, there is some wait-time to allow the brain activity to return to the baseline. This design affords a *pre-stimulus* and a *post-stimulus* period. In the pre-stimulus period, some or all of the interference source are expected to be present and in the post-stimulus period, statistically similar interference activity is expected along with the activity of interest. This assumption that the activity of interest, in the post-stimulus period, is additive opens up the possibility of interference suppression that improves upon previously developed methods.

One way to remove noise is to use variational Bayesian factor analysis model

(VBFA), which decomposes sensor data into factors and independent noise (representing the sensor/electronic noise). (Attias, 1999) Once the independent sensor noise is removed, source localization can be performed on the cleaned data. VBFA provides a regularized estimator of the post-stimulus correlation matrix. (Many source localization algorithms use the sensor correlation and not the raw data to perform localization.) An extension of VBFA, is stimulus-evoked factor analysis (SEFA); this method removes the interference present in the pre-stimulus period from the post-stimulus period, in addition to the independent sensor noise. First, VBFA is run on the pre-stimulus period, then factors can be learned in the post-stimulus period after the pre-stimulus activity is removed from the post-stimulus sensor data. Neither VBFA nor SEFA localizes activity, rather it provides a cleaned signal for the post-stimulus period on which source localization can be performed. There are also source localization algorithms that incorporate this idea of noise suppression when estimating the location and time courses of brain activity. One that combines ideas from SEFA and source localization is the Neurodynamic Stimulus-Evoked Factor Analysis Localization (NSEFALoc) algorithm.(Zumer et al., 2008). NSEFALoc uses SEFA in conjunction with temporal basis functions learned from the data to localize brain activity. The algorithm in (Zumer et al., 2007), presents a probabilistic modeling framework for MEG/EEG source localization that estimates Source Activity using Knowledge of Event Timing and Independence from Noise and Interference and is called SAKE-

TINI. SAKETINI is a scanning technique that estimates the brain activity at every candidate, after removing the activity present in the pre-stimulus period using VBFA. While the paper demonstrates that the algorithm is able to localize activity at low SNR levels and in the presence of interference, it is a computationally intensive algorithm as it scans through each candidate location and is untenable for researchers without access to parallel computing centers.

In Chapter Two of this thesis, we derive a novel source reconstruction algorithm, called *Champagne*, that is derived in an empirical Bayesian schema and incorporates deep theoretical ideas about sparse-source recovery from noisy, constrained measurements. Champagne improves upon existing methods of source reconstruction in terms of reconstruction accuracy, robustness, and computational efficiency. It is designed to estimate the locations and time courses of a small (sparse) set of flexible, dipolar sources that adequately explain the observed sensor data. Champagne is a tomographic algorithm and it achieves its sparsity through a sparse penalty function, as described in Chapter Two. Noise and interference suppression is integrated in the Champagne algorithm, which makes it particularly robust when the signal-to-noise ratio is low. Champagne achieves robustness by fitting a structured covariance model to the sensor data covariance.

1.7 Difficult M/EEG Data Sets

There are some source configurations that are harder to localize than others. This section will outline the features of the data sets for which it is notoriously hard to accurately localize brain activity.

Localizing brain activity in the face of source correlations is one of the greatest challenges for source localization algorithms. Correlated sources refer to regions of brain activity that have almost precisely the same time course. Auditory responses are an example of where the correlated source problem arises with real M/EEG data. Correlated sources are a problem for multiple reasons. First, if sources are highly correlated and are located bilaterally, as is the case with auditory sources, then the projection to the sensors will be similar to having one strong source at the center of the head. Second, some source localization algorithms, such as beamformers, assume orthogonality in the source activities, making it difficult for them to localize activity that is far from orthogonal (or uncorrelated). These correlations make the problem of source localization more challenging, but are also meaningful for understanding the way that the brain transfers information from one region to another, known as *functional connectivity*. There has been much attention paid to this issue in the development of new source localization algorithms and the modification of existing algorithms (Dalal et al., 2006; Gross et al., 2001; Mosher and Leahy, 1998; Sekihara et al., 2001).

Deep sources present a similar challenge to source localization. It is accepted that sources deeper than the cortex are difficult to uncover with M/EEG imaging (Pascual-Marqui, 2002; Lin et al., 2006a). Deep sources have low signal-to-noise ratios (SNR) as they are farther away from the sensors. Likewise, their projection to the sensors is less unique than shallow sources because they are (more-or-less) equally close to all the sensors. These two characteristics make localizing deep sources challenging. The stronger SNR of the shallow sources can cause localization algorithms to be biased towards these sources (and ignore the deep ones). To address this, column normalizing the lead field is a standard practice in the source localization procedure (Pascual-Marqui, 2002), but this normalization does not always succeed in localizing the deep sources. Sensitivity to deep sources is important since there are medial brain structures, such as the basal ganglia, the amydgala, and hippocampus, which are equally important to understand as more lateral cortical regions.

Sources that are extended over many voxels are also hard to localize with M/EEG imaging. When the lead field is calculated, a resolution is chosen for the voxel grid. This resolution is somewhat arbitrary and does not always reflect the true resolution of activity. In the case of cortical activations, the extent of the patch of cortex with activity can be different from one region to another, or one subject to another. It is difficult to recover the true extent of the activation, especially when the lead field resolution (or voxel size) is smaller than the activations. There are modifications to source localization algorithms to account for the varying extent of cortical sources, including incorporating basis functions that allow for different resolutions of the lead field (Limpiti et al., 2006; Jerbi et al., 2002; Lin et al., 2006a; Yetik et al., 2005).

Other challenging source configurations are sources that are close together, especially true for correlated sources, and conditions where there are many sources in different regions of the brain. These conditions have not been extensively tested or addressed in the literature. When developing source localization algorithms, it is important to keep these challenging source configurations in mind.

In Chapter Three, we present results from both simulated and real MEG data that more rigorously characterize Champagne's performance and better delineate both its strengths and limits. For the simulated data, we have selected a number of challenging source configurations: a large number of distinct sources, deep sources, correlated sources and clusters of sources. We also present the results on simulated data using benchmark algorithms, MVAB, sLORETA/dSPM, and MCE. In addition to investigating these challenging source configurations with simulated data, we also present an analysis of the effect of lead field modeling errors and we present results from simulated EEG and we investigate the effect of subsampling the number of EEG sensors.

We also use real data sets to test Champagne's performance as compared to the benchmark algorithms: somatosensory-evoked field (SEF), auditory-evoked field (AEF), audio-visual integration task, and face-processing data, both MEG and EEG recordings. Taken together, the results on simulated and real data give a rather complete picture of Champagne's source localization capabilities. Ultimately an exhaustive performance analysis of a source localization algorithm serves two distinct purposes. The first being to motivate the use of a new algorithm by showing head-to-head performance with commonly-used methods. The second is to fully describe when an algorithm works and of equal importance, when it does not work. The conditions under which an algorithm fails motivates further development of source localization algorithms to advance beyond these shortcomings.

1.8 Statistical Thresholding

The statistical analyses of the source images from M/EEG imaging typically proceed with voxel-level statistics. At each voxel, a statistical test is used to determine if an effect of interest is present. These statistics then go into forming a *statistic image*, which provides a visualization of the statistic values. Unless there is a hypothesis for regions of the brain that will contain the experimental effect, every voxel has to be tested. Traditionally, a parametric statistical test is used, some common methods include t-tests, F-tests, paired t-tests, ANOVA, correlation, linear regression, multiple regression, and ANCOVA. All of these tests assume that the data are normally distributed with a mean parameterized by a general linear model. The estimated parameters for the specific test are then contrasted to produce a t-statistic, which then can be assessed for statistical significance (i.e. evidence against the null hypothesis) given an assumed distribution. This process yields p-values, which is the probability that the value would exceed that observed under the null hypothesis. The test *level*, usually denoted by α, is the accepted risk of the test or the probability of committing a Type I error. A Type I error reflects falsely rejecting the null hypothesis, i.e. errantly finding significance. Any p-values less than α lead to a rejection of the null hypothesis.

1.8.1 Correcting for Multiple Comparisons

In the case of M/EEG brain imaging, there are typically 5,000 to 15,000 voxels, which results in that many statistical tests. Therefore, the risk for committing Type I error is high. If the *level* is $\alpha = 5\%$ and there are 10,000 voxels, there could be 500 voxels that have Type I errors, which could greatly impact the interpretability of the results. As such, there are methods to correct for Type I errors when there are multiple comparisons. The Bonferroni correction (Bonferoni, 1935) and false discovery rate (FDR) (Benjamini and Hochberg, 1995) are two commonly-used methods to correct for Type I errors. There corrections are overly-conservative and they can remove a result or trend that is present in the uncorrected statistical values.

1.8.2 Nonparametric Statistics

While conventional, the statistical testing described has two main issues. The first is the assumption of normally distributed random variables. Under the central limit theorem, this assumption would hold, but in brain imaging the number of samples does not approach those needed to assure Gaussianity. Secondly, correcting for multiple comparisons with a method such as, the Bonferroni correction or FDR, is often overly stringent. Nonparametric statistics provide a solution to at least the first of these issues and can be formalized in a way to address the second. Nonparametric statistics do not assume a distribution for the variable. Rather, they use the data to obtain a null distribution over the data, from which, significance can be tested. Permutation tests are one type of nonparametric test; they have been increasingly used as computational power increases. The principle behind permutation testing is not complicated. Essentially the data are permuted, for instance across conditions, in a pseudo-random fashion and the test statistics are computed for each permutation of the data. If there is no difference between conditions, the test statistics should not be greatly altered by the permutations. Permutation testing methods use the data to generate a permutation distribution for every voxel. From this distribution, p-values can be calculated for the t-statistics from the original (or unpermuted) data.

These p-values, however, are not corrected for multiple comparisons as there are still many statistical tests being performed in parallel. The same methods that correct for multiple comparisons with parametric statistics can be applied to the p-values obtained from nonparametric statistics, but applying these correction techniques to nonparametric statistics also can be overly stringent. The maximal statistic approach to nonparametric statistics automatically corrects for the multiple comparisons problem. For each permutation, only the maximum t-statistic from every permutation is saved and a distribution of the maximal statistic is formed. Then a threshold for the statistic image can be obtained given a level α. Nonparametric statistics have been applied to neuroimaging data in a variety of studies and variants of the method have been used to improve the sensitivity of the tests (Nichols and Holmes, 2001; Singh et al., 2003; Sekihara et al., 2005; Pantazis et al., 2005; Chau et al., 2004; Dalal et al., 2008). These various implementations of nonparametric statistics differ in the way they generate the permutations, i.e. by exchanging conditions, exchanging subject groups, or exchanging time windows of the data, the formulation of the test statistic, and how the permutation distribution is determined, by the maximal statistic or otherwise.

A challenge of statistical thresholding is an appropriate means of arriving at a threshold for a single subject. Often statistics use an average across subjects, but in functional brain imaging, this means the MRIs of the individual subjects must be spatially normalized. Spatial normalization, while a widely-used technique, is not always ideal because it relies on the automatic segmentation of the MRI. If there is an error in the segmentation, the averaged results can be corrupted. As such, it is often preferable to look at the results on an individual subject level, but arriving at a principled threshold for a single subject's data is not a prescribed science.

In Chapter Four, we explore permutation testing methods for statistically thresholding real MEG data. Our method is valid for thresholding the results of only one subject, or could be applied to a condition contrast or average across subjects. First, we demonstrate Champagne's localization capabilities on data that is not trial-averaged and then three methods for performing the nonparametric statistics are formalized and applied to AEF, AV, SEF, and face-processing data sets. We have tailored the statistical thresholding procedure to be applicable to sparse algorithms, a issue not addressed in the literature.

Chapter 2

A Novel Bayesian Algorithm for Neural Source Reconstruction: Mathematical formulation [1]

2.1 Overview

In this chapter, a novel source localization algorithm, called *Champagne*, is derived and its theoretical properties are exposed. Champagne is an empirical Bayesian method that incorporates noise suppression with the localization of source locations and the estimation of the source time courses applicable to either MEG and EEG data. Champagne is essentially a structured maximum likelihood covariance estimate, the empirical sensor covariance is decomposed into a noise

---

[1] Portions of this chapter originally appeared as an article in (Wipf et al., 2010)

and signal component. The noise component is estimated from the sensor data recorded before the stimulus, referred to as the pre-stimulus period, and then the noise is subtracted from the empirical sensor covariance, leaving the signal-of-interest component. This cleaned sensor covariance is used to estimate the source locations and time courses from the data collected after the stimulus, referred to as the post-stimulus period. Champagne is a sparse algorithm; only a small fraction of the voxels have nonzero activity. The sparsity arises from a sparse penalty function, as opposed to imposing a sparse prior distribution on the sources. Champagne is robust to highly correlated sources, a feature that allows Champagne to outperform many commonly-used source localization algorithms.

2.2 The Champagne Algorithm

The Champagne algorithm relies on segmenting the data into pre- and post-stimulus periods, learning the statistics of the background activity from the pre-stimulus period, and then applying the statistics of the background activity to the post-stimulus data to uncover the stimulus-evoked activity. The underlying assumption is that noise and non-stimulus-locked brain activity present in the pre-stimulus period continues into the post-stimulus period, where the stimulus-evoked activity is linearly superimposed on top of the pre-stimulus activity.

We model post-stimulus sensor data (B_{post}) as:

$$B_{post} = \sum_{r=1}^{d_s} L_r S_r + \mathcal{E}, \tag{2.1}$$

where $B_{post} \in \mathbb{R}^{d_b \times d_t}$, where d_b equals the number of sensors and d_t is the number of time points at which measurements are made in the post-stimulus period. $L_r \in \mathbb{R}^{d_b \times d_c}$ is the lead field matrix in d_c orientations for the r-th voxel. Each unknown source $S_r \in \mathbb{R}^{d_c \times d_t}$ is a d_c-dimensional neural current-dipole source at d_t time points, projecting from the i-th voxel. There are d_s voxels under consideration. $\mathcal{E} \in \mathbb{R}^{d_b \times d_b}$ is a *noise-plus-interference* factor that is learned from the pre-stimulus period using Stimulus-Evoked Factor Analysis (SEFA)Nagarajan et al. (2007). Learning $\mathcal{E}$ is the first step of the Champagne algorithm. as discussed in Section 2.2.2. The second step to the source localization process is to estimate hyperparameters Γ that govern a statistical model of the post-stimulus data.

We can fully define the probability distribution of the data conditioned on the sources:

$$p(B_{post}|S) \propto \exp\left(-\frac{1}{2}\left\|B_{post} - \sum_{r=1}^{d_s} L_r S_r\right\|^2_{\Sigma_\epsilon^{-1}}\right), \tag{2.2}$$

where Σ_ϵ^{-1} is learned using SEFA from the pre-stimulus data (described in the Learning the *Noise-plus-Interference* Factor section below) and $\|X\|_W$ denotes the weighted matrix norm $\sqrt{\text{trace}[X^T W X]}$.

We assume the following for the source prior on S:

$$p\left(S|\Gamma\right) \propto \exp\left(-\frac{1}{2}\text{trace}\left[\sum_{r=1}^{d_s} S_r^T \Gamma_r^{-1} S_r\right]\right). \tag{2.3}$$

This is equivalent to applying independently, at each time point, a zero-mean Gaussian distribution with covariance Γ_r to each source S_r. We define Γ to be the $d_s d_c \times d_s d_c$ block-diagonal matrix formed by ordering each Γ_r along the diagonal of an otherwise zero-valued matrix. If the lead field has only one orientation (scalar/ orientation-constrained lead field), Γ reduces to a diagonal matrix.

Since Γ is unknown, we can find a suitable approximation $\hat{\Gamma} \approx \Gamma$ by integrating out the sources S of the joint distribution $p(S, B|\Gamma) \propto p(B|S)p(S|\Gamma)$ and then minimizing the cost function:

$$\mathcal{L}(\Gamma) \triangleq -2\log p(B|\Gamma) \equiv \text{trace}\left[C_b \Sigma_b^{-1}\right] + \log\left|\Sigma_b\right|, \tag{2.4}$$

where $C_b \triangleq d_t^{-1} B B^T$ is the empirical covariance and Σ_b is the data model covariance, $\Sigma_b = \Sigma_\epsilon + L\Gamma L^T$.

The first term of (2.4) is a measure of the dissimilarity between the empirical data covariance C_b and the model data covariance Σ_b; in general, this factor encourages Γ to be large because it is convex and nonincreasing in Γ (in a simplified scalar case, this is akin to minimizing $1/x$ with respect to x, which of course naturally favors x being large). The second term provides a regularizing or sparsifying effect, penalizing a measure of the volume formed by the model covariance Σ_b. The determinant of a matrix is equal to the product of its eigenvalues, a well-known volumetric measure. Since the volume of any high dimensional space is more effectively reduced by collapsing individual dimensions as close to zero as possible (as opposed to incrementally reducing all dimensions isometrically), this penalty term promotes a model covariance that is maximally degenerate (or non-spherical), which pushes elements of Γ to exactly zero (resulting in hyperparameter sparsity).

2.2.1 Estimating the Hyperparameters (Γ)

We can minimize this cost function (2.4) with respect to Γ by employing the expectation-maximization (EM) algorithm. The E-step of the EM algorithm corresponds to computing the posterior probability $p(S|B_{post})$, which is Gaussian and be expressed as follows, given Bayes Rule, as a function of time:

$$p(S(t)|B_{post}(t)) = \frac{p(B_{post}(t)|S(t))p(S(t))}{p(B_{post}(t))}$$

Considering the log of the posterior probability, we obtain:

$$log\, p(S(t)|B_{post}(t)) = log\, p(B_{post}(t)|S(t)) + log\, p(S(t)) - log\, p(B_{post}(t))$$

Since the posterior probability $p(S(t)|B_{post}(t))$ is Gaussian, we can find an expression for its mean, $\bar{S}(t)$, by taking the first gradient with respect to $S(t)$, setting it equal to zero and solving for $S(t)$:

$$\begin{aligned}
\frac{\partial}{\partial S(t)} log\ p(S(t)|B_{post}(t)) = \frac{\partial}{\partial S(t)} log\ p(B_{post}(t)|S(t)) + \frac{\partial}{\partial S(t)} log\ p(S) = 0 \\
\frac{\partial}{\partial S(t)}(-\frac{1}{2}(B_{post}(t) - LS(t))^T\Sigma_\epsilon(B_{post}(t) - LS(t)) - \frac{1}{2}S(t)^T\Gamma S(t) = 0 \\
(L^T\Sigma_\epsilon L)S(t) - L^T\Sigma_\epsilon B_{post}(t) + \Gamma S(t) = 0
\end{aligned}$$

where solving for $S(t)$ results $\bar{S}(t)$, the mean of the posterior:

$$\bar{S}(t) \quad = \quad L^T\Sigma_\epsilon(L^T\Sigma_\epsilon L + \Gamma)^{-1}B_{post}(t)$$

The M-step corresponds to updating the parameters that maximize the averaged complete data likelihood $\bar{\ell}$ defined as:

$$\bar{\ell} = \int dS\ \ p(S|B_{post}(t))\ \ \ell(B_{post}(t), S(t))$$

Maximizing the log-likelihood is the same as maximizing the likelihood, thus the parameter update rules for the M-step are derived as:

$$\begin{aligned}
log\ \bar{\ell} \quad &= \quad E_{p(S(t)|B_{post}(t))}\left\{\sum_{t=1}^{d_t}[log\ p(B_{post}(t)|S(t)) + log\ p(S(t))]\right\} \\
\frac{\partial}{\partial\Gamma} log\ \bar{\ell} \quad &= \quad \frac{\partial}{\partial\Gamma}E_{p(S(t)|B_{post}(t))}\left\{\sum_{t=1}^{d_t}\frac{1}{2} \mid \Gamma \mid + \frac{1}{2}S(t)^T\Gamma S(t)\right\}
\end{aligned}$$

Setting the gradient equal to zero:

$$\begin{aligned} E_{p(S(t)|B_{post}(t))}\left\{\sum_{n=1}^{N}\frac{1}{\Gamma}+S(t)S(t)^T\right\} &= 0 \\ \frac{d_t}{\mid\Gamma\mid}+d_t\Gamma^{-1}+\sum_{t=1}^{d_t}\bar{S}(t)\bar{S}(t)^T &= 0 \end{aligned}$$

results in the following update rule for the source precision (Γ):

$$\Gamma=\frac{d_t}{R_{SS}}$$

where

$$R_{SS}=d_t\Gamma^{-1}+\sum_{t=1}^{d_t}\bar{S}(t)\bar{S}(t)^T$$

The EM iterations consist of first computing the mean of the posterior distribution, $\bar{S}(t)$ and then maximizing the log-likelihood by computing the new parameter values Γ.

The EM implementation (and other generic methods) are exceedingly slow when d_s is large. Instead, as described in Wipf et al. (2009, 2010), we utilize an alternative optimization procedure that expands upon ideas from Sato et al. (2004); Wipf et al. (2007), handles arbitrary/unknown dipole-source orientations, and converges quickly. In Figure 2.1, we demonstrate the convergence rates of the EM implementation versus the proposed method described below. Figure 2.1 displays the reduction in the Champagne cost function $\mathcal{L}(\Gamma)$ as a function of the iteration number. The convergence rates are quite disparate, the proposed method converges after 20 iterations while the EM implementation is still decreasing the cost function after 100 iterations. Consistent with previous observations, the EM updates are considerably slower in reducing the cost. While the detailed rationale for this performance discrepancy is beyond the scope of this chapter, ultimately it is a consequence of the different underlying bounds being used to form auxiliary functions. EM leads to slower convergence because it is effectively using a *much* looser bound around zero than the bound described in Section 2.2.1 and therefore fails to fully penalize redundant or superfluous components. This prevents the associated hyperparameters from going to zero very quickly, drastically slowing the convergence rate. More details on this topic can be found in Wipf and Nagarajan (2010).

To begin, we note that $\mathcal{L}(\Gamma)$ only depends on the data B through the $d_b \times d_b$ sample correlation matrix C_b. Therefore, to reduce the computational burden, we replace B with a matrix $\widetilde{B} \in \mathbb{R}^{d_b \times \text{rank}(B)}$ such that $\widetilde{B}\widetilde{B}^T = C_b$. This removes any per-iteration dependency on d_t, which can potentially be large, without altering that actual cost function. It also implies that, for purposes of computing Γ, the number of columns of S is reduced to match $\text{rank}(B)$. We now re-express the cost function $\mathcal{L}(\Gamma)$ in an alternative form leading to convenient update rules and, by

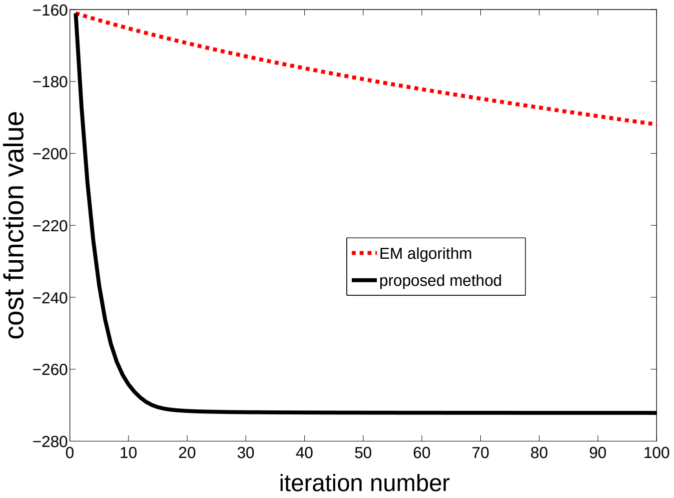

Figure 2.1 Convergence rate: EM implementation compared to upper-bounding hyperplane implementation (proposed method).

construction, a proof that $\mathcal{L}\left(\Gamma^{(k+1)}\right) \leq \mathcal{L}\left(\Gamma^{(k)}\right)$ at each iteration.

The procedure we will use involves constructing auxiliary functions using sets of hyperplanes. First, the log-determinant term of $\mathcal{L}(\Gamma)$ is a concave function of Γ and so it can be expressed as a minimum over upper-bounding hyperplanes via :

$$\log|\Sigma_b| = \min_Z \left[\sum_{i=1}^{d_s} \operatorname{trace}\left(Z_i^T \Gamma_i\right) - h^*(Z) \right], \tag{2.5}$$

where $Z \triangleq \left[Z_1^T, \ldots, Z_{d_s}^T\right]^T$ is a matrix of auxiliary variables that differentiates each hyperplane and $h^*(Z)$ is the concave conjugate of $\log|\Sigma_b|$. While $h^*(Z)$ is unavailable in closed form, for our purposes below, we will never actually have to compute this function. Next, the data fit term is a concave function of Γ^{-1} and so it can be also be expressed using similar methodology as:

$$\operatorname{trace}\left[C_b\Sigma_b^{-1}\right] = \min_X \left[\left\|\widetilde{B} - \sum_{i=1}^{d_s} L_i X_i\right\|^2_{\Sigma_\epsilon^{-1}} + \sum_{i=1}^{d_s} \|X_i\|^2_{\Gamma_i^{-1}}\right], \tag{2.6}$$

where $X \triangleq \left[X_1^T, \ldots, X_{d_s}^T\right]^T$ is a matrix of auxiliary variables as before. Note that in this case, the implicit concave conjugate function exists in closed form.

Dropping the minimizations and combining terms from (2.5) and (2.6) leads to the modified cost function:

$$\mathcal{L}(\Gamma, X, Z) = \left\|\widetilde{B} - \sum_{i=1}^{d_s} L_i X_i\right\|^2_{\Sigma_\epsilon^{-1}} + \sum_{i=1}^{d_s} \left[\|X_i\|^2_{\Gamma_i^{-1}} + \operatorname{trace}\left(Z_i^T\Gamma_i\right)\right] - h^*(Z), \tag{2.7}$$

where by construction $\mathcal{L}(\Gamma) = \min_X \min_Z \mathcal{L}(\Gamma, X, Z)$. It is straightforward to show that if $\{\hat{\Gamma}, \hat{X}, \hat{Z}\}$ is a local (global) minimum to $\mathcal{L}(\Gamma, X, Z)$, then $\hat{\Gamma}$ is a local (global) minimum to $\mathcal{L}(\Gamma)$.

Since direct optimization of $\mathcal{L}(\Gamma)$ may be difficult, we can instead iteratively optimize $\mathcal{L}(\Gamma, X, Z)$ via coordinate descent over Γ, X, and Z. In each case, when two are held fixed, the third can be globally minimized in closed form. This ensures that each cycle will reduce $\mathcal{L}(\Gamma, X, Z)$, but more importantly, will reduce $\mathcal{L}(\Gamma)$ (or leave it unchanged if a fixed-point or limit cycle is reached). The associated update rules from this process are as follows:

$$X_r^{\text{new}} \rightarrow \Gamma_r L_r^T \Sigma_b^{-1} \widetilde{B} \tag{2.8}$$

$$Z_r^{\text{new}} \to \nabla_{\Gamma_r} \log |\Sigma_b| = L_r^T \Sigma_b^{-1} L_r \tag{2.9}$$

$$\Gamma_r^{\text{new}} \to Z_r^{-1/2} \left(Z_r^{1/2} X_r X_r^T Z_r^{1/2} \right)^{1/2} Z_r^{-1/2}, \tag{2.10}$$

where Γ_r comprise the blocks of the block-diagonal matrix of hyperparameters Γ.

Each Γ_r is initialized with an identity matrix plus/minus a small random number, $O(1e-5)$. We found that this was the most robust initialization, as opposed to initializing with the source power of another algorithm, such as MVAB. When using a vector lead field, as opposed to a scalar/orientation-constrained lead field, a $dc \times dc$ covariance matrix is learned for each source. This covariance can be thought of as describing a noisy or unfixed source orientation.

In summary, the Champagne algorithm estimates Γ by iterating between (2.8), (2.9), and (2.10), and with each pass we are theoretically guaranteed to reduce (or leave unchanged) $\mathcal{L}(\Gamma)$.

The source time courses can be calculated using the posterior source distribution $p(S|B_{post}, \Gamma) \propto p(B_{post}|S)p(S|\Gamma)$, which is Gaussian. To estimate the source time courses, we choose the source posterior mean, i.e. the mean of $p(S|B_{post}, \Gamma)$), given by:

$$\hat{s}_r(t) = \Gamma_r L_r^T \left(\Sigma_\epsilon + L\Gamma L^T\right)^{-1} B_{post}(t), \tag{2.11}$$

where $\hat{s}_r(t) \in \mathbb{R}^{d_c \times 1}$ (a short vector across lead field components).

The sparsity in the hyperparameters, as discussed above, results in sparsity in the source time courses. As seen in Equation (2.11), we can see that if $\Gamma_r = 0$ then the source time course $\hat{s}_r(t)$ will also be zero.

The Champagne algorithm typically converges in 75 to 100 iterations. The lead field we used for the majority of the simulations had approximately 5500 voxels, which results in the estimation of approximately 16,500 hyperparameters, and takes approximately 10 minutes to run on a lead field of that size (and trial-averaged data) on an Intel(R) Core(TM)2 Quad CPU @ 3.00GHz, with 8GB of memory. The number of hyperparameters Champagne estimates is dependent of the number of lead field orientations. For a scalar (or fixed orientation) lead field, d_s hyperparameters are estimated. For a 2 or 3 component lead field, $3 \times d_s$ or $6 \times d_s$ unknown hyperparameters are estimated, respectively. (Each precision matrix Γ_r is a symmetric matrix.) We have plotted the magnitude of the hyperparameters over 1, 5, 10, 20, 50, and 100 iterations to show the pruning in Figure 2.2. For this example, after only 20 iterations the hyperparamters have been vastly pruned and do not change significantly over the next 80 iterations. If the value of a hyperparameter falls below a certain threshold (a parameter that is set in the algorithm and is usually $1e-8$), that voxel is pruned from the lead field. This pruning speeds up the time per iteration, as over time there are fewer and fewer voxels to consider.

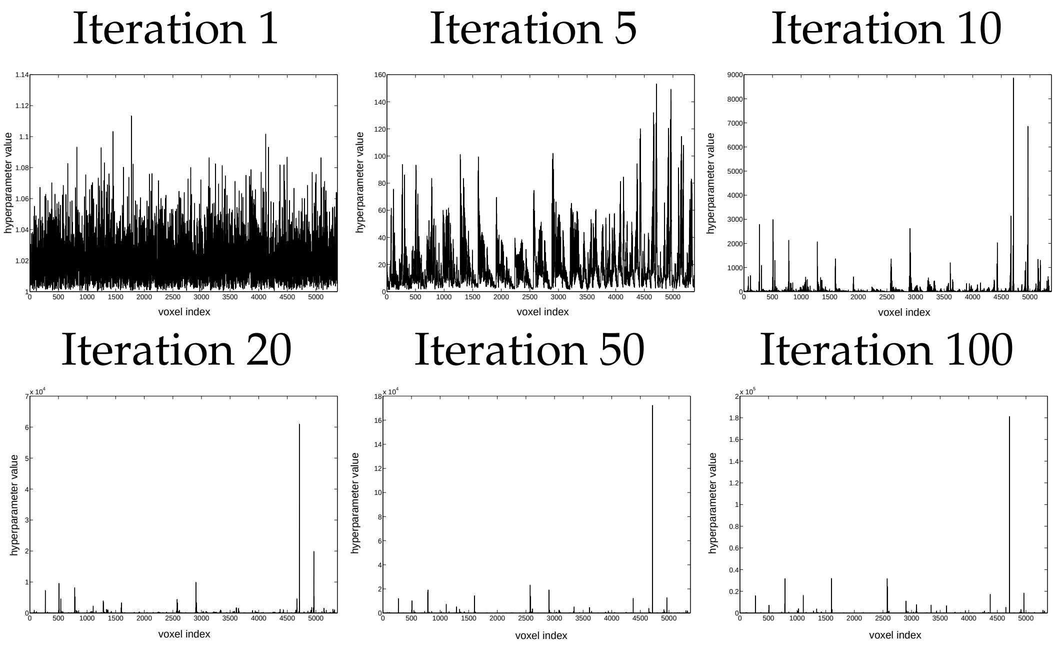

Figure 2.2 Demonstration of the pruning of the hyperparameters over iterations by Champagne. This figures shows the initialization (Iteration 1) and the subsequent pruning after 5, 10. 20, 50, and 100 iterations.

2.2.2 Learning the *Noise-plus-Interference* Factor (Σ_ϵ)

The learning procedure described in the previous section boils down to fitting a structured maximum likelihood covariance estimate $\Sigma_b = \Sigma_\epsilon + L\Gamma L^T$ to the data covariance C_b. The idea here is that $L\Gamma L^T$ will reflect the brain signals of interest while Σ_ϵ will capture all interfering factors, e.g., spontaneous brain activity, sensor noise, muscle artifacts, etc. Since Σ_ϵ is unknown, it must somehow be estimated or otherwise accounted for. Given access to pre-stimulus data (i.e., data assumed to have no signal/sources of interest), a factor analysis model provides a powerful means of decomposing a data covariance matrix C_b into signal and interference components.

There are many possible ways to learn Σ_ϵ; we have found that the Stimulus-Evoked Factor Analysis (SEFA) model is the most robust Zumer et al. (2007). Essentially SEFA decomposes the empirical sensor covariance as follows:

$$C_b \approx \Lambda + AA^T + DD^T, \tag{2.12}$$

where $A \in \mathbb{R}^{d_b \times d_e}$ represents a matrix of learned interference factors, Λ is a diagonal noise matrix, and $D \in \mathbb{R}^{d_b \times d_f}$ represents signal factors. Both the number of interference factors d_e and the number of signal factors d_f are learned from the data via a variational Bayesian factor analysis procedure. Using a generalized form of the expectation maximization algorithm, the method attempts to find a small number of factors that adequately explains the observed sensor data covariance during both the pre- and post-stimulus periods. The pre-stimulus is modeled with a covariance restricted to the terms $\Lambda + AA^T$, while the post-stimulus covariance, which contains the signal information DD^T we wish to localize, is expressed additively as in (4.3). Then, we can set $\Sigma_\epsilon \to \Lambda + AA^T$. Λ is learned first from the pre-stimulus period and then is updated using the post-stimulus period data; learning Λ accounts for nonstationarity in the sensor noise between the pre-and post-stimulus periods. Note that SEFA does not use the lead field L to model the activity, rather A and D are mixing matrices that do not have any physical (or physiological) interpretation and are simply statistical models used to suppress pre-stimulus activity from the post-stimulus sensor covariance.

2.3 Discussion

This chapter derives a novel empirical Bayesian algorithm, Champagne, for M/EEG source reconstruction that readily handles multiple correlated sources with unknown orientations, a situation that commonly arises even with simple imaging tasks. Based on a principled cost function and fast, convergent update rules, this procedure displays significant theoretical and empirical advantages over many existing methods. In the following chapters of this thesis, we will demonstrate Champagne's performance on simulated and real data sets. We also investigate alternative methods to estimate the noise-plus-interference factor (Σ_ϵ), including when there is not pre-stimulus sensor data.

Chapter 3

Performance Evaluation of the Champagne Algorithm on Simulated and Real M/EEG Data [1]

3.1 Overview

In this chapter, we present results from both simulated and real MEG data that more rigorously characterize Champagne's performance and better delineate both its strengths and limits. For the simulated data, we have selected a number of challenging source configurations: a large number of distinct sources, deep sources, and clusters of sources. We also present the results on simulated data us-

---

[1] Portions of this chapter originally appeared as an ISBI conference paper (Owen et al., 2009) and as an article in (Wipf et al., 2010)

ing benchmark algorithms, minimum-variance adaptive beamforming (MVAB), dynamic statistical parametric mapping and standardized low resolution brain electromagnetic tomography (sLORETA/dSPM), and minimum current estimate (MCE). In addition to investigating these challenging source configurations with simulated data, we also present an analysis of the effect of lead field modeling errors. In addition, we present results from simulated EEG and we investigate the effect of subsampling the number of EEG sensors. We use six real data sets to test Champagne's performance as compared to the benchmark algorithms: somatosensory-evoked field (SEF), auditory-evoked field (AEF), audio-visual integration task, a face-processing data, with both MEG and EEG recordings, and resting-state data form an epilepsy patient. Taken together, the results on simulated and real data give a rather complete picture of Champagne's source localization capabilities as we have explored the full parameter space for the simulations and have applied the algorithm to several real data sets. Ultimately an exhaustive performance analysis of a source localization algorithm serves two distinct purposes. The first being to motivate the use of a new algorithm by showing head-to-head performance with commonly-used methods. The second is to fully describe when an algorithm works and of equal importance, when it does not work. The conditions under which an algorithm fails motivates further development of source localization algorithms to advance beyond these shortcomings.

3.2 Benchmark Source Localization Algorithms

We chose to test the Champagne algorithm against four representative source localization algorithms from the literature: an adaptive spatial filtering method, minimum-variance adaptive beamforming (MVAB) , two non-adaptive spatial filtering methods, standardized low resolution brain electromagnetic tomography (sLORETA), and dynamic statistical parametric mapping (dSPM), and a version of minimum current estimation (MCE). For all the algorithms, we localize sources only in the post-stimulus period.

MVAB

Minimum variance adaptive beamformer (MVAB) spatial filters (Sekihara and Nagarajan, 2008) rely on the assumption that sources are temporally uncorrelated with each other and that the number of active sources is less than the number of sensors (regardless of how many voxels are included in the source reconstruction volume). The minimum-variance methods find the weight that minimizes the output power $\hat{s}_r^T \hat{s}_r$, where $\hat{s}_r = w_r^T B_{post}$ and therefore, minimizing the output power is equivalent to minimizing $w_r^T R w_r$, where $R = B_{post} B_{post}^T$ is the covariance of the sensor data. This minimization can be done subject to various constraints; using a unit-gain constraint is a widely-used form. The output power should be minimized since it usually has large contributions from noise and source activity at voxels other than r.

However, the weight should also strive to maximize the pass-through of the source power at the current voxel. Using a Lagrange multiplier with the unit-gain constraint $w_r^T L^r = 1$ yields the unit-gain-constraint minimum-variance adaptive beamformer. The source-time course ($\hat{s}_r(t)$) at a particular voxel (r) and time point (t) for MVAB is:

$$\hat{s}_r(t) = \left(L_r^T R^{-1} L_r\right)^{-1} L_r^T R^{-1} B_{post}(t), \tag{3.1}$$

sLORETA/dSPM

The source localization algorithms sLORETA and dSPM belong to the general family of minimum norm estimates (MNE). They are non-adaptive, meaning that they do not consider the sensor data in obtaining the spatial filter that will be applied to the data in order to obtain the source estimates. The non-adaptive property gives sLORETA and dSPM robustness to correlated sources that often impeded algorithms, like MVAB. If the sensor noise is ignored, the simplest initial guess for the weight would be the inverse of the forward field matrix. However, non-square matrices cannot be inverted; instead a pseudo-inverse can approximate a left- or right-sided inverse. The source estimate $\hat{s}_r(t)$ that minimizes the linear least-squares error is the one that uses the Moore-Penrose pseudo-inverse as the weight:

$$\hat{s}_r(t) = L_r^+ B_{post}(t) = L_r^T G^{-1} B_{post}(t) \tag{3.2}$$

where L_r indicates the lead field for voxel r and the $^+$ indicates the pseudo-inverse. G is called the Gram matrix, $G = LL^T$; it gives an indication of the spatial correlation or overlap between the sensitivity profile of the sensors. Since neighboring sensors have similar sensitivities, this Gram matrix is too close to singular to invert stably. Tikhonov regularization is needed to invert G, which involves adding a small amount of noise to the diagonal of the Gram matrix prior to inversion: $G^{-1} = (LL^T + \sigma\lambda I)^{-1}$. $\sigma\lambda$ regularizes the inverse; σ is a scalar that is chosen empirically, ranging from $[1e-2 : 1e2]$, and λ is the maximum eigenvalue of the data covariance matrix (R), which helps get the regularization factor in the right range. The Gram matrix is dependent on all the voxels that could contain source activity, thus, the source activity at a particular voxel is dependent on the forward field for the whole volume. The weights used in the adaptive methods, such as MVAB, only rely on the forward field for the voxel of interest.

dSPM, proposed in Dale et al. (2000), modifies this weight by normalizing it to ensure a uniform noise spatial distribution. This solution can also be obtained by finding the weight that minimizes $w^T G w$ subject to the unit-noise constraint $w^T w = 1$. Hence, the weight-normalized minimum-norm solution for dSPM is:

$$\hat{s}_r(t) = \left(tr\left\{L_r^T G^{-2} L_r\right\}\right)^{-\frac{1}{2}} L_r^T G^{-1} B_{post}(t). \tag{3.3}$$

sLORETA, proposed in (Pascual-Marqui, 2002), modifies the above weight-normalized form, where the source time course estimate is:

$$\hat{s}_r(t) = \left(L_r^T G^{-1} L_r\right)^{-\frac{1}{2}} L_r^T G^{-1} B_{post}(t), \tag{3.4}$$

The sLORETA weight is equivalent to normalizing the min-norm solution by the square root of the resolution kernel and, alternatively, can be derived by finding the weight that minimizes $w^T G w$ subject to the unit-total-leakage constraint $w^T G w = 1$.

We find that SL and dSPM yield almost identical results, thus we present their results together in this paper.

MCE

The MCE algorithm was first proposed in (Uutela et al., 1999) and can be formulated by the following. Recall that the MEG signals can be modeled by the linear model

$$B_{post} = \sum_{r=1}^{d_s} L_r S_r + \mathcal{E} \tag{3.5}$$

The MCE algorithm consists of minimizing the norm of the sources $||S||$ with constraint that $B_{post} = \sum_{r=1}^{d_s} L_r S_r$. The variants of the MCE algorithm differ in the selection of the norm, the constraints placed upon the minimization, and the errors allowed in the constraint.

We used a version of MCE that is specially tailored to handle multiple time-points and unconstrained source orientations (Wipf et al., 2009). This method extends standard MCE by applying a ℓ_2 penalty across time. In this version there is an ℓ_1-norm over space and an ℓ_2-norm over time, sometimes called an $\ell_{1,2}$-norm in signal processing. The source orientation components are also included within the ℓ_2 penalty. Similar to Champagne, MCE favors sparse/compact source reconstructions.

The MCE cost function is expressed as:

$$\hat{s}(t) = \arg\min_{\hat{s}} \left[\left\| B - \sum_{r=1}^{d_s} L_r \hat{s}_r \right\|^2_{\Sigma_\epsilon^{-1}} + \sum_{r=1}^{d_s} ||\hat{s}_r||_{\mathcal{F}} \right], \tag{3.6}$$

where $||\hat{s}_r||_{\mathcal{F}} = \sqrt{tr(\hat{s}_r^T \hat{s}_r)}$ and Σ_ϵ^{-1} is either set to be the inverse noise covariance learned with SEFA from the pre-stimulus period (as in Champagne) or it is set to be $\Sigma_\epsilon^{-1} = \sigma I$, where σ is a scalar that is empirically selected and I is an identity matrix of size $d_s \times d_s$.

And the estimate for the source time course ($\hat{s}_r(t)$) at a particular voxel (r) and time point (t) is obtained by iterating the following equations:

$$\hat{s}_r^{(k+1)}(t) \rightarrow \alpha_r^{(k)} L_r^T \left(\Sigma_\epsilon + \sum_{r=1}^{d_s} \alpha_r^{(k)} L_r L_r^T \right)^{-1} B_{post}(t), \tag{3.7}$$

where $\alpha_r^{(k)} = \sqrt{\frac{1}{d_c d_t} \sum_{t=1}^{d_t} \|\hat{s}(r,t)^{(k)}\|_2^2}$ and k is the iteration number.

3.3 Assessing Localization Accuracy

In order to evaluate performance, we used two features: localization accuracy and time course estimation accuracy. To assess localization accuracy, we used the A' metric (Snodgrass and Corwin, 1988; Darvas et al., 2004). When assessing the localization accuracy, it is important to take into account both the number of hits (sources that were correctly localized) and the presence of false positives, or spurious localizations. A principled way to take these two features into account is the ROC (receiver-operator characteristic) method, which is a technique that balances hit rate (H_R) versus false-positive rate (F_R). Specifically, we used the A' metric, which is a way to approximate the area under the ROC using one H_R/F_R pair. The A′ values range from 0 to 1, with 0 indicating all false positives and 1 indicating all hits. We determined H_R and F_R at each SNIR level and each algorithm in the following way. For each simulation, we calculated all the local peaks in the power map; the power map is a three-dimensional image of the power of the source time course at each voxel location, A local peak was defined as a voxel that is greater than its 20 three-dimensional nearest neighbors and was at least 10 percent of the maximum activation of the image. (This thresholding at 10 percent is designed to filter out any spurious peaks or ripples in the image that are much weaker than the maximum peak.)

