

UCSF

UC San Francisco Electronic Theses and Dissertations

Title

Probabilistic Methods for Neural Source Reconstruction from MEG Data

Permalink

<https://escholarship.org/uc/item/9ft9m8ht>

Author

Zumer, Johanna Margarete

Publication Date

2007-11-15

Peer reviewed|Thesis/dissertation

Probabilistic Methods for Neural Source Reconstruction from MEG Data

by

Johanna Margarete Zumer

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

BIOENGINEERING

in the

GRADUATE DIVISIONS

of the

UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

and the

UNIVERSITY OF CALIFORNIA, BERKELEY

UMI Number: 3289309



UMI Microform 3289309

Copyright 2008 by ProQuest Information and Learning Company.
All rights reserved. This microform edition is protected against
unauthorized copying under Title 17, United States Code.

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

Copyright © 2007

by

Johanna Margarete Zumer

Abstract

Probabilistic Methods for Neural Source Reconstruction from MEG Data

by

Johanna Margarete Zumer

The field of brain imaging has exploded in the past two decades due to new technological developments mostly on the hardware/acquisition end. Consequently, neuroscientists have been excited to study all aspects of brain function and the general public have been excited to hear what results come from such studies. Magnetoencephalography (MEG) is a brain imaging method which passively detects the naturally occurring magnetic fields outside the head resulting from neural cells' activity in the brain. The MEG sensors record direct neural activity with millisecond resolution. However, there is inherently no unique solution to determining where exactly in the brain the neural activity was located that produced the brainwaves measured by the sensors. With certain general assumptions, a reasonable estimate of the location can be made. However, measurement noise and other artifacts, including heartbeat and eyeblinks, swamp the signals of interest, making localization nearly impossible. In this dissertation, two novel methods are proposed which improve neural source estimate by removing the effects of noise and interference. These new methods are tested against standard meth-

ods for both simulated and real data and show improved performance. These methods are further tested on real data obtained from primates with the ultimate goal of using electrophysiological data from these primates to compare the MEG localization with the true location. Finally, an example is shown of one way to combine neural activity measured by MEG with a method for measuring white-matter anatomical connections obtained by diffusion tensor imaging.

Acknowledgments

I would like to acknowledge many people for providing support and inspiration through the many stages of this dissertation work.

Most importantly, I would like to thank my Ph.D. advisor, Dr. Srikantan Nagarajan, for outstanding guidance. His extensive knowledge in both statistical signal processing and neuroscience is awe-inspiring and has been instrumental in guiding my research interests. I appreciate his nudging me into studying probabilistic graphical models, which I would not have pursued nearly as much otherwise, but are a very useful and general tool. I also wish to thank him for his generosity in sending me to many international science conferences at which I have presented my work and established contacts, both critical to starting a career in academia.

I would next like to thank Dr. Elizabeth Disbrow for taking the time to be a secondary advisor, for providing guidance in learning about somatosensory cortex, and for teaching me how to perform MEG and MRI experiments in monkeys, along with the many details of veterinary and anesthesia knowledge associated

with the experiments. Dr. Hagai Attias, an unofficial secondary advisor, has been very patient with me while tutoring me on probabilistic graphical models. Dr. Kensuke Sekihara, another unofficial secondary advisor, has provided much knowledge and advice about MEG inverse methods.

I greatly appreciated the instruction, advice and support from my qualifying exam committee: Dr. Sarah Nelson, Dr. Mark D'Esposito, Dr. Bob Knight, Dr. Tom Ferree and Dr. Heidi Kirsch. In addition to chairing the exam committee, Dr. Nelson provided guidance when I was having trouble finding an advisor and suggested I consider Dr. Nagarajan's lab. In addition to advice while I prepared for the qualifying exam, Dr. Ferree provided excellent background knowledge on the physical principles of MEG and EEG in his class on Electromagnetic Neuroimaging.

I also greatly appreciate the guidance of my dissertation committee: Dr. Sri Nagarajan (chair), Dr. Mark D'Esposito, Dr. Heidi Kirsch, Dr. Liz Disbrow, and Dr. John Kornak, especially as my deadline to finish approached.

Many members and affiliates of Dr. Nagarajan's lab have made my experiences and education both fun and exciting. I couldn't imagine my graduate years without Sarang Dalal, fellow Bioengineering graduate student in the lab, office-mate, and great friend. That he always seemed at least one step ahead of me in many ideas or projects certainly pushed me to learn more and try harder. I can't thank him enough for all my questions he's answered, including the ones I should

have researched on Google instead of spending his time, and especially the ones I couldn't have found a solution easily elsewhere. His excitement and foresight in developing the NUTMEG software package was contagious and further helped me to learn programming skills.

Kenny Hild, a postdoc and officemate, also contributed much to my intellectual growth and enjoyable years in the lab, especially through the many early-morning broad-topic conversations and statistics tutorials we shared. Technicians Anne Findlay and Susanne Homna have provided invaluable help in experiment setup, data acquisition and data processing, especially with all aspects of the monkey experiments. I would like to thank the anonymous human and monkey subjects who gave their time for the data collection.

Other lab members and affiliates have all provided useful and interesting guidance, assistance and/or conversation: Drs. Leah Krubitzer, Roland Henry, Dave McGonigle, Leighton Hinkley, Ilana Hairston, Virginie Van Wassenhove, Adrian Guggisberg, Zhao Zhu, Marianne Lowenthal, David Wipf, Julia Owen, John Houde, Jeff Padberg, and Jeff Berman. Bioengineering students from other labs that have helped me along the way are: Drs. Felice Sun, Dan Handwerker, and Stephen David.

Funding for my initial two years of graduate school was through the NIH training grant provided to the UCSF/Berkeley Joint Graduate Group in Bioengineering. Funding for projects in Dr. Nagarajan's lab was provided by R01 NS44590,

R01 DC4855 and R01 DC6435.