After scanning through the voxel grid and locating all the local peak locations, we tested whether each local peak was within ten millimeters of a true source location. If a particular peak was within 10mm of a seeded source, that peak gets labeled as a hit and if there was not a seeded source within 10mm, that peak gets labeled as a false positive. H_R is then calculated by dividing the number of hits by the true number of seeded sources. (A 8mm voxel grid was used for the simulations, so a source would have to be in the true seeded location or in an adjacent voxel in order to be labeled a hit.) Determining F_R is more tenuous. There is not a clear maximum number of possible false positives, as there is with hits. We empirically determined the maximum number of false positives for each algorithm, for each given experiment. Since the spatial filtering techniques (MVAB, sLORETA, and dSPM) are inherently more smooth than the sparse solutions obtained from Champagne and MCE, the maximum number of false positives was determined empirically across all SNIR (signal to noise plus interference ratio) levels and 50 simulations for each algorithm. Then, the number of false positives was divided by this maximum false positive number in order to calculate F_R for each simulation and SNIR level. Lastly, the A′ metric was calculated for each

H_R/F_R pair with the following equations:

$$A' = \begin{cases} \frac{1}{2} + \frac{(H_R - F_R)(1 + H_R - F_R)}{4H_R(1-F_R)} & \text{for } H_R \geq F_R \\ \\ \frac{1}{2} + \frac{(F_R - H_R)(1 + F_R - H_R)}{4F_R(1-H_R)} & \text{for } F_R > H_R \end{cases} \tag{3.8}$$

To assess the accuracy of the time course estimates, we used the correlation coefficient between the true and estimated time courses. Note that we only computed the time course correlation if a local peak is deemed a hit. For a particular simulation, we average the correlation coefficients to obtain one number that reflects the time course reconstruction. The correlation coefficient also ranges from 0 to 1, with 1 implying perfect time course estimation. The equation used to obtain the average correlation coefficient ($\bar{R}$) between the true source time course $S_i^{true}(t)$ and the estimated source time course $S_i^{est}(t)$ for the i^{th} correctly localized source is as follows:

$$\bar{R} = \frac{1}{N} \sum_{i=1}^{n} \frac{C(S_i^{true}(t), S_i^{est}(t))}{\sqrt{C(S_i^{true}(t), S_i^{true}(t)) C(S_i^{est}(t), S_i^{est}(t))}} \tag{3.9}$$

where $C(x_i, x_j)$ is the covariance of x_i and x_j and N is the total number of correctly localized sources.

From this procedure, we obtain three useful metrics: hit rate (H_R), A′ metric, and average correlation coefficient ($\bar{R}$). We have developed a metric that captures both the accuracy of the location and the time courses of the algorithms, which we call the *Aggregate Performance*, (AP). It combines the A', $\bar{R}$, and H_R, in the following equation is:

$$AP = \frac{1}{2}(A' + H_R\bar{R}) \tag{3.10}$$

We use the H_R as a weight for $\bar{R}$ since we only compute the correlation coefficient of the sources that are correctly localized. AP also ranges from 0 to 1. For this paper, we decided to use an AP value of 0.75 as the cutoff for a successful localization.

3.4 Generating Simulated Data

The simulated data in this chapter were generated by simulating dipole sources, either fixed in orientation or with some variation in orientation. We seeded the voxel locations with damped sinusoidal time courses (except where noted) and then projecting the voxel activity to the sensors with the lead field. The brain volume was segmented into 8mm voxels and a two-orientation ($d_c = 2$) forward lead field was calculated using a single spherical-shell model (Sarvas, 1987) implemented in NUTMEG (Dalal et al., 2004) unless where otherwise noted. The volume of interest (VOI) is volumetric, as opposed to cortically constrained, and results in lead field with approximately 4,500 voxels. The data time course was partitioned into pre- and post-stimulus periods. The pre-stimulus period

(270 samples) contained only sensor noise and interfering brain activity and in the post-stimulus period (450 samples), the activity of interest, or the stimulus-evoked activity, was superimposed on the noise and interference present in the pre-stimulus period. The noise plus interference activity consisted of the resting-state sensor recordings (except where noted) collected from a human subject presumed to have only spontaneous activity (i.e., non-stimulus evoked sources) and sensor noise. Each source location was seeded with a distinct time course of activity, described in more detail below, and the sources were only present for half of the post-stimulus period (225 samples). The voxel activity was projected to the sensors through the lead field and the noise was added to achieve as desired signal to noise ratio, also described below.

The locations for the sources were chosen with a variety of constraints, as detailed in the experimental sections below, to test Champagne's performance on challenging source configurations. We could adjust both the signal-to-noise-plus-interference ratio (SNIR), the correlations between the different voxel time courses (inter-dipole), and the correlation between the orientations of the dipolar sources (intra-dipole). For our purposes, SNIR is defined as:

$$\text{SNIR} \triangleq 20 \log \frac{\|LS\|_{\mathcal{F}}}{\|\mathcal{E}\|_{\mathcal{F}}}, \tag{3.11}$$

where L is the lead field matrix, S are the source time courses, $\mathcal{E}$ is the noise or non-stimulus evoked activity, see () and $\mathcal{F}$ denotes the Froebenius norm or the

ℓ_2-norm.

We simulated data to replicate the signal-to-noise ratios of real stimulus-evoked data (trial-averaged data). We chose 10dB and 0dB as the upper and lower bound for typical trial-averaged, real MEG data.

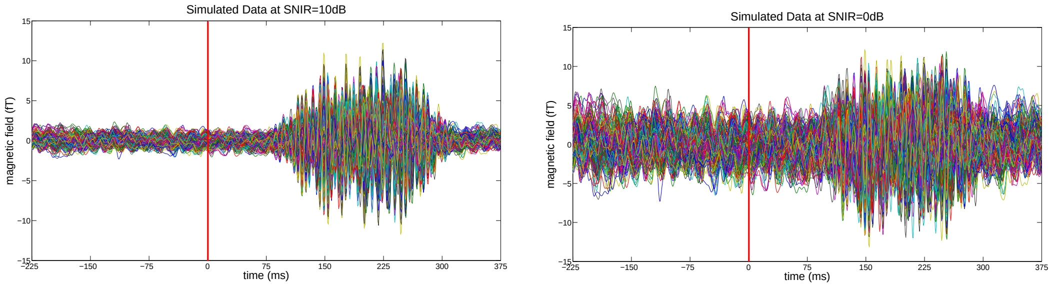

Figure 3.1 Example of a simulated data at 10dB and 0dB. The red line indicates the start of the post-stimulus period (or the stimulus onset) at 0ms. There are 5 sources seeded in this example and the noise is real resting-state data.

When simulating data, we can set the inter-dipole correlations (the amount of correlation between the dipoles seeded throughout the brain) and the intra-dipole correlations (the correlation between the orientations of the lead field for each individual dipole). The inter-dipole correlations allow us to test the algorithms' performance when disparate areas of the brain have correlated activity. The intra-dipole correlations allow us to model deviations in the dipole's orientation as a correlation of 1 would imply a fixed orientation and a correlation of 0 would model a dipole that is fully rotating and not stable in orientation. Another interpretation of the intra-dipole correlation is that the brain volume within a single voxel might contain two or more current sources, and these sources might be aligned in different directions. This would be true of voxels that fall on a 90 degree fold of the cortical surface. We impose intra- and inter-dipolar correlations (for a two-orientation lead field) for n dipoles in the following manner.

First we generate the time courses for the two orientations of each dipole: $S_n^1(t)$, the time course of the n_{th} dipole and the first orientation and $S_n^2(t)$, the time course of the n_{th} dipole and the second orientation. Then we correlate the two orientations of each individual dipole. If the intra-dipolar correlation coefficient is α_{intra}, then $S_n^2(t)$ is defined to be:

$$S_n^2(t) \rightarrow \alpha_{intra} S_n^1(t) + \sqrt{1-\alpha_{intra}^2} S_n^2(t) \tag{3.12}$$

After we set the intra-dipolar correlation, we correlate all the dipoles (n in number) to the first dipole, $S_1(t)$. $S_n(t)$ is a vector formed by concatenating the time courses for each orientation for the nth dipole: $S_n(t) = \begin{bmatrix} S_n^1(t) \\ S_n^2(t) \end{bmatrix}$, for a two-component lead field. If the inter-dipolar correlation coefficient is α_{inter}, then $S_n(t)$ (for $n \neq 1$) is defined to be:

$$S_n(t) \rightarrow \begin{bmatrix} \alpha_{inter} & 0 \\ 0 & \alpha_{inter} \end{bmatrix} S_1(t) + \begin{bmatrix} \sqrt{1-\alpha_{inter}^2} & 0 \\ 0 & \sqrt{1-\alpha_{inter}^2} \end{bmatrix} S_n(t) \tag{3.13}$$

Two examples of the simulated sensor data at 0dB and 10dB are shown in

Figure 3.1. The red line indicates the start of the post-stimulus period (effectively the stimulus onset) at 0ms. In this example, there are 5 sources seeded throughout the brain.

3.5 Performance on Difficult Source Configurations

In two previous publications (Wipf et al., 2009, 2010), Champagne was demonstrated to be robust to correlated sources and capable of localizing a large number (up to ten) of dipoles. Here its performance has been further tested to explore the limits of its ability to reconstruct complex simulated brain activity. We also compare Champagne's performance to existing, commonly-used source localization algorithms (as described above). The results we obtained using simulated data are presented in two forms. First, we show the plots of mean A', $\bar{R}$, H_R, and/or our aggregate performance metric, AP. For each configuration, the results are averaged over 50 simulations and we have plotted these averaged results with standard error bars. These plots give a feel for overall performance as the sources are randomly seeded (with some constraints). For some of the experiments, we also show examples of the localization results from single simulations, which complement our aggregate results. We compute the source power at every voxel and project the activity to the surface of a rendered MNI-template brain. These plots contain a projection of the true source power, called *ground truth*.

3.5.1 Discriminating Two Sources

We examined the minimum distance at which two sources are able to be resolved. The spacing between voxels on our grid was 8mm, thus we tested the localization accuracy when the distance between the two sources (inter-dipole distance) was 16, 24, 32, or 40mm. The locations of the two sources were chosen randomly with the constraint that the minimum distance from the center of the head was 35mm (since deeper sources are harder to localize). (The maximum distance from the center of the head in the volume of interest (VOI) is 65mm.) We aggregated 50 simulations at SNIR levels of 0 and 10dB. We averaged the results for each inter-dipolar distance and SNIR level and computed the standard error. The sources in these simulations had an inter-dipole correlation coefficient of 0.5 and an intra-dipole correlation coefficient of 0.25.

The results from this experiment are presented in Figure 3.2. These results serve two purposes; the first being to demonstrate Champagne's ability to resolve two correlated sources as we move them farther apart. The second is to give some intuition about our novel performance metric, AP and described above in Section 3.3. The majority of the results in the subsequent experiments with simulated data will show only plots of AP, which we believe captures the two aspects of performance, localization and time course reconstruction, in one number.

The first column of Figure 3.2 shows the hit rate (H_R) plotted against inter-dipolar distance at 10dB and 0dB. The H_R results show that at 10dB Champagne is able to localize two sources at any inter-dipole distance and at 0dB Champagne is able to distinguish two sources if there is 32mm (or 3 voxels width) between the sources. MVAB does well at 10dB as long as there is at least 24mm between the sources and at 0dB, however, MVAB is unable to resolve the two sources at any inter-dipole distances. MCE does fairly well at both noise levels, although the hit rate at best is approximately 0.7 at 10dB and 0.6 at 0dB (at 40mm separation). SL/dSPM has difficulty localizing the two sources at all inter-dipole distances at both noise levels. A' is computed from the H_R and F_R, so naturally, the A' plots shown in the second row are similar to the H_R plots. In general, if the F_R is low, the A' value will be higher than the H_R value (this is true even if the H_R is poor). The average correlation coefficient ($\bar{R}$) between the seeded and reconstructed time courses are shown in the third column. We only compute the correlation coefficient for sources that are counted as hits, so our results reflect how accurately the algorithms reconstruct the seeded time courses when the localization is successful. Champagne is able to reconstruct the time courses at all inter-dipole distances at 10dB and at 0dB, Champagne is successful when there is at least 32mm between the sources (similar to the H_R/A' plots). MCE is able to reconstruct the time courses at both noise levels with 32mm or 40mm distance between the sources, but the $\bar{R}$ values drop off for 16mm and 24mm inter-dipole spacing. Both SL/dSPM and MVAB have difficulty reconstructing the time courses at both SNIR levels and all inter-dipole distances.

The three values presented in the first three columns of Figure 3.2, H_R, A', and $\bar{R}$, are combined to compute Aggregate Performance (AP) shown in the fourth column. For Champagne, the AP results look very similar to the plots in the previous three columns since Champagne tended to perform well across metrics (or perform poorly as is the case at 0dB and the smaller inter-dipole distances). Likewise, MCE shows similar performance across the first three metrics (or columns) and yields similar trends in the AP metric. MCE performs most similarly to Champagne at large inter-dipole distances and performs like MVAB at smaller inter-dipole distances at 10dB. At 0db, however, Champagne and MCE have similar trends across the inter-dipole distances, but Champagne performs better than MCE at the lower SNIR level. The AP results for MVAB at 10dB demonstrate that MVAB was penalized for having poor time course reconstruction in spite of having good localization results. SL/dSPM does poorly across all of the first three metrics at both noise levels and thus, does poorly in terms of the AP metric.

3.5.2 Detecting Multiple Sources: Vector Lead Field

In order to examine Champagne's ability to localize many individual sources, we performed extensive simulations with randomly seeded sources.

We used a volumetric (two-component) lead field computed in NUTMEG (as described above). We randomly seeded 3 to 30 sources throughout the brain. The locations for the sources were chosen so that there was some minimum

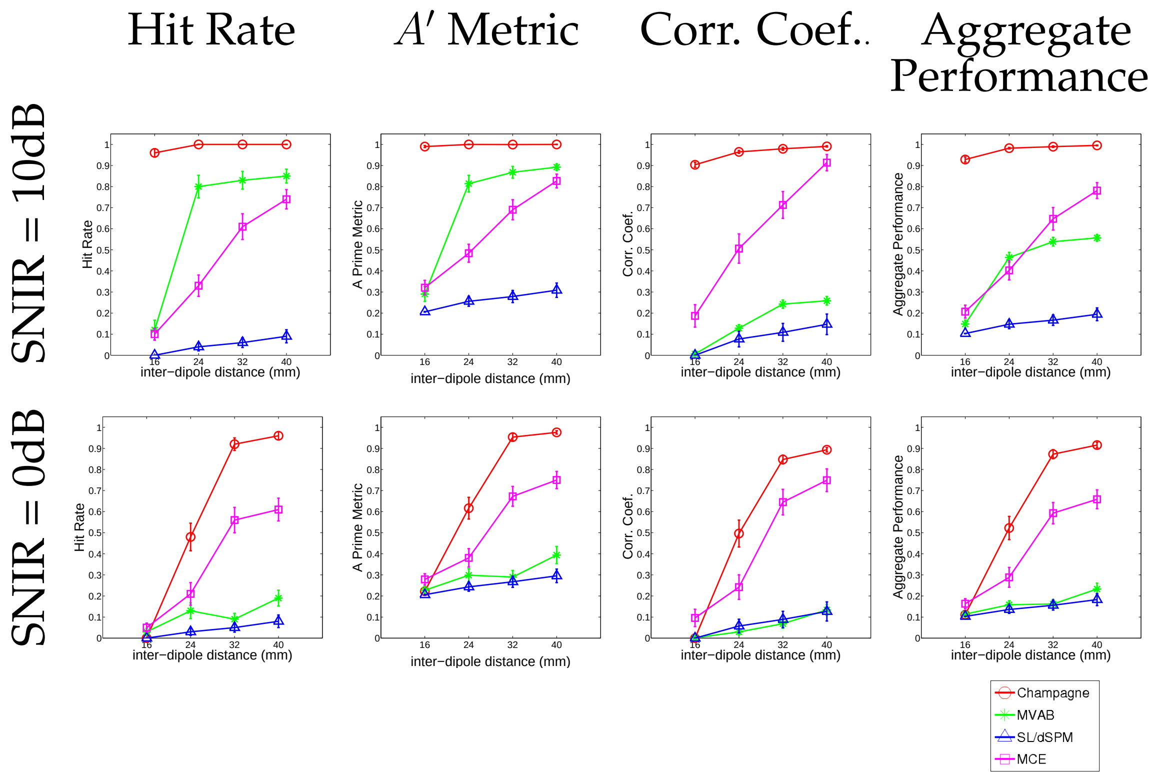

Figure 3.2 Simulations with two dipoles seeded with inter-dipole distances of: 16mm, 24mm, 32mm, or 40mm. The inter-dipole correlation is 0.5 and the intra-dipole correlation is 0.25. Components: (a) A' metric, (b) Hit Rate (H_R) and (c) Average Time-Course Correlation Coefficient ($\bar{R}$) used to compute (d) Aggregate Performance Metric (AP). The following equation is used: $AP = \frac{1}{2}(A' + H_R\bar{R})$. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

distance between sources (at least 10mm) and a minimum distance from the center of the head (at least 35mm). We ran 50 simulations of randomly (located) seeded sources at SNIR levels of 0dB and 10dB. The sources in these simulations had an inter-dipole correlation coefficient of 0.5 and an intra-dipole correlation coefficient of 0.25.

The results from the experiment with the vector lead field are shown in Figure 3.3. In this figure, we plot number of sources versus AP at SNIR levels of 10dB and 0dB (inset plot). At both SNIR levels, Champagne outperforms the other source localization algorithms. Champagne is able to accurately reconstruct up to 10 sources at 10dB and up to 5 sources at 0dB. For both SNIR levels and across number of sources, the three other source localization algorithms perform at almost the same level. MCE performs better than MVAB and SL/dSPM at 3 and 5 sources at both SNIR levels, but at more than 5 sources performance is fairly similar for the benchmark algorithms. We also show single simulation examples at 10dB and 0dB for both 5 (Figure 3.4) and 10 sources (Figure 3.5). The 5 sources examples demonstrate that Champagne is able to recover 5 sources at both SNIR levels. Both MCE and MVAB are able to recover most of the 5 sources, but also have some false positives and blur around the sources. At 0dB, MVAB and MCE do not successfully localize the sources, there is a peak at the center of the head (which gets projected to the surface) and one successful source for MVAB. SL/dSPM did not localize any of the 5 sources at either SNIR level. These results show a peak at the center of the head and while this peak does extend out to where the sources were seeded, there are no distinct peaks at the seeded locations. The 10 sources examples demonstrate that Champagne is able to recover 9 out of 10 sources at a SNIR of 10dB. At 0dB, Champagne is able to recover over half the sources, but also has some false positives, including a peak at the center of the head. Both MVAB and MCE are able to localize some of the 10 sources at 10dB, but at 0dB these algorithms are unable to resolve more than one source. As in the 5 source example, SL/dSPM is not able to recover any of the 10 sources at 10dB or 0dB; there is only a peak at the center of the head.

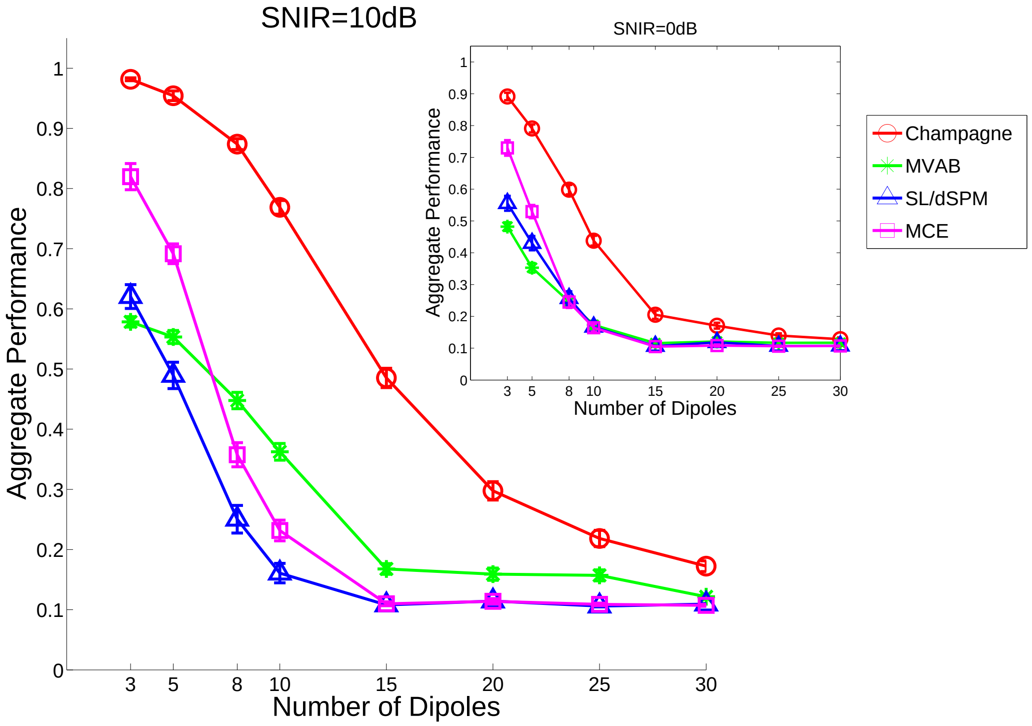

Figure 3.3 Results from simulations testing localization on multiple dipoles (3 to 30) with a volumetric lead field (2 components). AP is plotted against number of dipoles for SNIR=10dB and SNIR=0dB (inset plot). The inter-dipole correlation is 0.5 and the intra-dipole correlation is 0.25. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

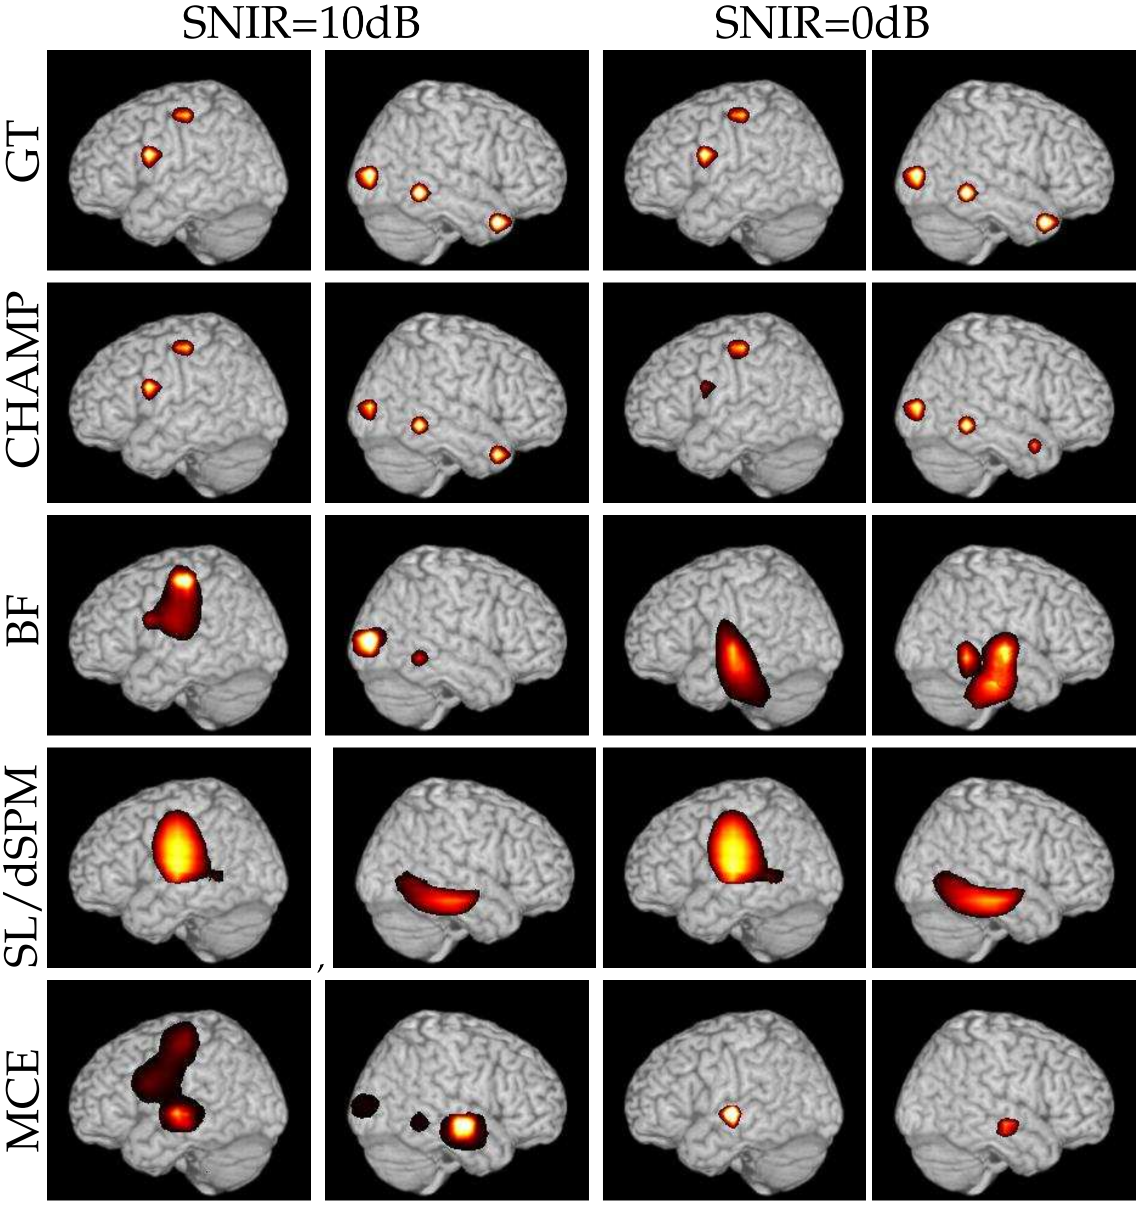

Figure 3.4 A single example of the localization results with the vector lead field for 5 dipoles at SNIR=10dB (right columns) and SNIR = 0dB (left columns). The ground truth (GT) location of sources are shown for comparison, first row. The results with Champagne (CHAMP) are shown in the second row and the comparison algorithms, BF, SL/dSPM, and MCE are shown in the subsequent rows. We project the source power to the surface of a template brain.

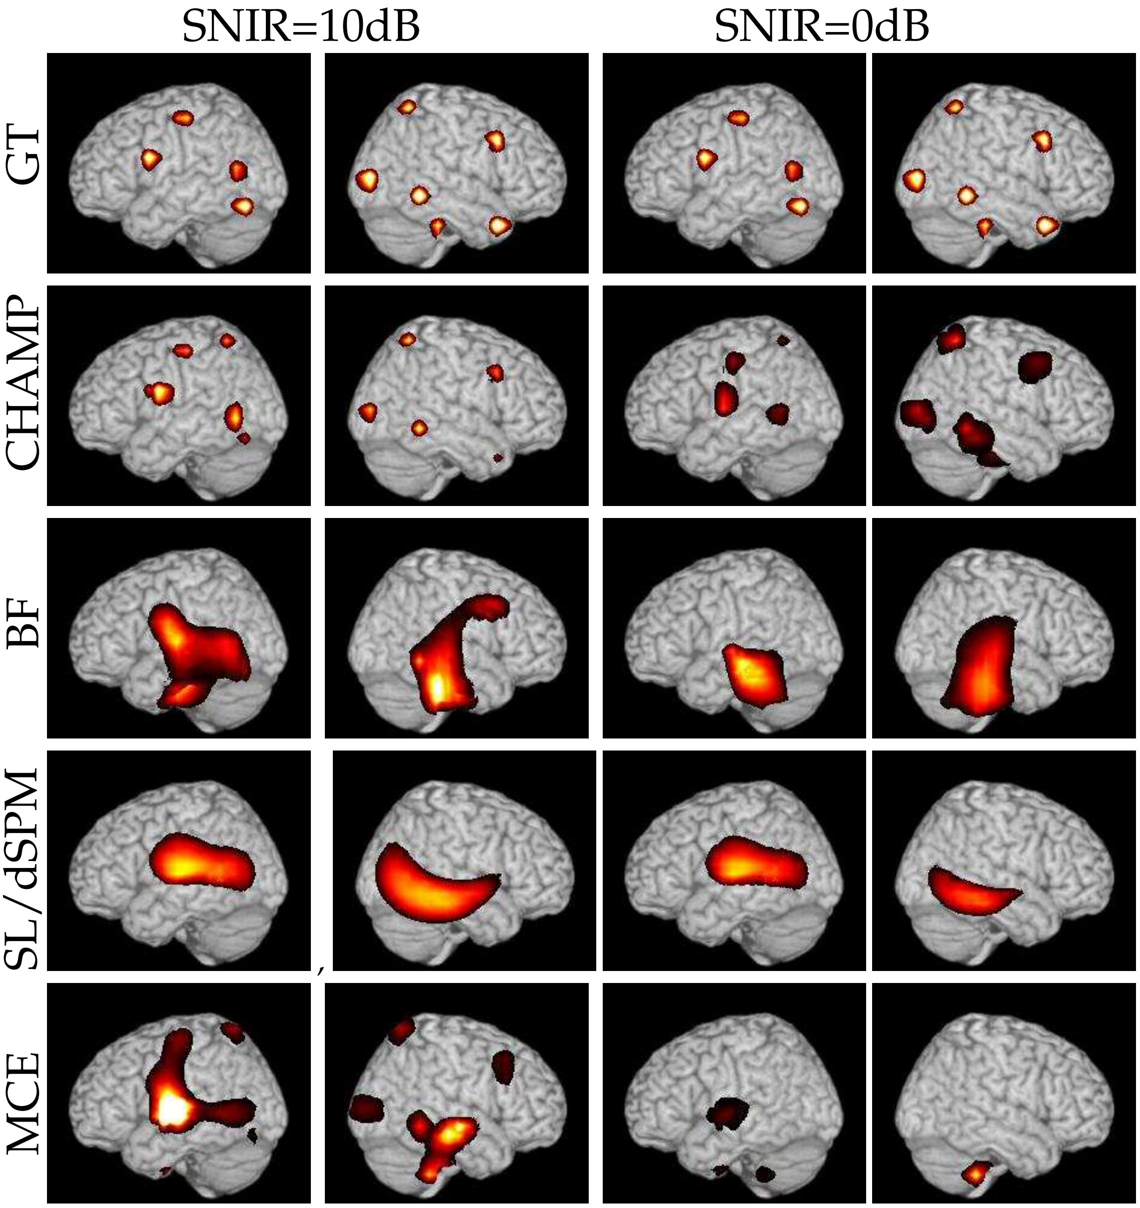

Figure 3.5 A single example of the localization results with the vector lead field for 10 dipoles at SNIR=10dB (right columns) and SNIR = 0dB (left columns). The ground truth (GT) location of sources are shown for comparison, first row. The results with Champagne (CHAMP) are shown in the second row and the comparison algorithms, BF, SL/dSPM, and MCE are shown in the subsequent rows. We project the source power to the surface of a template brain.

3.5.3 Detecting Multiple Sources: High SNIR

The above experiment, in Section 3.5.2, was repeated at a SNIR of 20dB for Champagne; the results of this experiment compared to the results for 0dB and 10dB are shown in Figure 3.6. As expected, Champagne is able to localize more sources at 20dB than at 10dB. At 20dB, Champagne is able to localize 15 to 20 dipoles. Also, at 20dB, it can almost perfectly localize up to 10 dipoles and the drop off in performance does not begin until 15 dipoles.

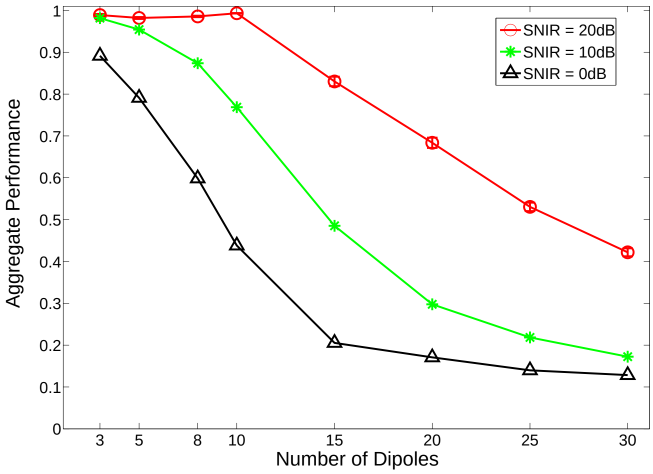

Figure 3.6 Comparison of localization results at SNIR levels: 0dB, 10dB, and 20dB. The inter-dipole correlation is 0.5 and the intra-dipole correlation is 0.25. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

3.5.4 Detecting Multiple Sources: No Noise

The experiment in Section 3.5.2 was repeated again with high SNIR, 100 dB, and with Gaussian time courses in lieu of the damped sinusoids. The generative model for the Champagne algorithm assumes Gaussian source time course and localization is easier with this modification to the way the simulated data were generated. The results from this experiment can be found in Figure 3.7. In this figure, we can observe the improved performance when there is virtually no noise and the time courses more closely fit the model. Champagne is able to localize up to 70 sources under these conditions.

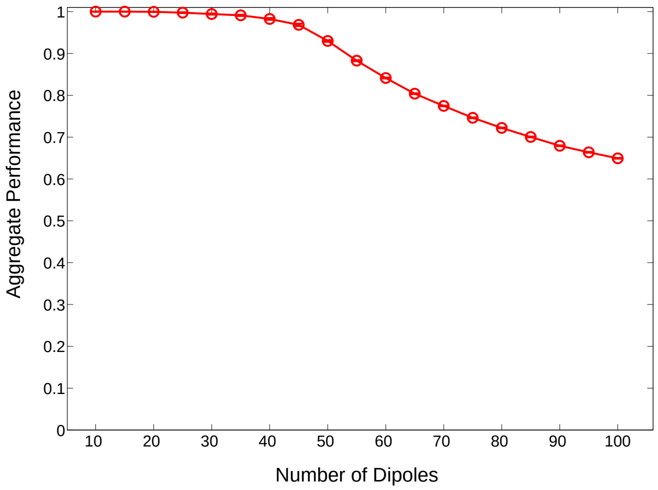

Figure 3.7 Simulation at SNIR=100dB with Gaussian time courses and Gaussian noise model. The inter-dipole correlation is 0.5 and the intra-dipole correlation is 0.25. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

3.5.5 Detecting Multiple Sources: Inter-dipole Correlations and Source Time Courses

We tested Champagne on four conditions to assess the effect of the inter-dipole correlation and the type of source time course. We used inter-dipole correlations, α_{inter}, of 0 and 0.5 and Gaussian and damped sinusoidal time courses. The damped sinusoidal time courses and $\alpha_{inter} = 0.5$ mirror the conditions used in Section 3.5.2 and is the lower bound on performance for this experiment. The results from the two conditions, sinusoidal time courses and $\alpha_{inter} = 0.5$ and Gaussian time courses and $\alpha_{inter} = 0$, are similar to each other and perform better than the damped sinusoidal time courses and $\alpha_{inter} = 0.5$ condition. The Gaussian time courses and $\alpha_{inter} = 0$ condition performs the best of the four conditions and Champagne is able to localize up to 15 dipoles under these conditions, which is in contrast to 10 dipoles for the other conditions.

3.5.6 Detecting Multiple Sources: Scalar Lead Field

Many research groups prefer to use a cortically- and orientation-constrained leadfield for M/EEG source reconstructions. The reasoning behind this preference is that the pyramidal cells in the cortex that give rise to the majority of the M/EEG signal are mostly oriented perpendicular to the cortical surface. When the current dipoles are all assumed to be oriented normal to the cortical surface, the lead field simplified to have only one orientation. For the simulations with the scalar lead field, we randomly seeded 2 to 100 dipoles throughout cortex; the minimum distance between any two sources is 10mm. We selected the coarse grid spacing in SPM, which results in approximately 5,000 voxels. The voxels are not placed in a regular grid, but in general the voxels are spaced approximately 5mm

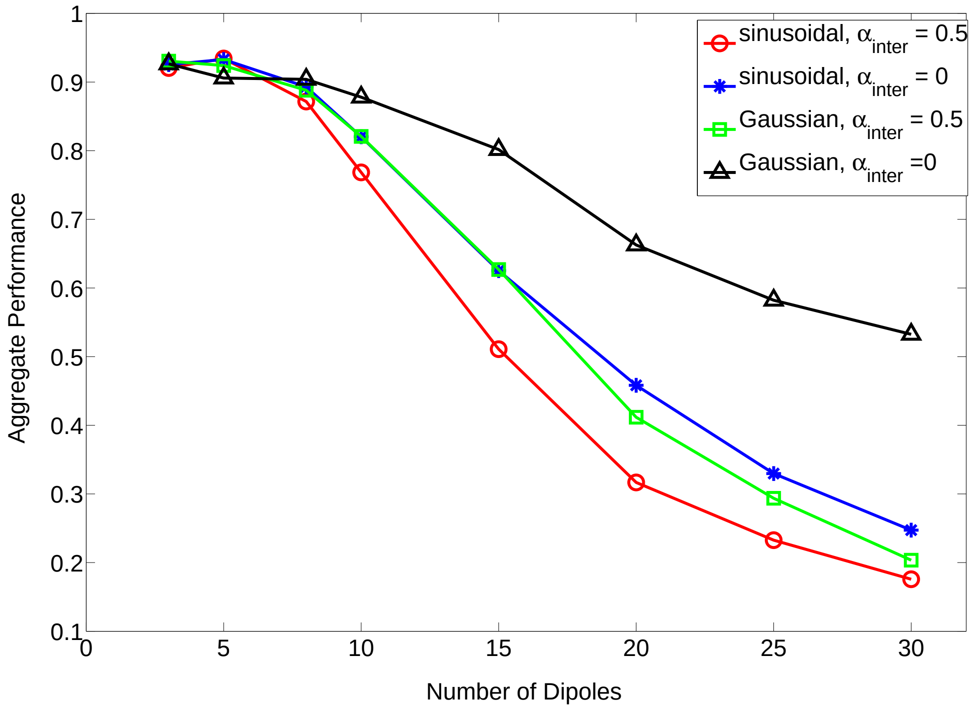

Figure 3.8 Simulations at 10dB with different inter-dipole correlations, $\alpha_{inter} = 0$ and $\alpha_{inter} = 0.5$, and types of source time courses, Gaussian and damped sinusoidal. The intra-dipole correlation is 0.25. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

apart. The sources in these simulations had an inter-dipole correlation coefficient of 0.5.

The results from this experiment with a scalar lead field are shown in Figure 3.9. Aggregate performance is plotted against number of dipoles at SNIR levels of 10dB and 0dB (inset plot). At both SNIR levels, Champagne outperforms the three other algorithms at both SNIR levels. The only exception is that MCE does better than Champagne for 2 dipoles at 10dB. At 10dB, Champagne is able to localize up to 25 dipoles, while the other three algorithms have a performance drop off at 5 dipoles. At 0dB, Champagne is able to localize up to 10 dipoles, while the other algorithms are not able to accurately localize 2 dipoles (MVAB performs better than SL/dSPM and MCE at 0dB). The single simulation example at 10dB, Figure 3.10, demonstrates that Champagne is able to localize 19 of the 20 dipoles seeded, while the other algorithms are only able to localize at most 6 (MVAB) and at least 2 (SL/dSPM).

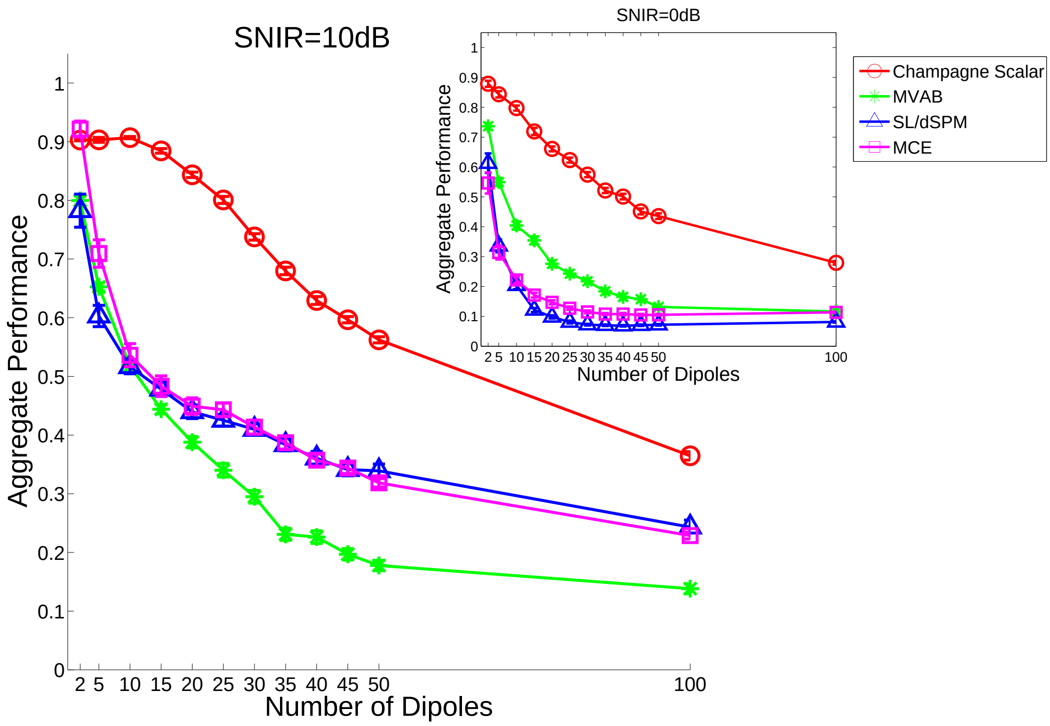

Figure 3.9 Results from simulations to test localization on multiple dipoles (2 to 100) with a scalar leadfield (1 component). Aggregate Performance (AP) is plotted against number of dipoles for SNIR=10dB and SNIR=0dB (inset plot) The inter-dipole correlation is 0.5. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

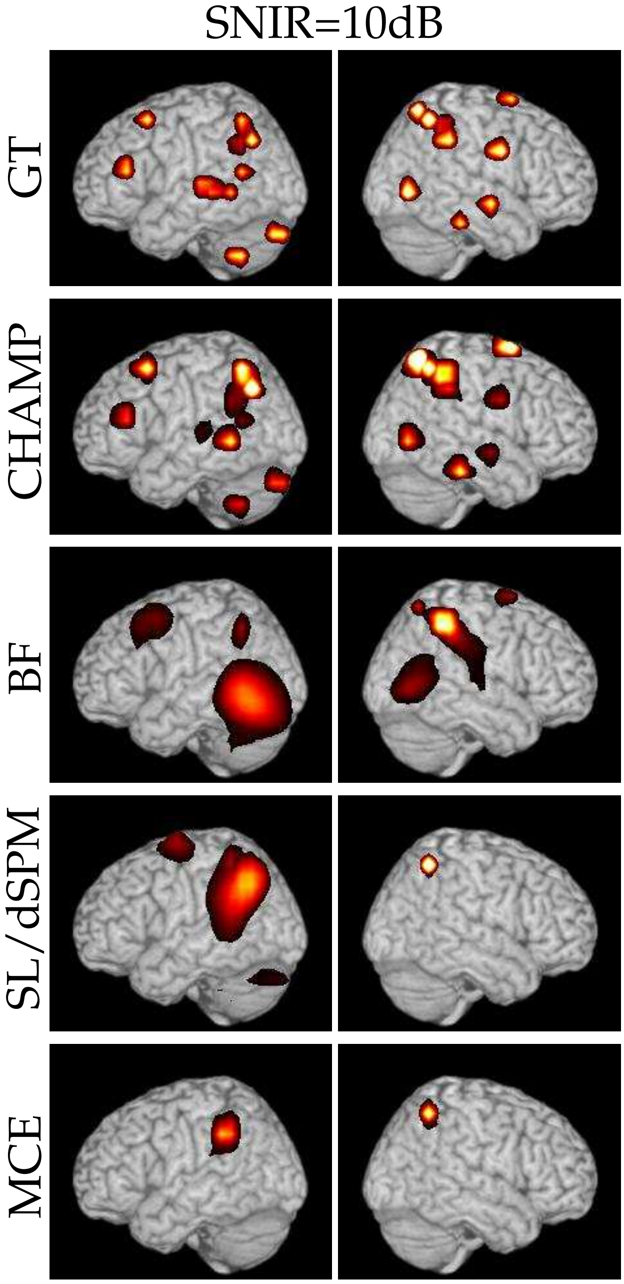

Figure 3.10 A single example of the localization results with the scalar lead field for 20 dipoles at SNIR=10dB. The ground truth (GT) location of sources are shown for comparison, first row. The results with Champagne (CHAMP) are shown in the second row and the comparison algorithms, BF, SL/dSPM, and MCE are shown in the subsequent rows. We project the source power to the surface of a template brain.

3.5.7 Rotating versus Fixed Dipole Model

Altering the correlation between the lead field components, called the intra-dipole correlation, changes the degree that a source's orientation rotates. A correlation of one indicates a fixed orientation and a correlation of zero indicates a freely rotating dipole. While there isn't much evidence for how much the approximate current dipoles in the brain rotate over the time course of activity, it is most likely somewhere between 0 and 1 (and not exactly equal to zero or one). We wanted to investigate the effect of the intra-dipole correlation (α_{intra}) on Champagne's ability to localize 10, 15, and 20 sources. We chose intra-dipole correlations of 0.25, 0.75, and 1. If the orientation was truly fixed, then the most probable orientation for the sources would be normal to the cortical surface (the pyramidal cells in the cortex that give rise to the majority of the M/EEG signal are mostly orientated perpendicular to the surface of the cortex). Many research groups choose to use a cortically- and orientation-constrained lead field for M/EEG source reconstructions, called a scalar lead field. Scalar lead fields have only one orientation per voxel and the orientation is set to be normal to the surface. In addition to testing the vector lead field with varying intra-dipole correlations (as described above), we also tested the scalar lead field when the intra-dipole correlation was 0.25 and 1. The inter-dipole correlation is 0.5 for all conditions. In all these cases a vector lead field was used for the forward model and we varied the lead field used for solving the inverse problem. The lead field used for the inverse is referred to as the *inverse model*.

AP is plotted against five conditions: 1) $\alpha_{intra} = 1$ & vector inverse model , 2) $\alpha_{intra} = 0.75$ & vector inverse model, 3) $\alpha_{intra} = 0.25$ & vector inverse model , 4) $\alpha_{intra} = 1$ & scalar inverse model, and 5) $\alpha_{intra} = 0.25$ & scalar inverse model. Instead of using a scalar lead field from another software package, which would have a different grid than the vector lead field from NUTMEG, we transformed the vector lead field into a scalar lead field by assuming an optimal orientation for every voxel. With this assumed orientation, we can transform the vector lead field into a scalar lead field with the same voxel grid.

The optimal orientation for a particular voxel, r, is obtained with the following equation:

$$\eta_{opt}(r) = v_{min}\left\{L^T(r)R^{-1}L(r)\right\} \tag{3.14}$$

where $v_{min}\{\}$ is the eigenvector corresponding to the minimum value of the matrix in the curly braces, $L(r)$ is the forward lead field from voxel r and R^{-1} is the inverse covariance of the sensor data.(Sekihara and Nagarajan, 2008). With this optimal orientation for every voxel, we can transform the vector lead field into a scalar lead field with the same voxel grid. (Theoretically, we could use any arbitrary orientation matrix to transform the vector lead field into a scalar lead field since we were working with simulated data.)

The results from this experiment are shown in Figure 3.7. AP for Champagne is plotted for the five conditions described above for 10, 15, and 20 sources. The performance using a vector inverse model (the first three conditions) demonstrate that the presence of intra-dipole correlations makes the localization problem more difficult and that the weaker the correlation, the easier it is to localize multiple sources. The results using a scalar inverse model show that if the sources' orientations are truly fixed, it is advantageous to use a scalar lead field for the inverse problem. The performance on 15 and 20 sources is better with the scalar lead field than the performance obtained using the vector lead field if the orientation is fixed. On the other hand, if the orientation is rotating and a scalar inverse model is used for the reconstruction the performance is drastically reduced as compared to the experiments where a vector inverse model is used.

3.5.8 Effect of Lead Field Errors

We wanted to investigate the effect of orientation errors on the results obtained with the scalar lead field from the previous experiment. When using a scalar lead field, an assumption is being made about the orientation of the current sources in the brain. To test the effect of errors between the true and estimated orientation, we randomly rotated the orientation of each voxel after simulating the data and before performing source localization with Champagne. The maximum perturbation to the orientation ranged from 0 to $\frac{\pi}{4}$. Not every voxel's orientation was rotated by the same amount, rather every voxel was rotated by a randomly

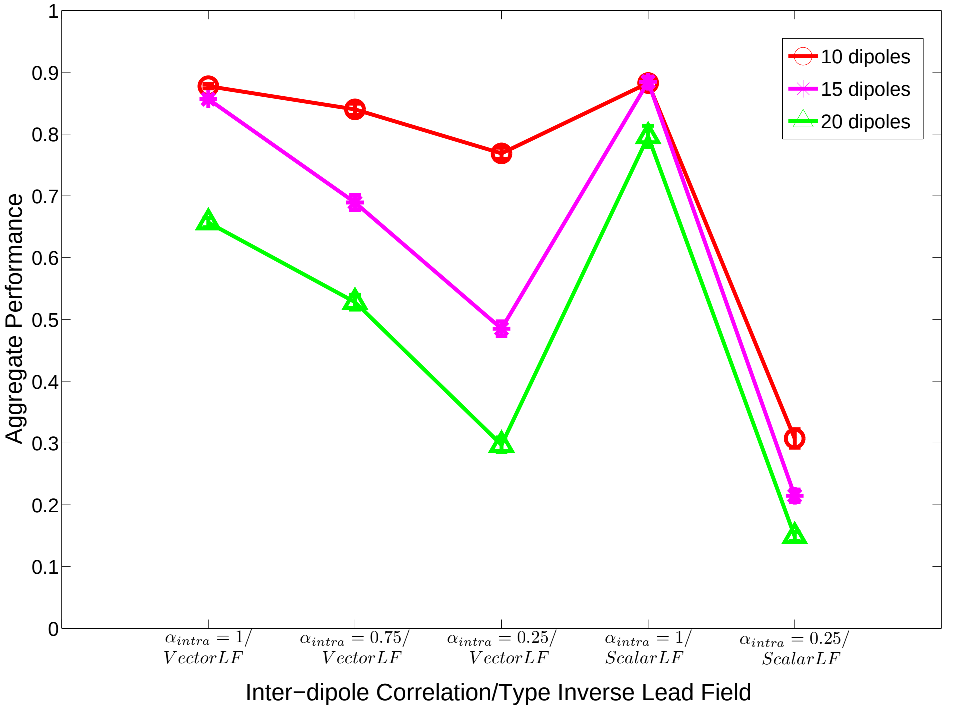

Figure 3.11 Exploration of lead field errors with 10, 15, and 20 dipoles (depicted by color) by using one forward model and a different inverse model. AP is plotted against five conditions: 1) $\alpha_{intra} = 1$ & vector inverse model , 2) $\alpha_{intra} = 0.75$ & vector inverse model, 3) $\alpha_{intra} = 0.25$ & vector inverse model , 4) $\alpha_{intra} = 1$ & scalar inverse model, and 5) $\alpha_{intra} = 0.25$ & scalar inverse model. The inter-dipole correlation is 0.5 for all conditions. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

generated angle between zero and the maximum perturbation angle. We chose to run these simulations on 10, 15 and 20 sources since this is the regime for which there is a discrepancy in performance for the scalar and vector lead field.

The results from this experiment are shown in Figure 3.12. AP is plotted against maximum orientation error for 10, 15, and 20 sources. The trend in performance across the perturbations is consistent for 10, 15 and 20 sources. The performance drops approximately 10% immediately with the smallest maximum perturbation, $\frac{\pi}{64}$. The performance stays the same for maximum perturbations of $\frac{\pi}{32}$ and $\frac{\pi}{16}$, and then starts falling off more rapidly for maximum perturbations of $\frac{\pi}{8}$ and $\frac{\pi}{4}$.

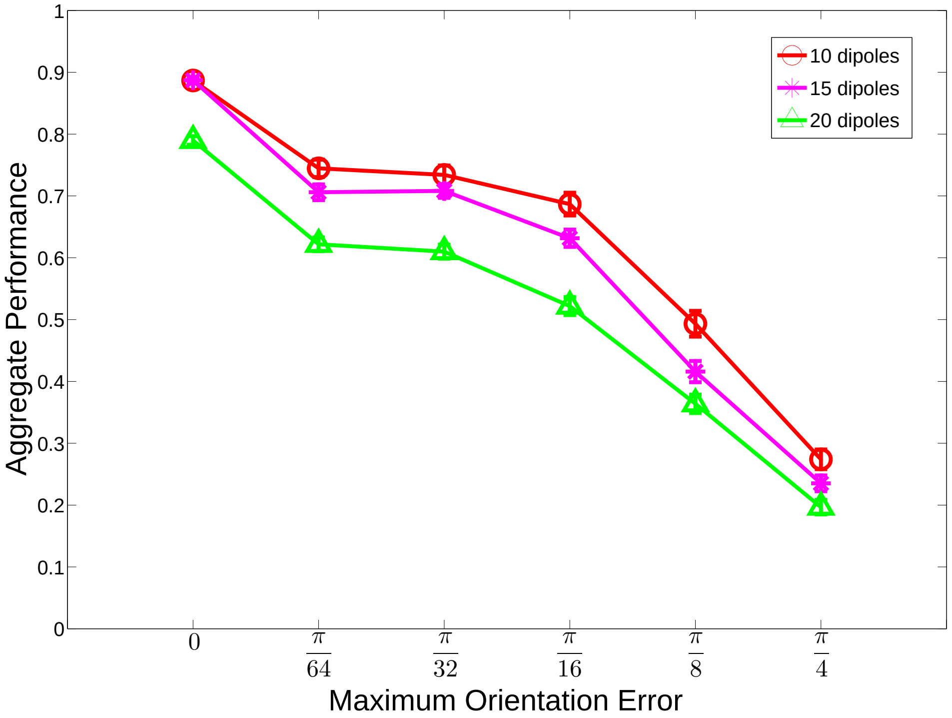

Figure 3.12 Localization results for 10, 15, and 20 dipoles (depicted by color) with orientation errors introduced to the scalar lead field. AP is plotted against maximum perturbation angle. Each voxel's orientation vector is perturbed by a random angle from 0 to the maximum angle (shown on the x-axis). The inter-dipole correlation is 0.5 (and the intra-dipole correlation is 0.25 for the vector lead field). The results are averaged over 50 simulations at each data point and the error bars show the standard error.

We conducted an additional experiment to assess the effect of lead field errors on localization accuracy with Champagne. We used a vector lead field to simulate data at a SNIR of 10dB. Then we shifted the lead field grid diagonally the distance of about half the diagonal of a voxel (5mm) and recomputed the lead field for this grid of source locations. We solved the inverse problem with Champagne using the lead field from the shifted grid. We tested Champagne with 3 to 30 sources at 10dB and found that the AP metric was reduced, at worst, 10% for each number of sources. The results are shown in Figure 3.13. This result supports the evidence from the previous experiment that Champagne is robust to errors in the lead field modeling.

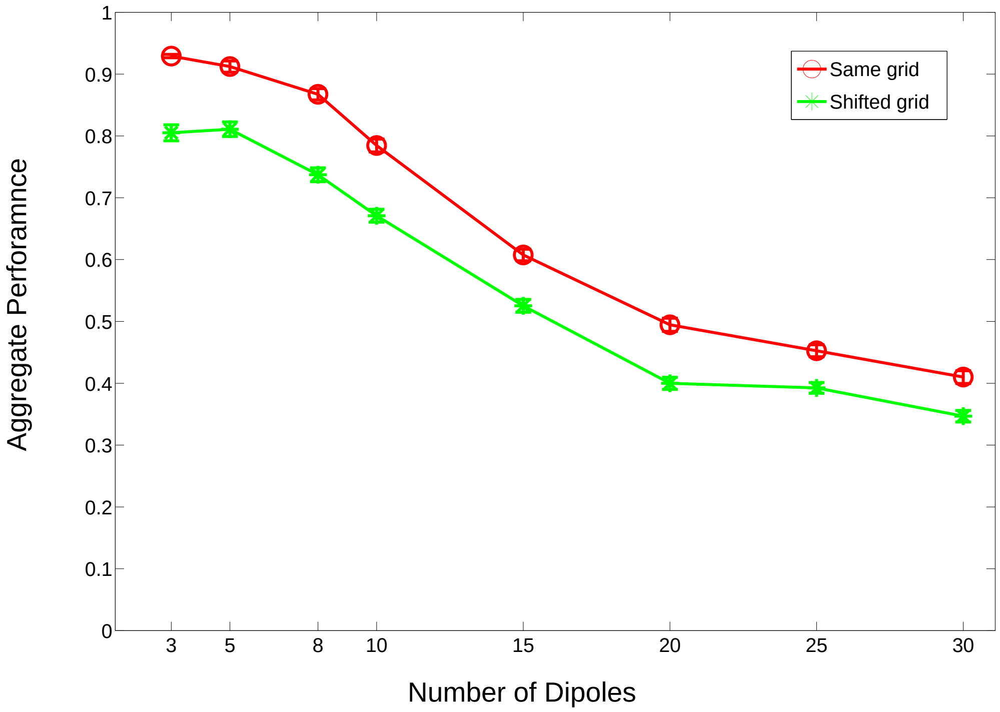

Figure 3.13 Simulations with un-shifted and shifted voxel grid at SNIR of 10dB. The inter-dipole correlation is 0.5 and the intra-dipole correlation is 0.25. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

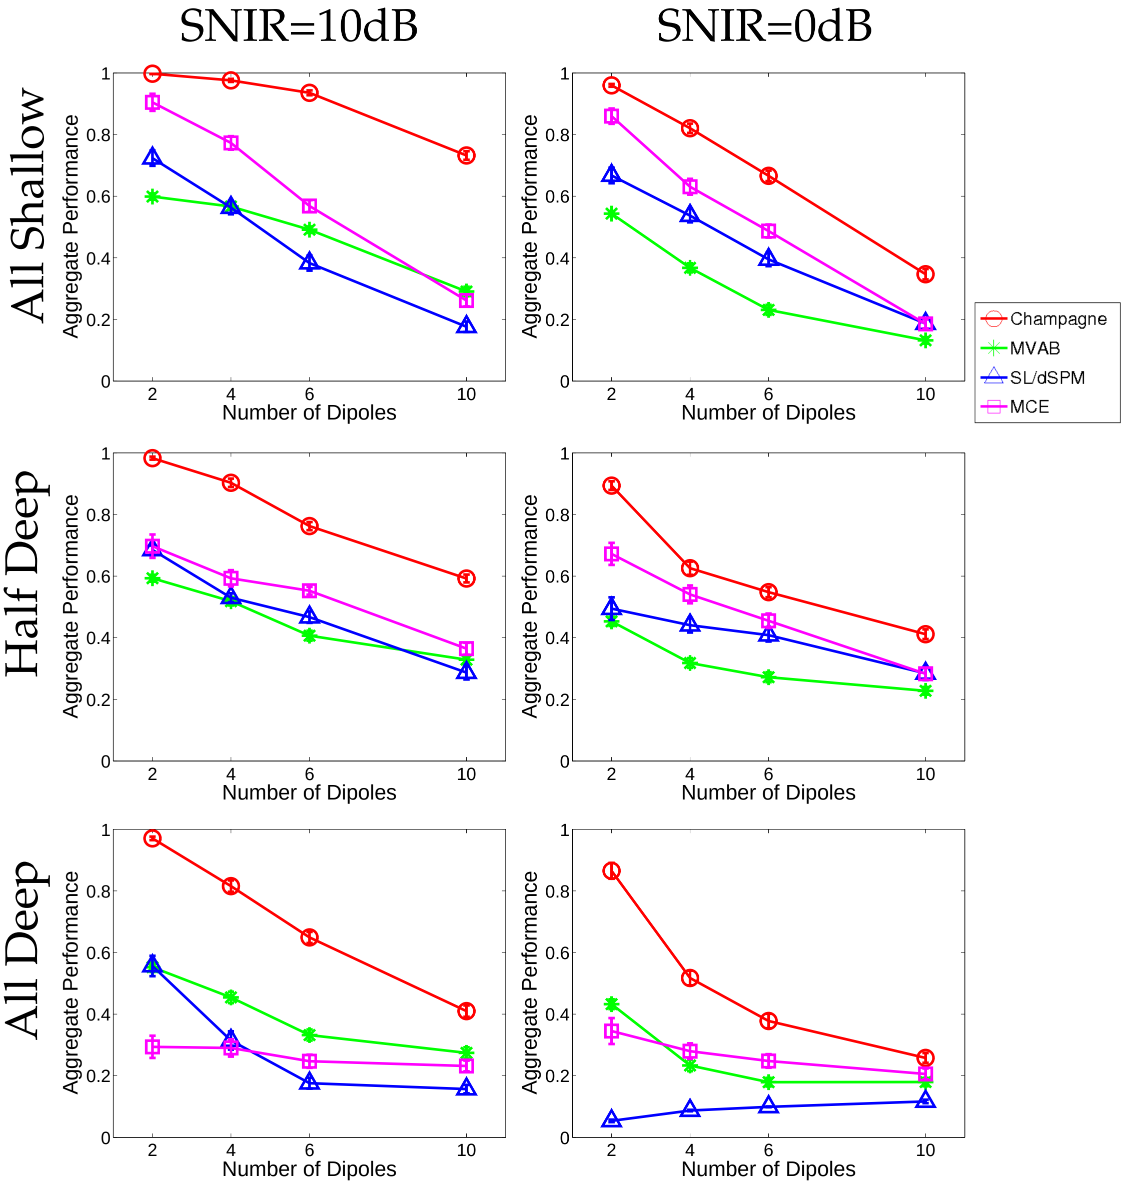

Figure 3.14 Localization results in the presence of deep sources at SNIR=10dB (left column) and SNIR=0dB (right column). AP is plotted against total number of sources in each panel. There were three conditions (corresponding to the rows of the figure): 1) no deep sources (or all shallow sources), 2) half deep sources and half shallow sources, 3) all deep sources. The inter-dipole correlation is 0.5 and the intra-dipole correlation is 0.25. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

3.5.9 Deep Sources

To assess Champagne's ability to localize deep sources, which are typically hard to localize with MEG, we constructed three conditions to compare performance with: 1) no deep sources (or all shallow sources), 2) half deep sources and half shallow sources, 3) all deep sources. A deep source was defined as less than 35mm from the center of the head and a shallow source was defined as above, at least 35mm from the center of the head. These conditions were designed to test sensitivity to deep sources, when there are only deep sources and when there is a combination of deep and shallow sources. The configurations with only shallow sources were included to give a basis for comparison. The placement of the sources was random (within the distance constraints) and we aggregated 50 simulations for SNIR levels of 0 and 10dB. The sources in these simulations had an inter-dipole correlation coefficient of 0.5 and an intra-dipole correlation coefficient of 0.25. The total number of sources was 2, 4, 6 and 10, where each total number of sources had three conditions associated with it. For example, for 4 total sources, there was one condition where there were 4 shallow sources, one condition where there were 2 shallow and 2 deep sources, and one condition where there were 4 deep sources. We chose the maximum number of sources to be 10 because that is the largest number of sources that Champagne was able to localize in the *Detecting Multiple Dipoles* experiment.

The results for this experiment are shown in Figure 3.14. We have plotted

AP against total number of sources (2, 4, 6, or 10) for the three source configurations described above (rows of Figure 3.14), at 0 and 10dB (columns of Figure 3.14). Across all the source configurations and noise levels, Champagne outperforms the other source localization algorithms across all numbers of sources. As we saw in the *Detecting Multiple Dipoles* experiment, Champagne is able to reconstruct 10 sources at 10dB and 4 to 6 sources at 0dB. MCE consistently does better than MVAB and SL/dSPM for the conditions where there are all shallow sources or half shallow (and half deep) sources. The presence of deep sources (whether they constitute half or all the sources) degrades performance for all the source localization methods. This decrease in performance is especially marked for MCE at both noise levels and SL/dSPM at 0dB.

3.5.10 Clusters

Given Champagne's sparsity, we tested its ability to localize distributed activity by simulating clusters of sources. We seeded 5, 10, or 15 clusters each with 10 sources in each cluster. These cluster sizes correspond to 50, 100, and 150 voxels having non-zero activity. The placement of the cluster center was random and the clusters consisted of sources seeded in 9 nearest neighboring voxels. We aggregated 50 simulations for SNIR levels of 0 and 10dB. The source time courses within each cluster had an inter-dipole correlation coefficient of 0.8 and an intra-dipole correlation coefficient of 0.25. The multiple clusters were correlated with a

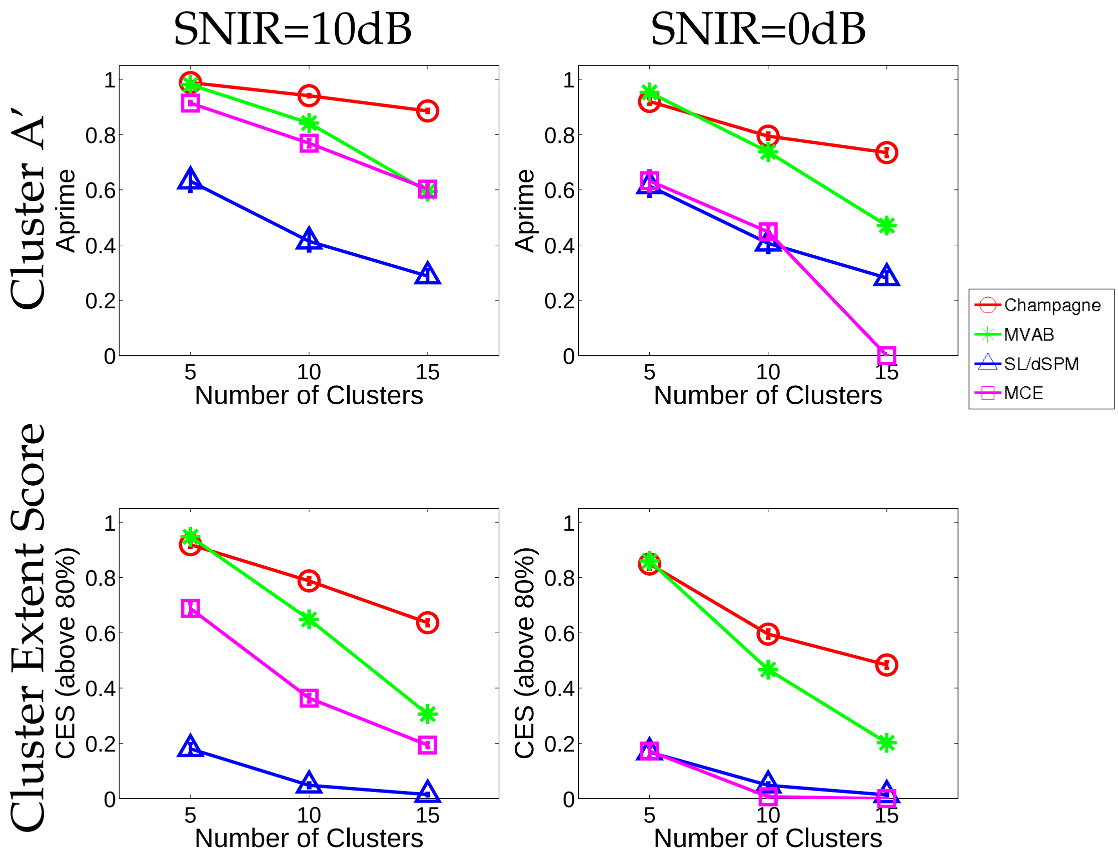

Figure 3.15 Performance on 5, 10 and 15 clusters of dipoles (each cluster has 10 contiguous active voxels) using a vector lead field. Cluster A' is plotted against number of clusters in the top row, showing the localization accuracy and Cluster Extent Score (CES) is plotted against number of clusters in the bottom row, showing the algorithms' ability to correctly model the extent of the clusters. The inter-dipole correlation within the cluster is 0.8, the inter-dipole correlation between clusters is 0.5, and the intra-dipole correlation is 0.25. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

correlation coefficient of 0.5. We made the correlations within the clusters higher than between clusters because nearby voxels are more plausibly correlated than voxels at a distance. For the clusters, we are both interested in whether the cluster is localized and whether the extent of the cluster is accurately reconstructed. To assess the localization of the clusters, we use the A' metric. The A' metric is cal-

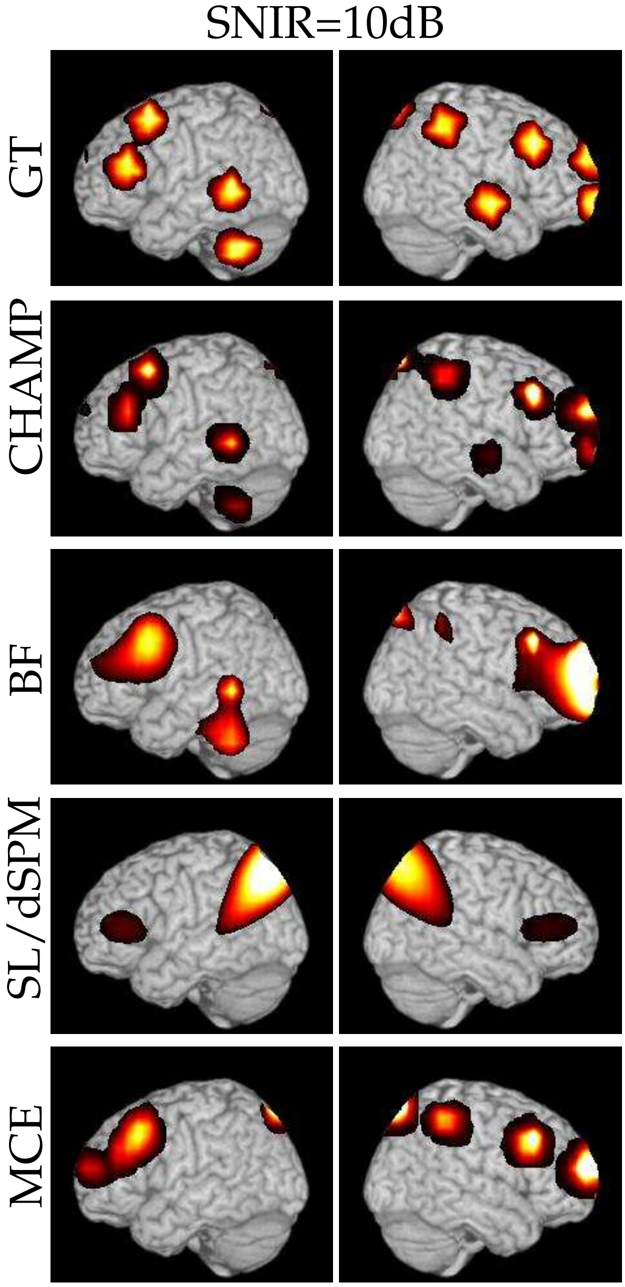

Figure 3.16 A single example of the localization results for 10 clusters (each with 10 dipoles) at SNIR = 10dB with the vector lead field. The ground truth (GT) location of the clusters are shown for comparison, first row. The results with Champagne (CHAMP) are shown in the second row and the comparison algorithms, BF, SL/dSPM, and MCE are shown in the subsequent rows. We project the source power to the surface of a template brain.

culated for the clusters by testing if there is a local peak within the known extent of the cluster. A local peak that is not within the extent of any cluster is deemed a false positive. To assess the accuracy of the extent of the clusters, we use our knowledge of where the 10 sources that constitute each cluster were seeded to calculate the fraction of the seeded voxels with activity in the 80th percentile of all the voxels. We call this fraction the *Cluster Extent Score* (CES).

The results from this experiment are shown in Figure 3.15. In the first row, we show A' metric results for 10dB and 0dB. In the second row, we show the CES results for 10dB and 0dB. For 5 clusters, Champagne and MVAB perform equally for both metrics at both SNIR levels, but at 10 and 15 clusters Champagne performs better than MVAB and the other algorithms both in terms of localizing the clusters and reconstructing their extent. MCE does a fairly good job at localizing the clusters at 10dB as seen in the A' plot, but is not able to reconstruct the extent of the clusters as seen in the CES plot. At 0dB, MCE and SL/dSPM perform similarly in terms of A' and CES and in general, SL/dSPM is not successful in localizing or reconstructing the extent of the clusters. A single-simulation experiment with 10 clusters is presented in Figure 3.16. Champagne is able to localize all 10 clusters in this example. MVAB and MCE are able to localize a fair number of clusters (7 and 5 respectively). SL/dSPM is only able to localize one cluster (the cluster in the occipital lobe).

3.5.11 EEG Simulations: Vector Lead Field

The forward model for EEG accounts for the conductivity of the biological tissue(s) through which the electric fields travel. It is generally thought that the forward model for EEG is harder to correctly calculate and thus, source localization with EEG data is more difficult that with MEG data. EEG is more prevalent because it is less expensive and portable. The vector EEG lead field is on a 5mm grid and has approximately 6,000 voxels. The vector lead field has three components for every voxel (as compared to two components for the vector MEG lead field) and was calculated using the BESA software (www.besa.de). We used re-montaged lead fields for 257, 129, 81, 33, and 19 sensors at 0 and 10dB. The sources in these simulations had an inter-dipole correlation coefficient of 0.5 and the sources had an intra-dipole correlation coefficient of 0.25. We did not have any resting-state data from these subjects, so Gaussian random noise was used for the noise-plus-interference.

The results from this experiment with simulated EEG data are shown in Figure 3.17. These data imply that localization with EEG is significantly harder than with MEG. Even at 257 sensors and at 10dB, Champagne is only barely able to localize 3 sources. At all numbers of sensors and SNIR levels, all the algorithms do not meet the $AP = 0.75$ criterion for a successful performance. Although, these results do demonstrate that Champagne is better able to localize sources with EEG at the higher number of sensors (257, 129, and 81 sensors). MVAB does almost as well as Champagne for most conditions, but SL/dSPM and MCE are not able to localize 3 source with any number of sources or at any SNIR level.

3.5.12 EEG Simulations: Scalar Lead Field

In addition to the experiment with the vector lead field, we tested Champagne on simulated EEG data using a scalar, cortically-constrained lead field computed in SPM (http://www.fil.ion.ucl.ac.uk/spm) by selecting the coarse resolution. This results in approximately 5,000 voxels at 10mm spacing. We repeated the *Detecting Multiple Sources* experiment with this lead field. In addition to running on the maximum number of sensors (128), we also investigated the effect of subsampling the number of sensors on the ability to localize sources. We subsampled the 128 sensor lead field to 64, 32, and 16 sensors. Some EEG researchers use as few as 10 electrodes for a standard 10-20 montage or 16 sensors for clinical EEG systems. As such, we wanted to test Champagne's localization performance when the number of sensors is greatly reduced. In order to decrease the number of sensors, the EEG channels are grouped in 4 groups depending on their location. For each of the subsampled lead fields, we individually subsample each of the four subsections in order to preserve coverage. We tested the lead fields at both 10dB and 0dB. We aggregated 50 simulations at each number of sensors and sources and SNIR level. The sources in these simulations had an inter-dipole correlation coefficient of 0.5.

The results from this experiment are shown in Figure 3.18. The left column shows the results at 10dB and the results at 0dB are shown in the right column. The results obtained with the full 128-sensor lead field are shown in the top row, and the number of sensors in the lead field decrease in the subsequent rows. At almost every SNIR level, number of sources and sensors, Champagne outperforms the other algorithms. In the one source case, only MCE is able to localize that source better than Champagne with 128 and 64 sensors at both 10dB and 0dB. MCE's performance after one source drops off quickly in these plots. In general, Champagne's performance is the same for 128 and 64 sensors, we do not see a degradation in performance until 32 sensors. This is true for the other localization algorithms as well. At 10dB and 32 sensors, Champagne is able to localize 8 sources and at 0dB and 32 sensors, Champagne is able to localize 3 sources. At only 16 sensors, Champagne is able to localize up to 5 sources at 10dB and 3 at 0dB. Overall, MVAB performs better than MCE and SL/dSPM and while MCE does well on one source, SL/dSPM and MCE perform similarly across all number of sources.

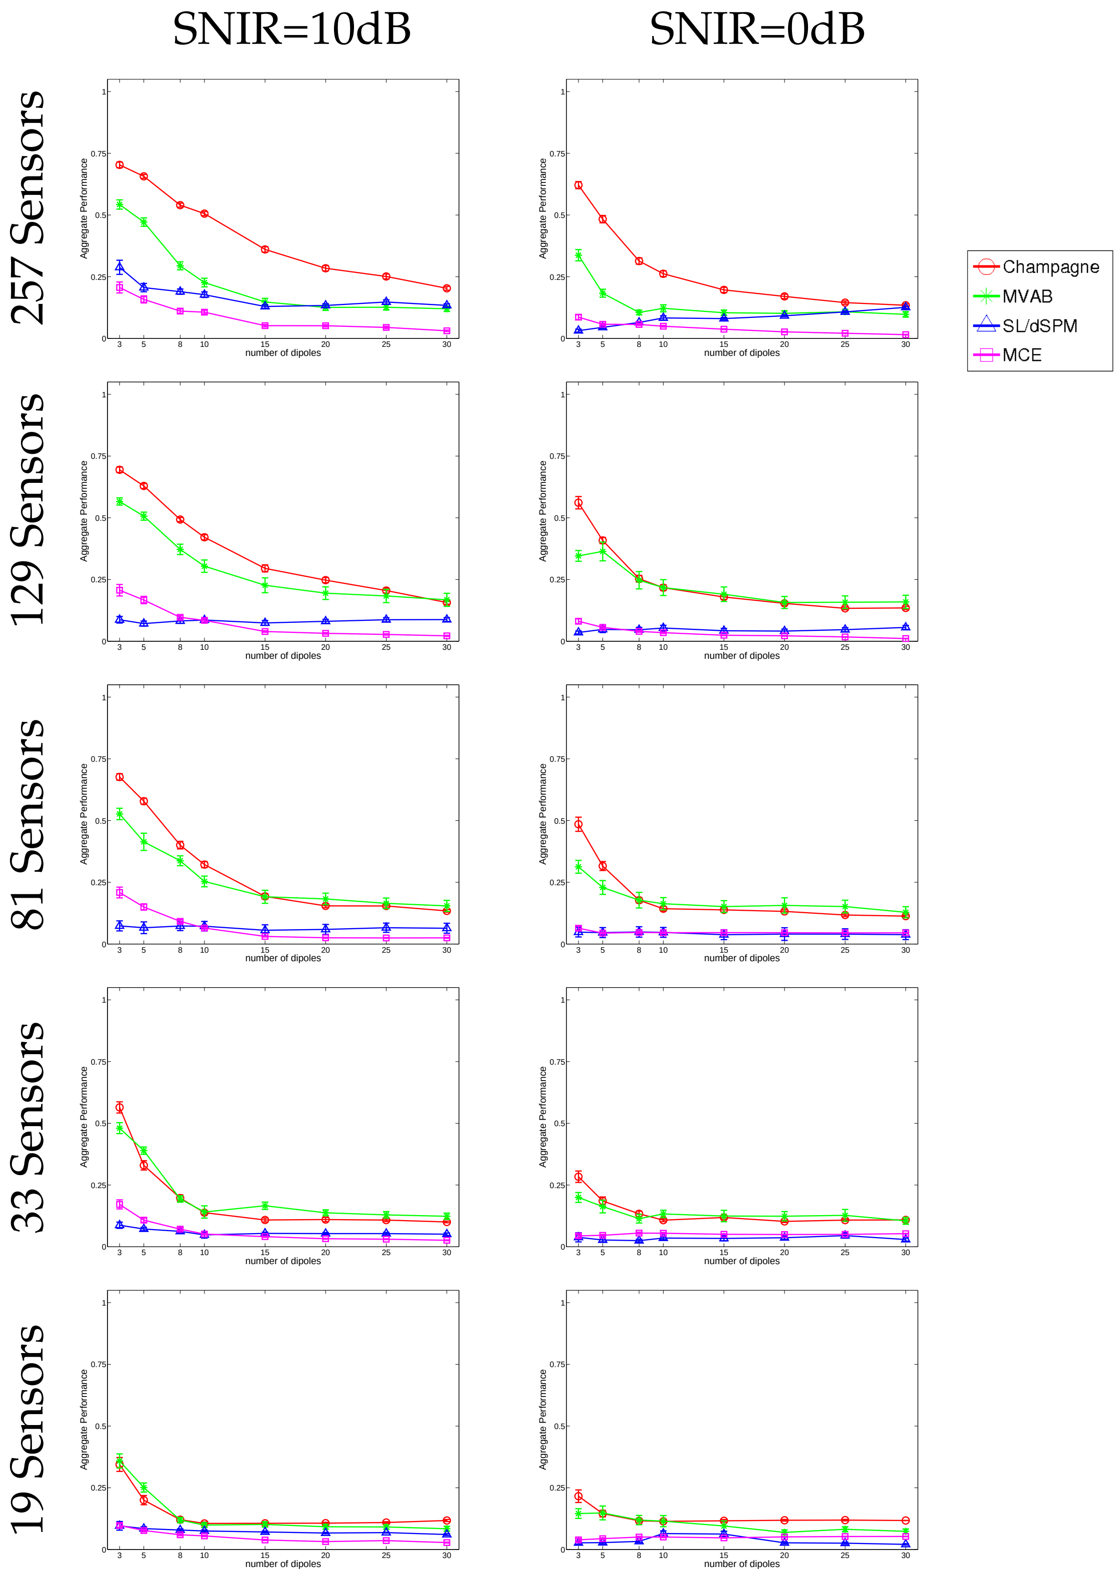

Figure 3.17 Simulations with a vector EEG leadfield with 257, 129, 81, 33, and 19 sensors at 0dB and 10dB.

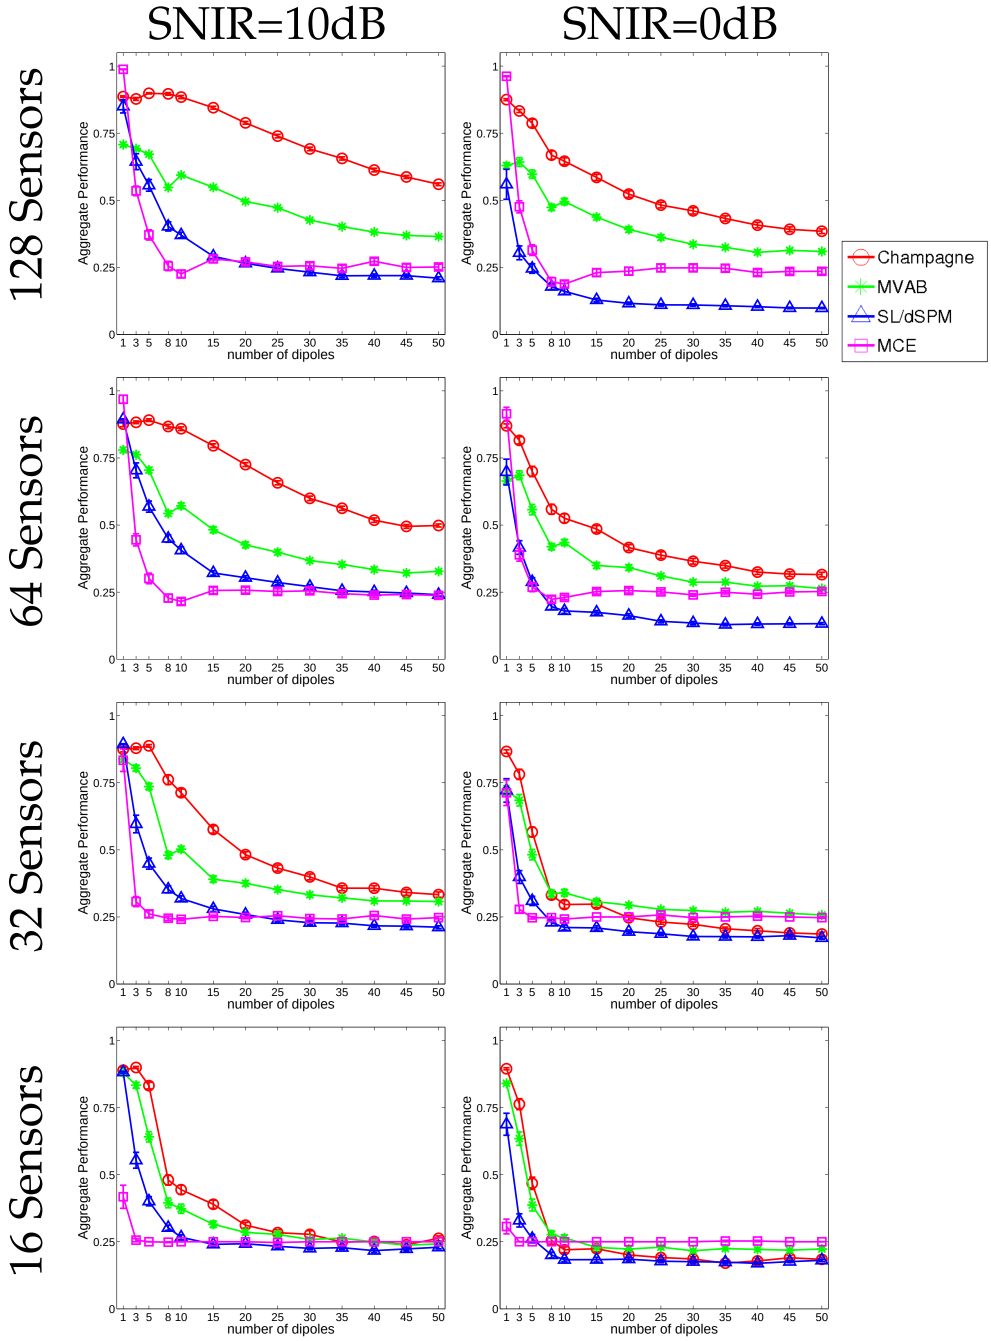

Figure 3.18 Results for EEG simulations with a scalar EEG lead field, AP is plotted against number of dipoles (1 to 50). Also shown is the effect of downsampling the number sensors on performance (shown in each row). The inter-dipole correlation between clusters is 0.5. The results are averaged over 50 simulations at each data point and the error bars show the standard error.

3.5.13 Functional Connectivity

The brain integrates information across multiple brain regions during the performance of a task. Traditionally in the literature, *functional connectivity* describes the correlations between brain areas and *effective connectivity* describes the direction of information flow. (Here, for simplicity, we will only use the term functional connectivity, even when direction is being assessed.) Functional connectivity analyses with MEG depend largely on reliably estimating the voxel time courses. Without good estimates of brain activity, functional connectivity analyses cannot shed light on the interactions and transfer of information in the brain.

Methods of Assessing Functional Connectivity

There are two classes of functional connectivity metrics: bivariate and multivariate metrics. Bivariate methods are metrics that produce a scalar measure of the interdependency between two time series. In the case of neuroimaging data, these values are calculated between every voxel or electrode voltage time series in a pair-wise fashion. These methods can either be symmetric (non-directional) or they can be asymmetric, providing information about the directionality of the interaction, i.e. region A drives activity in region B or vice verse. Multivariate methods incorporate multiple areas (more than two) and multiple time points or lags into the connectivity measures. They are both asymmetric measures and provide information about directionality. In general, multivariate metrics are more involved to compute, especially as the number of areas and lags increase, but they give a more comprehensive picture of the brain interactions.

We have assessed two bivariate metrics, *coherence* and *imaginary coherence* and one multivariate metric, *multivariate autoregression* (MVAR).

The coherence function, $S_{x,y}$, gives the correlation between two time series, $S_x(t)$ and $S_y(t)$ as a function of frequency:

$$F_{x,y}(f) = \mathscr{F}\left\{C_{x,y}(t)\right\} = S_x(f)S_y^*(f) \tag{3.15}$$

Conventionally, the normalized cross spectrum is calculated, where $F_{xx}(f)$ and $F_{yy}(f)$ are the auto-cross spectrums of $x(t)$ and $y(t)$ respectively:

$$C_{xy}(f) = \frac{\mid F_{x,y}(f) \mid}{\mid (F_{xx}(f) \mid\mid F_{yy}(f) \mid)^{\frac{1}{2}}} \tag{3.16}$$

There are volume conduction artifacts in MEG recordings; these artifacts arise from sources blurring across neighboring sensors, as such spurious interactions can be inferred from the sensor data. This blurring occurs at a zero-time lag as the electromagnetic signals travel nearly at the speed of light. There are methods that attempt to isolate the zero-time lag interactions form the non-zero time lag interactions, one of which is imaginary coherence (IC). IC is defined as:

$$IC = imag\left\{C_{x,y}(f)\right\} \tag{3.17}$$

The MVAR model assumes that current values of a time series, $S_1(t)$, is dependent on both previous values of itself, $S_1(t-k)$ and previous values of other time series, $S_2(t-k)$, $S_3(t-k)$ up to $S_n(t-k)$. In brain imaging this is described as a particular brain area's activity being dependent on previous activity in that area and also dependent on previous activity in other regions. The MVAR model is formalized in the following equation:

$$\begin{pmatrix} S_1(t) \\ S_2(t) \\ . \\ . \\ S_n(t) \end{pmatrix} = \sum_{k=1}^{K} A_k \begin{pmatrix} S_1(t-k) \\ S_2(t-k) \\ . \\ . \\ S_n(t-k) \end{pmatrix} + \varepsilon(t) \tag{3.18}$$

where n is the number of nodes in the network, K is the number of delays in the MVAR model, A_k is the mixing matrix at a given delay k, and $\varepsilon(t)$ is the additive noise in the model.

Functional Connectivity

As previously stated, without reliable estimates of brain activity, functional connectivity analyses cannot provide insights into the brain. In the following experiments, we wanted to show that Champagne is a robust source localization algorithm to use when networks of brain activity are expected through simulated data. In the first experiment, we used damped sinusoid time series for the voxel activity and Gaussian noise, instead of the real-brain noise described above. We used a SNIR of 5dB and an intra-dipole correlation of 0.5. We simulated a network of 7 nodes (or voxels), where there were two networks, one left and one right, and two deep ″common source nodes″. The general network structure is depicted in Figure 3.19. We first localized the sources using Champagne, MVAB, MCE, and sLORETA and then assessed the connectivity from the estimated sources. The inter-dipole correlations, α_{intra}, are depicted in the diagram found in Figure 3.24(a) where the color of the lines between the sources denotes the strength of correlation, with red being high and blue being weak (see colorbar in Figure 3.24). The line type indicates whether the mixing was instantaneous (dashed) or non-instantaneous (solid). The ″common source″ nodes were added to simulate the effect of instantaneous correlations on the metrics. The voxels in the left and right networks were all instantaneously coupled with the common source voxels, but the coupling was a different strength for each of the five nodes. We used the pair-wise coherence measure to reconstruct the correlations (shown by the color of the lines) and pair-wise imaginary coherence to determine whether the correlations are instantaneous or non-instantaneous (shown by dashed versus solid lines).