Prior to my arrival in graduate school, several professors encouraged and developed my interest in brain imaging as mentors: Dr. Todd Constable (Yale summer project), Dr. John Gore (Yale undergraduate thesis) and Dr. Armando Manduca (Mayo Clinic summer project). Going even further back to high school, I strongly thank teachers Mike Andrews and Jack Marino for excellent training and encouragement in math and physics.

I wish to thank my roommates and fellow Bioengineering students, Janine Lupo and Jill Ulrich, whom I lived with for five and four years, respectively. Their friendship has meant so much to me, from sympathizing on graduate student issues, to going for runs through Golden Gate Park, going out on weekends, and sharing many laughs.

Of course I am extremely thankful to my parents Andrew and Sally for continued support, advice and encouragement, especially when I felt I didn't need it. My sister Asta has always provided a good laugh and a fresh perspective, and kept me on my toes. I also wish to thank my aunts Dorothy and Theresa for their kindness and support through the years.

Lastly, I am beyond grateful to my boyfriend Ben Inglis for the many flavors of support he has provided over the past four years. He has advised me through many major decisions about my work and life, down to the very minor details of phrasing of emails. He encouraged and helped me while I studied for my qual-

ifying exam and was by my side when I gave oral presentations at international conferences. He has proofread this dissertation and drafts of other papers. He helped fix the broken C key on my laptop, upon which most of this dissertation has been written, and reminded me not to split infinitives. He has been my best friend, confidant and source of relaxation.

Part of the text of this dissertation has appeared in accepted publications. Parts of Chapters 2 and 3 on the SAKETINI method were first presented in a NIPS conference paper (Zumer et al., 2007a) and later a full paper in NeuroImage (Zumer et al., 2007b). Parts of Chapters 2 and 3 on the NSEFALoc method were presented in an ISBI conference paper (Zumer et al., 2006). All of these publications were co-authored by my research director Dr. Srikantan Nagarajan, as well as Dr. Hagai Attias and Dr. Kensuke Sekihara.

Contents

Acknowledgments	v
List of Tables	xiii
List of Figures	xiv
1 Introduction	1
1.1 Magnetoencephalography: Overview	1
1.2 MEG versus EEG	4
1.3 Neural basis of MEG/EEG	6
1.4 Magnetoencephalography: Forward Solution	8
1.5 Magnetoencephalography: Inverse Problem	9
1.5.1 Non-adaptive weights including sLORETA	12
1.5.2 Adaptive weights including MVAB	15
1.5.3 Noise and interference	18
1.6 Background Mathematics	20
1.6.1 Probabilistic methods for the inverse problem	22
1.7 Graphical Models	23
2 Mathematics of Probabilistic Models	25
2.1 Overview	27
2.2 Factor Analysis	27
2.2.1 Dimension reduction and denoising methods	27
2.2.2 Expectation-Maximization	31
2.2.3 Factor Analysis estimates	32
2.2.4 Variational Bayesian Expectation-Maximization	33
2.2.5 Variational Bayesian Factor Analysis (VBFA)	36
2.3 Stimulus Evoked Factor Analysis (SEFA)	41
2.4 SAKETINI	48
2.4.1 SAKETINI: Full model overview	49

2.4.2	SAKETINI: Separation of background sources learned from pre-stimulus data	51
2.4.3	SAKETINI: Localizing evoked factors learned from post-stimulus data	53
2.4.4	Relationship of SAKETINI to MVAB	59
2.5	NSEFALoc	60
2.5.1	NSEFALoc1: Spatial weighting of temporal basis functions	62
2.5.2	NSEFALoc2: Spatial weighting of temporal basis functions	67
2.6	Conclusion	71
3	Performance in Simulations and Real Data	72
3.1	Overview	72
3.2	MEG Simulations	74
3.2.1	Localization accuracy of a single source	77
3.2.2	Multiple sources, including correlated sources	78
3.2.3	Simulation: average performance	83
3.2.4	Simulation: effects of number of sensors and time points	86
3.2.5	Model Order and Initialization	90
3.3	Real MEG Data	97
3.3.1	Bilateral auditory data	97
3.3.2	High noise somatosensory data	99
3.4	Discussion	100
4	Comparison of Methods and Consistency of Results	104
4.1	Overview	104
4.2	Somatosensory and Auditory Evoked MEG data in Human and Nonhuman Primates	106
4.2.1	Data acquisition and processing	106
4.2.2	Sensor data results	114
4.3	Comparison of Methods on Full Data	117
4.3.1	Analysis of method comparison on full data	117
4.3.2	Results and discussion of method comparison on full data	121
4.4	Comparison of Methods on Focused Data	126
4.4.1	Analysis of method comparison on focused data	126
4.4.2	Results and discussion of method comparison on focused data	129
4.5	Conclusion	134
5	Separability of functionally-defined regions and networks through analysis of MEG data and MEG-seeded DTI tractography	136
5.1	Overview	136
5.2	MEG latencies across stimulation type and species	138
5.2.1	MEG latencies in humans	139

5.2.2	MEG latencies in monkeys	140
5.2.3	Discussion of MEG latencies	141
5.3	Spatial Separability of MEG MVAB Results	145
5.3.1	Spatial Separability of human MEG MVAB results	145
5.3.2	Spatial separability of monkey MEG MVAB results	147
5.3.3	Discussion of MEG MVAB results and neuroanatomical variation	152
5.4	DTI Tractography using MEG MVAB Localization	155
5.4.1	Tractography overview	155
5.4.2	DTI acquisition and processing	158
5.4.3	DTI results	161
5.4.4	Discussion of MEG-seeded tractography	169
5.5	Conclusion	175
6	Conclusion	177
6.1	Overview	177
6.2	Future Directions	178
6.2.1	Direct extensions of models	178
6.2.2	Improved model order selection	179
6.2.3	Time-frequency analysis	180
6.2.4	Improved head modeling and forward field errors	182
6.2.5	Spatial smoothness and tomographic reconstruction	184
6.2.6	Combine MEG data with other modalities	184
6.2.7	Networks and causal models	186
6.2.8	Relating scales of measurement	187
A	Appendix of Useful Matrix Relationships	189

List of Tables

5.1	MEG latencies in monkeys and humans	141
-----	---	-----

List of Figures

2.1	SAKETINI graphical model	51
2.2	NSEFALoc graphical model	63
3.1	Localization error of all methods for a single source	76
3.2	Performance of all methods in example simulations	80
3.3	Performance of all methods in example simulations with correlated sources	81
3.4	Localization and time course estimation accuracy for all methods	86
3.5	All methods' performance as a function of number of data sensors	88
3.6	All methods' performance as a function of number of data time points	90
3.7	Effect of hyperparameter dimension selection for SEFA	92
3.8	Effect of type of basis functions for NSEFALoc	93
3.9	Effect of hyperparameters and dimension choice on SAKETINI and NSEFALoc2	95
3.10	All methods' performance as a function of dimension reduction	96
3.11	All methods applied to real auditory evoked data	98
3.12	All methods applied to high noise real somatosensory data	101
4.1	Averaged MEG sensor data from human subject 4	116
4.2	Averaged MEG sensor data from monkey 1	118
4.3	A' value for human MEG data under digit, lip and auditory stimulation	122
4.4	Digit localization results from all methods on subject H4	124
4.5	A' value for monkey MEG data under digit, lip and auditory stimulation	125
4.6	Comparing A' for all 275 versus reduced 37 sensors in two monkeys	130
4.7	Percentage threshold comparison	131
4.8	Peak maximum comparison	132
4.9	Single Cluster Index comparison	133
4.10	Average localization errors of all methods	134
5.1	Beamformer source localization of human MEG dataset	146
5.2	Average MVAB localization results in four humans	148

5.3	Beamformer source localization of MEG dataset from monkey 1	149
5.4	Average MVAB localization results in four monkeys	153
5.5	DTI tractography results in humans	162
5.6	DTI tractography results in monkeys	165
5.7	DTI tractography summary in monkeys and humans	168

Chapter 1

Introduction

1.1 Magnetoencephalography: Overview

The development and use of noninvasive methods for imaging brain function have grown rapidly in the past few decades. Magnetoencephalography (MEG) is one such method that detects the tiny magnetic fields outside a person's head that naturally occur as a result of neural activity. Together with electroencephalography (EEG), which measures the corresponding voltage potentials at the scalp due to similar neural activity, these techniques are an effective way to directly measure and map neural function. The tiny magnetic fields measured by MEG are roughly one billion times smaller than the earth's magnetic field ($\sim 50\mu\text{T}$). These magnetic fields are produced by a simple concept in physics: electric current flowing through a wire produces a magnetic field circling around the wire, in the direction of the *right-hand rule*; likewise, ionic current flowing through a stimulated

neuron produces magnetic fields circling the cell.

Several important types of neural cells in the brain, specifically pyramidal neurons, are structurally shaped as branching trees. The neuron is excited by other neurons via a small current flowing inward to the cell starting from the far branches. This current flows down the branches (dendrites) and through the trunk (apical dendrite) to the soma (cell body), at which point the cell computes if the signal is strong enough to be worth passing on to other cells. The current flowing through one neuron would produce a magnetic field roughly 250fT near the neuron (Blagoev et al., 2007), but only 0.008aT outside the head since the strength of the magnetic field weakens as the inverse of the distance squared. SQUIDS (Superconducting Quantum Interference Devices), the very sensitive detectors used in MEG, have inherent noise around $10\text{fT}/\sqrt{\text{Hz}}$. It so happens that pyramidal neurons are arranged with their trunks in parallel throughout all areas of cortex, so that the magnetic fields produced by current flowing through them add, rather than cancel if they were arranged randomly. Additionally, neurons nearby each other tend to have similar stimulus preferences and functions so that they tend to be simultaneously active. This spatial and temporal summation of activity over many neurons in a small patch (roughly $1\text{-}8\text{mm}^3$) produce magnetic fields approximately 10-100fT outside the head: strong enough to be measured by SQUIDS.

Mapping the spatiotemporal neural activity of the entire brain is an impor-

tant problem in basic neuroscience research. It is also clinically important for patients with brain tumors and epilepsy, both in localizing regions important for cognitive function and for identifying epileptogenic brain regions. Such brain mapping procedures are useful to guide neurosurgical planning, navigation, and resection.

Many noninvasive techniques have emerged for functional brain mapping, such as functional magnetic resonance imaging (fMRI) and MEG/EEG. Although fMRI is the most popular method for functional brain imaging with high spatial resolution, it suffers from poor temporal resolution since it measures blood oxygenation level-dependent (BOLD) signals with fluctuations in the order of seconds. These BOLD signals are also indirect measures of neural activity which might not accurately reflect neural activity, especially in regions of altered vasculature. However, dynamic neuronal activity has fluctuations in the submillisecond time-scale that can only be directly measured with MEG and/or EEG.

This dissertation focuses on methods improving the ability to estimate the underlying neural activity from measured MEG data. The rest of this first chapter provides an introduction to MEG, including the signal propagation from the neurons to the MEG sensors (forward solution), common methods for estimating neural activity from the sensor data (inverse problem), and common problems for these methods including noise in the data. Chapter 2 presents two new algorithms for estimating neural activity which incorporate removal of

noise/interference. These methods are developed within a probabilistic graphical model framework, which is also introduced in Chapter 1. In Chapter 3, the performance of these new probabilistic methods is examined in many simulated datasets and some example real datasets. They are shown to be an improvement over existing methods, in estimation of both neural source localization and neural source dynamics.