In a second experiment, we increased the coupling between the common-sources and all the voxels relative to the strength of coupling within the networks in order to test the robustness to the common-source interference. In addition to applying the pair-wise metrics, we also tested the performance of the Multivariate Autoregressive (MVAR) model on the 5 node network. Resolving the functional connectivity is complicated by the correlations introduced by the common sources, When we compute the mixing matrix, A, we did not include the common sources; in real experiments we do not have access to the artifacts and sources of noise that introduce instantaneous correlations in the data. In other words, we simulated the data with all seven nodes, but then used only the five nodes of the two networks of interest to perform the MVAR analysis. Champagne was the only source localization algorithm tested for this analysis.

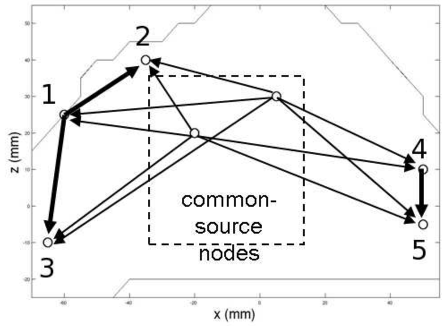

Figure 3.19 General connectivity schema for simulations: This figure shows the general schema for the networks used in the functional connectivity experiments.

Figure 3.20 show the source reconstruction results from Champagne (a), MCE (b), MVAB (c), and sLORETA (d) respectively. The white and black circles mark the true locations of the sources and the surface plot shows the maximum intensity projection of the power of the source estimate at every voxel, illustrating the inferred location of the sources. While MCE came close to uncovering all 5 nodes of the two networks, Champagne was the only algorithm able to resolve the location and time-courses of the two networks entirely. (Both common source nodes were uncovered with Champagne, but one was below the threshold of the image in Figure 3.20 (a).)

The functional connectivity results of the first experiment are depicted in Figure 3.24 and the second experiment are depicted in Figure 3.22. As described above, in the first experiment we used the coherence measure to reconstruct the correlations (shown by the color of the lines) and imaginary coherence to determine which correlations are instantaneous or non-instantaneous (shown by dashed versus solid lines). The similarity of the ground-truth (a) and Champagne (b) plots demonstrates that these two quantities can be used in conjunction to uncover the strength and lags (instantaneous vs. non-instantaneous) of interactions in a network of brain areas. The common sources are not shown to confound the connectivity results with Champagne. MCE (c), MVAB (d) and sLORETA (not shown) showed an over-estimation of the connectivity and fail to reconstruct the ground-truth connectivity. We decided to proceed with the connectivity analysis with all the algorithms, regardless of success, because it is common practice to do region-of-interest analyses.

In the second experiment, we increased the coupling of the common-sources to investigate the resilience of the functional connectivity methods to strong common source component. We found that the pair-wise metrics were not able to uncover the two networks (with any algorithm) due to this increased common-source coupling (not shown). MVAR, on the other hand, is able to uncover the networks of interest by providing accurate information on the direction of the interactions (causality) and the time lag of these interactions, Figure 3.22.

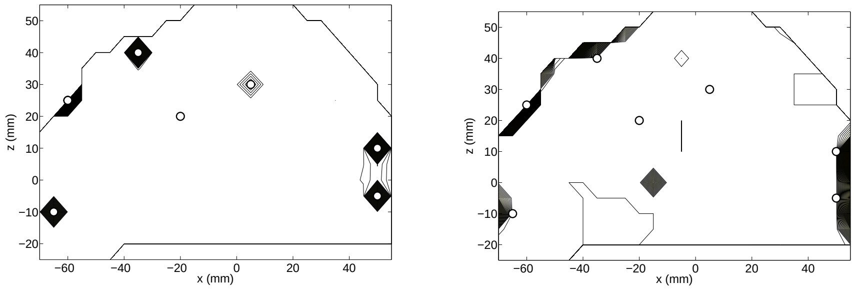

(a) Champagne (b) MCE

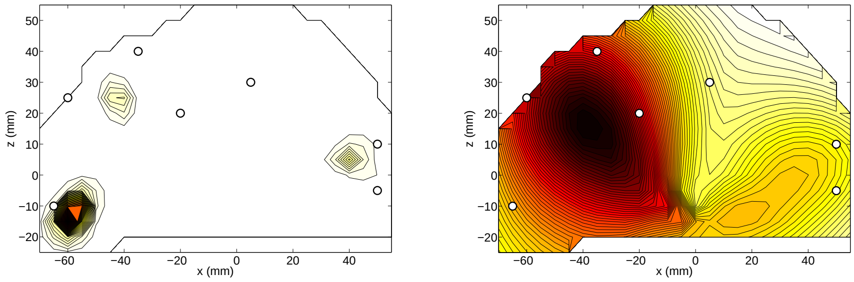

(c) MVAB (d) SL

Figure 3.20 Source Localization: Source localization results for (a)Champagne, (b) MCE, (c) MVAB, and (d) sLORETA. The white circles show the seeded location of the sources and the surface plot shows the estimated location of the sources.

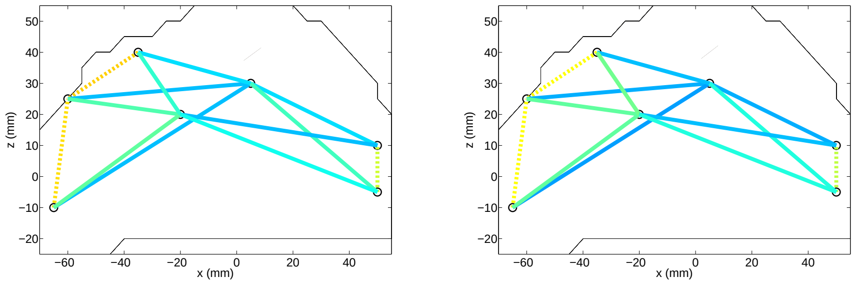

(a) Ground Truth Connectivity (b) Champagne

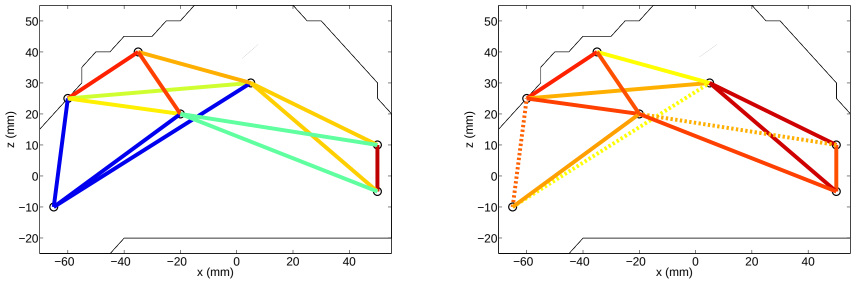

(c) MCE (d) MVAB

Figure 3.21 Functional Connectivity: (a) Ground-truth functional connectivity between sources and ”common sources”. Reconstructed networks using (b) Champagne, (c) MCE and (d) MVAB. The color (see (e)) shows the strength of coupling and the line type shows the lag of integration.(solid for instantaneous, dashed for non-instantaneous).

The results from the second functional connectivity experiment prompted a more extensive investigation of Champagne ability to resolve functional networks of differing sizes, i.e. number of nodes. We simulated networks with 2, 3, and 5 nodes, shown in Figure 3.23. These networks each have a specific causal flow; some nodes drive the activity in other nodes. The time courses were generated by randomly seeding each node with a random, Gaussian time course, $x_i(t)$ and then correlating the nodes though a mixing matrix (A) as defined in 3.5.13.

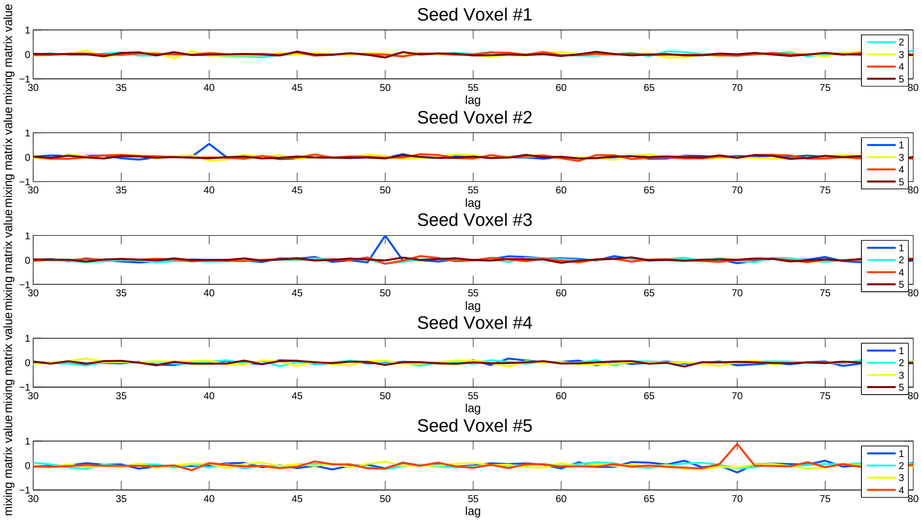

(a)

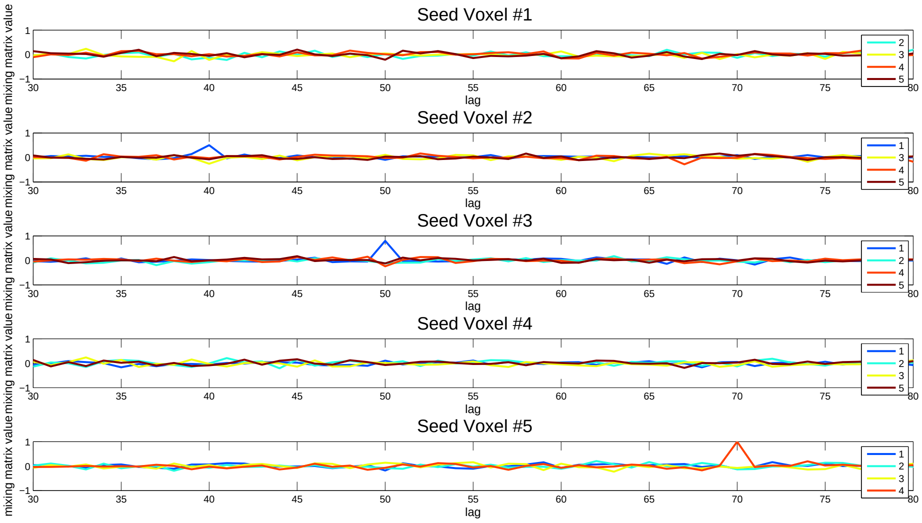

(b)

Figure 3.22 Functional Connectivity: Plot of mixing matrix obtained from voxel time-course reconstructions using Champagne, shown in (b). The reconstructed connectivity found with Champagne is nearly identical to the true connectivity shown in (a). This plot shows that voxel 1 cause voxels 2 and 3, and that voxels 4 causes voxels 5. The values here are normalized to the maximum.

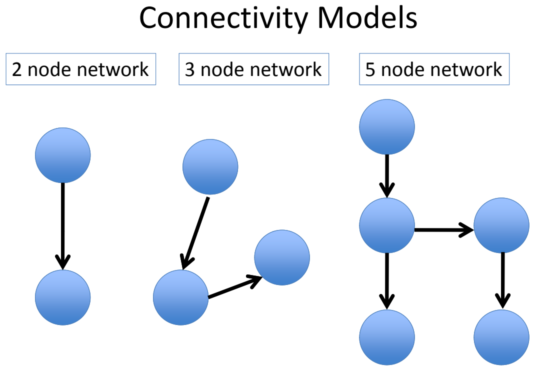

Figure 3.23 Networks with 2, 3, or 5 nodes.

The mixing matrix, A reflects the structure shown in Figure 3.23. For example, for a two node network, the following is an example of a mixing matrix used to simulate the structure:

$$A = \begin{pmatrix} 0.3 & 0.1 \\ 0.5 & 0.3 \end{pmatrix}$$

The diagonal of this matrix has the couplings between a particular node and itself, whereas the off diagonal terms are the couplings between nodes. This A matrix demonstrates that the second node (second row of A) is driven by the first node (first row of A), but the first node is only weakly caused by the first.

For each number of nodes, we simulated a total of 50 runs by randomly seeding the location of the nodes and perturbing the MVAR mixing matrix used to generate the source time courses at SNIR levels of 10dB and 0dB. Since not all possible mixing matrices are stable, we chose a stable mixing matrix for each of the three structures shown in 3.23 and then perturbed the values to test the robustness of the MVAR model to uncover the structure of the network. After seeding the voxel time courses, we project to the sensors and the reconstruct the voxel time courses with Champagne. Once the voxel time courses are obtained, we use the arfit package for MATLAB (http://www.gps.caltech.edu/ tapio/arfit) to obtain an estimation of the mixing matrix, $\tilde{A}$.

In order to test the success of both the source localization and the estimation of the network, we used three metrics. We used the A' metric to test the localization accuracy and the correlation coefficient R to assess the estimation of the time courses. We added a third metric, the multidimensional correlation coefficient between the true mixing matrix, A and the estimated mixing matrix, $\tilde{A}$. This is similar to the time course correlation coefficient, R, but it measures the similarity between two matrices as opposed to two vectors.

In Figure 3.24, we present the results from this experiment. The A' and R results, seen in Figure 3.24 (a) and (b) respectively, demonstrate the Champagne is able to localize activity and reconstruct the time courses at both SNIR levels. The A matrix correlation coefficient results show that even though the localization and time course reconstruction are near perfect, the A matrices estimated from the time courses are only reliable to 2 and 3 node networks at 0dB and even at 10dB, the A matrices are not correctly estimated for the 5 node network (the correlation coefficient was less than 0.75).

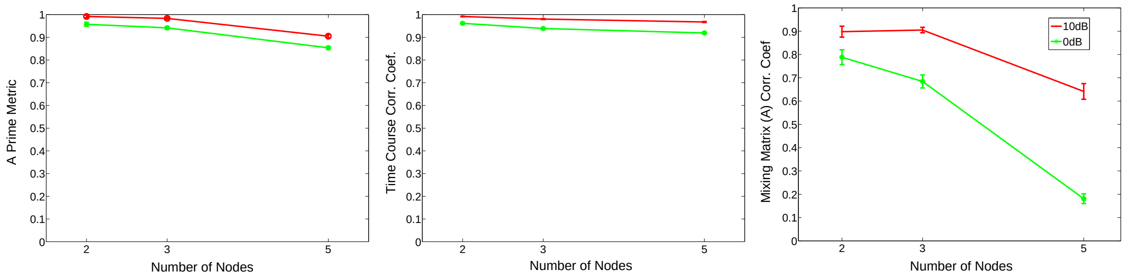

(a) (b) (c)

Figure 3.24 Results from simulations with 2, 3 and 5 nodes at 10dB (red) and 0dB (green). We show the A' metric in (a) and R in (b) to show the localization and time course reconstruction results. In (c), we show the mixing matrix correlation coefficient, which demonstrates the ability to reconstruct the connectivity of the network from the estimated time courses. The results are averaged over 50 simulations with standard error bars.

3.6 Performance on Real Data

We used a number of real data sets to examine Champagne's performance on real MEG and EEG data. All MEG data was acquired in the Biomagnetic Imaging Laboratory at UCSF with a 275-channel CTF Omega 2000 whole-head MEG system from VSM MedTech (Coquitlam, BC, Canada) with a 1200 Hz sampling rate. The lead field for each subject was calculated in NUTMEG (Dalal et al., 2004) using a single-sphere head model (two-orientation lead field) and a 8mm voxel grid. The data was digitally filtered from 1 to 160Hz to remove artifacts and the DC offset was removed. The EEG data was downloaded from the SPM website (http://www.fil.ion.ucl.ac.uk/spm/data/mmfaces) and the lead field was calculated in SPM8 using the coarse resolution. Note that all MRI images are shown in the neurological convention, i.e. right is on the right.

3.6.1 Auditory Evoked Field

We analyzed the neural responses of seven subjects to Auditory Evoked Field (AEF) stimulus, which was elicited with single 600ms duration tones (1 kHz) presented binaurally. The data was averaged across 120 trials (after the trials were time-aligned to the stimulus). The pre-stimulus window was selected to be -100ms to 5ms and the post-stimulus time window was selected to be 5ms to 250ms, where 0ms is the onset of the tone. We applied Champagne, MVAB, SL/dSPM, and MCE to these data sets to give a comparison for performance across algorithms.

The results from analyzing the AEF data from seven subjects are shown in Figure 3.25. In each subplot, we show the power at each voxel in a 50ms to 75ms window around the M100 peak. Champagne is able to localize bilateral auditory activity for all seven subjects (shown in the first column of Figure 3.25). The activity is in Heschel's gyrus, which is the location of primary auditory cortex. SL/dSPM is able to localize bilateral auditory activity in five of the seven subjects. In these five subjects, the activations are diffuse and in most cases biased to one side. MVAB is only able to localize bilateral activity in one subject (Subject 5) and in the other six cases it localizes the activity to the center of the head. MCE is only able to localize the auditory activity on one side in most subjects. In Subjects 4 and 6, MCE is able to localize bilateral activity, but the activity is more lateral than Heschel's gyrus. The MCE algorithm favors voxels on the edge of the voxel grid, and often does not accurately localize cortical areas.

3.6.2 Audio-Visual Task

We analyzed a data set designed to examine the integration of auditory and visual information. We presented single 35ms duration tones (1 kHz) simultaneous with a visual stimulus. The visual stimulus consisted of a white cross at the center of a black monitor screen. The data was averaged across 100 trials (after the trials were time-aligned to the stimulus). The pre-stimulus window was se-

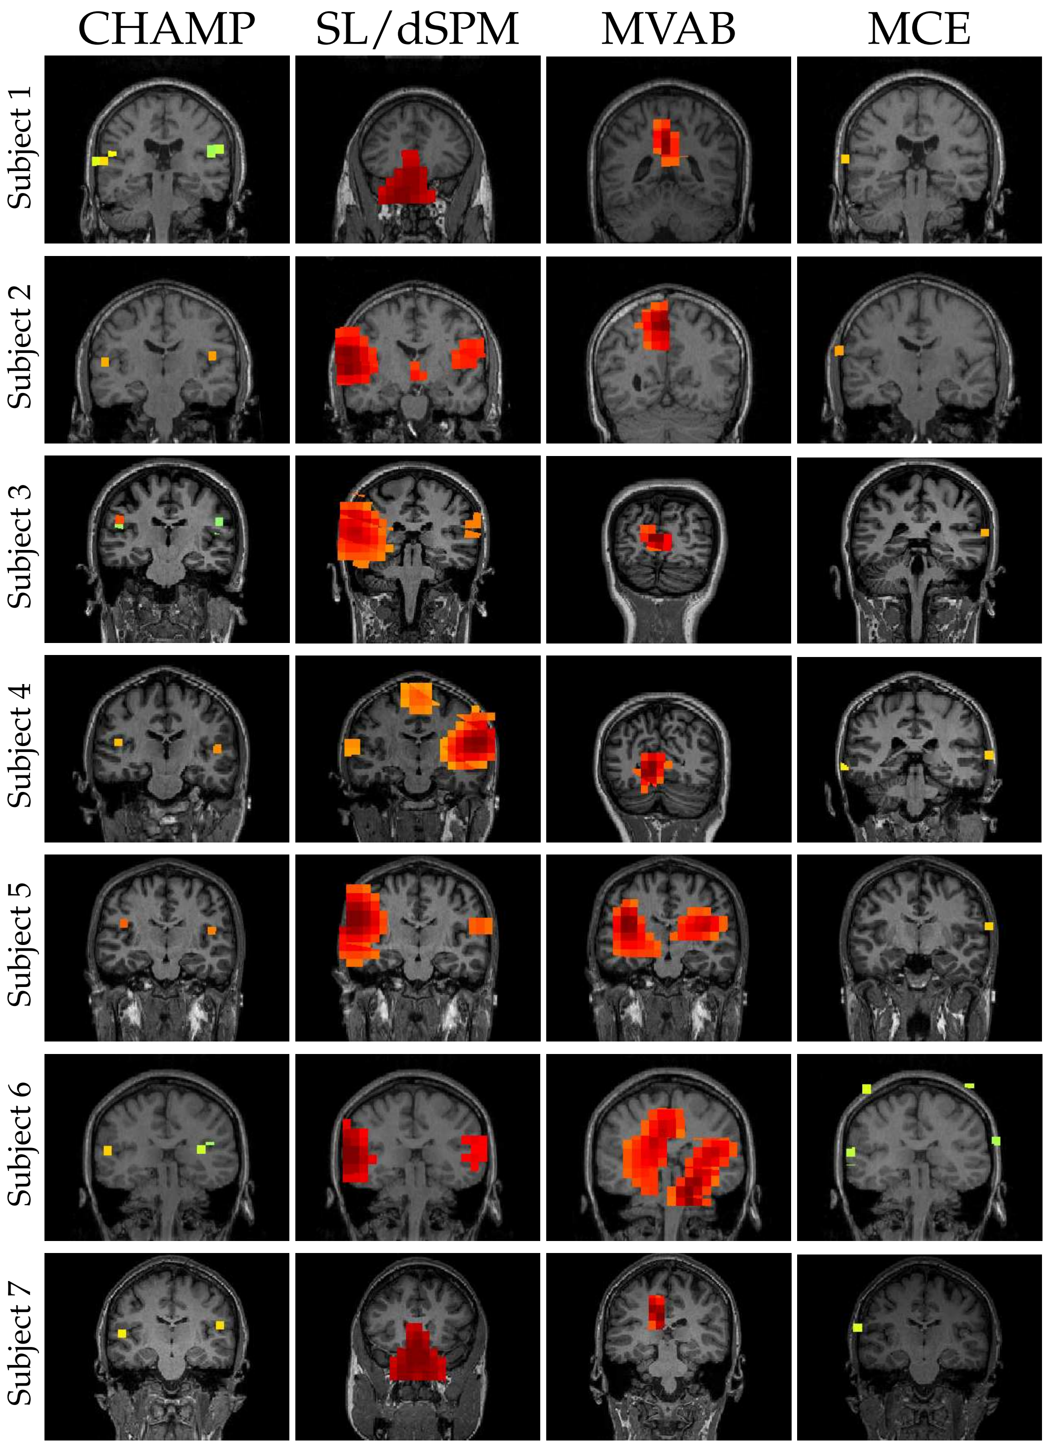

Figure 3.25 Auditory evoked field (AEF) results for 7 subjects. The results from using Champagne are shown in the right-most column and the results from using the benchmark algorithms are shown the other three columns.

lected to be -100ms to 5ms and the post-stimulus time window was selected to be 5ms to 450ms, where 0ms is the onset of the simultaneous auditory and visual stimulation.

Champagne's results from the audio-visual task are presented in Figure 3.26. In the first and second rows, we show the brain activations associated with the auditory stimulus. Champagne is able to localize bilateral auditory activity in Heschel's gyrus in the window around the M100 peak, shown in Figure 3.26(a) and (c). The time courses for the left and right auditory sources are shown in Figure 3.26(b) and (d), along with the window used around the M100 peak. For all the plots presented in the following sections, the crosshairs in the plot to the left show the location of the voxel whose time course is plotted to the right. The two auditory sources had the maximum power in the window around the M100 peak. Second, we present the early visual response in Figure 3.26(e) and (f). Champagne is able to localize a source in medial, occipital gyrus with a peak around 150ms. We plot the power in the window around this peak in Figure 3.26(e) and the time course of the source marked with the crosshairs in Figure 3.26(f). The voxels with the maximum power in this window are still found in the auditory areas (seen in Figure 3.26(a) and (c)) as the M100 peak bleeds into the time window used to localize the visual activation (around 150ms). Using a later time window shown in Figure 3.26(g) we can localize a later visual response with a time course (Figure 3.26(h)) that has power extending past 200ms.

We also ran the benchmark algorithms on this data set. The results from SL/dSPM in Figure 3.27, demonstrate that SL/dSPM was able to localize a source in the left auditory cortex (3.27 (a)) and a visual activation (3.27 (e)). The peaks for the auditory and visual sources, see Figure 3.27 (b) and (f), have peaks around 100ms and 150ms respectively. SL/dSPM also localizes a large source along the midline in the precuneous (Figure 3.27 (c)), but this source has the same time course at the auditory source, see Figure 3.27 (d). We can see that the localizations provided by MVAB and SL/dSPM are diffuse and do not capture the auditory and visual activity that Champagne is able to localize. The results from MCE are presented in Figure 3.28; these results show that MCE is able to localize bilateral auditory activations (Figure 3.28 (a), (c)) and some visual activity (Figure 3.28 (e)). The time courses for the auditory and visual sources are shown in Figure 3.28 (b), (d), and (f). These time courses look most similar to those estimated with Champagne and are able to provide a clear distinction between the auditory activity at 100ms and the visual activity around 150ms. As seen in the AEF results, the MCE algorithm favors voxels on the edge of the voxel grid and often, does not accurately localize cortical areas, but the activations seem to be over reasonable locations for activity given the task. The results from MVAB, in Figure 3.29, show that MVAB can localize auditory activity Figure (3.29 (a) left) and a visual activation (Figure 3.29 (c)). The time courses, Figure 3.29 (b) and (d), for these sources (shown by the crosshairs) do not contain useful information. When the power was taken around the M100 and around 150ms, the localization results were in the same locations, but with less power.

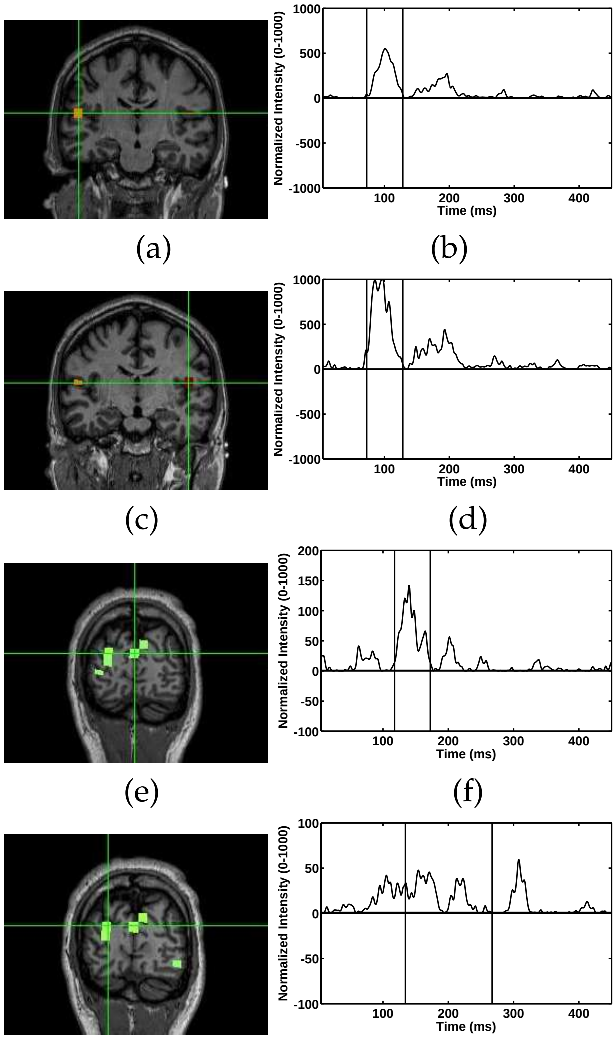

(a) (b)

(c) (d)

(e) (f)

(g) (h)

Figure 3.26 Champagne localizes a bilateral auditory response at 100ms after the simultaneous presentation of tones and a visual stimulus. Champagne localized bilateral auditory activity, (a) and (c), with time courses shown in (b) and (d). Champagne localizes an early visual response at 150ms after the simultaneous presentation of tones and a visual stimulus. The time course in (f) corresponds to the location indicated by the crosshairs in the coronal sections (e). Champagne localizes a later visual response later than 150ms after simultaneous presentation of tones and a visual stimulus. The time course in (h) corresponds to the location indicated by the crosshairs in the coronal sections (g).

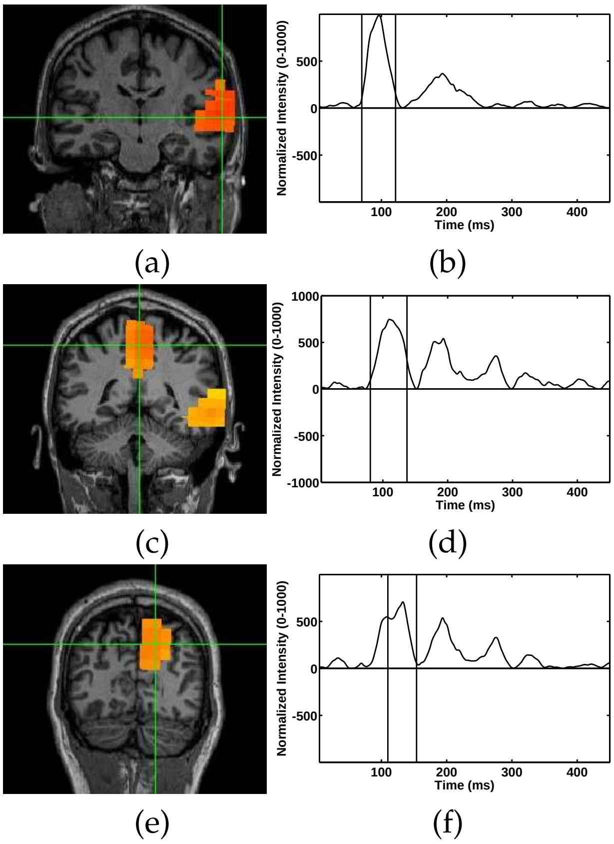

Figure 3.27 SL localizes a source in left auditory cortex (a) and a source in visual cortex (e). SL also localizes a source in the center of head that does not have functional significance (c). The time courses in (b), (d), and (f) correspond to the crosshair locations in (a), (c). and (e), respectively.

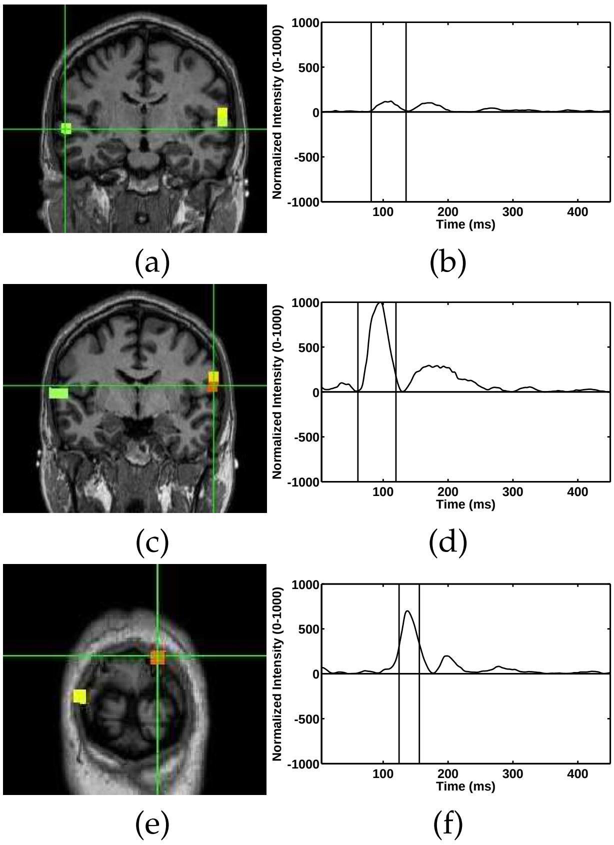

Figure 3.28 MCE localizes bilateral activity over auditory cortex (a) and a source over visual cortex (b,c). MCE favors the voxels on the edge of the voxel grid and does not successfully localize cortical areas, but does show activity above functionally significant areas.

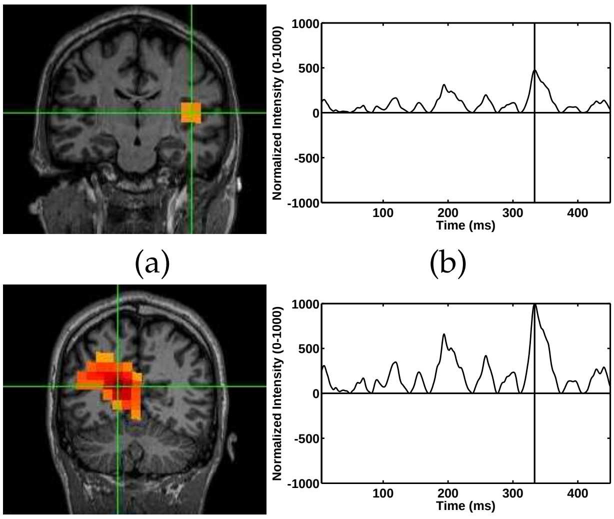

(a) (b)

(c) (d)

Figure 3.29 MVAB localizes auditory activity on the left (a) and one source in a visual area (c) for the audio-visual data. The time courses in (b) and (d) correspond to the crosshair locations in (a) and (c), respectively.

3.6.3 Face-Processing Task

MEG

We analyzed a MEG data set from a subject in which faces and scrambled faces were presented in a random order with an interstimulus interval of 1 second. We applied Champagne to this dataset. For the face stimulus, the pre-stimulus window was selected to be -200ms to 5ms and the post-stimulus time window was selected to be from 5ms to 450ms, where 0ms is the appearance of the visual stimulus. In addition to running Champagne on the face processing data, we also constructed a contrast dataset to localize the brain areas more active when the faces were presented than when scrambled faces were presented. The pre-stimulus period of this data set consisted of the trial-averaged post-stimulus period (5ms to 450ms) of the scrambled-face data and the post-stimulus period consisted of the trial-averaged post-stimulus period (5ms to 450ms) of the face data.

The face processing results are shown in Figures 3.30 and 3.32. First we examine the brain areas involved in the processing of the faces. Figure 3.30(a) shows an early visual response to the presentation of the face visual stimulus in medial occipital activation with time course in Figure 3.30(b). Figure 3.30(c) shows a later visual response more lateral to the early response in occipital cortex with the time course in Figure 3.30(d). In Figure 3.30 (e) and (g) we show a bilateral activation in the fusiform gyrus with time courses shown in 3.30 (f) and (h) that show peaks around 170ms. In the contrast condition, results in Figure 3.32, we can see a bilateral activation in the fusiform gyrus Figure 3.32 (a) and (c) with a peak around 170ms (time courses shown in Figure(b) and (d)). The activations in the fusiform gyrus are the maximum in the window shown in Figure(b) and (d), around 170ms. The results concur with those from the literature (Kanwisher et al., 1997). An area in fusiform gyrus (on the ventral surface of the occipital lobe), called fusiform face area (FFA), has been shown to be an area preferentially activated by the presentation of faces versus other visual stimuli, such as scrambled faces. In Figure 3.31 the ventral surface of the brain is shown; Champagne is able to localize sources on the ventral surface with peaks around 170ms.

The benchmark algorithms were also applied to the face processing data set.

The results from MVAB are shown in Figure 3.33. MVAB fails to localize any activity of interest, rather it localizes one source close to the center of the head Figure 3.33(a) and does not estimate a meaningful time course for this activity (Figure 3.33 (b)). SL/dSPM is more successful on this data set, results shown in Figure 3.34. A visual source is localized (Figure 3.34 (a)), in addition to bilateral sources near to fusiform gyrus Figure 3.34 (c) and (e). The time course for the visual source (Figure 3.34 (b)) shows a peak at 100ms. The time courses for the sources in fusiform gyrus (Figure 3.34 (d),(f)) show peaks at 170ms, but the source on the left has a larger peak at 100ms. These time courses do not help to separate activity in the visual sources and fusiform gyrus as clearly as the time courses obtained from Champagne. SL/dSPM also localizes activity near the center of the head, which is not functionally relevant, as seen in Figure 3.34 (c) and (e). MCE is able to localize a visual source (Figure 3.35 (a)) with a time course (Figure 3.35 (b)) that has a peak around 100ms. It is also able to localize bilateral sources near fusiform gyrus, as seen in Figure 3.35 (c) and (e). The time courses for these sources, shown in Figure 3.35 (d) and (f), have peaks around 170ms, but the left fusiform source does not have its maximum at 170ms. Since these benchmark algorithms do not account for the pre-stimulus period in their generative models, we did not run the face versus scrambled contrast.

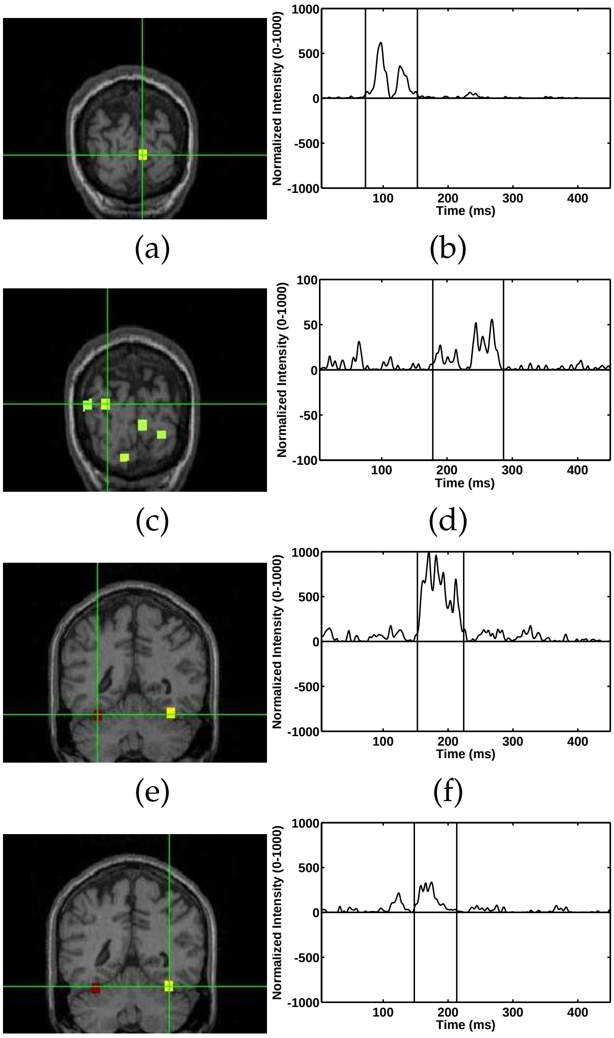

(a) (b)

(c) (d)

(e) (f)

(g) (h)

Figure 3.30 Champagne localizes an early visual response 100ms after the presentation of a face stimulus, shown in (a). The time course for this source is shown in (b) and corresponds to the location indicated by the crosshairs in the coronal section (a). Champagne localizes a later visual response around 200ms after the presentation of a face stimulus, seen in (c), with the time course shown in (d). Champagne localizes a bilateral activation in fusiform gyrus that is thought to be in FFA, shown in (e) and (g). The peak for these sources is around 170ms after the presentation of a face stimulus, time courses shown in (f) and (h).

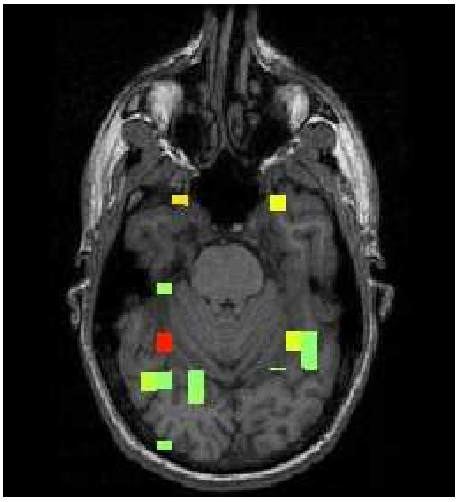

Figure 3.31 A view of the ventral surface of the brain. Champagne localizes many sources on the ventral surface of the occipital and temporal lobes with peaks around 170ms.

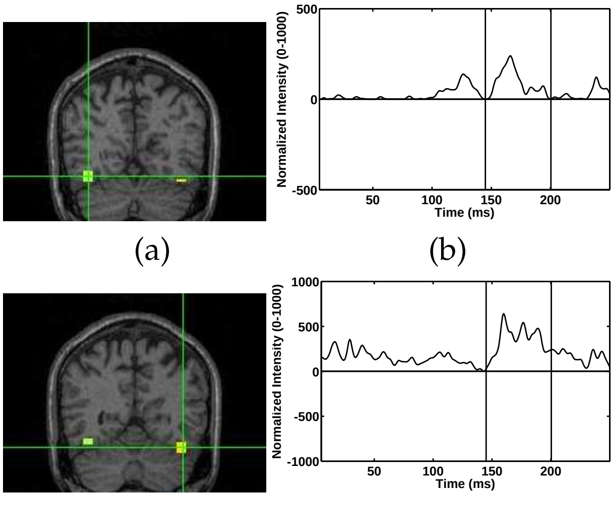

(a) (b)

(c) (d)

Figure 3.32 Champagne localizes a bilateral activation in fusiform face area (FFA) with the face versus scrambled data set, seen in (a) and (b). The peak is around 170ms after the presentation of a face stimulus, as seen in the time courses for these sources in (b) and (d).

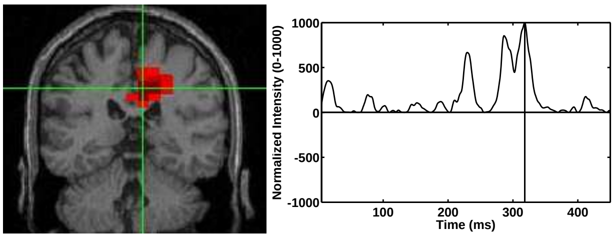

(a) (b)

Figure 3.33 MVAB localizes only one source near the center of the head, localization seen (a) and time course seen in (b).

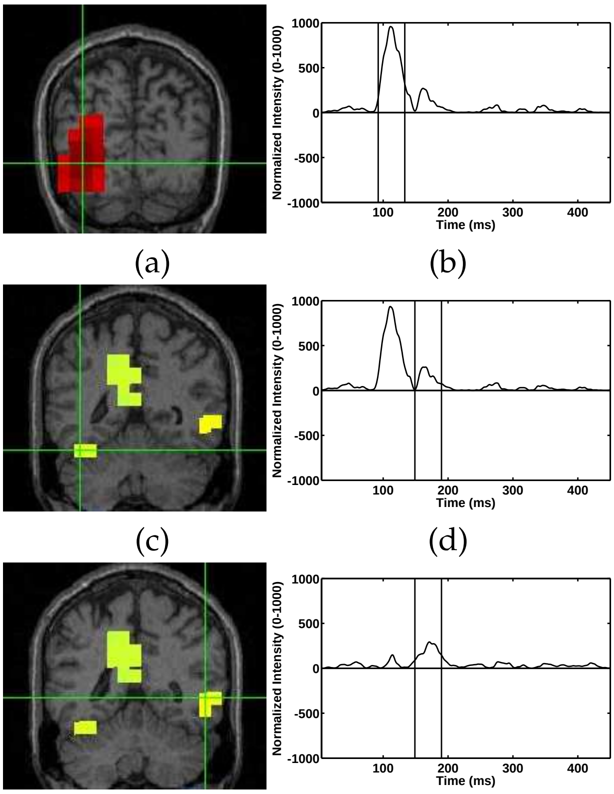

(a) (b)

(c) (d)

(e) (f)

Figure 3.34 SL is able to localize one source in a lateral visual area in middle occipital cortex (a), with the time course shown in (b). It is also able to localize bilateral activation in (or near) the fusiform gyrus (c,e), with time courses in (d,f).

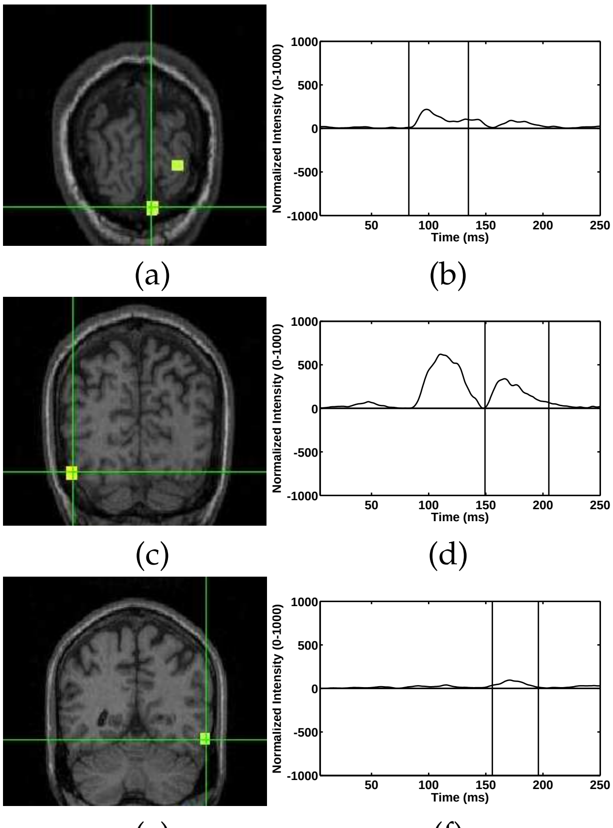

(a) (b)

(c) (d)

(e) (f)

Figure 3.35 MCE is able to localize a visual source (a) and bilateral activity in the fusiform gyrus, (c) and (e). The time courses for these sources are shown in (b), (d) and (f).

EEG

The EEG data were downloaded from the SPM website (http://www.fil.ion.ucl.ac.uk/spm/data/mmfaces/), which were acquired on a 128-channel ActiveTwo system at 2048 Hz. The paradigm involves randomized presentation of at least 86 faces and 86 scrambled faces, although we did not use the scrambled face data. We averaged the time-aligned the trials to the presentation of the face and created an averaged data set. The pre-stimulus window was selected to be -200ms to 5ms and the post-stimulus time window was selected to be 5ms to 250ms. For this real data set we found that using the three-component (vector) lead field in SPM was more robust than the orientation constrained lead field. The lead field was cortically constrained and we selected the coarse tessellation for our grid resolution, which resulted in approximately 5000 voxels.

The results from using Champagne on this EEG data set are shown in Figure 3.36. The time courses for the peak voxels are plotted; the arrows point from a particular voxel to its time course. In Figure 3.36(a), we see that Champagne is able to localize early visual areas that have a peak around 100ms. In Figure 3.36(b), the ventral surface of the brain is shown. There are a few activations in and around fusiform gyrus. These activations are larger in extent and have peaks around 170ms corresponding to the N170 seen in the sensor data. These results are consistent with those obtained in (Henson et al., 2010) using the same EEG data set.

The benchmark algorithms were applied to the EEG data set as well. The results are shown in Figure 3.37. The time courses for the peak voxels are plotted; the arrows point from a particular voxel to its time course. MVAB (Figure 3.37 (a),(b)) and SL/dSPM (Figure 3.37 (c),(d)) are able to localize sources at the occipital pole, an early visual processing area. They are not able to localize distinct sources on the ventral surface, as was Champagne. They also find sources in other areas of the brain that are not functionally relevant. MCE is able to localize a visual source and a source on the left ventral surface of the occipital lobe, seen in Figure 3.37 (e) and (f). With the exception of these two sources, all the remaining voxels are pruned to zero with MCE.

3.6.4 Localizing Spiking Activity in an Epileptic Patient

We used the resting-state MEG data from a patient with recurrent seizures to assess Champagne's ability to localize spiking, epileptic activity. This patient had previously underwent surgery to resect a large portion of her frontal lobe, but she still had seizures after recovering from the surgery. It was unclear whether the seizures were originating from the region around the lesion (perilesional) or from another part of the brain. We performed source localization on a single spike with Champagne and the benchmark algorithms to assess the source of her spiking activity. We used -200ms to -5ms for the pre-stimulus period and 0ms to 200ms

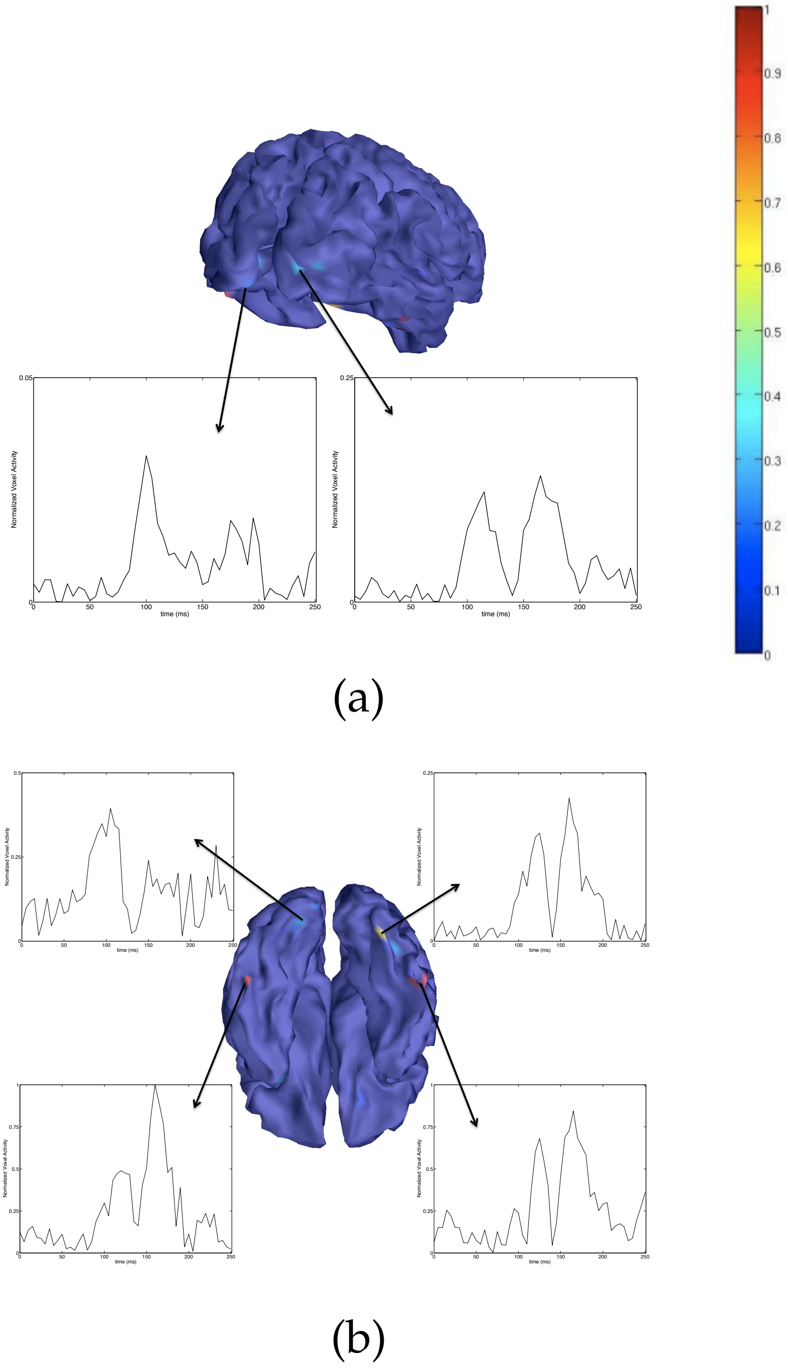

Figure 3.36 EEG Face Processing Data: (a) Shows two early visual responses in occipital cortex with the time courses. (b) Shows four ventral activations in (or near) fusiform face area (FFA) with time courses showing peaks around 170ms. (Ventral side of brain shown in (b), with right on the right.)

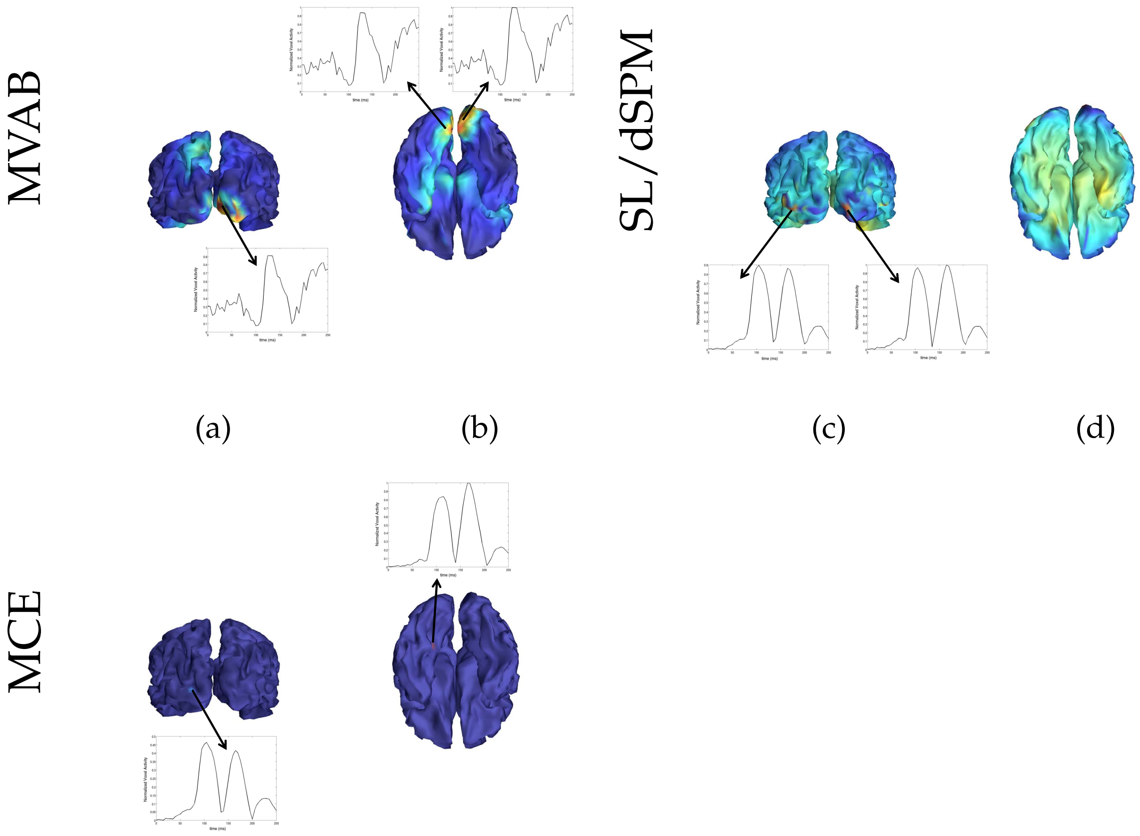

(a) (b) (c) (d)

(e) (f)

Figure 3.37 EEG Face Processing Data: Benchmark algorithms: BF, SL/dSPM, and MCE were applied to the data. BF, (a,b) is able to localize visual areas in the occipital pole. SL/dSPM, (c,d) is also able to localize occipital (visual) sources, particularly on the occipital pole. MCE, (e,f) is able to localize two sources, one in left occipital cortex and one on the ventral surface of the brain. (The occipital lobe is shown in (a), (c), and (e) and the ventral surface of brain shown in (b), (d), and (f) with right hemisphere on the right.

for the post-stimulus period, the spike occurred just after the 0ms point.

The results from running Champagne on the single spike are found in Figure 3.38. Champagne was able to localize the epileptic activity to a region around the lesion from the previous surgery. The MEG data was collected prior to a second surgery and we learned that removing perilesional brain matter in a second surgery has relieved the patient's symptoms. This anecdotal post-operative information indicates that Champagne is effective in localizing spiking activity with the data from one spike. The results from running MCE are presented in Figure 3.39. MCE localizes a similar region on the border of the lesion. The results from MVAB and SL are shown in Figure 3.40 and 3.41, respectively. The maximum voxel for MVAB is in the center of the head, Figure 3.40(a), but it is able to localize a weak perilesional peak, Figure 3.40(b). SL shows a maximum peak in the most anterior part of the frontal cortex, Figure 3.41(a), but is able to localize weaker, perilesional activity, Figure 3.41(b).

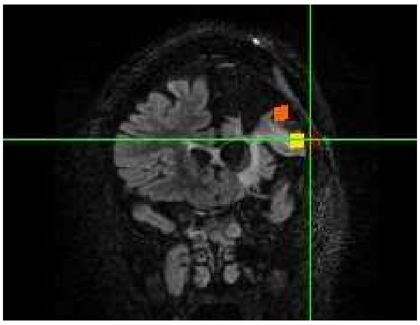

Figure 3.38 Localization results for Champagne on a single spike in a epileptic patient who had seizures even after a area of her frontal cortex was removed. Champagne is able to localize the spike to a region around the lesion, which is confirmed to be the source of the epileptic activity as reported by surgeons who performed a second surgery on this patient. The location of the crosshairs indicates the maximum voxel in the VOI.

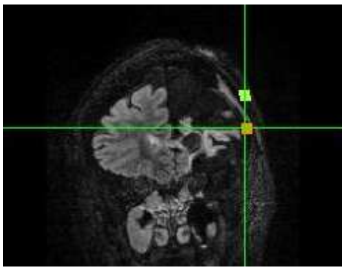

Figure 3.39 Localization results for MCE on a single spike in a epileptic patient. MCE is able to localize the spiking activity to a region on the border of the lesion.

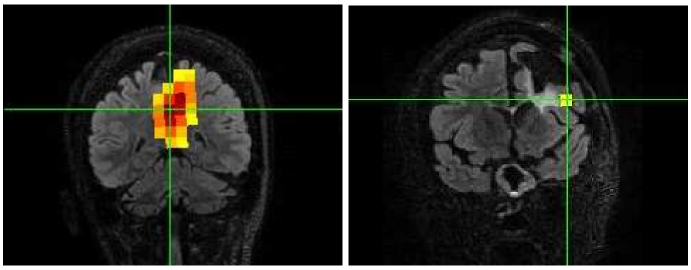

(a) (b)

Figure 3.40 Localization results for MVAB on a single spike in a epileptic patient. MVAB localizes the maximum for the spike to the center of the head (a), but also finds a peak perilesional (b).

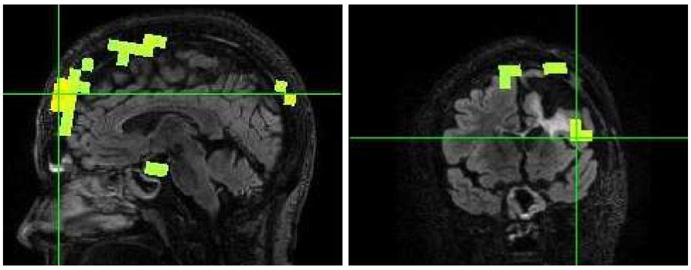

(a) (b)

Figure 3.41 Localization results for SL on a single spike in a epileptic patient. SL localizes the maximum for the spike to the an anterior region of the frontal cortex(a), but also finds a peak perilesional (b).

3.6.5 Somatosensory Evoked Field

We used two somatosensory evoked field (SEF) data. This paradigm should elicit an activation in the contralateral (left) somatosensory cortex when the right index finger (RD2) is stimulated. The stimulation is administered by air puffs with a pseudorandom interstimulus interval of 450 to 500ms. The data was averaged over 252 trials and the pre-stimulus period was selected to be -90ms to -5ms and the post-stimulus period was selected to be 5ms to 250ms.

SEF data is typically one of the easiest on which to perform source localization. The results from applying Champagne to this data set can be seen in Figure 4.2.2. Champagne is not able to localize the contralateral somatosensory cortex in this subject, despite the simplicity of the data; the other source localization algorithms are able to localize the contralateral somatosensory cortex. In order to fully explore why Champagne was unsuccessful on this data set, we expanded on this investigation by using unaveraged data from both the stimulation of RD2 and of the upper right lip. We concatenated the pre- and post-stimulus periods of each trial to form one long epoch for both the pre-stimulus and post-stimulus periods. For the pre-stimulus period, we took the window of data between -100 to -5ms from each trial and for the post-stimulus period, we took the window between 5ms to 200ms, where 0ms is the onset of the stimulus. We used both the full number of trials (252) and only the first five trials for the two data sets to examine the effect of drastically reducing the amount of data (and the signal-to-noise ratio) on somatosensory localization.

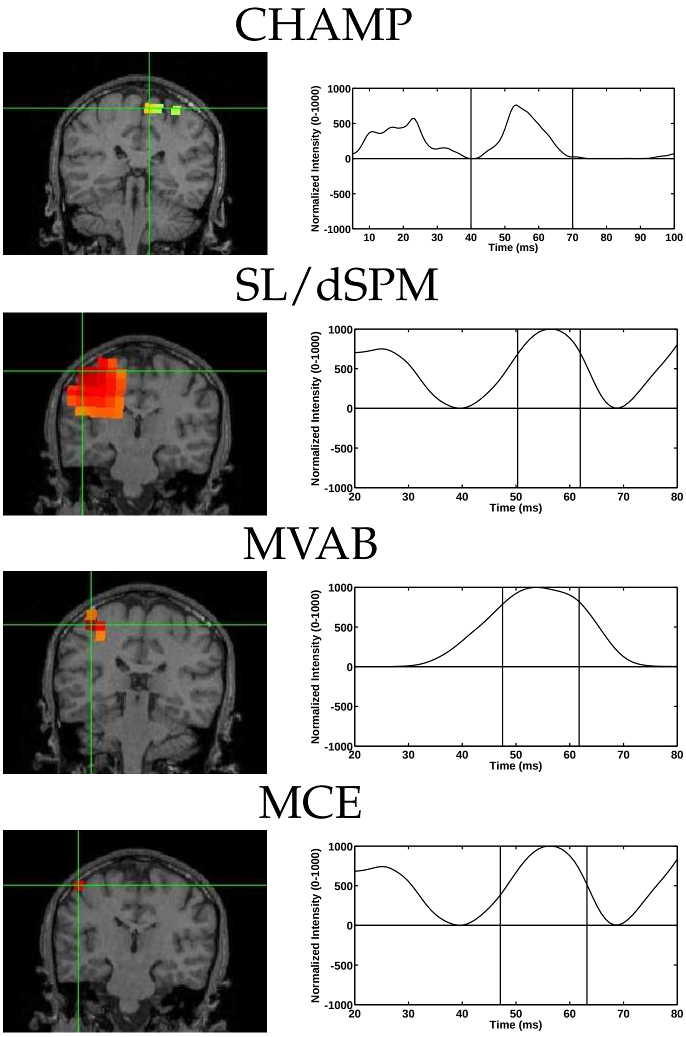

Figure 3.42 Localization results for Champagne, MVAB, SL/dSPM, and MCE for the SEF data. All the algorithms yield good localization and time course recovery except Champagne. Champagne mis-localize the activity to the ipsilateral hemisphere, as opposed to the contralateral hemisphere, as the other algorithms show.

The results from the unaveraged data experiment with SEF experiments are shown in Figure 3.43. For each plot, the power at each voxel was computed and the power map was projected to the surface of a MNI template brain. We present the results from 5 and 252 trials for both stimulation of RD2 and the upper right lip. The first and second rows show the results from stimulating RD2 using 252 trials and 5 trials, respectively, and the third and fourth rows show the results from stimulating the upper right lip using 252 trials and 5 trials, respectively. (We projected to the surface of a MNI template brain to show the location along the somatosensory strip.) Stimulation of RD2 elicits a left-lateralized activation in the left (contralateral) somatosensory cortex. Champagne is able to localize this source using both 252 trials and 5 trials. Stimulation of the right upper lip elicits a bilateral response in somatosensory cortex for both 252 and 5 trials. These results correspond to a vast literature on somatosensory localization.

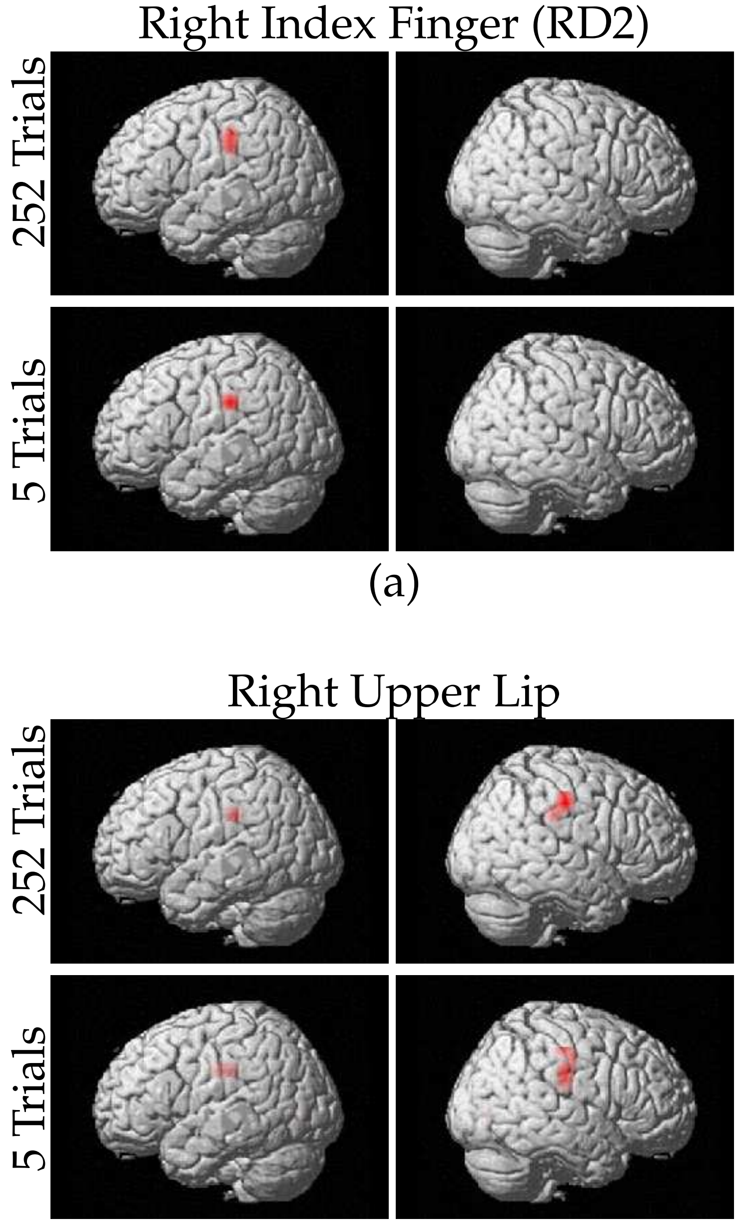

Figure 3.43 Somatosensory processing for stimulation of the right index finger (RD2) shown in (a) and the right upper lip shown in (b) as localized by Champagne. Results from using all the trials of data (252 trials) shown in upper rows of (a) and (b) and results from using 5 trials shown in lower rows of (a) and (b). The data was not trial-averaged for this data set.

3.7 Discussion

In this chapter, we present the results from a rigorous battery of performance tests of Champagne, our novel source-localization algorithm. For the most part, we have compared Champagne's performance to commonly-used source localization algorithms (MVAB, SL/dSPM, and MCE). The tests use simulated and real data; the simulated data explores performance with difficult source configurations and the real data demonstrates Champagne's ability localize real brain activity in the face of source correlations and real noise, interference, and signal-to-noise ratios.

The experiments with simulated data exemplify that Champagne provides robust localization and time course estimation with complex source configurations and noisy, correlated sensor data. At high SNIR (10dB), Champagne almost perfectly resolves two distinct, correlated sources even when there is only one voxel's width between the sources. The other algorithms need at least 32mm (or 3 voxels' width) between the sources in order to accurately distinguish two sources. Champagne is able to distinguish two neighboring, correlated sources, which could have real implications when analyzing real MEG data with brain areas activated in close proximity to one another. This is true of speech motor data as auditory sources are close to the mouth motor area. Auditory/tactile experiments also activate regions in close proximity; sources in auditory cortex and secondary somatosensory cortex are on either side of the lateral sulcus. In addition, visual areas, such as V1/V2/V3/V4 are close to one another and are typically hard to resolve.

At the SNIR level, 10dB, we found that Champagne is able to localize 10 sources when using a vector (two orientation) lead field when the intra-dipole correlation (α_{intra}) was 0.25 and more than 30 dipoles when using a scalar lead field. These numbers are well above the number of sources that the benchmark algorithms are able to localize. At 20dB and a vector lead field, Champagne is able to localize up to 15 sources, and is able to perfectly localize 10 sources. The type of source time course and inter-dipole correlation has a large effect on the number of sources Champagne can localize with simulations of over 10 dipoles. The majority of our simulations were done with $\alpha_{inter} = 0.5$ and damped sinusoidal sources, which is the most difficult configuration of correlations and source time courses to localize.

We did an additional experiment to investigate the effect of α_{intra} and in general, the localization results improved with increasing correlation between the dipole components. In this experiment, we also compared using a scalar lead field versus a vector lead field for localizing sources with a fixed orientation. We found a divergence in the performance at 20 sources; the scalar lead field is able to recover 20 sources more reliably than the vector lead field (see Figure). This discrepancy can be explained by the fact that when using a scalar inverse model with Champagne, only one hyperparameter is learned per voxel (Γ_r). When using a two-component lead field with Champagne, it learns three hyperparameters per voxel. (Remember that for a two-component lead field, each block of Γ is a symmetric matrix (size 2x2) and for a scalar lead field, Γ is a diagonal matrix.) It should also be noted that if the dipolar sources are jiggling and a scalar lead field is used for the inverse problem, the localization results are quite poor. From these results, we can conclude that using a scalar lead field for source localization is favorable with Champagne if your sources have a fixed orientation, but if your sources have variation in their orientation, it is better to use a vector lead field to allow for flexible orientation. For most of the simulations using a vector lead field, we have imposed an intra-dipole correlation of 0.25 on the dipole components. This weak correlation between the lead field components makes the source localization problem more difficult and our results show good performance on difficult simulated data.

These results with MEG data do not necessarily imply that EEG recordings would be as successful. Due to the blurring of the electric fields, resolving two sources with EEG data is more difficult. Likewise, due to the smearing of the electrical field, source localization with EEG data might be more sensitive to errors in the leadfield than the MEG simulations we have shown in this chapter.

The results for localizing 20 sources with a fixed orientation could lead one to believe that there is a benefit to using a orientation-constrained or a scalar lead field for solving the inverse problem. To investigate this question, we wanted to looked at how the performance with the scalar lead field was affected by errors in the orientation. When using the vector lead field, there is not an imposed orientation for the sources, rather the algorithm essentially learns the orientation from the data. For the scalar lead field, an orientation is assumed, so using the same lead field for the forward and inverse step corresponds to being able to perfectly calculate the orientation of every source. We know that this would be near impossible, so in order to better interpret the results with the scalar lead field, it was important to test the effect of errors in the orientations. We perturbed the orientation of the lead field that was used to solve the inverse problem by different angles (maximum perturbations ranged from $\frac{\pi}{64}$ to $\frac{\pi}{4}$). We found that even small errors in orientation decreased performance with the scalar lead field, leading to the conclusion that the results for 20 sources we obtained with the scalar lead field are very much dependent on perfectly modeling the orientation of the sources.

We conducted an additional experiment to assess modeling errors in the lead field that entailed a shift of the voxel grid. A diagonal shift close to the maximum shift did not significantly affect the localization results. This finding supports the other evidence that Champagne is robust to errors in the lead field modeling. This experiment and the orientation-error experiment were designed to address a confound of using simulated data. When simulating data one must use a forward model to create the data, but it is often argued that using that same forward model to solve the inverse problem does not accurately replicate actual source localization performance. (It is highly unlikely that we could calculate perfectly the lead field.) This confound is often called ”the inverse crime.” We have tried to address this confound by testing Champagne in less than ideal circumstances.

Our experiments on deep source and clusters also demonstrate that Champagne improves upon existing methods of source localization when there are deep sources and distributed activity. Deep sources are notoriously hard to localize since they are far away from the MEG/EEG sensors, and have low SNIR. In our experiments, we can control the SNIR, so with the all deep sources condition, the reduced SNIR is not the driving factor for localization performance. Another mitigating factor in localizing deep sources is that the contribution to the sensors is fairly uniform, as deep sources are approximately equidistant from all the sensors, whereas sources in the cortex (shallow sources) have a more unique contribution to the sensors (they are close to some sensors and far from others). Champagne is able to provide better localization of deep sources both when there is a mix of deep and shallow sources and only deep sources. Champagne produces the largest improvement when localizing only deep sources at both noise levels as compared to the other algorithms. The contribution of shallow sources is much greater to the sensors than that of deeper sources.

The localization results for all algorithms improve with the cluster simulations as compared to the single source simulations. In some ways, the clusters are easier to localize because there is more signal in a particular location as opposed to when the sources are spread around. But, localizing the clusters is not the only aspect of performance that we investigated, we are also concerned with how well the extent of the cluster is reconstructed. Localizing the individual clusters and capturing the cluster extent is done most accurately by Champagne at both SNIR levels. MVAB does nearly as well as Champagne; the localization of the clusters is most likely aided by the increase in the signal at the cluster locations (as described above) and accurately modeling the extent of the sources is likely aided by the inherent blur in beamformer solution. MCE does well at localization, but does not accurately model the extent, due to its sparsity profile. Thus, it is notable that Champagne is able to both get the location and the extent of the clusters despite its sparsity profile. Champagne is sparse, but not so sparse that it cannot model a total of 50 to 150 active voxels at one time. (We also tried clusters of smaller and larger extent, for which the results were similar.)

This paper is the first evaluation of Champagne's performance on EEG data. Using a scalar lead field, we found that Champagne performed similarly on the EEG data with 128 sensors to the MEG results with 275 sensors. The results were not drastically affected by reducing the number of sensors by half, i.e. going down to only 64 sensors. Performance is still decent at 32 and 16 sensors as compared to the benchmark algorithms, leading to the conclusion that Champagne is an effective source localization algorithm for use with only 16 to 32 sensors of EEG data when a sufficiently small number of sources are expected. Using a vector lead field for EEG source reconstruction proved to be more challenging. The maximum number of sources at any number of sensor or SNIR that Champagne was able to localize was 3. However, as compared to the other source localization algorithms, Champagne does provide an improvement for EEG source localization.

The sparse solution to the inverse problem obtained from Champagne is well suited for functional connectivity analyses as the number of active voxels is significantly smaller than with other techniques commonly used, such as MVAB and sLORETA. We found that Champagne is better able to localize and reconstruct the time courses with highly correlated brain activity as compared to the benchmark algorithms. Champagne is able to reconstruct the brain connectivity, even in the presence of common source coupling. We have demonstrated that in certain situations, MVAR outperforms coherence and IC in uncovering interactions and lags in a network of brain areas in simulation. Although the time courses were accurately reconstructed in many case with Champagne, the mixing matrices from the MVAR model were not reconstructed with a similar fidelity. Uncovering the networks of brain areas involved in a task is a difficult problem, one that boils down to more than accurately reconstructing the voxel time courses. But, it can be concluded that without accurate time courses, it would be impossible to uncover networks, so our results with Champagne demonstrate that it is an effective source localization technique to use for functional connectivity analyses.

Using simulated data to assess a source localization algorithm's performance is necessary as one must know the true locations and time courses of the sources in order to definitively recognize a successful localization. But, while simulated data can be generated in a realistic fashion, it is still artificial data. Also, the way in which data is simulated can be biased towards a particular algorithm. The experiments in this paper use simulated data that has been generated in a manner that best replicates key aspects of real MEG data. We have also tried to generate data that does not favor Champagne by making the following choices. We have chosen realistic signal-to-noise ratios and number of time points to best model true experimental data. We use real brain noise instead of Gaussian noise; the noise models in both SEFA and Champagne assume Gaussian distributions. The source time courses are damped sinusoids, while the Champagne model assumes Gaussian time courses. Also, the inter- and intra-dipole correlations make the source localization problem more difficult, but more closely model the complexity of brain activity. We found that Champagne is able to localize on the order of 80 sources when the source time courses are Gaussian, the SNIR level is 100dB, and the inter-dipoles correlation is zero. In general, we have found that α_{inter}, α_{intra}, SNIR and the type of source time course has an effect on the localization results at higher number of sources, i.e. over 10 sources.

The experiments with real data highlight Champagne's source localization abilities. It is harder to evaluate localization accuracy with real data since the ground truth is not known. For this reason, we have chosen real data sets that have well-established patterns of brain activity; the SEF, AEF, audio-visual, and face-processing data. We are able to localize SEF activations with only 5 trials, which has implications for shortening data collection time. SEF is often added to an experiment as a functional localizer. If the amount of trial collected could be cut down to 5-10 as opposed to 252 that would have a impact on the time subjects are in the MEG machine.

For each of the seven AEF data sets presented, Champagne is able to localize bilateral auditory activity in the physiologically accepted region. The other source localization algorithms are not able to provide equally focal, bilateral activations. Historically, AEF is a difficult data set to obtain accurate activations due to the highly correlated sources. Variants of beamforming have been developed in order to handle the correlated sources, such as coherent suppression beamformers and dual-core beamformers (Diwakar et al., 2010; Dalal et al., 2006). Champagne is able to localize correlated activity without any modifications and it provides a robust solution to this long-standing deficit in source localization. The results from the more complicated data sets, the face processing and the audio-visual task, demonstrate that Champagne is able to localize many distinct, functionally-relevant brain activations. The benchmark algorithms are able to localize some of the sources that Champagne uncovers, but they cannot localize every source nor do they provide focal peaks, as is the case with MVAB and SL/dSPM. A major point of divergence between Champagne and the benchmark algorithms is the estimation of the time courses. In general, Champagne is able to produce time courses that tease apart brain areas, while the other algorithms tend to have very similar time courses at all locations in the brain. The results from the EEG data set shows that Champagne is an effective source localization method for EEG data, localizing activity in both visual areas and in the fusiform gyrus. The benchmark algorithms have better success at localizing the visual activity than the face-specific activations in fusiform gyrus with the EEG data set. Champagne's ability to localize epileptic activity from one spike demonstrates that it could be an effective algorithm to use pre-operatively to determine the focus of seizures.

Chapter 4

Sensitivity to the Estimation of the Background Noise for the Champagne Algorithm

4.1 Overview

The Champagne algorithm improves upon the localization capabilities of benchmark algorithms due to a few, key features, as demonstrated in Chapter Three of this thesis. First, it is robust to correlated sources and yields a sparse solution. Second, it integrates noise suppression in the source localization process. In the version of Champagne derived in Chapter Two and applied to real and simulated data in Chapter Three, the pre-stimulus period is used to learn the background activity. The background activity is thought to consist of independent sensor noise (electronic noise), external sources of noise, such as power lines and physiological noise, and non-stimulus evoked brain activity. Champagne learns the statistics of this baseline activity and then suppresses that activity in the post-stimulus period. An underlying assumption of the Champagne algorithm is that the stimulus-evoked activity, in the post-stimulus period, is added on top of the background activity in the pre-stimulus period.

An alternative to suppressing the noise in an integrated fashion is split to the localization into two steps. First, the stimulus-evoked factor analysis (SEFA), discussed in Chapter One, can be applied to the pre- and post-stimulus period to "clean" the post-stimulus period. Then second, the learning procedure used in Champagne can be used on the cleaned post-stimulus data. We discuss this variant of the Champagne algorithm and demonstrate its performance on real and simulated data in this chapter.

Both the Champagne algorithm formulated in Chapter Two and the variant described above assume that the activity in the post-stimulus period is added on top of background activity in the pre-stimulus period. If this assumption does not hold, then the subtraction of the pre-stimulus activity from the post-stimulus period could have disastrous effects on localizing the activity in the post-stimulus period. Non-stationarity in the noise between the pre- and post-stimulus periods would result in a situation where the additive assumption breaks down. Non-stationarity could arise from a visual or auditory prompt for a paradigm, or from artifacts from movement or eye blinks that are stimulus locked. The later is particularly relevant for tasks that require an verbal response; it is traditionally difficult to obtain clean pre-stimulus data in this case. If there is non-stationarity in the noise between the pre- and post-stimulus period, another method for estimating the noise covariance, without using the pre-stimulus data, would enable Champagne's use on these data sets.

In addition to data sets with non-stationarity in the background noise, estimating the post-stimulus activity without using the pre-stimulus period is relevant for data where there was not a pre-stimulus period collected. To save time in the MEG machine, some data sets do not collect enough samples before the prompt to accurately estimate the background activity. Another advantage of localizing activity without using the pre-stimulus data would be Champagne's application to time-frequency analysis. Typically, a ratio of the power, in a particular frequency band, between an active and control period is used as an indication of neuronal activity. Using the pre-stimulus period to estimate the noise in the post-stimulus period can only detect positive chances in power (a result of the additive assumption). If we can run Champagne on the pre- and post-stimulus periods to localize activity in both periods, we can calculate the increases and decreases in spectral power.

It is unclear that Champagne is able to localize activity without using the pre-stimulus period for the noise estimation, so before we can apply Champagne to data analysis techniques like time-frequency analysis, we must establish that it works without access to a pre-stimulus baseline. We have developed three variants of Champagne that use the same learning procedure for the source locations with the post-stimulus data, but differ in the way that the noise covariance is estimated (without access to the pre-stimulus data). These variants are described below and are applied to simulated and real data sets in order to demonstrate their effectiveness.

4.2 Denoising the Post-Stimulus Data

As discussed in Chapter Two, the noise covariance is estimated by running SEFA on the pre- and post-stimulus data. SEFA partitions the sensor data covariance into three components:

$$C_b \approx \Lambda + AA^T + DD^T, \tag{4.1}$$

where AA^T is learned from the pre-stimulus data and $\Lambda + DD^T$ are learned from the post-stimulus period data. We then set the noise covariance to be $\Sigma_\epsilon \to \Lambda + AA^T$. In the version of Champagne derived in Chapter Two, we discard the DD^T. (The computation of D in the post-stimulus period is done to estimate a more accurate Λ to account for nonstationarity in the noise.) Alternatively, we can use the D to clean the post-stimulus sensor data. We then replace, C_b, the empirical post-stimulus sensor covariance, with $\tilde{C}_b \rightarrow DD^T$ and estimate the source time courses from this "cleaned" sensor covariance. One advantage of localizing with the cleaned covariance is that the source time course ($s_r(t)$) are more smooth than those obtained with the full data covariance.

When we use this cleaned sensor covariance, however, we cannot use our estimate for the noise, $\Lambda + AA^T$, because we have already removed this activity from the post-stimulus sensor covariance. As such, we need a new way to estimate the noise covariance, Σ_ϵ. We have found that setting the noise to be, $\Sigma_\epsilon \rightarrow \sigma I$, where σ is a scalar that is empirically selected, is effective. We refer to this method as *Cleaned Data*. The estimate for the noise covariance, i.e. setting σ, is a crucial step. If σ is set too high or low, it will interfere with the source localization accuracy.

If we consider the source time course estimate for Champagne, with some modifications, we can see the relation between the value of sigma and the SNR. The source time course estimate for Champagne:

$$\bar{S}(t) = L^T \Sigma_\epsilon (L^T \Sigma_\epsilon L + \Gamma)^{-1} B_{post}(t) \tag{4.2}$$

We can replace Γ and Σ_ϵ with:

$$\Gamma = \mu I$$

$$\Sigma_\epsilon = \omega I$$

where μ and ω are constants.

When plugged into the source time course equation and the constants, μ and ω, rearranged, we obtain the following equation:

$$\bar{S}(t) \propto L^T (L^T \frac{\omega}{\mu} I L + I)^{-1} B_{post}(t)$$

The fraction, $\frac{\omega}{\mu}$, scales with the inverse of the SNIR. As such, $\sigma \propto \frac{\omega}{\mu}$, gives a more principled way to select σ. The σ regularizes the inverse in Equation 4.2 and other source localization algorithms, such as MNE, also use the inverse of the SNIR to set the regularization (Lin et al., 2006b).

If the SNIR is not known, it can be estimated by the eigenspectrum of the sensor covariance. We plot the eigenspectrum of the sensor covariance for the 10, 0 , -2, and -10dB simulated data sets (the full data) in Figure 4.1. The ratio of the average of the eigenvalues that explain the top 95% of the variance and the average of the bottom 5% gives an estimate for the SNIR.

4.2.1 Simulated Data

We generated simulated data as described in , using sinusoidal time courses and the real brain noise. We seeded three sources in the brain and they had an

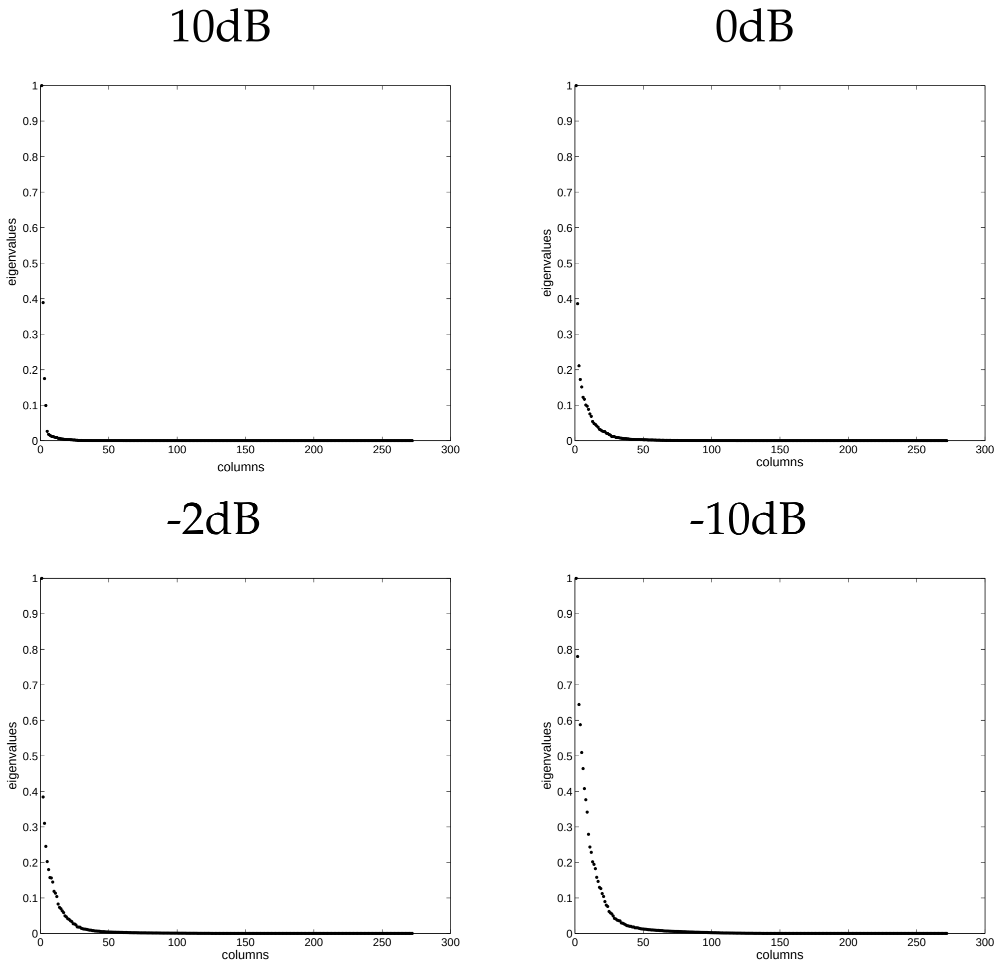

Figure 4.1 Eigenspectrum of simulated data at 10, 0, -2, and -10dB. These plots demonstrate that the SNIR of the data can be estimated from the eigenspectrum.

intra-dipole correlation of 0.75 and an inter-dipole correlation of 0.5. We used SNIR levels of 10, 0, -2, and -10dB. We ran the version of Champagne derived in Chapter Two and used in Chapter Three, which we refer to as *Full Data*, and the *Cleaned Data* version. For the Cleaned Data, we used $\sigma = 1$ for the 10 and 0dB data sets, but for -2 and -10dB we used $\sigma = 10$ and $\sigma = 100$, respectively. The comparison of the results for these two versions of the Champagne algorithm is found in Figure 4.2. The Full Data and Cleaned Data perform similarly for all the noise levels, and for 10, 0, and -2dB the activity all three sources are accurately localized, but for -10dB only the source on the left is localized.

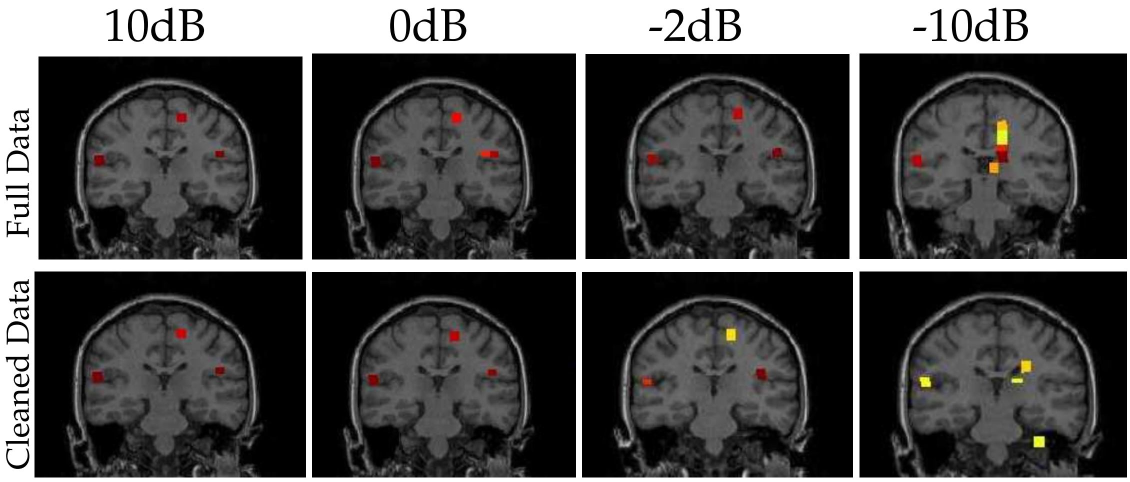

Figure 4.2 The *Full Data* and *Cleaned Data* versions of Champagne perform similarly for all the noise levels, and for 10, 0, and -2dB the activity all three sources are accurately localized, but for -10dB only the source on the left is localized.

4.2.2 Real Data

We have applied this method to the SEF, AV, and AEF data sets. The results are presented below in Figures 4.3, 4.4, and 4.5. In general, the *Cleaned Data* method provides similar source localization results to running Champagne on the full data covariance. One point of divergence is that the time courses are more smooth. The most substantial improvement is seen with the SEF data set. These data prove challenging for Champagne when the full data covariance is used (for trial-averaged data), see Figure in Chapter Three. The localization with this data set is greatly improved with running Champagne on the cleaned data covariance. This method while effective in localizing the SEF activation, does not estimate the noise covariance, Σ_ϵ, in a principled fashion and its effectiveness is somewhat sensitive to the choice of σ. For the SEF and AEF data we used $\sigma = 1$, but the AV data necessitated setting $\sigma = 0.01$.