Chapter 4 extends the method comparison to a larger set of real data in which the approximate neural location is known, multiple brain areas are active, and/or large artifacts are present. Data from healthy humans are examined, as well as data collected from anesthetized primates as part of a larger study developing a primate model for MEG validation studies. In Chapter 5, the feasibility of separating brain areas in primates using MEG is demonstrated. The spatial separability is further strengthened by the distinct neuroanatomical networks to which they are connected, as measured by diffusion tensor MRI. The relation of the monkey to human data is also discussed. Finally, Chapter 6 discusses future directions for the models proposed in Chapter 2, potential problems of the methods, and future directions for MEG research in general.

1.2 MEG versus EEG

Since MEG and EEG are measuring such similar phenomena, why should one be used over the other? The number of experiments using EEG rather than

MEG is much larger for several reasons. The technology for EEG is much simpler: electrodes plus conducting gel and amplifiers allowed the first recordings in the 1920s, while the very sensitive SQUID detectors enabling MEG were not invented until the 1960s. The cost of a commercial MEG system is roughly 50 times that of EEG. EEG is certainly portable, whereas the dewar of liquid helium enclosing the SQUIDs stays in the laboratory. Once the EEG electrodes are placed on the scalp, the subject can move their head as needed, whereas a subject in the MEG must hold still to within a few millimeters relative to sensor array.

Practical considerations aside, MEG and EEG have different sensitivities to the neural activity: EEG depends greatly on the estimate of tissue conductivity to account for the current distortion across tissue boundary layers (from brain to skull to scalp), whereas these conductivity values do not matter for MEG. Since EEG measures voltage potentials, these potentials always have to be relative to some reference electrode, the best placement of which is not clear; MEG does not have this reference sensor problem. EEG is more sensitive to a radially-oriented source than a tangential one at the same depth, which may be good or bad depending on the neural sources of interest. On the other hand, if the head is a perfect sphere, MEG is totally blind to radially-oriented sources; however, due to cortical surface folding and the fact that the head is not a perfect sphere, some near-radial sources can be detected. Mainly due to the lack of smearing of voltage potential across the different conductivity layers, MEG typically has a larger es-

estimated spatial precision in its estimate of source location (Leahy et al., 1998). Therefore, the source reconstruction methods presented here have focused on MEG data with the hope of later extension to the more difficult case of EEG source localization.

1.3 Neural basis of MEG/EEG

Changes in stimulus, attentional state, or goal-directed actions certainly alter neural activity. However, there is still debate as to what aspect of neural activity best represents the brain's way of encoding these changes. Traditionally, the rate of action potential spikes (the output of neurons in the axon) recorded with an electrode inside an animal brain was studied as the main correlate of neural activity. Local field potential recordings (also using an electrode inside the brain) represent synaptic- and/or dendritic-based electric field changes, more related to the input a neuron receives than the output decision it makes. Through the use of current source density recordings, field reversal changes when crossing over a source, and other studies (Okada et al., 1999), it is generally accepted that the local field potentials relate much more closely to the signals measured in MEG and EEG than spiking activity.

An electrode placed in a primary sensory cortex tuned to a particular stimulus feature (such as location of a touch on the skin) will record a burst of spiking activity and an increase in field potentials within milliseconds after that stimulus

is given. Likewise, this same stimulus will produce waveforms of similar time course in MEG data, although slightly different due to the greater spatial area of activity being summed. Unfortunately, these *stimulus evoked* field changes measured by the MEG sensors are as small as all the other ongoing brain activity unrelated to the stimulus. But since this ongoing brain activity has no temporal relation to the exact timing of the stimulus onset, an average of many trials of stimulus presentation will reveal the activity related to the stimulus since the other activity will average out towards zero, relative to a pre-stimulus baseline. Most MEG experiments have focused on studying the amplitudes and shapes of these evoked fields, and it is this paradigm of experiment that is the focus of the source reconstruction methods presented in this dissertation.

That the neural activity measured by EEG and MEG oscillates at certain frequencies has been known since the first EEG recordings in the 1920s. An increase in the so-called alpha frequency (roughly 8-12Hz) occurs when a subject closes his eyes. Higher frequencies tend to be more associated with active brain function, such as beta (12-30Hz) with motor actions and gamma (30-60Hz) (Pulvermüller et al., 1995) and high-gamma (60-200Hz) with higher cognitive processes (Crone et al., 2006). Lower frequencies such as theta (4-8Hz) and delta (below 4Hz) are related to drowsiness, sleep, and some clinical disorders. These discrete frequency boundaries are approximate and their relations to brain-state mentioned here are a broad generalization.

An interesting alternative to stimulus-evoked analysis is to study the changes in power of these various frequency bands that result from a stimulus, action or attentional-shift. One hypothesis is that the power at a given frequency band for individual neurons is constant, but that when many neurons synchronize their phase, the average power in that frequency band increases purely due to the individual phases changing, not an increase in the power/activity of each neuron. This phenomenon is termed *event-related (de)synchronization* (ERS/ERD). ERS/ERD studies have highlighted brain activity related to stimuli/actions that would not have been seen by an evoked trial-average method alone. Most likely, some combination of both power increases/decreases and phase changes underlie the brain signals recorded in MEG, and these changes are always relative to some “baseline” state which is certainly not rest/inactivity.

1.4 Magnetoencephalography: Forward Solution

The first step of computing what magnetic field change would be expected at a SQUID sensor caused by a dipole current at a given location in the brain is termed the *forward solution*. The magnetic field does not just decrease as a simple function of distance from the neural current source as it would if the current source were in isolation. Instead, the source is inside the brain, where its current is partially contained by cerebrospinal fluid and skull, and fully contained by the scalp. The head is therefore a conducting volume for which shape and location

from the sensors must be taken into account.