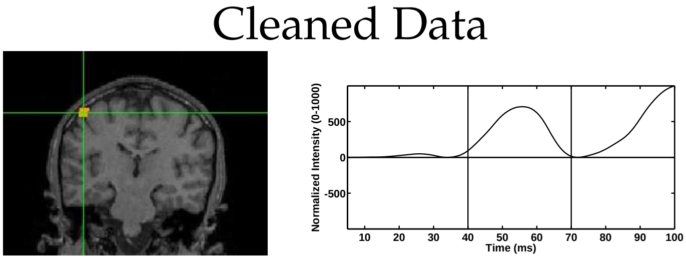

Figure 4.3 Champagne's results on the cleaned signal covariance for the SEF data. Champagne is able to accurately localize the contralateral somatosensory cortex when using the denoised signal.

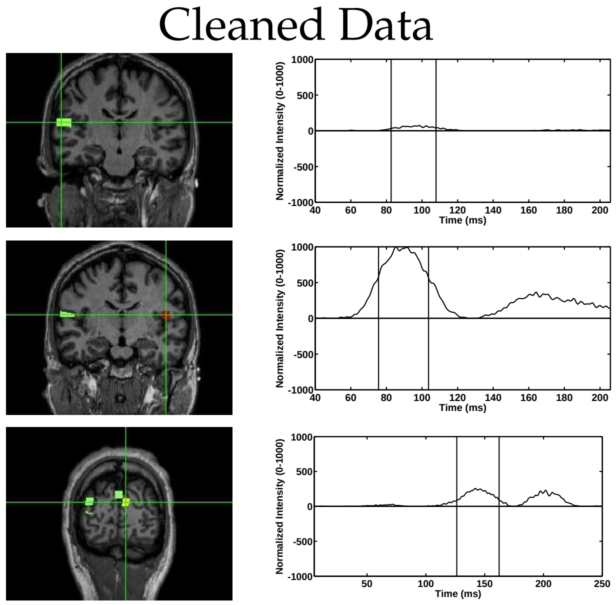
Cleaned Data

Figure 4.4 Champagne's results on the cleaned signal covariance for the AV data. Champagne is able to accurately localize auditory and visual activity when using the denoised signal.

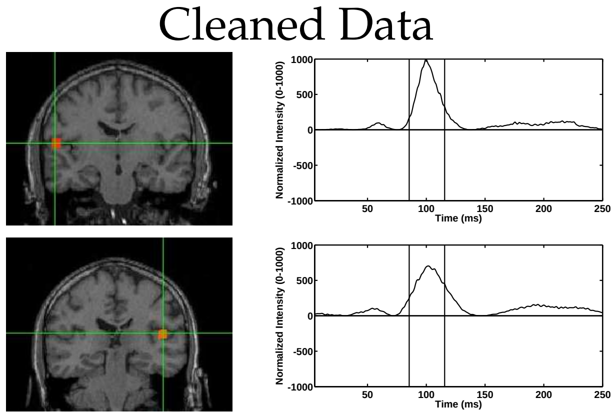
Cleaned Data

Figure 4.5 Champagne's results on the cleaned signal covariance for the AEF data. Champagne is able to accurately localize bilateral auditory activity when using the denoised signal.

4.3 Estimating the Noise Covariance without Pre-Stimulus Data

We have also developed methods to run Champagne on data sets without using the pre-stimulus period to estimate the noise covariance (or to clean the post-stimulus data). There are different options for estimating the noise covariance without pre-stimulus data; we describe three methods below and apply them to simulated and real data to test their efficacy.

First, we can use the same estimate for the noise covariance as in the previous section, $\Sigma_\epsilon \rightarrow \sigma I$, where σ is a scalar that is empirically selected. We use the full sensor covariance, however, not the cleaned covariance. We refer to this method as *Isotropic Noise*. This estimate for the noise does not have a true theoretical underpinning, but there is a relation between the setting of σ and the SNIR as discussed above.

Second, we can run Variational Bayesian Factor Analysis (VBFA) Nagarajan et al. (2007) on the post-stimulus sensor data. (VBFA is the first step of the SEFA algorithm, but SEFA runs VBFA on the pre-stimulus data.) VBFA partitions the sensor covariance into $\Lambda + AA^T$. We can then set the noise covariance to be $\Sigma_\epsilon \rightarrow \Lambda$ and then run Champagne on the post-stimulus data covariance. In some cases, we need to scale the noise covariance by a scalar, thus changing the equation to $\Sigma_\epsilon \rightarrow \sigma\Lambda$. This scaling is also dependent on the SNIR of the data, we find that with high SNIR it does not need to be scaled, but with low SNIR it needs to be multiplied by a $\sigma > 1$. We refer to this method as *Non-Isotropic, Diagonal Noise*.

The third method relies on the partitioning of the post-stimulus sensor covariance into a signal and a noise subspace with a singular value decomposition. The covariance can be modeled as:

$$C_b = E_S \Lambda_S E_S^T + E_N \Lambda_N E_N^T$$

where $\Lambda_S = [\lambda_1 \lambda_2 ... \lambda_Q]$ and $E_S = [e_1, e_2, ..., e_Q]$ and $\Lambda_S = [\lambda_Q + 1 \lambda_Q + 2 ... \lambda_M]$ and $E_N = [e_Q + 1, e_Q + 2, ..., e_M]$, λ_j and e_j are the jth eigenvalue and eigenvector, respectively, of C_b. We choose the Q largest eigenvalues (and the corresponding eigenvectors) as source components and the remaining $Q + 1$ through M eigenvalues as noise components. We choose the cutoff, Q, by plotting the eigenvalues The eigenspectrum) and selecting Q to be the point where the spectrum levels off near zero. Then, we can set $\Sigma_\epsilon \rightarrow E_N \Lambda_N E_N^T$. We refer to this method as *Non-Isotropic, Non-Diagonal Noise*. We could also make the noise estimate diagonal for this method, but we find that it performs similarly to the non-diagonal estimate.

4.3.1 Simulated Data

We used the same simulated data sets and SNIR levels in Section 4.2.1 to test the three variants of Champagne that do not require a pre-stimulus period, as described above.

The results from running the *Isotropic Noise* variant on the simulated data is shown in Figure 4.6. These results substantiate the analysis that the setting of σ should scale with the inverse of the SNIR. We chose $\sigma = 1$ for the 10 and 0 dB case. but $\sigma = 10$ yielded the best localization results for -2dB and $\sigma = 100$ for the -10dB data set. This variant of Champagne is able to localize all three sources at 10, 0, -2dB, but is only able to localize the source on the left at -10dB. These results are similar to those in Figure 4.2, demonstrating that as long as the σ scalar is set appropriately, localization is still accurate with an overly simplified estimate of the noise covariance.

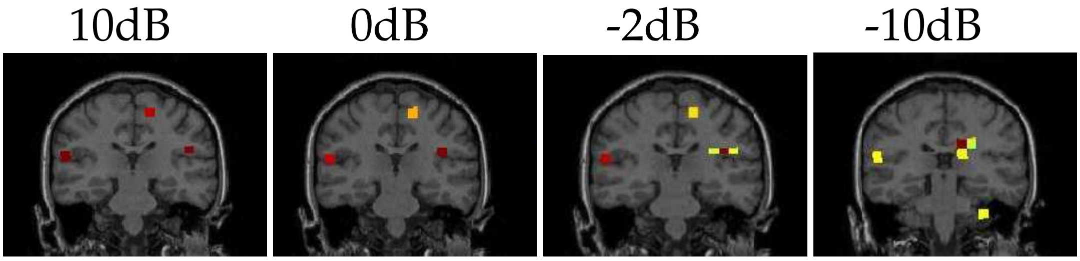

Figure 4.6 The *Isotropic Noise* variant of Champagne is able to localize all three sources at 10, 0, -2dB, but is only able to localize the source on the left at -10dB.

The results from running the *Non-Isotropic, Diagonal Noise* variant on the simulated data is shown in Figure 4.7. Here we used the independent noise term obtained from the post-stimulus period from running VBFA on the post-stimulus data; for the 10dB case we did not have to change the σ value, but for 0, -2, and -10dB we used $\sigma = 10$. This variant of Champagne is able to localize all three sources at 10, 0, -2dB, but is only able to localize the source on the left at -10dB. These results are similar to those in Figure 4.2 and Figure 4.6.

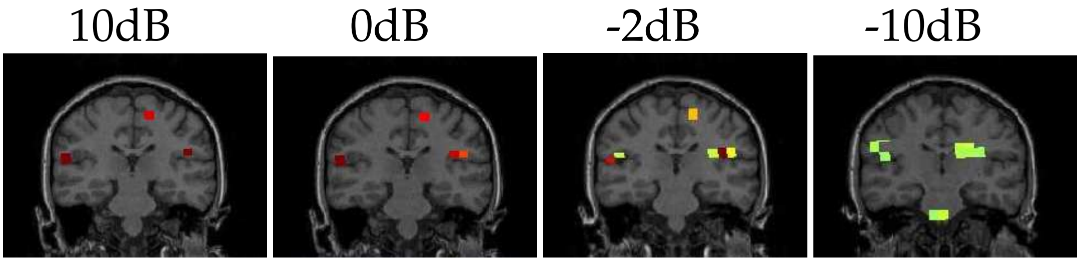

Figure 4.7 The *Non-Isotropic, Diagonal Noise* variant of Champagne is able to localize all three sources at 10, 0, -2dB, but is only able to localize the source on the left at -10dB.

The results from running the *Non-Isotropic, Non-Diagonal Noise* variant on the simulated data is shown in Figure 4.8. Here we used a different number of eigenvectors to partition the post-stimulus data into signal and noise subspaces. The 10dB data set was best localized when only one eigenvector was used for the signal subspace and the remaining components were used for the noise, whereas with the noisier datasets, more components were needed to partition the signal from the noise. The 0dB data set had the best localization with 4 components, the -2dB data set needed 5 components, and the -10dB data set was best with 10 components. We tried the localization for -2dB and the -10dB data set with various numbers of components and 5 and 10, respectively, worked the best. This variant of Champagne is able to localize all three sources at 10, 0, -2dB, but is not able to localize any of the three sources at -10dB.

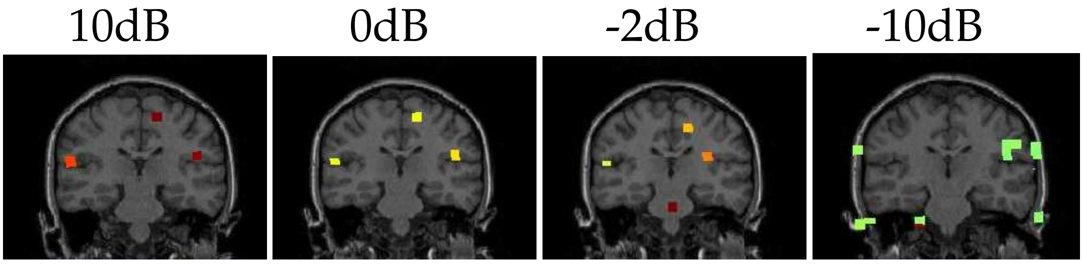

Figure 4.8 The *Non-Isotropic, Non-Diagonal Noise* variant of Champagne is able to localize all three sources at 10, 0, -2dB, but is not able to localize any of the three sources at -10dB.

4.3.2 Real Data

We applied these three methods to real data: the SEF, AV, and AEF data sets. The results from the SEF data set are presented in Figure 4.9. The localization with the *Isotropic Noise* method with $\sigma = 1$ is successful. Whereas, the *Non-Isotropic, Diagonal Noise* method with $\sigma = 1$ does not correctly localize the contralateral somatosensory cortex, while setting $\sigma = 100$ yields a successful localization. This demonstrates the importance of the σ parameter even when the noise is estimate using VBFA. The *Non-Isotropic, Non-Diagonal Noise* method using 5 components for the signal subspace does not accurately localize the somatosensory activity and looks like the results obtained from running Champagne on the full data covariance. When we reduced the signal subspace to being composed of only the first eigenvector, the localization of somatosensory cortex is accurate. The signal to noise ratio on this averaged data set is high, so only one component is needed.

All three methods work on the AV data set; they all provide similar localization results, but the time courses differ. The results from the two diagonal methods, *Isotropic Noise* and *Non-Isotropic Noise*, are shown in Figure 4.10 and the localizations and time courses are similar for these two methods. While the *Non-Isotropic, Non-Diagonal Noise* method does localize similar areas shown in Figure 4.11, the time courses are more smooth and cleaner than the previous methods. Also, the amplitudes of the auditory and visual sources are similar, as opposed to the other methods where there is disparity between the amplitudes of the activations.

Figure 4.12 shows the results from the three noise covariance estimation techniques applied to the AEF data set. They all localize bilateral auditory activity and the time courses are similar, but the time courses from the *Non-Isotropic, Diagonal Noise* method are not as smooth as those obtained from the *Non-Isotropic Noise* and *Non-Isotropic Noise* methods.

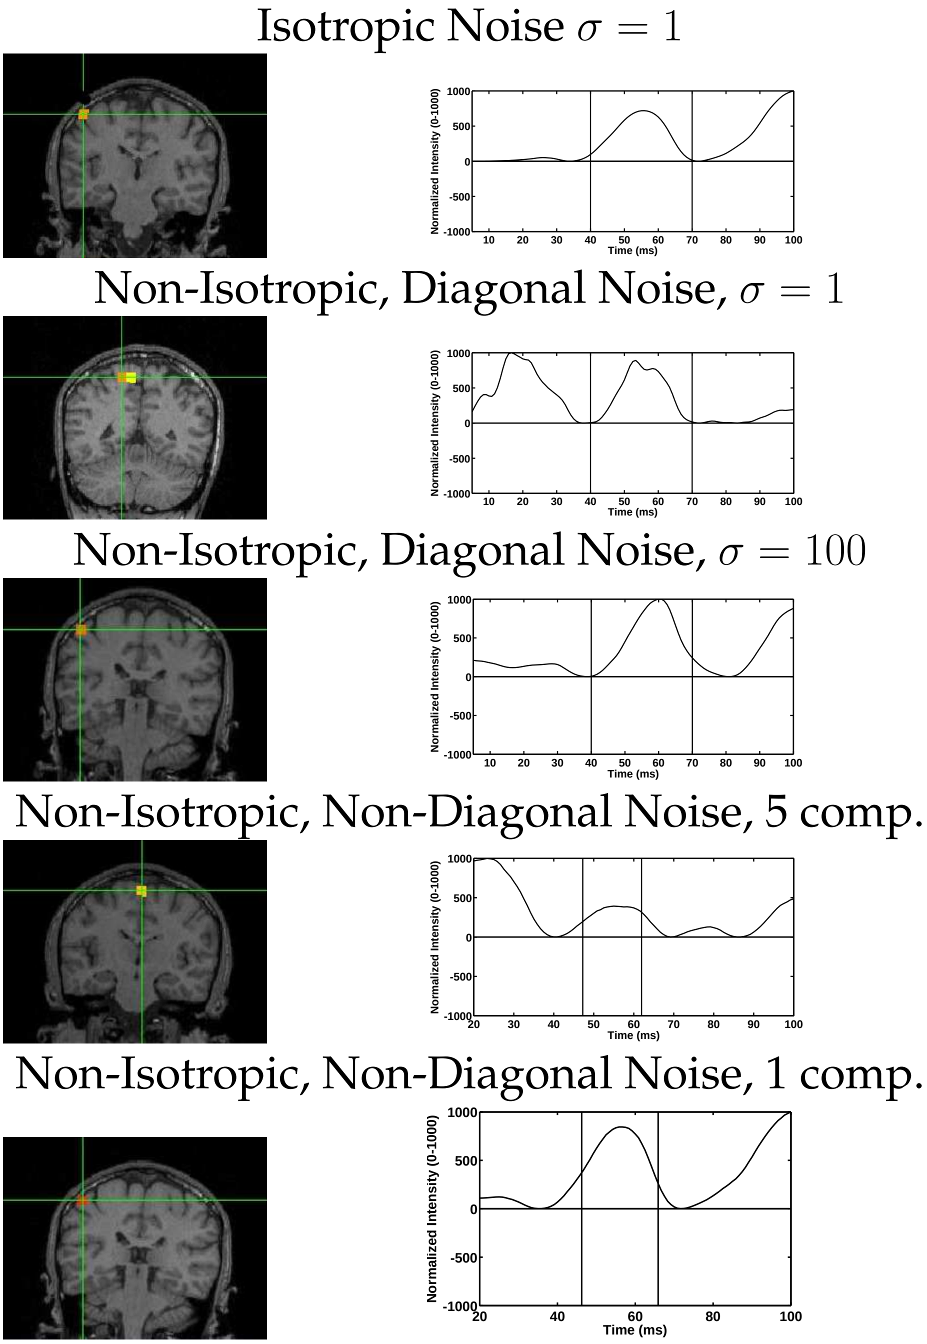
Isotropic Noise $\sigma = 1$

Non-Isotropic, Diagonal Noise, $\sigma = 1$

Non-Isotropic, Diagonal Noise, $\sigma = 100$

Non-Isotropic, Non-Diagonal Noise, 5 comp.

Non-Isotropic, Non-Diagonal Noise, 1 comp.

Figure 4.9 Results for Champagne run on the SEF data without access to the pre-stimulus period to learn the noise covariance. The results from using an isotropic noise estimate with $\sigma = 1$, a non-isotropic, diagonal noise estimate with $\sigma = 1$ and $\sigma = 100$, and a non-isotropic, non-diagonal noise estimate using 1 or 5 components for the signal subspace.

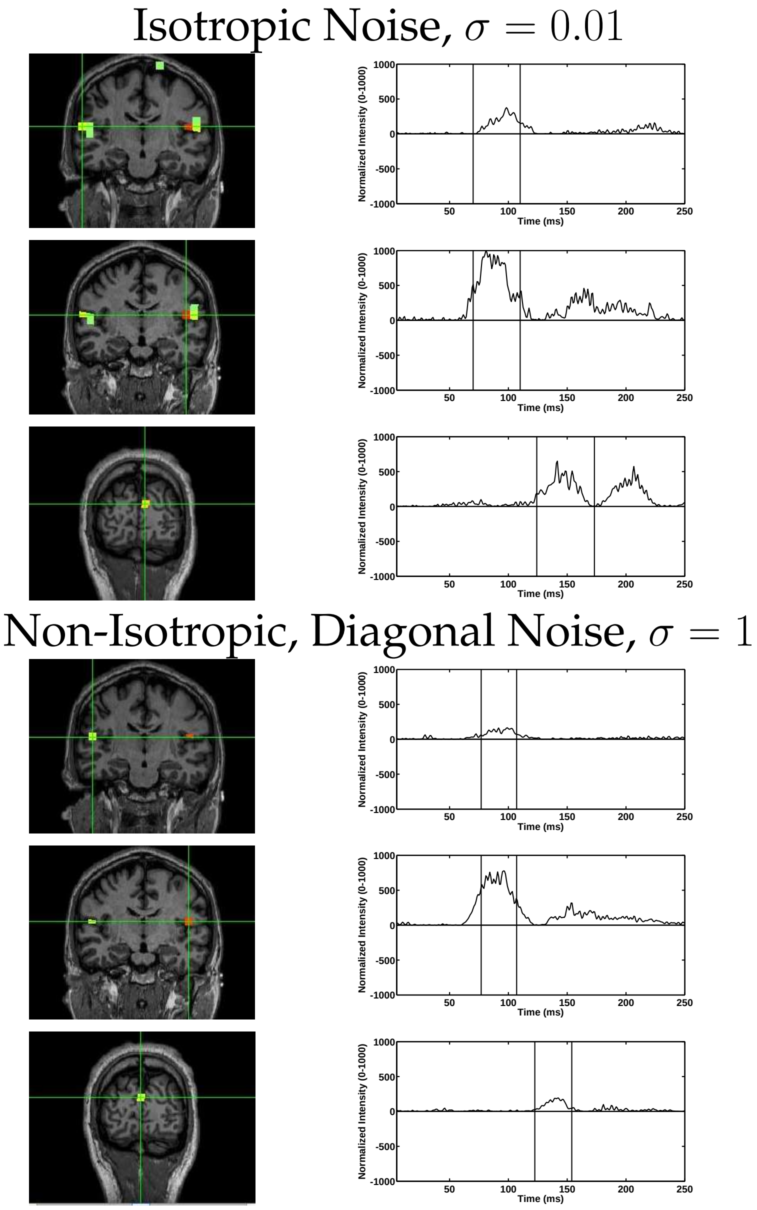

Figure 4.10 Results for Champagne run on the AV without access to the pre-stimulus period to learn the noise covariance from using the Isotropic Noise estimate with $\sigma = 0.01$ and the Non-Isotropic, Diagonal Noise estimate with $\sigma = 1$.

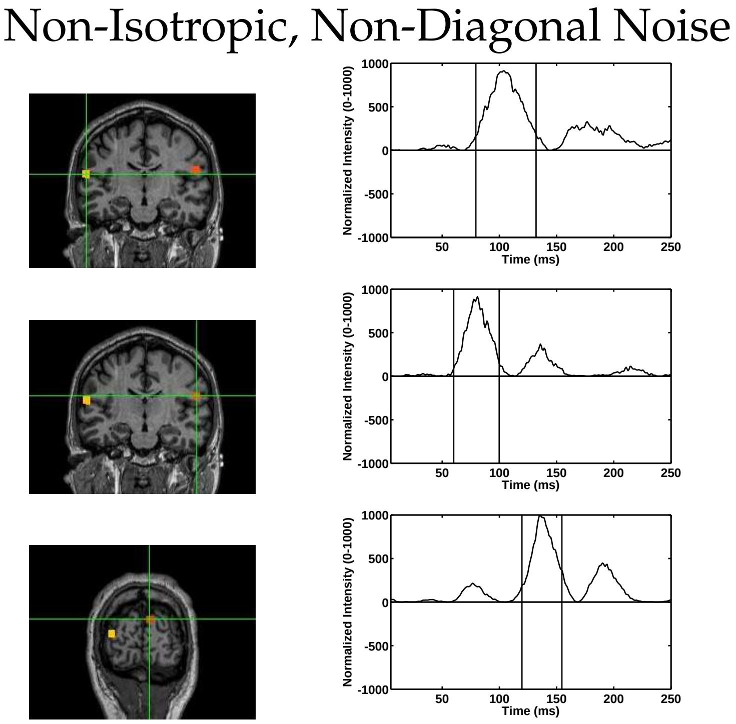

Figure 4.11 Results for Champagne run on the AV without access to the pre-stimulus period to learn the noise covariance from using the Non-isotropic, Non-diagonal Noise estimate with five components.

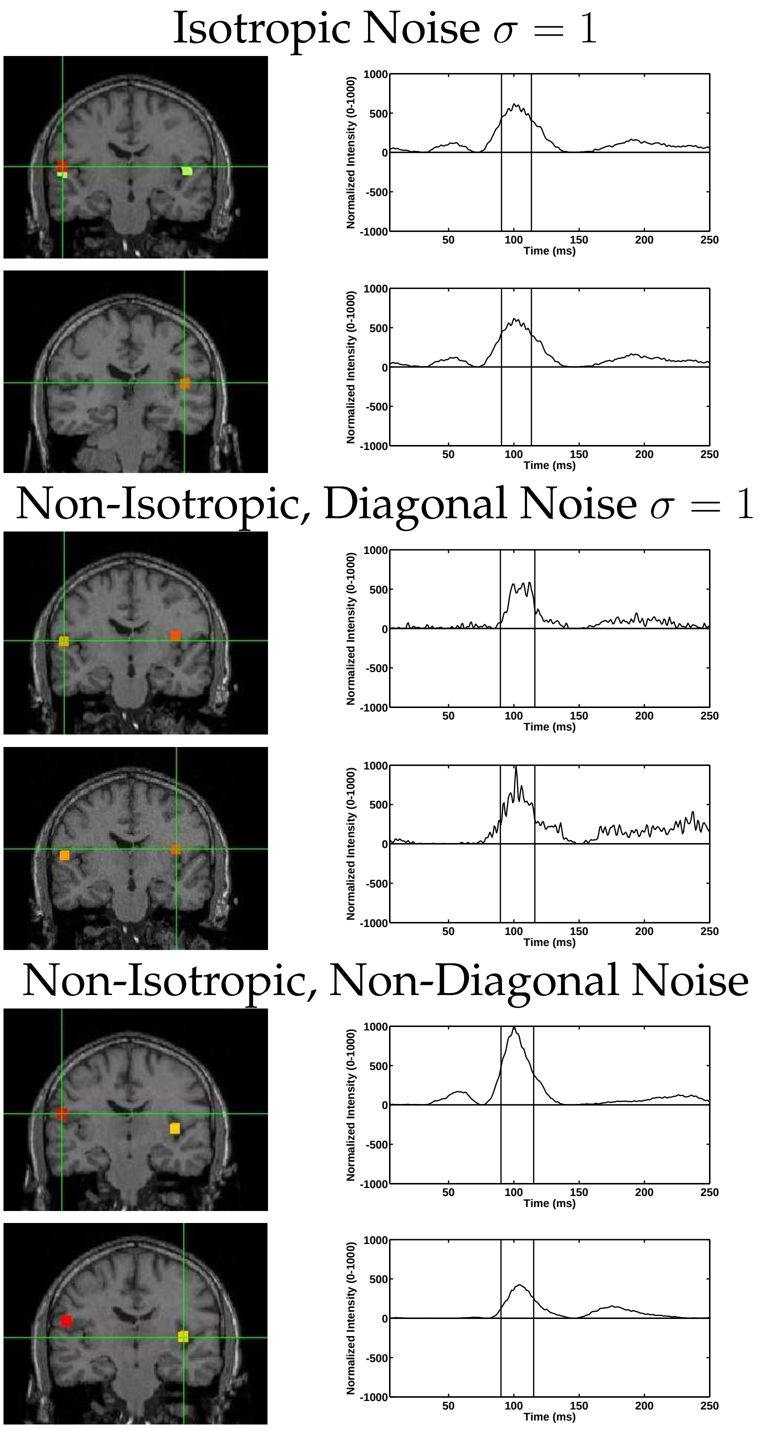

Isotropic Noise $\sigma = 1$

Non-Isotropic, Diagonal Noise $\sigma = 1$

Non-Isotropic, Non-Diagonal Noise

Figure 4.12 Results for Champagne run on the AEF data without access to the pre-stimulus period to learn the noise covariance. The results from using The Isotropic Noise estimate with $\sigma = 1$, the Non-Isotropic, Diagonal noise estimate with $\sigma = 1$, and the Non-Isotropic, Non-Diagonal Noise estimate with one component.

4.4 Discussion

The variants of the Champagne algorithm, which differ in the estimation of the noise covariance and in the case of the *Cleaned Data*, the post-stimulus data used, perform well on the simulated and real MEG data.

While we only tested on one simulated data set, we can infer from these results that the Champagne algorithm is robust beyond its ability to estimate the pre-stimulus noise covariance. Running Champagne on the "clean data" yielded similar results at all SNIR levels. Surprisingly, the *Isotropic Noise* estimate works as well as the original formulation and the *Cleaned Data* variant. While its performance does depend on the selection of σ, we have developed a principled way to set this scalar based on the SNIR of the post-stimulus data. Also, an inspection of Champagne's estimate from the source time courses yields some intuition for why this method works so well. The *Non-Isotropic Noise, Diagonal* variant also performs well on the simulated data, but does not perform better than the Isotropic Noise method at the low SNIR levels. And it also needs some scaling with a parameter σ, so there does not seem to be a clear benefit for this method over the simpler Isotropic Noise method. The *Non-Isotropic Noise, Non-Diagonal* method is also effective with the simulated data example, but like the previous two methods discussed, requires the setting of the number of eigenvectors to use for the signal versus the noise subspaces. We can see from these experiments that while all equally effective at localizing the three sources in this simulated data example, only the original Champagne method can localize activity without setting a parameter value. But, if there is not a baseline or pre-stimulus period, or it is too noisy, we can rely on these variants to localize activity with the Champagne algorithm.

The real data sets demonstrate the effectiveness of all the variants and emphasize the import of selecting the parameter, σ or the number of components to include in the signal subspace. In the case of the SEF data set, all the variants are able to localize the contralateral somatosensory cortex when the original Champagne algorithm was not able to localize this activity. For instance, when we reduced the signal subspace to being composed of only the first eigenvector, instead of the top five, for the Non-Isotropic, Non-Diagonal Noise variant, the localization of somatosensory cortex is accurate. The signal to noise ratio on this averaged data set is relatively high and could explain why it does not yield accurate results with the original formulation of Champagne. If the post-stimulus period is "too clean" then the estimate of the noise from the pre-stimulus period could throw off the localization in the post-stimulus period. Since the variants on the Champagne algorithm do not use the noise covariance learned from the pre-stimulus period directly, they are immune to the misleading effects it has on the localization. The other data sets, AV and AEF, demonstrate that all the variants are successful at localizing activity without learning the noise covariance from the pre-stimulus period.

In general, these variants open new possibilities for the type of data on which Champagne can be used, such as when there is not a pre-stimulus period or where the pre-stimulus period has a large artifact. These experiments also imply that Champagne could be used to assess power changes in the source time courses as compared to the pre-stimulus period, so-called time-frequency analysis, which would expand the types of neuroscience questions Champagne could help to answer.

Chapter 5

Statistical Thresholding with Permutation Testing

5.1 Overview

The statistical analyses of the source images from M/EEG imaging typically proceed with voxel-level statistics. At each voxel, a statistical test is used to determine if an effect of interest is present. Traditionally, a parametric statistical test is used, some common methods include t-tests, F-tests, paired t-tests, ANOVA, correlation, linear regression, multiple regression, and ANCOVA. All of these tests assume that the data are normally distributed with a mean parameterized by a general linear model. The estimated parameters for the specific test are then contrasted to produce a t-statistic, which then can be assessed for statistical significance (i.e. evidence against the null hypothesis) given an assumed distribution. This process yields p-values, which is the probability that the value would exceed that observed under the null hypothesis. The test *level*, usually denoted by α, is the accepted risk of the test or the probability of committing a Type I error. A Type I error reflects falsely rejecting the null hypothesis, i.e. incorrectly finding significance. Any p-values less than α lead to a rejection of the null hypothesis.

In the case of M/EEG brain imaging, there are typically 5,000 to 15,000 voxels, which results in that many statistical tests. Therefore, the risk for committing Type I error is high. If the *level* is $\alpha = 5\%$ and there are 10,000 voxels, there could be 500 voxels that have Type I errors, which could greatly impact the interpretability of the results. As such, there are methods to correct for Type I errors when there are multiple comparisons. The Bonferroni correction (Bonferoni, 1935) and false discovery rate (FDR), both implemented in a step-up (Benjamini and Hochberg, 1995) and a step-down procedure (Benjamini and Liu, 1999), are commonly-used methods to correct for Type I errors. These corrections have different sensitivities, but tend to be overly conservative and can remove a result or trend that is present in the uncorrected statistical values.

While conventional, the statistical testing described has two main issues. The first is the assumption of normally distributed random variables. Under the central limit theorem, this assumption would hold, but in brain imaging the number of samples does not approach those needed to assure Gaussianity. Secondly, correcting for multiple comparisons with a method such as, the Bonferroni correction or FDR, is often overly stringent. Nonparametric statistics provide a solution to at least the first of these issues and can be formalized in a way to address the second. Nonparametric statistics do not assume a distribution for the variable. Rather, they use the data to obtain a null distribution over the data, from which, significance can be tested. Permutation tests are one type of nonparametric test; they have been increasingly used as computational power betters. However, the statistical values obtained with permutation testing are not corrected for multiple comparisons as there are still many statistical tests being performed in parallel. The same methods that correct for multiple comparisons with parametric statistics can be applied to the p-values obtained from nonparametric statistics, but applying these correction techniques to nonparametric statistics also can be overly stringent. The maximal statistic approach to nonparametric statistics automatically corrects for the multiple comparisons problem.

Permutation methods have been applied extensively to neuroimaging data to find nonparametric statistical thresholds (Nichols and Holmes, 2001; Singh et al., 2003; Sekihara et al., 2005; Pantazis et al., 2005; Chau et al., 2004; Dalal et al., 2008). As discussed in Chapter One, the implementation of permutation testing differs across these references. In (Pantazis et al., 2005), the authors present an approach that creates surrogate data sets by randomly exchanging the pre- and post-stimulus periods. This way of generating surrogates differs from the conventional approach, i.e. exchanging the labels (usually condition or group labels) to test for a significance difference between two groups or conditions. Exchanging time windows of an individual data set opens the opportunity to test significance for a single subject's results. By collecting multiple trials of data, the mean power across trials (or time) can be tested for difference from zero using a test statistic, a pseudo t-value and the maximal statistic method corrects for multiple comparisons.

In this chapter, we apply this method described in (Pantazis et al., 2005) to real MEG data and we develop two variations on this method, which have been tailored to use on source localization results obtained from sparse algorithms. Sparse algorithms yield brain images that are already thresholded to a certain degree, but they can still contain some spurious peaks. The distribution of source power values obtained from a sparse algorithms, such as Champagne or MCE, is quite different than the distribution obtained from other algorithms, such as MVAB and SL. Sparse algorithms have only a handful of non-zero voxels, resulting in a distribution of power values that has a prominent peak at zero and a tail that drops off quickly. It follows that permutation methods used for non-sparse algorithms might not be suitable for thresholding sparse algorithms.

A challenge of statistical thresholding is an appropriate means of arriving at a threshold for a single subject. Often statistics use an average across subjects, but in functional brain imaging, this means the MRIs of the individual subjects must be spatially normalized. Spatial normalization, while a widely-used technique, is not always ideal because it relies on the automatic segmentation of the MRI. If there is an error in the segmentation, the averaged results can be corrupted. As such, it is often preferable to look at the results on an individual subject level, but arriving at principled threshold for a single subject's data is not a prescribed science.

Source localization techniques, such as SL and MVAB, often yield diffuse activations and the distribution of power values is much more continuous than the distribution obtained from a sparse algorithm. These diffuse activations can be hard to interpret as they cover many neighboring brain areas. If these results could be thresholded in a way that only very focal activations survive, the results obtained from these algorithms would be more interpretable. Typically, this is achieved by arbitrarily setting the threshold to some percentage of the maximum, such as 90%, but this threshold is heuristic and must be tuned to every data set. Alternatively, a statistical threshold could be obtained to stringently threshold the results. A nonparametric statistical approach to this issue would yield a threshold that is corrected for multiple comparisons (through the maximal statistic) and could be stringent enough to achieve a sparse-like image.

The work in this chapter seeks to answer three questions. First, can traditional nonparametric statistical thresholding methods be applied to the inverse solution obtained from sparse algorithms? Second, can nonparametric statistical thresholding reject spurious peaks obtained by sparse algorithms? And third, can brain images obtained from non-sparse algorithms resemble the sparse maps through stringent thresholding?

5.2 Source Localization with Unaveraged Data

We performed source localization on the unaveraged sensor data $B(t)$ by choosing a time window of interest in the pre-stimulus and post-stimulus periods from N trials. Then, we concatenated the pre-stimulus windows and the post-stimulus windows to form long pre- and post-stimulus periods. We compared Champagne's performance on unaveraged data to the performance of three benchmark algorithms: sLORETA (SL), minimum variance adaptive beamforming (MVAB), and minimum current estimate (MCE) performance. These algorithms are formalized in Section 3.2. The source localization algorithms were run on these concatenated pre- and post-stimulus periods. (We did not run dSPM for these experiments as our previous results demonstrate that SL and dSPM yield the same solution to the inverse problem.) All the source localization methods generate a spatial filter w such that:

$$s_r(t) = w_r B(t). \tag{5.1}$$

where r is the voxel index and t are the time points in the post-stimulus period.

The source time courses were averaged across trials to generate a time course $\bar{s}_r(t)$, for every voxel and then, we calculated the power P in a given time window ($t_2 \geq t \geq t_1$) across voxels:

$$P_r = \frac{1}{T} \sum_{t=t_1}^{t_2} \bar{s}_r(t)^2 \tag{5.2}$$

where T is the number of time points in the window.

5.3 Statistical Thresholding

Based on the method presented in (Pantazis et al., 2005), we can calculated a pseudo t-value t using the power in the pre- and post-stimulus periods:

$$t_r = \frac{P_r^{post} - P_r^{pre}}{\sigma_{\tilde{r}}^{pre}} \tag{5.3}$$

$$\tag{5.4}$$

where P_r^{post} and P_r^{pre} are the pre- and post-stimulus power averaged over trials and σ^{pre} is the standard deviation of the power in the pre-stimulus period, calculated by applying the weights to the pre-stimulus data and then calculating the power as with the post-stimulus power. Instead of using only the standard deviation from one voxel r, we pooled the standard deviation from 20 neighboring voxels $\tilde{r}$ using a Gaussian-weighted average.

The pseudo t-value can be distinguished from the t-value in that it pools the standard deviation across neighboring voxels. The pseudo t-statistic is used when there is a low degree of freedom and the estimate of the standard deviation is less reliable. It pools the standard deviation of neighboring voxels in order to increase the degrees of freedom; in effect it smoothes the standard deviation. Deviating from the approach in (Pantazis et al., 2005), an alternative to using the pseudo t-statistic is to use the post-stimulus power, P_r^{post}, as the statistic. Since Champagne and MCE are sparse, the estimate for the pre-stimulus standard deviation is zero or near zero in most voxels. This causes spurious peaks in the pseudo t-statistic values as the value statistic is large when the standard deviation is almost zero.

In order to obtain a statistical threshold for the true or original pseudo t-values or the post-stimulus power, generically referred to as Ψ^O, we used resampling methods and the maximal statistic to obtain the null distribution for the data. This method yields a threshold that is not subject to the assumptions of parametric statistics and is corrected for multiple comparisons. We developed two methods to create M surrogate data sets, B^m, where m is the permutation number. The first method consisted of exchanging the pre- and post-stimulus periods of a subset of the N trials (as done in (Pantazis et al., 2005)). There are 2^N possible surrogate data sets possible. The second method resampled the pre-stimulus data by randomly drawing N trials from the total trials available (greater than N). If there are approximately 100 trials total and $N = 30$, then there will be $\binom{100}{30}$ possible surrogate data sets. Both ways to create surrogate data sets yields millions of possible permutations, as such, we chose to subsample the surrogates by randomly creating 1000 surrogate data sets. To ensure normalization between the surrogates and the original data, we normalized the power of each surrogate to the power of the original sensor data.

We then applied the spatial filter weights obtained from the source localization procedure described in the previous section to each surrogate data set to obtain source time courses, which were averaged across trials to generate a trial-averaged time course for every voxel. For each permutation, we can then calculated either the pseudo t-value t_r^m or the power, P_r^m in the time window $(t_2 \geq t \geq t_1)$ across voxels, generically referred to as Ψ_r^m.

To employ the maximal statistic correction for both methods, we then took the maximum across voxels Ψ_r^m from each permutation:

$$\Psi^{max^m} = \max_r(\Psi_r^m), \tag{5.5}$$

and use the Ψ^{max^m} to estimate the null distribution of Ψ^O, referred to as $\hat{F}_{\Psi^O}$. Given a significance level of α, a statistical threshold, θ_α^{max}, can be set as the $c+1$ largest member of Ψ^{max^m}, where $c = \alpha M$ and c is rounded down if not an integer. Ψ^O can be thresholded by θ_α^{max}.

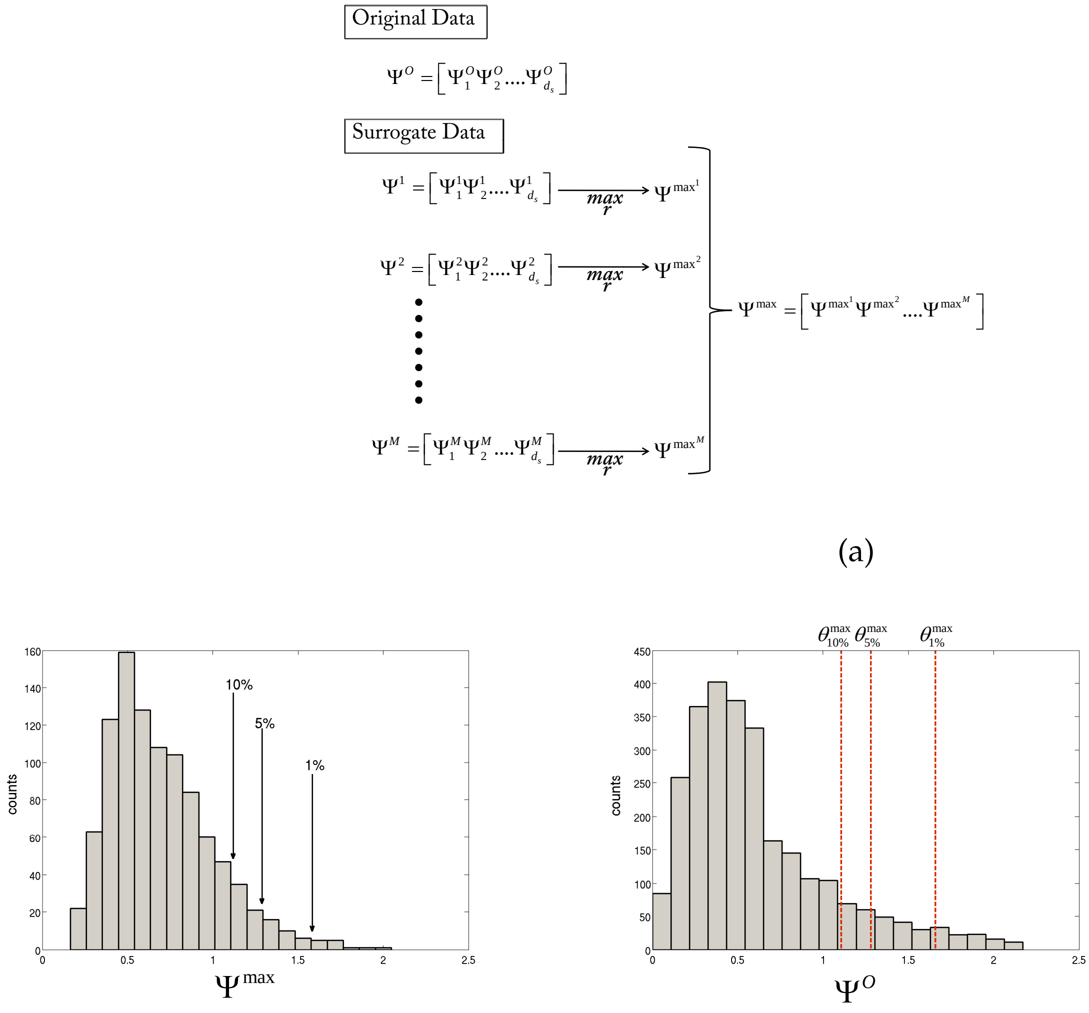

(a)

(b)

(c)

Figure 5.1 Cartoon illustrating the thresholding procedure. (a) The test statistic is calculated for every voxel for the unpermuted (original data), Ψ_r^O. Then, for each permutation of the data, Ψ_r^m, is computed. Finally, the maximum over r is taken to obtain Ψ^{max^m}. (b) A histogram of the maximal distribution, Ψ^{max^m}, with arrows pointing to the 1^{st}, 5^{th}, and 10^{th} percentiles, corresponding to $\alpha = 1\%$, $\alpha = 5\%$, and $\alpha = 10\%$ respectively. (c) A histogram of the original statistic, Ψ_r^O, with the $\theta_{1\%}^{max}$, $\theta_{5\%}^{max}$, and $\theta_{10\%}^{max}$, corresponding to the values obtained in (b).

We have developed three methods for performing the permutation testing, which differ in the way the surrogate data sets are generated and the test statistic used. Method 1 uses switching of random pre- and post-stimulus period to generate the surrogates and the pseudo t-value for the test statistic (Ψ). Method 2 uses switching of random pre- and post-stimulus period to generate the surrogates and the post-stimulus power for the test statistic (Ψ). Method 3 uses resampling of the pre-stimulus period to generate the surrogates and the post-stimulus power for the test statistic (Ψ). See Figure 5.2 for a table describing these methods.

Statistic

Surrogate Data	T-value	Post-Stim. Power
Pre-Post Switch	**Method 1**	**Method 2**
Resampled Pre-Stim.	----------	**Method 3**

Figure 5.2 Methods for thresholding

Alternatively, p-values can be calculated using the maximal statistics, which are corrected by multiple comparisons. For each Ψ^O, we counted the number of values that are greater in Ψ^{max^m} and divide by M. Then, we can use a threshold α to determine areas with significant activity.

A less conservative approach than the maximal statistic is to save more than just the maximum statistical value from every permutation. The maximal statistic can be driven by outliers; if there is one errant voxel (with high power) in each permutation, the threshold obtain for Ψ^O could be overly conservative. We propose saving the top n^{th} percentile of the statistic values from each permutation and using these values to estimate the null distribution of Ψ^O. Just as with the maximal statistic, we can then obtain $\theta_\alpha^{n\%}$ by taking the $c+1$ largest member of the distribution, where $c = \alpha n M$ and c is rounded down if not an integer. Then, Ψ^O can be thresholded by $\theta_\alpha^{n\%}$. We compared the threshold derived from the maximal statistics to using the top 1^{st} and 5^{th} percentiles.

5.4 Performance on Simulated Data

We first demonstrate Champagne's ability to localize sources with unaveraged data on simulated MEG data. We also applied the statistical thresholding procedure, Method 3, to the results obtained from the source localization for Champagne and the comparison algorithms, MCE, MVAB, and SL. We used the three source configuration from the previous chapter. We simulated 30 trials of data at two SNIR levels, 0dB and -5dB. The lower SNIR levels reflect that we are using unaveraged data, so the individual trials would have a lower SNIR than the "averaged" data used previously. We thresholded the results at $\alpha = 1\%$, $\alpha = 5\%$ and $\alpha = 10\%$ using the maximal statistic procedure.

The results from the simulations for all algorithms are presented in Figures 5.3 and 5.4. Champagne is able to localize the three sources at both SNIR levels with the unaveraged simulated data. The thresholding procedure at all confidence levels,$\alpha = 1\%$, $\alpha = 5\%$ and $\alpha = 10\%$, retains the three sources while thresholding out the non-zero activity surrounding the sources. The results from MCE are similar to those obtained with Champagne and the thresholding results show that the three sources survive the thresholding at all confidence levels. The source localization results with MVAB at both SNIR levels show that there are peaks in the correct locations but the results are quite blurred. The thresholding does not clean up the localization results with MVAB. At 0dB, SL is able to localize sources in the correct locations and the statistical thresholding cleans up the results such that distinct peaks can be seen. At -5dB, two of the three sources are localized, but the statistical thresholding does not clean up the image, even at $\alpha = 1\%$.

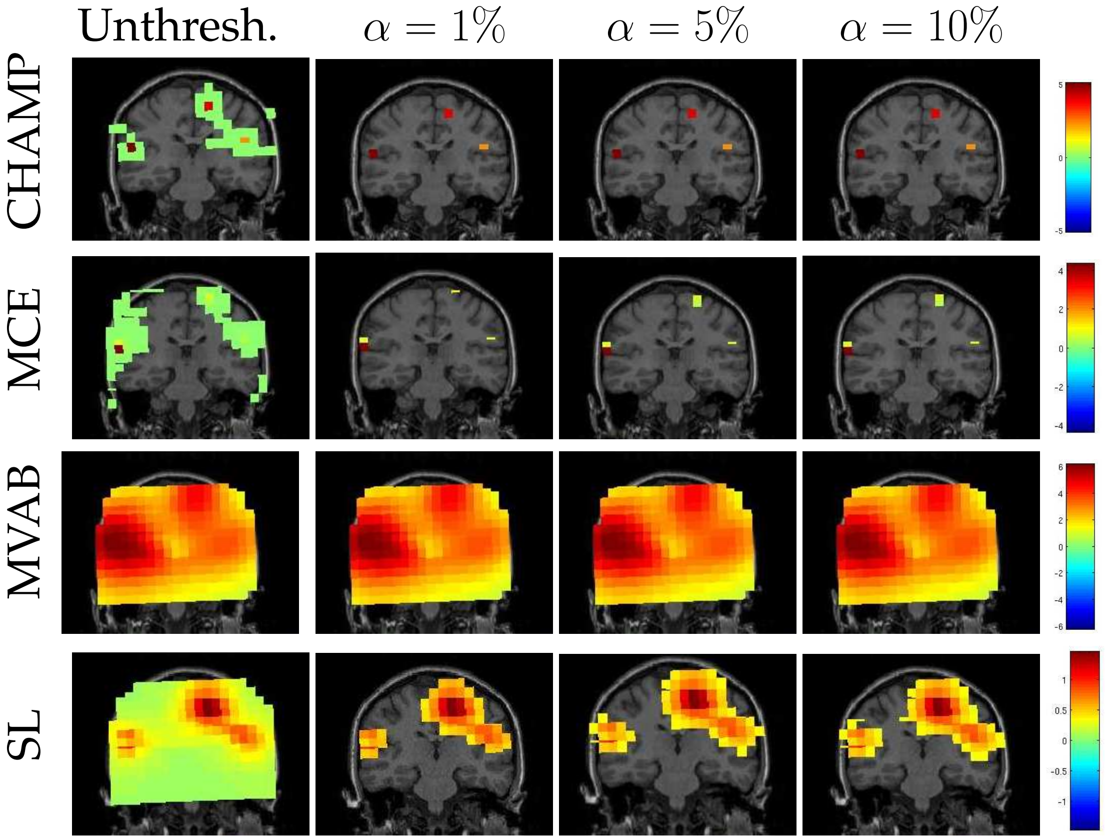

Figure 5.3 Simulated data example (0dB) with surrogates generated with resampling pre-stimulus data and computing post-stimulus power, Method 3. The unthresholded power is shown in the far left column (coronal slice). The power is thresholded with the maximal statistic at $\alpha = 1\%$, $\alpha = 5\%$ and $\alpha = 10\%$. Champagne, MCE, and SL benefit from the statistical thresholding, while MVAB does not.

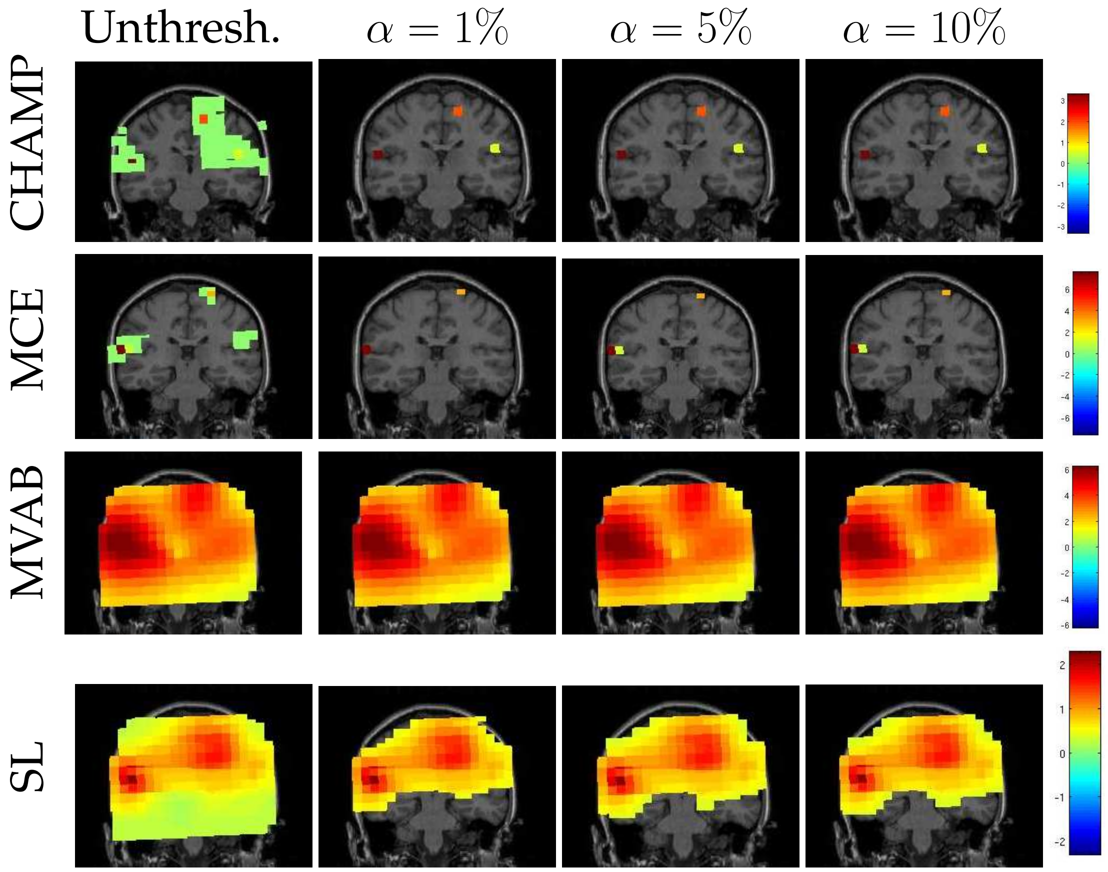

Figure 5.4 Simulated data example (-5dB) with surrogates generated with resampling pre-stimulus data and computing post-stimulus power, Method 3. The unthresholded power is shown in the far left column (coronal slice). The power is thresholded with the maximal statistic at $\alpha = 1\%$, $\alpha = 5\%$ and $\alpha = 10\%$. Champagne, MCE, and SL benefit from the statistical thresholding, while MVAB does not.

5.5 Performance on Real Data

We present the localization results with unaveraged data for four data sets, audio-visual (AV), somatosensory evoked field (SEF), face processing, and auditory evoked field (AEF) data sets. We ran Champagne on these data sets and compared its performance to MCE, SL, and MVAB. For all the overlays on the MRI presented here, we show the coronal (and axial) section that intersects the maximum voxel for the time window being investigated.

5.5.1 Audio-Visual Task

We analyzed a data set designed to examine the integration of auditory and visual information (the same audio-visual data set used in 3.6.2). We presented single 35ms duration tones (1 kHz) simultaneous with a visual stimulus. The visual stimulus consisted of a white cross at the center of a black monitor screen. The pre-stimulus period was the window from -100ms to -5ms and the post-stimulus window was taken to be 5ms to 250ms, where 0ms is the onset of the simultaneous tone and visual stimulus. We concatenated the pre-stimulus and post-stimulus periods for 30 trials, as described above. Then we computed the power in two windows, from 80ms to 140ms to capture the auditory activation and 100ms to 180ms to capture the visual activation. For this data set, we applied Method 1, Method 2, and Method 3 to obtain statistical thresholds for the results obtained from 80ms to 140ms with Champagne. We present the results from Method 3 below (the results from Method 1 and 2 can be found in the appendix to this chapter). Then, we applied only Method 3 to threshold the power map obtained from 80ms to 140ms with the other algorithms and from 100ms to 180ms with all algorithms.

We applied Method 3 to the 80ms to 140ms window, results shown in Figure 5.5, and compared the thresholding results across algorithms. For Champagne, the unthresholded post-stimulus power values are shown in the first column. We found that the thresholds obtained at $\alpha = 1\%$, $\alpha = 5\%$, and $\alpha = 10\%$ were sufficient to clean up the post-stimulus power maps, but not overly stringent, as both auditory cortical sources remained after threholding at all levels. We applied Method 3 to the benchmark algorithms, SL, and MCE, also shown in Figure 5.5. For MCE, the localization results show bilateral activity (the left source is dorsal to auditory cortex) and we found the amount of thresholding to be similar to Champagne's results. SL is able to localize bilateral activity that is diffuse. The thresholding at $\alpha = 1\%$ allows for distinguishing the left and right activations, while thresholding at $\alpha = 5\%$ and $\alpha = 10\%$ does create separation between the auditory activations. The localization for MVAB was unsuccessful with the unaveraged data, as demonstrated in Figure 5.9 and therefore, we did not perform the thresholding procedure for MVAB. The p-values obtained from the maximal statistic distribution (as described above in 5.3) are shown in Figure 5.6. We plot $1-p$ and threshold at 0.95, corresponding to $\alpha = 5\%$. The p-values for Champagne and MCE show similar results to those in Figure 5.5 at all levels and provide focal sources in the two auditory cortices, while the results with SL are diffuse and do not threshold the power values adequately.

For the window around the visual activation, 100ms to 180ms, we only applied Method 3 as this method strikes the best balance between stringency and allowing more than just the maximal voxel to pass the test. The results for the visual localization with Champagne, MCE, and SL are presented in Figure 5.7, with the unthresholded post-stimulus power values in the first column. Champagne is able to localize visual activity in this time window; thresholding at $\alpha = 1\%$ and $\alpha = 5\%$ allows activation in one visual area to pass, while thresholding $\alpha = 10\%$ allows a second visual activation (more medial) to pass. The unthresholded results from MCE show that there are activations in auditory areas (as in the 80ms to 140ms time window) and there is activation in the visual cortex, but these voxels do not have the maximum power in the time window (different from Champagne). Only the right auditory source passes to significance at all levels and the visual activations are thresholded out. SL also shows both visual and auditory activations in the unthresholded maps. At $\alpha = 1\%$ and $\alpha = 5\%$, the visual activation is distinguished from the auditory activation and at $\alpha = 10\%$ the threshold leaves the image blurred. The p-values computed from the maximal statistic distribution are shown in Figure 5.8. We plot $1 - p$ and threshold at 0.95, corresponding to $\alpha = 5\%$. Champagne has significant p-values in one visual area, while MCE has significant p-values in the same location as the auditory window results, Figure 5.6, and SL has a diffuse cluster of significant voxels showing that the p-values do not allow for more stringent thresholding with SL.

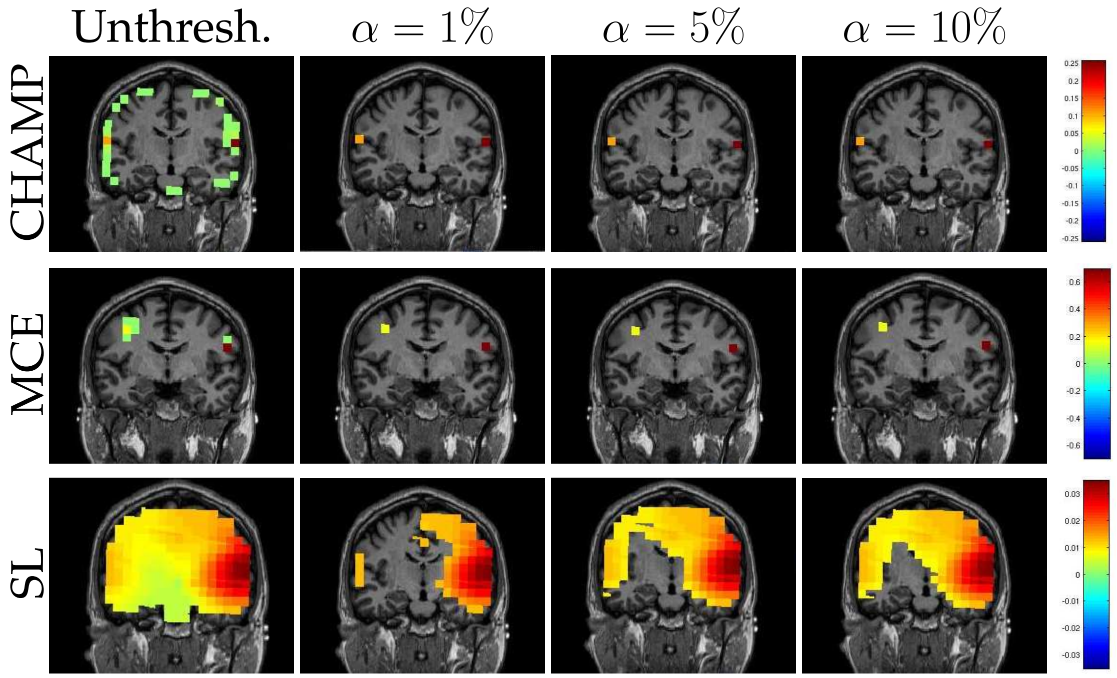

Figure 5.5 AV data with surrogates generated with resampling pre-stimulus data and computing post-stimulus power, Method 3. The unthresholded post-stimulus power values in the window from 80ms to 140ms are shown in the first column (coronal slice). The power is thresholded with the maximal statistic at $\alpha = 1\%$, $\alpha = 5\%$ and $\alpha = 10\%$. For Champagne, both auditory sources survive the threshold even at the most stringent level, $\alpha = 1\%$. For MCE, bilateral sources survive the threshold, but the localization is not in auditory cortex. SL localizes bilateral activity, but the thresholding, even at the most stringent level, does not create focal sources.

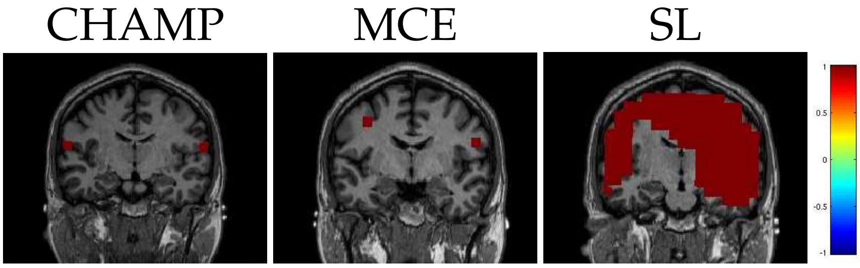

Figure 5.6 P-values computed from the maximal statistic distribution for the window around the auditory activation for the AV data. We plot (1-p) so that high p-values denote high significance. The images are thresholded at $1 - p \geq 0.95$. The p-values for Champagne and MCE threshold the activity similarly to those in Figure 5.5, whereas, the p-value for SL are equally, if not more, diffuse and not an effective means of thresholding.

5.5.2 Somatosensory Evoked Field

We used a somatosensory evoked field (SEF) data set (the same SEF data set used in 3.6.5). The stimulation is administered by air puffs with a pseudorandom interstimulus interval of 450 to 500ms. We concatenated the pre- and post-stimulus periods of each trial to form one long epoch for both the pre-stimulus and post-stimulus periods. For the pre-stimulus period, we took the window of data between -100 to -5ms from each trial and for the post-stimulus period, we took the window between 5ms to 200ms, where 0ms is the onset of the stimulus. We used the first ten trials of data. We calculated the source power in the window between 40ms and 80ms and applied Method 3 of the statistical thresholding procedures, and compared across algorithms.

In Figure 5.10, we present the unthresholded source power results along with the thresholded results for $\alpha = 1\%$, $\alpha = 5\%$, and $\alpha = 10\%$ using the maximal

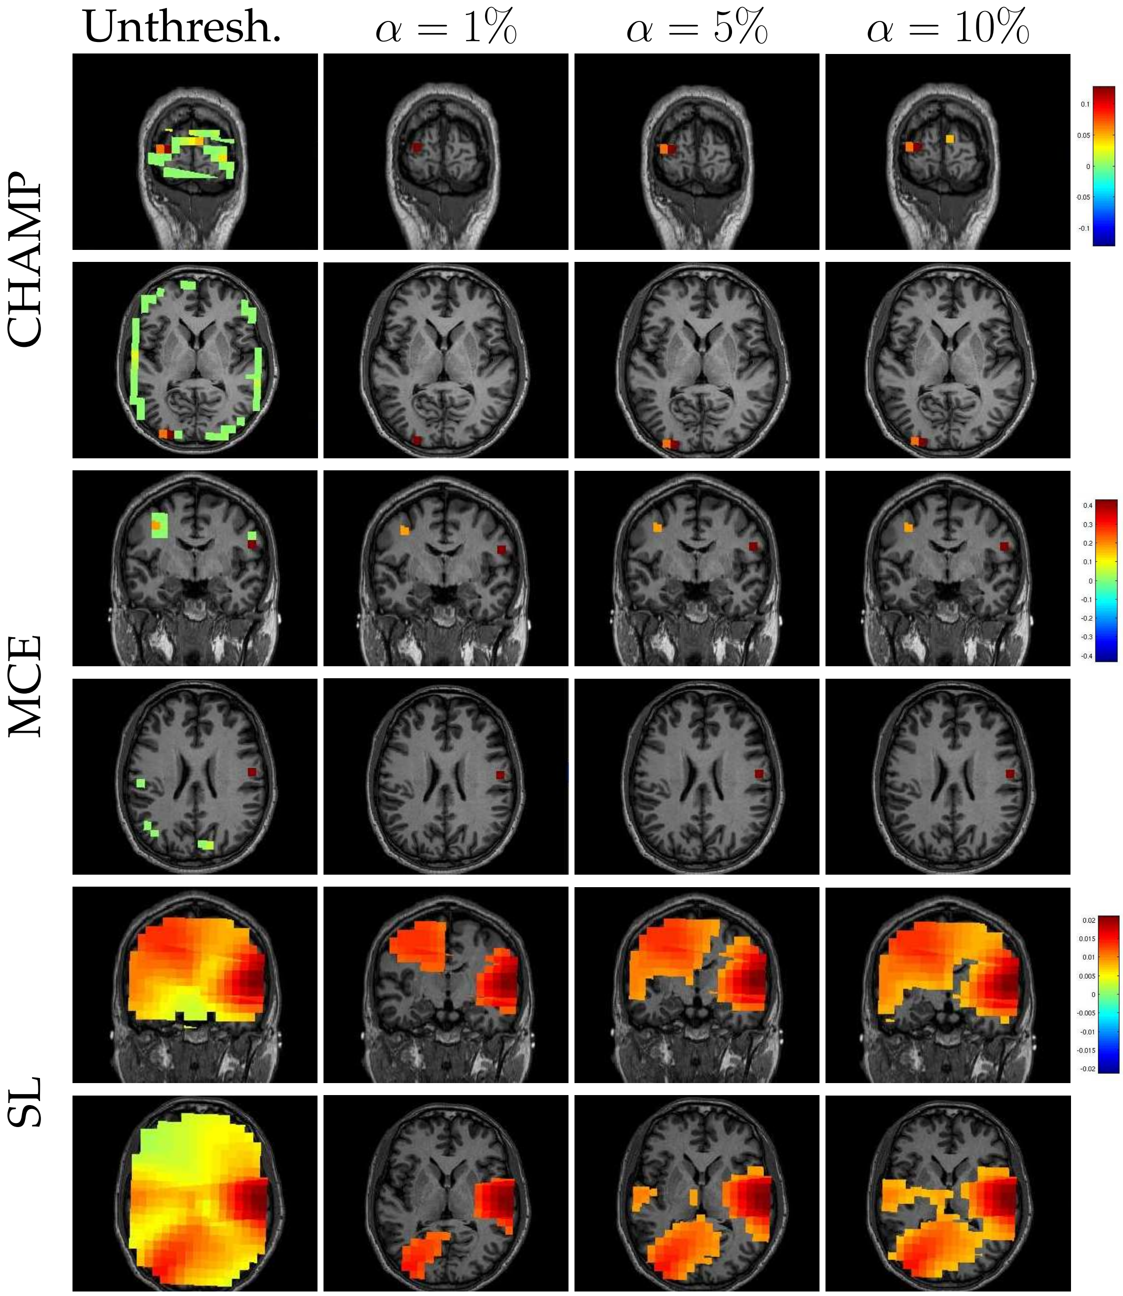

Figure 5.7 AV data with surrogates generated with resampling pre-stimulus data and computing post-stimulus power, Method 3. The unthresholded post-stimulus power values in the window from 100ms to 180ms are shown in the first column (coronal slice). The power is thresholded with the maximal statistic at $\alpha = 1\%$, $\alpha = 5\%$ and $\alpha = 10\%$. See text for description of results.

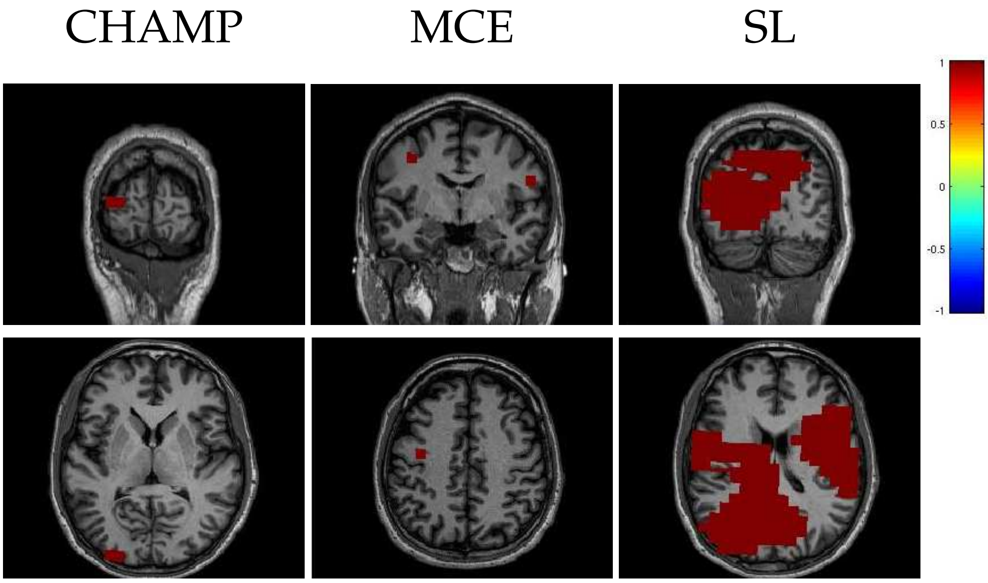

Figure 5.8 P-values computed from the maximal statistic distribution for the window around the visual activation for the AV data. We plot (1-p) so that high p-values denote high significance. The images are thresholded at $1 - p \geq 0.95$. The p-values for Champagne and MCE threshold the activity similarly to those in Figure 5.7 for $\alpha = 1\%$, whereas, the p-values for SL are equally, if not more, diffuse and not an effective means of thresholding.

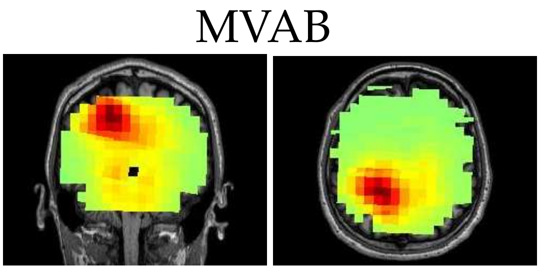

Figure 5.9 AV results for MVAB, the correlated auditory activity impedes MVAB's performance. The post-stimulus power values have the same distribution in both time windows.

statistic approach for all algorithms, Champagne, MCE, MVAB, and SL,. The unthresholded results from Champagne demonstrate that it is able to localize the contralateral somatosensory cortex, but there are voxels in functionally irrelevant areas that were not pruned. Thresholding at all three confidence levels cleans up the source power image, leaving only the source in the contralateral somatosensory cortex. The results from MCE are similar; the unthresholded power image shows that there is a source in somatosensory cortex, but there are also non-zero voxels in other brain areas. Thresholding at $\alpha = 1\%$ leaves only the source in somatosensory cortex and thresholding at $\alpha = 5\%$ and $\alpha = 10\%$ reveals another source nearby. The unthresholded results for MVAB and SL show that there is a peak in the contralateral somatosensory cortex and the thresholding at all levels cleans up the images to some degree. All threshold levels removes more of voxels for MVAB than SL, and the $\alpha = 1\%$ level with MVAB has similar sparisty to Champagne and MCE.

The p-value results for SEF are shown in Figure 5.11. We plot $1 - p$, as such, high values indicate high significance. We thresholded at 5% or $1 - p \geq 0.95$. For Champagne and MCE the p-values look very similar to the results in Figure 5.10. The p-values for MVAB look like the results for $\alpha = 5\%$ and the p-values for SL look like the results for $\alpha = 10\%$; the p-values do not provide more stringent thresholding for SL.

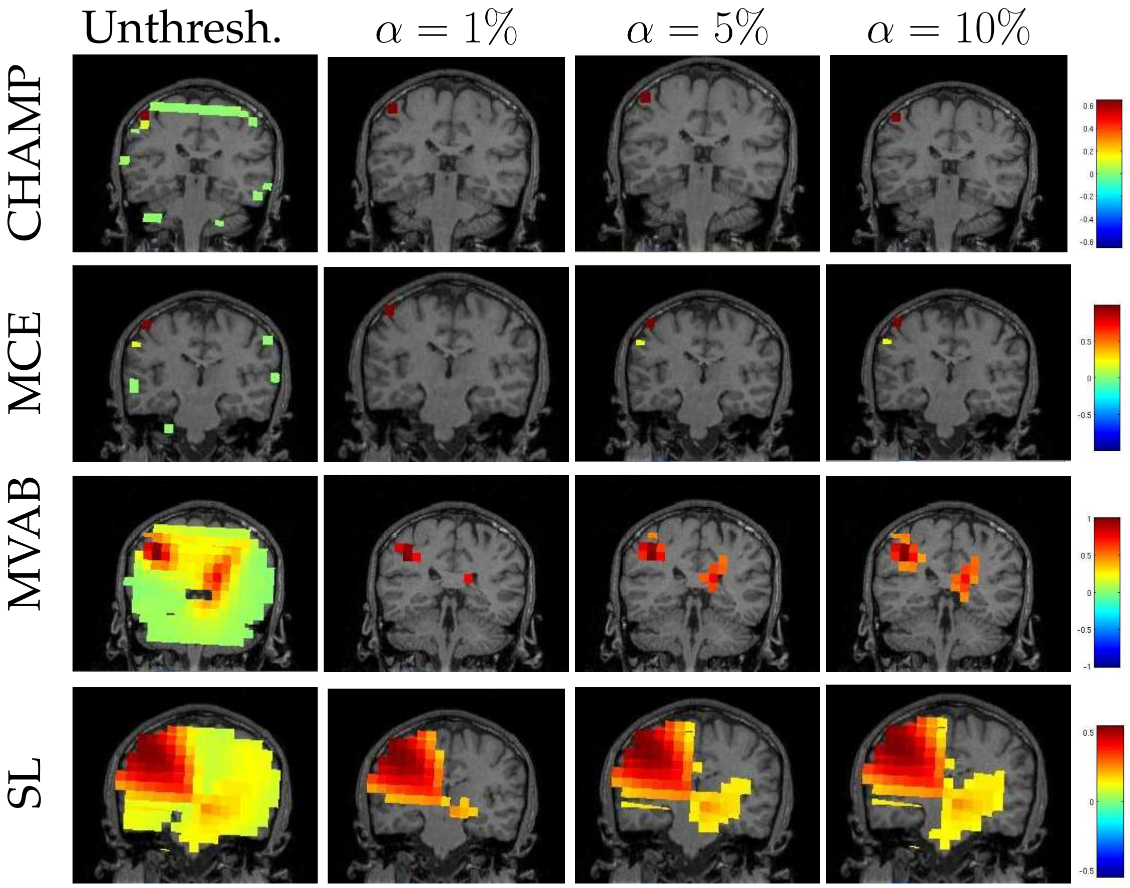

Figure 5.10 Somatosensory (SEF) data with surrogates generated with resampling pre-stimulus data and computing post-stimulus power, Method 3. The unthresholded post-stimulus power values in the window from 40ms to 80ms are shown in the first column (coronal slice). The power is thresholded with the maximal statistic at $\alpha = 1\%$, $\alpha = 5\%$ and $\alpha = 10\%$. For Champagne, MCE, and MVAB all levels of thresholding are able to uncover a focal source in the contralateral somatosensory cortex. SL, on the other hand, does not benefit from the thresholding as the activations are still diffuse even after stringent thresholding.

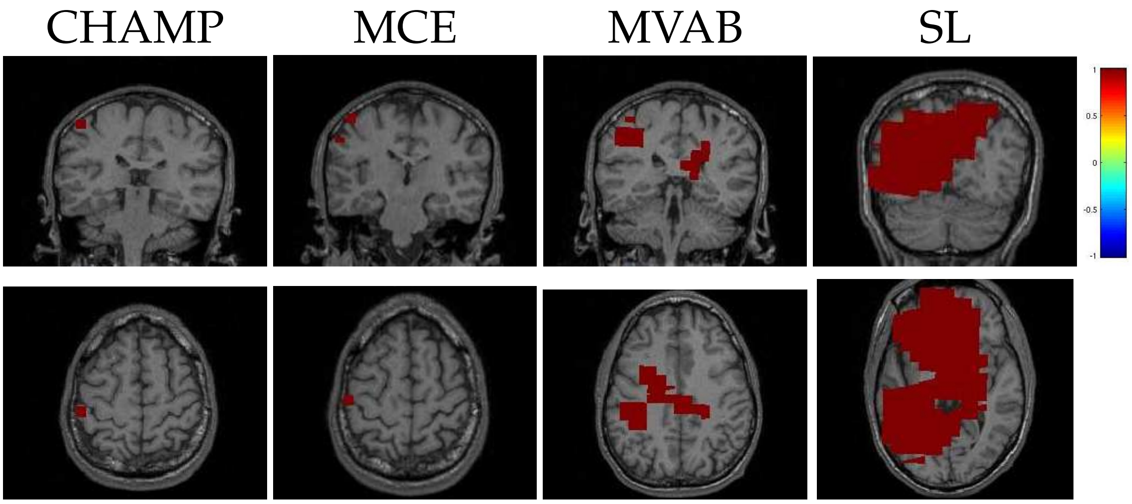

Figure 5.11 P-values for the SEF data computed from the maximal statistic distribution. We plot (1-p) so that high p-values denote high significance. The images are thresholded at $1 - p \geq 0.95$. The p-values for Champagne and MCE threshold the activity similarly to those in Figure 5.10 for $\alpha = 1\%$, whereas, the p-values for SL are equally, if not more, diffuse and not an effective means of thresholding.

5.5.3 Auditory Evoked Field

We analyzed an auditory evoked field (AEF) data set for which the subject was presented single 600ms duration tones (1 kHz) presented binaurally (the same AEF data set used in 3.6.1). We concatenated 35 trials for this data set, choosing the window from -90ms to -5ms as the pre-stimulus period and the window from 5ms to 200ms as the post-stimulus period from each trial. We then calculated the power in the window around the M100, the auditory response, from 90ms to 120ms. We applied Method 3 to obtain statistical threshold for these activations.

The results from the AEF data are shown in Figure 5.12. The first column displays the unthresholded results from the unaveraged data for Champagne, MCE, and SL. All three algorithms show bilateral activity in the time window around the auditory response. For Champagne, the thresholded results for all levels is the same. The thresholded results leave the bilateral auditory activity (the right activation can be seen in the axial slice). MCE also localizes bilateral activity (the left activation can be seen in the axial slice) and the statistical threshold at all levels, like Champagne, yields the same significant voxels at all levels. The statistical thresholding for SL is still quite liberal even at $\alpha = 1\%$ and the thresholding does not provide focal activations. The localization was not successful with MVAB, Figure 5.14, so we did not perform the thresholding with these results. The p-values are displayed in Figure 5.13. The p-values for Champagne differ slightly from the localization as compared to the post-stimulus power maps in Figure 5.12, as the two auditory activations are in the same plane with the p-values, but the results are equally successful. The p-values for MCE demonstrate the same brain activations as the post-stimulus power values and the thresholding for SL suffers with the p-values as the activations are more diffuse than the power map.

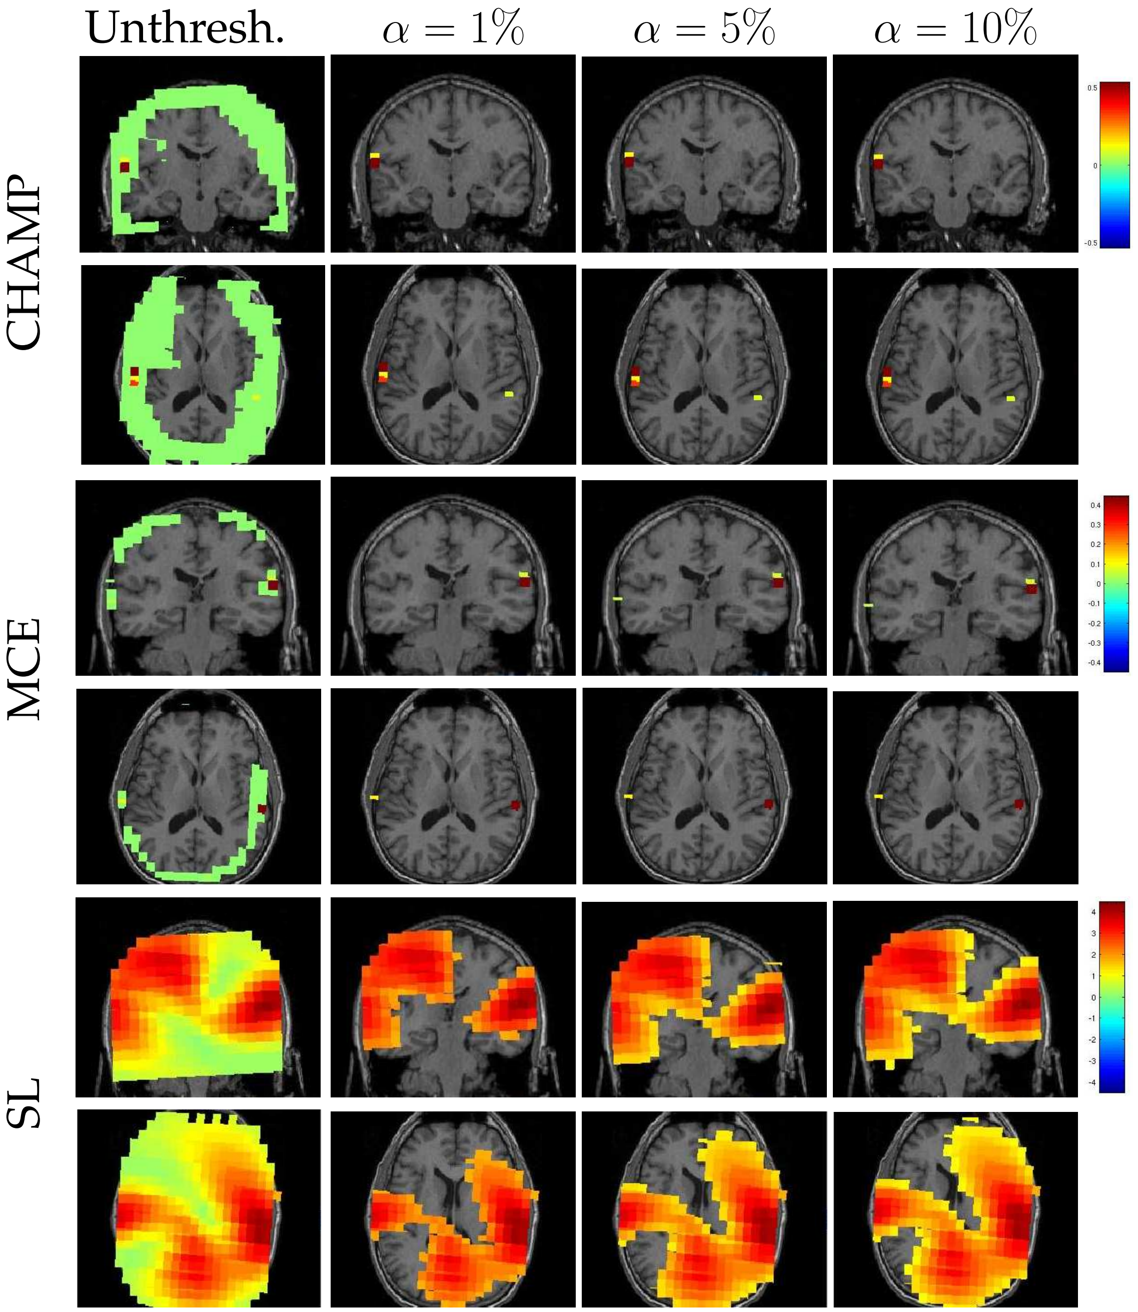

Figure 5.12 Auditory evoked field (AEF) data with surrogates generated with resampling pre-stimulus data and computing post-stimulus power, Method 3. The unthresholded post-stimulus power values in the window from 90ms to 120ms are shown in the first column (coronal slice). The power is thresholded with the maximal statistic at $\alpha = 1\%$, $\alpha = 5\%$ and $\alpha = 10\%$. See text for description of results.

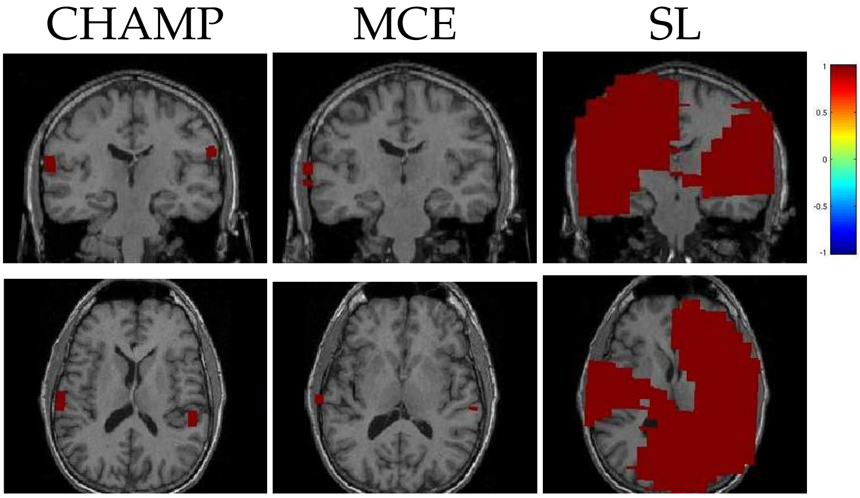

Figure 5.13 P-values for the AEF data computed from the maximal statistic distribution. We plot (1-p) so that high p-values denote high significance. The images are thresholded at $1 - p \geq 0.95$. The p-values for Champagne and MCE threshold the activity similarly to those in Figure 5.12 for $\alpha = 1\%$, whereas, the p-values for SL are equally, if not more, diffuse and not an effective means of thresholding.

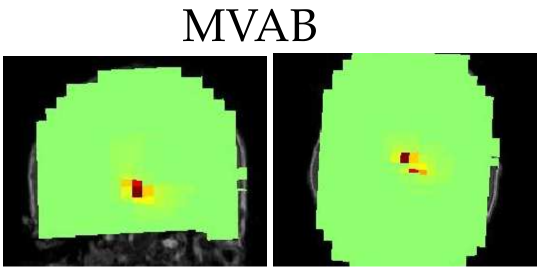

Figure 5.14 AEF results for MVAB, the correlated activity impedes MVAB's performance.

5.5.4 Face-Processing Task

We analyzed a data set from a subject in which faces were presented in a random order with an interstimulus interval of 1 second (the same face processing data set used in 3.6.3). The pre-stimulus window was selected to be -200ms to 5ms and the post-stimulus time window was selected to be from 5ms to 450ms, where 0ms is the appearance of the visual stimulus. We concatenated the pre- and post-stimulus periods from 25 trials for this data set. We used the time window around the M170, from 150ms to 180ms, to calculate the post-stimulus power and we applied Method 3 to obtain statistical thresholds.

The results from the face processing data are in Figure 5.15. The unthresholded results are shown in the first column for Champagne, MCE, and SL. The unthresholded results for Champagne reveal activations in the fusiform gyrus on both the right and left. Thresholding at all levels removes the activation on the right and leaves only the fusiform activation on the left. MCE has a similar pattern of activity, but the thresholding does not eliminate the activity on the right. At the same time, there is anterior activity that is not functionally relevant that also passes the significance test. SL localizes bilateral fusiform activations, but only the activity in the left fusiform area survives thresholding at all levels. It seems that for this data, as evidenced by Champagne and SL, that the left activation is stronger in the fusiform gyrus. MVAB is unsuccessful on this data set, thus the thresholding is not performed, see Figure 5.17. The p-values are displayed in

Figure 5.16. The p-values for Champagne and MCE threshold the activity similarly to those in Figure 5.15 for $\alpha = 5\%$, whereas, the p-values for SL are equally, if not more, diffuse and not an effective means of thresholding.