The $K \times Q$ forward lead field matrix F^r represents the physical (and linear) relationship between a dipole source at voxel location r for each dipole orientation $q = 1 : Q$ along a coordinate basis and its influence on sensor $k = 1 : K$ (Sarvas, 1987). The lead field F^r is calculated from knowing the geometry of the source location to the sensor location, as well as assuming the type of conducting medium in which the source lies. The human head is most commonly approximated as a single sphere volume conductor for MEG data; for EEG data, a more detailed model taking into account the tissue conductance boundaries is usually needed. In the most general case, $Q = 3$ for all three possible directions of coordinate bases of a source dipole. In the case of the single-shell sphere, the radial component of source dipoles contribute nothing to MEG sensors, thus $Q = 2$ if a spherical coordinate system is used. If there is knowledge of a subject's cortical anatomy, the source may be constrained to be perpendicular to the gray matter surface, thus $Q = 1$. Throughout the rest of this dissertation, the single-shell model with $Q = 2$ is used for both simulations and real data from MEG.

1.5 Magnetoencephalography: Inverse Problem

The calculation of the magnetic field changes at a sensor due to a specific neural current via knowing the forward field is a straightforward computation given the head geometry. (Obtaining the exact head geometry in practice, however, is

not trivial). Regardless, the *inverse problem* of computing the neural current, given knowledge of the external magnetic fields even with the most precise knowledge of head geometry, is ill-posed: no unique solution is guaranteed. Additional assumptions must be made to compute a single solution.

Continuing the notation above, the mathematical formulation of MEG data can be written:

$$y_n = \sum_r F^r s_{rn} + v_n \quad (1.1)$$

Here the subscript n denotes the time index ranging from $1 : N$, y_n is the MEG sensor data of dimension $K \times 1$, the forward field F^r is as described above, s_{rn} is the source activity at voxel r and dimension $Q \times 1$, and finally the sensor noise v_n is dimension $K \times 1$. The sum over r indicates that all voxels with nonzero source activity contribute to the sensor data. To be most precise, the sum should really be an integral over the whole conducting volume, but in practice this is the same in the limit that the voxel spacing becomes small.

Historically, the most common source reconstruction method is the equivalent current dipole (ECD) method. ECD assumes that there are few point locations with active neural currents and solves for the best location and orientation of these few dipoles by minimizing the error between the projected estimate and the data. This is a nonlinear search over the parameter space and is highly over-determined since the number of sensors is usually 50-300. When there is only one true active source, ECD usually performs well. However, for two or more sources,

the results can be highly dependent on initial location and on the total number of dipoles being fit. Often, trying to fit three or more dipoles will fit noise rather than true source locations.

Source localization algorithms can be broadly classified as parametric or tomographic, making different assumptions to overcome the ill-posed inverse problem. Parametric methods, including ECD, assume knowledge about the number of sources and their approximate locations. A single dipolar source can be localized well, but ECD techniques poorly describe multiple sources or sources with large spatial extent. Parametric methods could also involve multipole, line source, and/or patch reconstructions (Schmidt et al., 1999; Yetik et al., 2005; Jerbi et al., 2002). Many parametric methods involve computation of noise and/or signal covariances from data where several Bayesian approaches have been reported (Mosher and Leahy, 1998; Jun et al., 2005, 2006; Huizenga et al., 2002; Waldorp et al., 2002; de Munck et al., 2004; Bijma et al., 2005).

Tomographic methods reconstruct an estimate of source activity at grid points across the whole brain, either simultaneously or through scanning. Tomographic methods could also be classified as adaptive (data-dependent) or non-adaptive (data-independent). Of many tomographic algorithms, the adaptive beamformer has been shown to have the best spatial resolution and zero localization bias (Sekihara et al., 2005, 2002a; Küçükaltun-Yildirim et al., 2006); however, it is sensitive to temporally correlated sources.

Therefore, the main alternative source reconstruction methods estimate the source activity at each voxel across the whole brain, termed distributive source reconstruction methods. When the brain volume is discretized to a reasonable level (usually between 3-5mm edge cubes), the total number of voxels R can range from 5,000-20,000. This is now a highly under-determined system; however, if the number of active sources is less than the number of sensors, then the problem returns to being over-determined.

The basic equation for distributed source reconstructions is to solve for a weight w that can be applied to the sensor data y to produce the estimated source activity \hat{s} at a particular voxel r :

$$\hat{s}_{rn} = w_r^T y_n \quad (1.2)$$

The next two sections discuss the two main ways to compute the weight: non-adaptive weights depend solely on the forward field and do not change with the data recorded, and adaptive weights are a function of both the forward field and the data. This background material is summarized from work published by Dr. Sekihara and Dr. Nagarajan (Sekihara et al., 2005, 2002a,b, 2001), as well as outlined in their upcoming textbook to be published by Springer-Verlag in 2008.

1.5.1 Non-adaptive weights including sLORETA

If 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, but rather may have a left- or right-sided inverse, usually approximated by a pseudo-inverse. Similar to solving for the best ECD discussed above, the source estimate \hat{s} that minimizes the linear least-squares error is the one that uses the Moore-Penrose pseudo-inverse as the weight:

$$\hat{s}_{rn} = F^{r+} y_n = F^{rT} [\mathbf{F}\mathbf{F}^T]^{-1} y_n \quad (1.3)$$

Here, the $+$ indicates the pseudo-inverse, and \mathbf{F} indicates the full composite forward field for all voxels (size $K \times QR$). The quantity $\mathbf{F}\mathbf{F}^T$ is also called the Gram matrix G and gives an indication of the spatial correlation or overlap between the sensitivity profile of the sensors. Since nearby sensors have quite similar sensitivities, this Gram matrix is close to singular and so is not stable to invert. Hence, Tikhonov regularization is needed, which involves adding a small amount of noise to the diagonal of the Gram matrix prior to inversion. This regularization can be arrived at in an alternate way, by minimizing a sum of squared-error plus the norm of the source weighted by a Lagrange multiplier, termed the *min-norm* solution.