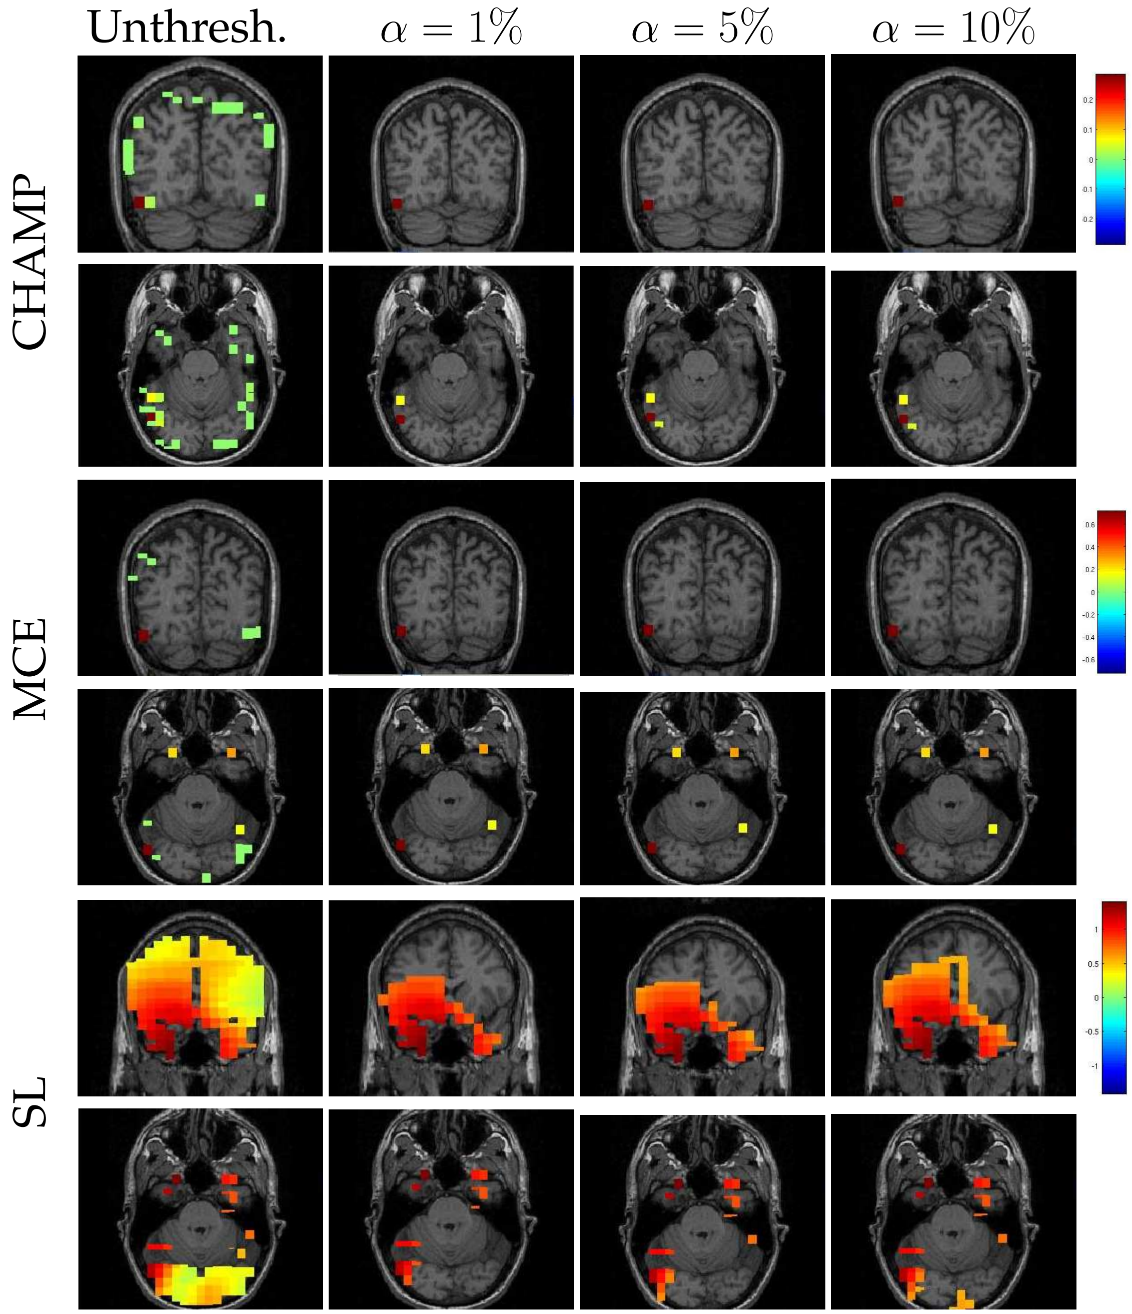

Figure 5.15 Face processing data with surrogates generated with resampling pre-stimulus data and computing post-stimulus power, Method 3. The unthresholded post-stimulus power values in the window around the M170, from 150 to 180ms, are shown in the first column (coronal slice). The power is thresholded with the maximal statistic at $\alpha = 1\%$, $\alpha = 5\%$ and $\alpha = 10\%$. For Champagne, MCE, and SL are all able to localize a source in the left fusiform gyrus. The statistical thresholding helps to clean up the Champagne and MCE results, while the SL results benefit only slightly from thresholding.

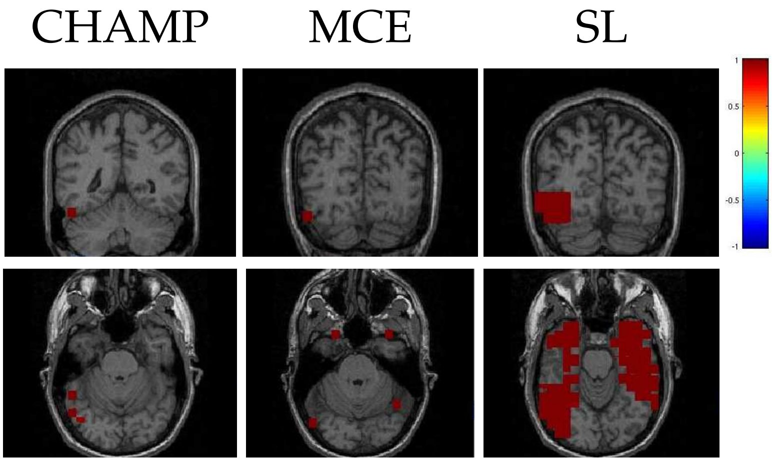

Figure 5.16 P-values for the face data computed from the maximal statistic distribution. We plot (1-p) so that high p-values denote high significance. The images are thresholded at $1 - p \geq 0.95$. The p-values for Champagne and MCE threshold the activity similarly to those in Figure 5.15 for $\alpha = 5\%$, whereas, the p-values for SL are equally, if not more, diffuse and not an effective means of thresholding.

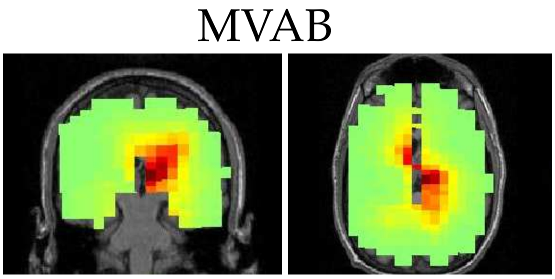

Figure 5.17 MVAB results on face processing data set; it is unable to localize any functionally relevant activity for the face data set.

5.5.5 Alternative to Maximal Statistic

In Figure 5.23, we show the audio-visual data with surrogates generated with pre- and post-stimulus switching and statistical thresholding with the alternative to the maximal statistic with the t-values. The unthresholded post-stimulus power values are shown in the first column (coronal slice). The power is thresholded by keeping the top 5^{th} percentile of the surrogate statistic and thresholding with $\alpha = 5\%$, denoted as $\theta_{5\%}^{5^{th}}$. Both the left and right auditory sources survive the threshold, but there are spurious peaks introduced from the computation of the t-values seen at the base of the head. Even with less stringent thresholding than the maximal statistic, computing t-values for sparse images is problematic due to the difficulties estimating the variance.

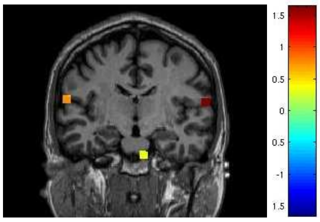

Figure 5.18 Audio-visual data with surrogates generated with pre- and post-stimulus switching and statistical thresholding with the alternative to the maximal statistic with the t-values. The unthresholded post-stimulus power values are shown in the first column (coronal slice). The power is thresholded by keeping the top 5^{th} percentile of the surrogate statistic and thresholding with $\alpha = 5\%$. Both the left and right auditory sources survive the threshold, but there are spurious peaks introduced from the computation of the t-values.

We also applied this alternative to the maximal statistic, the n^{th} percentile threshold, to the SEF and AV results using the post-stimulus power (Method 3).

This threshold is obtained by keeping the 1^{st} and 5^{th} percentile, denoted as $\theta_{\alpha}^{1^{st}}$ and $\theta_{\alpha}^{5^{th}}$ respectively. In Figure 5.19, we show the unthresholded results for the SEF data set along with the thresholded results using maximal statistic at $\alpha = 5\%$ and the 1^{st} and 5^{th} percentile results, also at $\alpha = 5\%$. As compared to the maximal statistic threshold, the percentile thresholds are less stringent and allow an activation in the ipsilateral somatosensory cortex to pass the significance threshold. In Figure 5.20, we show the unthresholded results for the AV data set along with the thresholded results using maximal statistic at $\alpha = 5\%$ and the 1^{st} and 5^{th} percentile results, also at $\alpha = 5\%$. As compared to the maximal statistic threshold, the n^{th} percentile thresholds are less stringent and allow more of the visual activity present in the unthresholded map to pass the significance threshold.

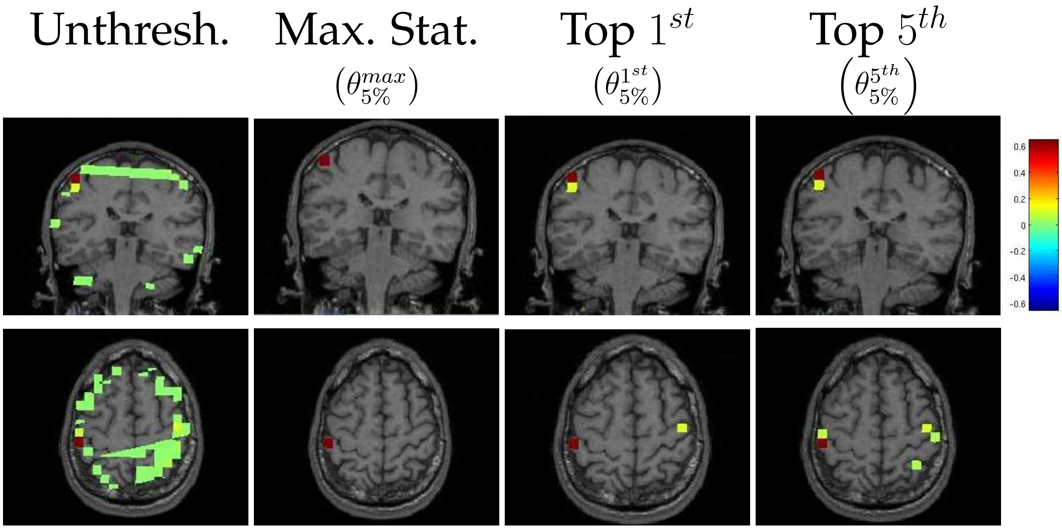

Figure 5.19 SEF results: These results compared the threshold at $\alpha = 5\%$ for the maximal statistic method as well as the method we have developed where we keep the top 1^{st} or 5^{th} percentile of each surrogate data sets.

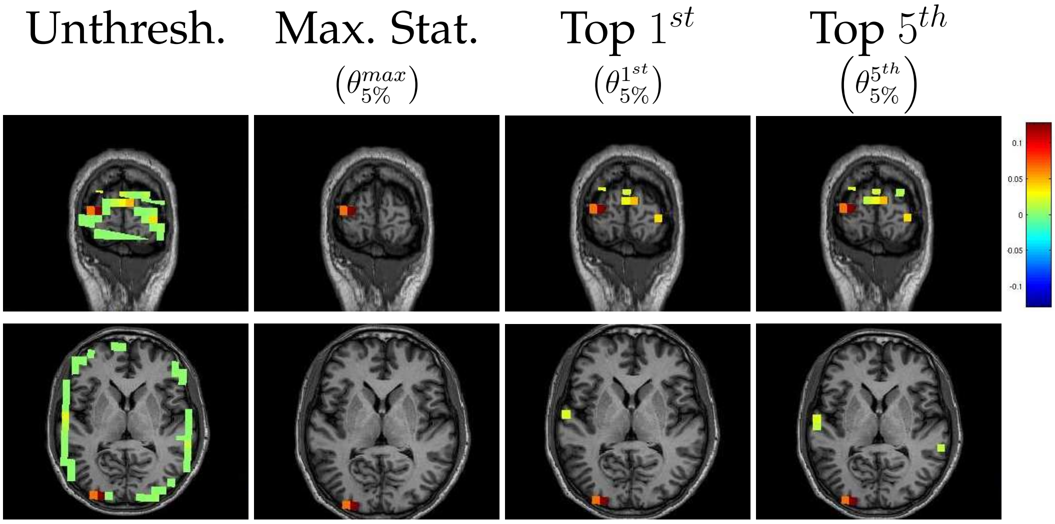

Figure 5.20 AV results: These results compared the threshold at $\alpha = 5\%$ for the maximal statistic method as well as the method we have developed where we keep the top 1^{st} or 5^{th} percentile of each surrogate data sets.

5.6 Discussion

In this chapter we have demonstrated that Champagne and the benchmark algorithms are able to localize activity with sensor data that has not been averaged across trials. The simulated data demonstrates that the permutation testing method for statistical thresholding developed in this chapter effectively removes activity for voxels not seeded with activity with Champagne and MCE. The thresholding is able to clean up the results with SL, but does not yield sparsity, and the results from MVAB do not benefit from the thresholding procedure. The simulations, while only one case, give evidence that the nonparametric thresholding procedure we have developed is compatible with sparse reconstructions.

In the case of the SEF, AEF, and AV data sets, Champagne is able to localize similar brain activity as compared to running the algorithm on trial-averaged data (see Section 3.6). The benchmark algorithms, MCE, SL, and MVAB, are also able to localize similar activity, although MVAB suffers more than the other algorithms without trial averaging. The face processing data proves to be more challenging for the algorithms when the data are not trial averaged. Champagne is able to localize the fusiform activation at 170ms, but does not localize the earlier visual activity that Champagne, run on trial-averaged data, was able to localize. When looking at the averaged sensor data for this data set, the peak at 170ms is the dominant peak, so it is not hard to believe that the visual activity would be more difficult to localize. Being able to localize brain activity on unaveraged MEG data has important implications for examining high frequency activity. This activity is averaged out of the data when an average is taken across trials as the high frequency activity is not phase-locked to the stimulus. These results with unaveraged data imply that Champagne could be used to localize power increases in the high frequency range, and if methods from 4.3 are implemented for unaveraged data, power decreases in the high frequencies could be detected as well.

We investigated methods to statistically threshold the results from non-trial averaged data. We designed three nonparametric, permutation-testing techniques referred to as Method 1, Method 2 and Method 3. All three methods employ the maximal statistic approach, but differ in the way they generate the surrogate data sets and the test statistic used. Method 1 consisted of creating surrogate data sets by exchanging the pre- and post-stimulus periods of random trails. Then a pseudo t-value was computed for the real data and all of the surrogate data sets for each voxel. Method 2, also, consisted of creating surrogate data sets with switching pre- and post-stimulus periods, but it used the post-stimulus power as opposed to the pseudo t-value as the statistic. These two methods proved to be overly stringent for Champagne when coupled with the maximal statistic. Typically only the maximum voxel survived the threshold, which eliminates functionally relevant activity seen in the unthresholded images. As such, we devised a third method of thresholding, called Method 3. In this method, the surrogates are generated by resampling the pre-stimulus period and the post-stimulus power is used as the statistic. We found that this method is less stringent and effectively thresholds out some of the non-zero activity from Champagne's activation maps, while leaving functionally relevant peaks in the activation. Confidence intervals of 1%, 5%, and 10% were used; in general all confidence intervals effectively thresholded the results. When this method was applied to MCE, similar thresholding effects were observed. When MVAB was able to localize activity, it also benefited from the statistical thresholding with Method 3. However, SL was overly diffuse, even after thresholding at 1%, leading to the conclusion that sparse-like solutions cannot always be obtained with stringent statistical thresholding of a non-sparse algorithm, such as SL.

In addition to using the maximal statistic threshold on the post-stimulus power values (Method 3), we also computed p-values from the distribution of the maximal statistic. These p-values, with a cutoff of $p = 0.05$, closely resembled the thresholded results. The p-values are a standardized approach and like the maximal statistic threshold, are corrected for multiple comparisons. Lastly, we designed a variant to the maximal statistic in which we saved the top percentile from each permutation or surrogate, as opposed to only the maximum. This method proved to be less stringent than the maximal statistic and in some instances, allowed functionally relevant activity, which were thresholded away by the maximal statistic, to survive; such was the case for the SEF data set and the visual activations in the AV data set.

The widely-used statistical thresholding techniques, such as Method 1 and 2 in this chapter, were not designed with sparse algorithms in mind. If we look at the sparsity profile of an algorithm like SL as compared to Champagne, the histogram of the post-stimulus power values across voxels is drastically different in shape. SL has a more or less smooth histogram, where Champagne has many voxels with little to no power and only a small subset with high power. The difference between the highest power value for Champagne and the second highest power value is large, see Figure 5.21 for a comparison of the histograms. As such, even when we resample the pre-stimulus period to create surrogate data sets, this distribution of power values persists. If only the maximum statistic is saved for the null distribution, the threshold obtained can be driven by spurious voxels.

Saving more than just the maximum from each permutation, protects the threshold from the spurious, high-powered voxels, which are more prevalent in sparse algorithms.

Yet another approach to statistical thresholding with resampling techniques is to run Champagne on each surrogate, i.e. learning different weights for each surrogate and then proceeding with the maximal statistic procedure. We found that relearning the weights for each surrogate data set was too noisy and only the maximum voxels passed the maximal statistic threshold. Fixing the weights also guarantees that the scaling of the voxel time courses will be consistent across surrogates, as we fix both the sensor power across trials and the weights. This further protects against the maximal statistic threshold being driven by spurious voxels, an effect that Champagne and other sparse algorithms are inherently prone, as discussed above. In addition, we chose to apply the weights learned from the unpermuted data to the surrogate data sets, instead of adopting this approach because running Champagne on unaveraged data is more computationally expensive than running it on trial averaged data (the time to compute each iteration scales with the number of time points in the post-stimulus period). If we recomputed the weights for each surrogate data set, the process of obtaining the maximal statistic threshold would take hours, perhaps days longer. We would like to try the *Top* 5^{th} and *Top* 1^{st} percentile approaches to the value obtained by relearning the weights for each permutation. This might protect from the spurious voxels we found to drive the maximal statistic approach.

Our approach of creating the surrogate data sets by resampling of the pre-stimulus period, Method 3, is an unconventional methodology. Exchanging the pre- and post-stimulus periods, Method 1 and Method 2, is the usual practice, as presented in (Pantazis et al., 2005). In this paper, the authors use a linear inverse method to localize the sources, which is an non-adaptive method. Champagne is an adaptive method, meaning that the weights are dependent on the data, while non-adaptive methods, such as SL and other minimum-norm algorithms, do not factor in the data when calculating the weights. The adaptive and sparse nature of Champagne makes the conventional method of generating surrogates problematic. Champagne prunes the majority of voxels to be zero, which essentially remove these voxels from the VOI. When these sparse weights are then applied to new data, the locations in the brain where there is non-zero activity is highly constrained, see Figure 5.21. If the surrogates are generated by switching the pre- and post-stimulus period, the surrogate post-stimulus periods will contain some of the same signal as the original data. This has more of an effect on Champagne and other adaptive methods than it does on the non-adaptive methods. Thus, we think it is a justifiable approach to use only pre-stimulus data for the surrogates. When we do this, we are assessing the source power obtained on data that we assume has no signal of interest, which is the assumption under the null hypothesis.

Method 3 also diverges from conventional methods in that we use the post-stimulus power as opposed to a pseudo t-value (or other statistic). When we apply Champagne's sparse weights to the pre-stimulus period in order to obtain an estimate for the variance, used in the pseudo t-value calculation, we only obtain non-zero variance in a small subset of the voxels. Usually the variance is pooled across neighboring voxels to protect from spurious values driving the t-values, but in Champagne's case pooling the variance does not have a smoothing effect. Thus, we found it was more stable to use the post-stimulus power values as our measure. It should be noted that Champagne subtracts the baseline from the post-stimulus source estimates and thereby the power values obtained are effectively the subtraction of the pre- and post-stimulus power.

We have developed methods for statistically thresholding single-subject brain-activity maps. These methods are designed for sparse algorithms and are able to produce statistical thresholds which preserve functionally relevant activity, while removing spurious voxels that do not get pruned away during source localization. They effectively address the multiple comparisons problem by obtaining statistical thresholds that are inherently corrected and this method provides an alternative to spatial normalization and averaging across subjects, a common approach to assessing statistic significance.

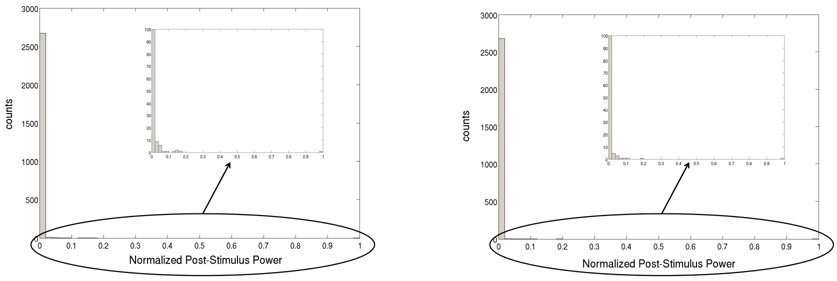

(a)Champagne (b)MCE

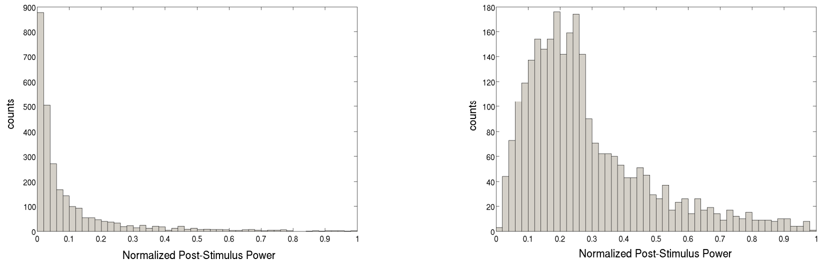

(c) MVAB (d) SL

Figure 5.21 Histograms of the post-stimulus power to illustrate the difference between sparse algorithms, Champagne and MCE, and non-sparse algorithms, MVAB and SL.

5.7 Appendix

We present the results from thresholding the AV data around the auditory activation with Method 1 and 2 in this appendix. The results from thresholding with Method 1 are shown in Figure 5.22. In the first column, the unthresholded pseudo t-values are shown. We thresholded the pseudo t-values at $\alpha = 1\%$, $\alpha = 5\%$, and $\alpha = 10\%$. At all levels, only the right auditory cortical source survives. In Figure 5.23, the results from Method 2 are shown. The unthresholded post-stimulus power values are displayed in the first column. As with Method 1, all levels of thresholding with Method 2 only leave the right auditory source. Comparing the unthresholded t-values in Figure 5.22 and the post-stimulus power values in 5.23, one can notice that while the location of non-zero values are the same, the amplitudes are different in some voxels. The t-values are computed by dividing by the standard deviation, which when small compared to the numerator (the difference between pre- and post-stimulus power), causes the t-values to be artificially large. In Figure 5.24, we show the standard deviation values (smoothed over 20 neighboring voxel as described above in (5.3)) used in the t-value calculation. We can see that voxels with low post-stimulus power in Figure 5.23 can also have low standard deviation in Figure 5.24, thus resulting in errantly high t-values.

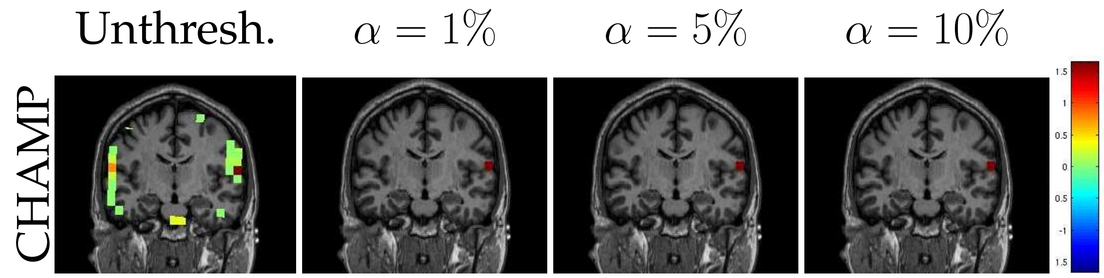

Figure 5.22 Audio-visual (AV) data with surrogates generated with pre- and post-stimulus switching and calculating t-values, Method 1. The unthresholded t-values are shown in the first column (coronal slice). The t-values are computed from the post-stimulus power in the window from 80ms to 140ms. These results are thresholded with the maximal statistic at $\alpha = 1\%$, $\alpha = 5\%$ and $\alpha = 10\%$. Only the left auditory source survives the threshold even at the least stringent level, $\alpha = 10\%$.

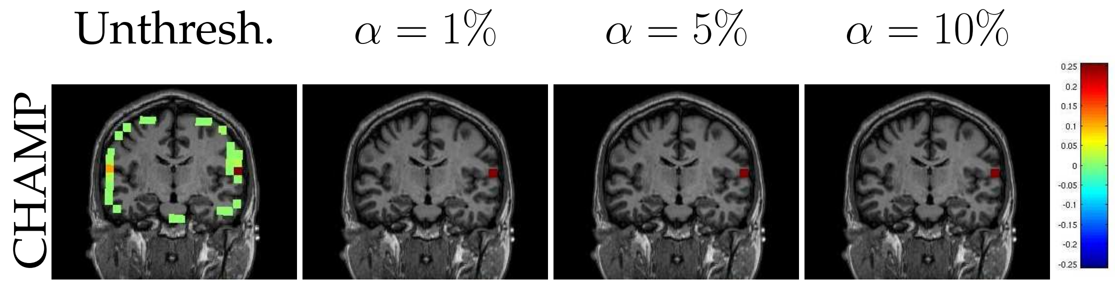

Figure 5.23 AV data with surrogates generated with pre- and post-stimulus switching and calculating post-stimulus power, Method 2. The unthresholded post-stimulus power values in the window from 80ms to 140ms are shown in the first column (coronal slice). The power is thresholded with the maximal statistic at $\alpha = 1\%$, $\alpha = 5\%$ and $\alpha = 10\%$. Only the left auditory source survives the threshold even at the least stringent level, $\alpha = 10\%$.

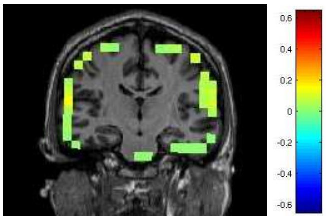

Figure 5.24 Plot of standard deviation of the pre-stimulus period for the AV data set, showing that this term in the denominator introduces spurious peaks and lessens that true peaks in the image.

Chapter 6

Conclusions

6.1 Overview

In Chapter Two of this thesis, we presented a novel source localization algorithm, called Champagne. Although Champagne is embedded in deep theoretical ideas, the resulting algorithm essentially iterates between four simple steps. We believe that the results presented in Chapter Three of this thesis demonstrate that Champagne makes a significant breakthrough in the reconstruction of brain activity with M/EEG data. Both simulated and real data sets were used to test Champagne in comparison to commonly-used source localization algorithms. In Chapter Four, we have investigated variants on the Champagne algorithm, including methods for applying the Champagne algorithm to data sets that do not contain a pre-stimulus period. Champagne is able to provide reliable estimates of the source locations and time courses with these types of data sets, demonstrating the robustness of the algorithm. In addition to developing and testing the Champagne algorithm, in Chapter Five, we have developed methods for statistically thresholding the results obtained from Champagne and other source localization algorithms. These methods provide thresholds that are corrected for multiple comparisons and are designed to be effective on sparse, adaptive source localization techniques. Taken as a whole, the work presented in this thesis improves upon the analysis of M/EEG data as it provides robust methods for localizing activity and versatile statistical thresholding techniques.

6.2 Future Directions

6.2.1 Comparisons and Extensions of Champagne Algorithm

Champagne's generative model is related to the multiple sparse priors model (MSP) (Friston et al., 2008). Both are covariance component analysis algorithms with some key differences. First, while MSP assumes a set of hundreds of covariance components that represent activity correlated across voxels, Champagne assumes thousands of covariance components representing uncorrelated activity in each voxel. Second, Champagne uses flat hyperpriors for variances and computes MAP estimates for variances with convergent update rules, whereas MSP assumes a Gamma-distribution for hyperpriors for variances, and computes a posterior distribution under a Laplace approximation. Finally, the noise is simultaneously estimated in the MSP algorithm, rather than separately estimated with the pre-stimulus data, as in Champagne. Given both the similarities and points of divergence, we plan to investigate the comparison of the MSP model to Champagne in future research.

Additionally, Champagne does not make assumptions about the smoothness of sources, nor does it consider uncertainties in the lead field. We are currently investigating the use of spatial- and temporal-smoothness priors, either in the form of basis functions, as in (Zumer et al., 2008), or in the form of auto-regressive smoothness priors. Incorporating spatial priors and better noise models that also model spatiotemporal correlations in the background noise could potentially improve performance. These extensions of the Champagne method hold promise for improving upon an already robust source localization algorithm. We investigated the effect of lead field errors in Chapter Three and found that a small amount of error is permissible. There are source localization algorithms that account for errors or uncertainties in the lead field (Stahlhut et al., 2010). Incorporating these lead field uncertainties into the Champagne algorithm is an extension we seek to investigate.

In addition to using Champagne to localize activity, we plan to use Champagne in the analysis of functional connectivity. Champagne, as discussed in Chapter Three, is well suited for functional connectivity as it is robust to correlated sources (which is expected when areas are functionally connected) and is a sparse algorithm, selection of the nodes is done by the algorithm and not heuristically by the researcher. To date, we have only applied Champagne to the simulated networks of brain regions (see 3.5.13) and some preliminary investigations with a MEG motor data set (Sekihara et al., 2010). These preliminary investigations imply that Champagne would be an effective means for analyzing the functional connectivity in complex tasks. Also, we would like to extend Champagne to time-frequency analysis, the first step of which is to be able to localize activity without access to a pre-stimulus period. In Section 4.3, we outline the possible ways to obtain source estimated without estimating the noise from the pre-stimulus period. This technique will be extended to the analysis of MEG in the frequency domain, similar to the technique in (Dalal et al., 2008).

6.2.2 Applications to MEG Data Sets

Champagne is being applied to a number of clinical data sets, especially those data sets where traditional source localization algorithms, such as sLORETA or beamforming have failed. Some clinical populations that we are currently investigating with Champagne are auditory processing in patients with schizophrenia and auditory and somatosensory integration in children with sensory processing disorders, such as autism. The children with sensory processing disorders present a challenge to source localization because they tend to move while in the MEG scanner, creating artifacts in the data. Our first attempts to apply Champagne has proven to be successful as the auditory and somatosensory activations are localized with Champagne, while other source localization algorithms fail. In the future, we plan to continue using Champagne on difficult data sets, and we plan to incorporate the statistical procedure into the analysis of these data sets.

Bibliography

Attias, H., 1999. Inferring parameters and structure of latent variable models by variational bayes, in: Proc. 15th Conf. Uncert. Art. Intell., pp. 21–30.

Baryshnikov, B., Van Veen, B., Wakai, R., 2004. Maximum-likelihood estimation of low-rank signals for multiepoch MEG/EEG analysis. Biomedical Engineering, IEEE Transactions on 51, 1981–1993.

Benjamini, Y., Hochberg, Y., 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. Journal of the Royal Statistical Society, Series B (Methodological) 57, 289–300.

Benjamini, Y., Liu, W., 1999. A step-down multiple hypothesis testing procedure that controls the false discovery rate under independence. Journal of Statistical Planning and Inference 82, 163–170.

Bijma, F., de Munck, J.C., Heethaar, R.M., 2005. The spatiotemporal MEG covariance matrix modeled as a sum of Kronecker products. NeuroImage 27, 402–415.

Bonferoni, C., 1935. Il calcolo delle assicurazioni su gruppi di teste. Studi in Onore del Professore Salvatore Ortu Carboni , 13–60.

Chau, W., McIntosh, A., Robinson, S., schulz, M., Pantev, C., 2004. Improving permutation test power for group analysis of spatially filtered MEG data. Neuroimage 23, 983–996.

Dalal, S., Guggisberga, A., Edwards, E., Sekihara, K., Findlay, A., Canolty, R., Berger, M., Knight, R., Barbaro, N., Kirsch, H., Nagarajan, S., 2008. Five-dimensional neuroimaging: Localization of the timefrequency dynamics of cortical activity. Neuroimage 40, 1686–1700.

Dalal, S., Sekihara, K., Nagarajan, S., 2006. Modified beamformers for coherent source region suppression. IEEE Trans. Biomed. Eng. 53, 1357–1363.

Dalal, S., Zumer, J., Agrawal, V., Hild, K., Sekihara, K., Nagarajan, S., 2004. NUTMEG: A neuromagnetic source reconstruction toolbox. Neurol. Clin. Neurophysiol. 52, 2004–2052.

Dale, A., Liu, A., Fischl, B., Buckner, R., Belliveau, J., Lewine, J., Halgren, E., 2000. Dynamic statistical parametric mapping: Combining fMRI and MEG for high-resolution imaging of cortical activity. Neuron 26, 55–67.

Darvas, F., Pantazis, D., Kucukaltun-Yildirim, E., Leahy, R.M., 2004. Mapping human brain function with MEG and EEG: methods and validation. NeuroImage 23 Suppl 1, S289–S299.

De Munck, J.C., Huizenga, H.M., Waldorp, L.J., Heethaar, R.M., 2002. Estimating stationary dipoles from MEG/EEG data contaminated with spatially and temporally correlated background noise. IEEE Trans. Signal Processing 50, 1565–1572.

Diwakar, M., Huanga, M., Srinivasand, R., Harringtona, D., Robbc, A., Angelesc, A., Muzzattia, L., Pakdamana, R., Songa, T., Theilmanna, R., Leea, R., 2010. Dual-Core beamformer for obtaining highly correlated neuronal networks in MEG. Neuroimage 54, 253–263.

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 MEG/EEG inverse problem. NeuroImage 39, 1104–20.

Gerven, M.v., Cseke, B., Oostenveld, R., Heskes, T., 2009. Bayesian source localization with the multivariate laplace prior. Advances in Neural Information Processing Systems 22.

Gorodnitsky, I., Rao, B., 1997. Sparse signal reconstruction from limited data using FOCUSS: A re-weighted minimum norm algorithm. IEEE Trans Signal Processing 45, 600–616.

Gross, J., Kujala, J., Hamalainen, M., Timmermann, L., Schnitzler, A., Salmelin, R., 2001. Dynamic imaging of coherent sources: Studying neural interactions in the human brain. Proc Natl Acad Sci U S A 98, 694–699.

Hämäläinen, M., Ilmoniemi, R., 1994. Interpreting magnetic fields of the brain: Minimum norm estimates. Med. Biol. Eng. Comput 32, 35–42.

Haufe, S., Nikulinc, V., Ziehea, A., Mllera, K., Nolte, G., 2008. Group imaging of task-related changes in cortical synchronisation using nonparametric permutation testing. Neuroimage 42, 726–738.

Henson, R., Flandin, G., Friston, K., Mattoutet, J., 2010. A parametric empirical Bayesian framework for fMRI-constrained MEG/EEG source reconstruction. Human Brain Mapping 31, 15121531.

Huizenga, H.M., de Munck, J.C., Waldorp, L.J., Grasman, R.P.P.P., 2002. Spatiotemporal EEG/MEG source analysis based on a parametric noise covariance model. IEEE Trans. Biomed. Eng. 49, 533–539.

Jerbi, K., Mosher, J.C., Baillet, S., Leahy, R.M., 2002. On MEG forward modelling using multipolar expansions. Phys. Med. Biol. 47, 523–555.

Jun, S.C., George, J.S., Paré-Blagoev, J., Plis, S.M., Ranken, D.M., Schmidt, D.M., Wood, C.C., 2005. Spatiotemporal Bayesian inference dipole analysis for MEG neuroimaging data. NeuroImage 28, 84–98.

Jun, S.C., George, J.S., Plis, S.M., Ranken, D.M., Schmidt, D.M., Wood, C.C., 2006a. Improving source detection and separation in a spatiotemporal Bayesian inference dipole analysis. Phys. Med. Biol. 51, 2395–2414.

Jun, S.C., Plis, S.M., Ranken, D.M., Schmidt, D.M., 2006b. Spatiotemporal noise covariance estimation from limited empirical magnetoencephalographic data. Phys. Med. Biol. 51, 5549–5564.

Kandel, E., Schwartz, J., Jessel, T. (Eds.), 2000. Principles of Neural Science. McGraw-Hill.

Kanwisher, N., McDermott, J., Chun, M., 1997. The fusiform face area: A module in human extrastriate cortex specialized for face perception. The Journal of Neuroscience 17, 4302–431.

Kiebel, S., Daunizeau, J., Phillips, C., Friston, K., 2008. Variational Bayesian inversion of the equivalent current dipole model in EEG/MEG. Neuroimage 39, 728–741.

Leahy, R.M., Mosher, J.C., Spencer, M.E., Huang, M.X., Lewine, J.D., 1998. A study of dipole localization accuracy for MEG and EEG using a human skull phantom. Electroencephalogr Clin. Neurophysiol 107, 159–73.

Limpiti, T., Van Veen, B., Wakai, R., 2006. Cortical patch basis model for spatially extended neural activity. Biomedical Engineering, IEEE Transactions on 53, 1740–1754.

Lin, F., Belliveau, J., Dale, A., Hmlinen, M., 2006a. Distributed current estimates using cortical orientation constraints. Biomedical Engineering, IEEE Transactions on 27, 1–13.

Lin, F., Witzela, T., Ahlfors, S., Stufflebeam, S., Belliveau, J., Hmlinen, M., 2006b. Assessing and improving the spatial accuracy in MEG source localization by depth-weighted minimum-norm estimates. Neuroimage 31, 160–171.

Lu, Z.L., Kaufman, L., 2003. Magnetic source imaging of the human brain. Lawrence Erlbaum Associates.

Lutkenhoner, B., 1998. Dipole separability in a neuromagnetic source analysis. Biomedical Engineering, IEEE Transactions on 45, 572–581.

Mattout, J., Phillips, C., Penny, W., Rugg, M., Friston, K., 2006. MEG source localization under multiple constraints: an extended Bayesian framework. NeuroImage 30, 753–767.

Mosher, J., Leahy, R., Lewis, P., 1999. EEG and MEG: forward solutions for inverse methods. Biomedical Engineering, IEEE Transactions on 46, 245 –259.

Mosher, J., Lewis, P., Leahy, R., 1992. Multiple dipole modeling and localization from spatiotemporal MEG data. IEEE Trans. Biomed. Eng. 39, 541–557.

Mosher, J.C., Leahy, R.M., 1998. Recursive MUSIC: a framework for EEG and MEG source localization. IEEE Trans. Biomed. Eng. 45, 1342–1354.

de Munck, J.C., Bijma, F., Gaura, P., Sieluzycki, C.A., Branco, M.I., Heethaar, R.M., 2004. A maximum-likelihood estimator for trial-to-trial variations in noisy MEG/EEG data sets. IEEE Trans. Biomed. Eng. 51, 2123–2128.

Nagarajan, S., Attias, H., Hild, K., Sekihara, K., 2007. A probabilistic algorithm for robust interference suppression in bioelectromagnetic sensor data. Stat. in Med. 26, 3886–3910.

Nagarajan, S., Attias, H., Sekihara, K., Hild II, K., 2006. Partitioned factor analysis for interference suppression and source extraction, in: International Workshop on Independent Component Analysis and Signal Separation, p. 189197.

Nichols, T., Holmes, A., 2001. Nonparametric permutation tests for functional neuroimaging data: A primer with examples. Human Brain Mapping 15, 1–25.

Nummenmaa, A., Auranen, T., Hämäläinen, M., I., J., Lampinen, J., Sams, M., Vehtari, A., 2007. Hierarchical Bayesian estimates of distributed MEG sources: Theoretical aspects and comparison of variational and MCMC methods. Neuroimage 35, 669–685.

Ossadtchi, A., Baillet, S., Mosher, J., Thyerlei, D., Sutherling, W., Leahy, R., 2004. Automated interictal spike detection and source localization in magnetoencephalography using independent components analysis and spatio-temporal clustering. Clinical Neurophysiology 115, 508–522.

Owen, J., Wipf, D., Attias, H., Sekihara, K., Nagarajan, S., 2009. Robust methods for reconstructing brain activity and functional connectivity between brain sources with MEG/EEG data., in: ISBI'09, pp. 1271–1274.

Pantazis, D., Nichols, T., Baillet, S., Leahy, R., 2005. A comparison of random field theory and permutation methods for the statistical analyses of meg data. Neuroimage 25, 383–394.

Pascual-Marqui, R., 2002. Standardized low-resolution brain electromagnetic tomography (sLORETA): technical details. Meth. Find. Exp. Clin. Pharmacol. 24 Suppl D, 5–12.

Phillips, C., Mattout, J., Rugg, M., Maquet, P., Friston, K., 2005. An empirical Bayesian solution to the source reconstruction problem in EEG. NeuroImage 24, 997–1011.

Sahani, M., Nagarajan, S., 2004. Reconstructing MEG sources with unknown correlations. Advances in Neural Information Processing Systems 16.

Sarvas, J., 1987. Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem. Phys. Med. Biol. 32, 11–22.

Sato, M., Yoshiokaa, T., Kajiharac, S., Toyamac, K., Godad, N., Doyaa, K., Kawatoa, M., 2004. Hierarchical Bayesian estimation for MEG inverse problem. NeuroImage 23, 806–826.

Schmidt, D.M., George, J.S., Wood, C.C., 1999. Bayesian inference applied to the electromagnetic inverse problem. Human Brain Mapping 7, 195–212.

Sekihara, K., Hild, K., Dalal, S., Nagarajan, S., 2008. Performance of prewhitening beamforming in MEG dual experimental conditions. Biomedical Engineering, IEEE Transactions on 55, 1112=1121.

Sekihara, K., Nagarajan, S., 2008. Adaptive Spatial Filters for Electromagnetic Brain Imaging. Springer. 1st edition.

Sekihara, K., Nagarajan, S., Poeppel, D., Marantz, A., Miyashita, Y., 2001. Reconstructing spatio-temporal activities of neural sources using an meg vector beamformer technique. Biomedical Engineering, IEEE Transactions one 48, 760–771.

Sekihara, K., Owen, J., Attias, H., Wipf, D., Nagarajan, S., 2010. Estimating directions of information flow between cortical activities using Phase-Slope Index. 17th International Conference on Biomagnetism Advances in Biomagnetism Biomag2010 IFMBE Proceedings 28, 1199–202.

Sekihara, K., Poeppel, D., Marantz, A., Koizumi, H., Miyashita, Y., 1997. Noise covariance incorporated MEG-MUSIC algorithm: a method for multiple-dipole estimation tolerant of the influence of background brain activity. Biomedical Engineering, IEEE Transactions one 44, 839–847.

Sekihara, K., Sahanib, M., Nagarajan, S., 2005. A simple nonparametric statistical thresholding for meg spatial-filter source reconstruction images. Neuroimage 27, 368–376.

Sekihara, K., Takeuchi, F., Kuriki, S., Koizumi, H., 1994. Reduction of brain noise influence in evoked neuromagnetic source localization using noise spatial correlation. Phys. Med. Biol. 39, 937–946.

Singh, K., Barnes, G., Hillebrand, A., 2003. Group imaging of task-related changes in cortical synchronisation using nonparametric permutation testing. Neuroimage 19, 1589–1601.

Snodgrass, J., Corwin, J., 1988. Pragmatics of measuring recognition memory: applications to dementia and amnesia. J. Exp. Psych.: General 117, 34–50.

Stahlhut, C., Mrup, M., Winther, O., Hansen, L., 2010. Simultaneous EEG source and forward model reconstruction (SOFOMORE) using a hierarchical Bayesian approach. J Sign Process Syst .

Tuomisto, T., Hari, R., Katila, T., Poutanen, T., Varpula, T., 1983. Studies of auditory evoked magnetic and electric responses: Modality specificity and modelling. Il Nuovo Cimento D 2, 471–483.

Ungan P, B.E., 1976. Comparison of wiener filtering and selective averaging of evoked potentials. Electroencephalography and Clinical Neurophysiology 40, 516–520.

Uutela, K., Hämäläinen, M., Salmelin, R., 1998. Global optimization in the localization of neuromagnetic sources. IEEE Trans. Biomed. Eng. 45, 716–723.

Uutela, K., Hämäläinen, M., Somersalo, E., 1999. Visualization of magnetoencephalographic data using minimum current estimates. Neuroimage 10, 173–180.

Vrba, J., Robinson, S., 2001. Signal processing in magnetoencephalography. Methods 25, 249–271.

Waldorp, L., Huizenga, H., Dolan, C., Molenaar, P., 2001. Estimated generalized least squares electromagnetic source analysis based on a parametric noise covariance model [EEG/MEG]. Biomedical Engineering, IEEE Transactions on 6, 737–741.

Wipf, D., Nagarajan, S., 2010. Iterative reweighted l1 and l2 methods for finding sparse solutions. Selected Topics in Signal Processing, IEEE Journal of 4, 317–329.

Wipf, D., Owen, J., Attias, H., Sekihara, K., Nagarajan, S., 2009. Estimating the location and orientation of complex, correlated neural activity using MEG. Advances in Neural Information Processing Systems 21, 1777–1784.

Wipf, D., Owen, J., Attias, H., Sekihara, K., Nagarajan, S., 2010. Robust Bayesian estimation of the location, orientation, and time course of multiple correlated neural sources using MEG. Neuroimage 49, 641–655.

Wipf, D., Ramirez, R., Palmer, J., Makeig, S., Rao, B., 2007. Analysis of empirical Bayesian methods for neuroelectromagnetic source localization. Advances in Neural Information Processing Systems 19, 1505–1512.

Yetik, I.S., Nehorai, A., Muravchik, C.H., Haueisen, J., 2005. Line-source modeling and estimation with magnetoencephalography. IEEE Trans. Biomed. Eng. 52, 839–851.

Zumer, J., Attias, H., Sekihara, K., Nagarajan, S., 2007. A probabilistic algorithm integrating source localization and noise suppression for MEG and EEG data. Neuroimage 37, 102–115.

Zumer, J., Attias, H., Sekihara, K., Nagarajan, S., 2008. Probabilistic algorithms for MEG/EEG source reconstructions using temporal basis functions learned from data. Neuroimage 41, 924–940.

Publishing Agreement

It is the policy of the University to encourage the distribution of all theses and dissertations. Copies of all UCSF theses and dissertations will be routed to the library via the Graduate Division. The library will make all theses and dissertations accessible to the public and will preserve these to the best of their abilities, in perpetuity.

I hereby grant permission to the Graduate Division of the University of California, San Francisco to release copies of my thesis or dissertation to the Campus Library to provide access and preservation, in whole or in part, in perpetuity.

9/19/11

Author Signature Date