$$\hat{s}_{rn} = F^{rT} [\mathbf{F}\mathbf{F}^T + \epsilon I]^{-1} y_n \quad (1.4)$$

If the source covariance Γ and sensor covariance Υ are assumed to be known, they can be included in the minimization function to further improve the estimate:

$$\hat{s}_{rn} = (\Gamma F^T)^r [\mathbf{F}\Gamma\mathbf{F}^T + \epsilon\Upsilon]^{-1} y_n \quad (1.5)$$

If Γ and Υ are determined before observing the data, this method remains non-adaptive. But if they are estimated using the current dataset, then this method appears more like an adaptive method. All of the above three equations can be expressed as:

$$\hat{s}_{rn} = \tilde{F}^{rT} \tilde{G}^{-1} y_n; \quad w_r^T = \tilde{F}^{rT} \tilde{G}^{-1} \quad (1.6)$$

As a side note, all of these min-norm variations that use the Gram matrix depend on its inclusion of all possible voxels that could contain source activity. Thus, even if the source activity at only a particular voxel is of interest, the forward field for the whole volume must be computed. It will be seen that the weights used in the adaptive methods presented in the next subsection and probabilistic methods presented in Chapter 2 only rely on the forward field for the voxel of interest.

Dale et al. (2000) proposed to modify 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 filter is:

$$w_r^T = \frac{F^{rT} G^{-1}}{\sqrt{\text{tr}(F^{rT} G^{-2} F^r)}} \quad (1.7)$$

Pascual-Marqui (2002) proposed a modification to the above weight-normalized form, called sLORETA (standardized low-resolution electromagnetic tomography), where the weight matrix is:

$$w_r^T = \frac{F^{rT} G^{-1}}{\sqrt{F^{rT} G^{-1} F^r}} \quad (1.8)$$

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 $w^T G w = 1$.

Both sLORETA and the weight-normalized min-norm filter can be modified so that the Gram matrix is regularized. A large amount of regularization leads to a simplified form approximating the spatial-matched filter $w_r^T = \frac{F^r}{\|F^r\|}$.

1.5.2 Adaptive weights including MVAB

The class of minimum variance adaptive beamformer (MVAB) filters rely on assumptions 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 $w_r^T R_{yy} w_r$ subject to various constraints, where R_{yy} is the data covariance matrix. 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 $w_r^T F^r = 1$ yields the unit-gain-constraint minimum-variance adaptive

beamformer (UGC-MVAB):

$$w_r^T = \frac{F^{rT} R_{yy}^{-1}}{F^{rT} R_{yy}^{-1} F^r} \quad (1.9)$$

While the unit-gain constraint seems most sensible at first, the gain for a source at the sphere-center will be infinite for a forward model computed using a single-sphere approximation since the forward field at the sphere center is zero. Thus, the norm of the forward field can be used as a constraint $w_r^T F^r = \|F^r\|$ to form the array-gain-constrained MVAB (AGC-MVAB):

$$w_r^T = \frac{\tilde{F}^{rT} R_{yy}^{-1}}{\tilde{F}^{rT} R_{yy}^{-1} \tilde{F}^r} \quad (1.10)$$

where $\tilde{F}^r = f^r / \|F^r\|$, so it also can be called the forward-field-normalized MVAB.

The source power estimate is $\hat{s}_r n = (\tilde{F}^{rT} R_{yy}^{-1} \tilde{F}^r)^{-1}$. The unit-noise-gain constraint of $w_r^T w_r = 1$ can also be used to give the UNG-MVAB:

$$w_r^T = \frac{F^{rT} R_{yy}^{-1}}{\sqrt{(F^{rT} R_{yy}^{-2} F^r)}} \quad (1.11)$$

If sources are temporally uncorrelated, then the MVAB estimate at a true source location will have zero influence from other sources and non-source locations will have zero power (except for very nearby sources whose resolution kernel has significant spread at that location). However, if two sources at locations p and q are temporally correlated with correlation value μ , the source estimate and power at p are:

$$\hat{s}_p = s_p - (\sigma_p / \sigma_q) \mu s_q \quad (1.12)$$

$$\hat{s}_p^2 = (1 - \mu^2) \sigma_p^2 \quad (1.13)$$

where σ is the true source power. For more than two correlated sources, the source estimate is:

$$\hat{s}_p = s_p + \sum_R \frac{R_s^{-1}(pr)}{R_s^{-1}(pp)} s_r \quad (1.14)$$

This reduction in source power and cross-talk of source time series estimate can be very problematic for certain real data situations, such as bilateral auditory cortex activation. Furthermore, the issue of the difference between "scalar" MVAB (weight computed for single known source orientation) or "vector" MVAB (for estimating either 2 or 3 directions) has been ignored until now. Solving for the x, y, z component of a given source means that those three components are 100% correlated with each other, assuming the source does not rotate over time or is not composed of multiple sources spatially overlapping.

These correlated source problems can be minimized by explicitly nulling the influence of other voxels onto the current voxel in computation of the weight. The linear-constrained minimum variance beamformer uses the additional constraint of $w_r^T l_q = 0$ for all known locations $q \neq r$ that are correlated with r . If the number of voxels q is large, a signal-space reduction can be used to estimate the main contribution from the region (Dalal et al., 2006). The vector MVAB imposes $w_{r_x}^T l_{r_y} = 0$ and $w_{r_x}^T l_{r_z} = 0$ to compute the x-component and similar constraints are used for the other orientations.

One final modification to the adaptive methods will be discussed here. Since the number of active sources can be reasonably assumed to be less than the num-

ber of sensors, the data covariance matrix can be partitioned into a "signal-space" portion and "noise-space" portion by using singular-value decomposition (SVD):

$$R_{yy} = E_s^T \Lambda_s E_s + E_n^T \Lambda_n E_n \quad (1.15)$$

where Λ_s are the eigenvalues associated with P active signals and Λ_n are eigenvalues associated with $K - P$ noise components. The lead field at true source locations is orthogonal to the space spanned by the noise eigenvectors. Thus, only the signal subspace of the data covariance remains if the weight is premultiplied by the projection $E_s E_s^T$:

$$w_s^T = w_r^T E_s E_s^T = \frac{F^T E_s^T \Lambda_s^{-1} E_s}{\varepsilon} \quad (1.16)$$

where ε corresponds to the various denominators used in Eqs. 1.9, 1.10, and 1.11. Furthermore, even if there are slight estimation errors in the forward field due to incorrect head model, the use of only the signal-space portion further reduces the contribution of the error and the data noise. It is this eigenspace beamformer with normalized forward-field that is referred to simply as "MVAB" for the rest of this dissertation, and whose performance is compared to the new methods proposed in Chapter 2.

1.5.3 Noise and interference

All existing methods for brain source localization are hampered by the many sources of noise present in MEG/EEG data. The magnitudes of the stimulus-evoked neural sources are on the order of noise on a single trial, and so typically

50-200 averaged trials are needed in order to clearly distinguish the sources above noise. This limits the type of cognitive questions that can be answered, and is prohibitive for examining processes such as learning that can occur over just a few trials. Acquiring sufficient trials for successful averaging is time-consuming, making it difficult for a subject or patient to hold still or pay attention through the duration of the experiment.

Gaussian thermal noise or Gaussian electrical noise are present at the MEG or EEG sensors themselves. Background room interference such as from powerlines and electronic equipment can be problematic. Biological noise such as heartbeat, eyeblink, or other muscle artifact can also be present. Ongoing brain activity itself, including the drowsy-state alpha (~ 10 Hz) rhythm, can drown out evoked brain sources. Finally, many localization algorithms have difficulty in separating neural sources of interest that have temporally overlapping activity (Sekihara et al., 2002a).

The basic min-norm method is biased towards the sensor array even with no noise, while min-variance methods have no bias. The min-norm method sLORETA has no localization bias when there is one active source and no noise, but multiple sources or added sensor noise can cause localization bias (Sekihara et al., 2005). The AGC-MVAB and UNG-MVAB methods can be shown to have no localization bias even when the input SNR is very low (Sekihara et al., 2005). Examples of performance in simulation and real examples of the MVAB

and sLORETA will be shown in Chapters 3 and 4.

1.6 Background Mathematics

A few mathematical rules and definitions used in Chapter 2 are presented here. The values that a variable can take are often assumed to be drawn from some distribution. For continuous variables, the Gaussian distribution (also known as the normal distribution or the “bell curve”) is the most common one assumed for many reasons. In any situation where random specific instances of the variable (known as a sample) tend to cluster around one specific value and large deviations from this mean are rare, then a Gaussian distribution is a good approximation, especially when the number of samples is large (due to the law of large numbers). The form of a Gaussian distribution is described by the probability that a specific value of the variable will occur, given the mean μ and precision Γ :

$$p(x_n) = \mathcal{N}(x_n|\mu, \Gamma) = \frac{|\Gamma|^{1/2}}{2\pi} \exp \left[-\frac{1}{2}(x_n - \mu)^T \Gamma (x_n - \mu) \right] \quad (1.17)$$

Like any probability distribution, the area under the curve for all values of x_n is normalized to one. The *precision* is the inverse covariance and is used throughout this thesis rather than the more commonly specified covariance. The above expression is the multivariate Gaussian and describes the distribution for a vector x_n . Hence the precision is, in general, a full-rank matrix if the individual elements of the vector x_n covary with each other.

Note there are only two parameters needed to specify the distribution: the mean and precision. For a Gaussian distribution, the estimated mean and covariance are the *sufficient statistics* (quantities that, once known, the parameters μ and Γ give no new information). They are given by:

$$\bar{x} = \frac{1}{N} \sum_{n=1}^N x_n \quad (1.18)$$

$$R_{xx} = E[(x_n - \bar{x})(x_n - \bar{x})^T] \quad (1.19)$$

In general, the mean of a quantity approximated by a Gaussian distribution can be estimated by the mean of the samples measured. However, it is often the case that the data that is measured is an indirect result of the underlying quantity of interest. This is just another way of stating the inverse problem. This can be described in a probabilistic framework through conditional probabilities and Bayes rule. Bayes rule relates the probabilities of two different variables with each other. First, think about the MEG data y that could be measured from any subject, in any head position, doing any cognitive task: the expected mean of this data is termed the *a priori* probability of y . Now, think about the expected mean of the data y given that the head position and functional task (and thereby location of expected neural activity) are known. If there is a direct/systematic way in which the value of one quantity influences/causes/predicts the value of another, this conditional probability is termed the likelihood. In the case of MEG, the true source activity at a location r causes specific changes in measured external magnetic fields specified by the forward field. Thus, the form of the likelihood distribution $p(y|x)$

is different from that of the prior $p(y)$. Since the forward solution relates x to y , then $p(y|x)$ is easy to compute. Before the data was measured, some prior guess exists as to what source activity might be present, given by the *prior distribution* $p(x)$. But really of most interest is the unknown/unmeasured quantity x given the measured data y , the *posterior distribution*, given by Bayes rule:

$$p(x|y) = \frac{p(y|x)p(x)}{p(y)} = \frac{p(y|x)p(x)}{\int dx p(x)p(x,y)} \propto p(y|x)p(x) \quad (1.20)$$

A useful concept when using Bayes rule is that of a *conjugate prior*. For a given likelihood distribution, a conjugate prior distribution is one which gives the same form of distribution for the posterior distribution (but with different parameter values). For a Gaussian likelihood distribution, the conjugate prior distribution is Gaussian. The mean of the posterior Gaussian is the weighted average of the likelihood and prior means, weighted proportionally by their certainty (precision).

1.6.1 Probabilistic methods for the inverse problem

Several tomographic algorithms use a Bayesian framework to explicitly incorporate prior assumptions about source configuration in the model (Sato et al., 2004; Phillips et al., 2005; Mattout et al., 2006). Wipf et al. (2007) provide a unifying framework relating methods through their choice in specification of a neural source prior, whether fixed or learned from the data.

Jun et al. (2005) create a multi-dipole model that explicitly computes the posterior for number of dipoles, location, and orientation. They use prior terms for

the temporal correlations of the sources, chosen based on the data. A prior on location or orientation can include information from fMRI or structural MRI, although the article demonstrates use of a uniform prior. The prior on number of dipoles roughly scales with the inverse of the number. They acknowledge the problem that the noise covariance has many parameters to estimate, so they use a Wishart prior distribution and estimate the noise posterior. They also integrate over the time courses to obtain the posterior for location, orientation and number of dipoles.

1.7 Graphical Models

Graphical models are a powerful and flexible tool for relating known quantities with unknown quantities in a probabilistic way. A node on the graph represents a variable or parameter, known or unknown, and lines are drawn between nodes that probabilistically depend on one another. In general, for a group of variables, the joint probability of all variables taking on some specific set of values assumes dependence of all variables on all other variables. However, a graphical model can easily indicate which quantities are independent, thus simplifying the joint probability distribution. Graphs are particularly useful for indicating conditional dependencies: whether A is independent of B given knowledge about C . The specific rules for whether two nodes are conditionally independent are different for directed versus undirected graphs. In an undirected graph, the line

connecting two nodes is nondirectional. However, the lines connecting nodes in directed graphs are directed, indicated by an arrowhead on the graph and the meaning of the connection is changed depending on which way the arrow points. In a directed graph, *parent* nodes point to *children* nodes, indicating the children have a probabilistic dependence on the parents; knowledge of the value of a parent node gives a greater-than-chance guess at the value of the child node. In undirected graphs, *neighboring* nodes have probabilistic dependencies on each other.

Throughout this work, directed graphs are used. A quantity in a circle is unknown, but a quantity in a square is known either by measurement or direct computation. A big box surrounds variables (quantities dependent on time, either measured or unknown), while parameters and hyperparameters, which specify distributions, are outside this box. When using Bayes rule, there is always a question of how to specify or choose the prior distribution of an unknown quantity. Often, this can be helped by assuming that the prior distribution at this level of the problem is really a posterior distribution of some higher level of the problem, and can be computed using Bayes rule and hyperparameters (parameters over parameters).

Chapter 2

Probabilistic Methods for Neural Source Reconstruction: Mathematical formulation ¹

Two distinct methods for neural source reconstruction which incorporate denoising have been developed. Both are novel probabilistic models that estimate source activity while suppressing the effect of interference and noise sources. Both rely on knowledge of stimulus/event timing to use pre-stimulus data to estimate the noise and then localize the evoked sources in post-stimulus data. The source image map is the likelihood of a dipole source at each voxel, and the source time course is also estimated.

¹Portions of this chapter originally appeared as an article in *NeuroImage* (Zumer et al., 2007b), NIPS conference paper (Zumer et al., 2007a), and ISBI conference paper (Zumer et al., 2006)

The first model (SAKETINI) estimates contributions to sensor data from evoked sources, interference sources and sensor noise using Bayesian methods and by exploiting knowledge about their timing and spatial covariance properties. Full posterior distributions are computed rather than just the MAP estimates. The second model (NSEFALoc) uses a set of temporal basis functions learned from the data. The data is assumed to be a sum of evoked neural sources plus other interference sources and sensor noise. First, a set of denoised temporal basis functions are computed from the data which describe the dynamics of all evoked sources. Next, the spatial weighting of these temporal basis functions is computed at each voxel to provide a spatiotemporal map of activity across the brain. One subtype (NSEFALoc1) models the activity outside a particular voxel by a full-rank covariance matrix and estimates unknown quantities by maximizing the likelihood. The other subtype (NSEFALoc2) parameterizes activity outside the voxel of interest as a linear mixture of a set of unknown Gaussian factors plus Gaussian sensor noise and estimates all unknown quantities using an Expectation-Maximization (EM) algorithm.

The mathematical background, framework and derivations of these methods are presented in this chapter. The performance of these methods in simulated and real data is presented in Chapter 3. An extended comparison on real data is further provided in Chapter 4.

2.1 Overview

Noise in MEG and EEG data is typically reduced by a variety of preprocessing algorithms before being used by source localization algorithms. Simple forms of preprocessing include filtering out frequency bands not containing a brain signal of interest. Additionally and more recently, ICA algorithms have been used to remove artefactual components, such as eyeblinks (Jung et al., 2000). More sophisticated techniques have also recently been developed using graphical models for preprocessing prior to source localization (Nagarajan et al., 2005, 2006). Therefore, current algorithms for source localization from MEG and EEG data typically use a two stage procedure - the first for noise/interference removal and the second for source localization.

2.2 Factor Analysis

2.2.1 Dimension reduction and denoising methods

MEG and EEG data often are acquired with many (~ 100) channels while the number of underlying neural sources of interest remains the same for a given task or stimulus. The number of activated brain areas with signal strength large enough to be measured by MEG/EEG and distinguished as separate entities is usually under 10; thus much signal-space reduction can be performed on the many channels of sensor data. Many methods exist for data dimension reduc-

tion; the two most popular for MEG/EEG data are PCA (principal components analysis) and ICA (independent components analysis). Factor analysis (FA) is discussed as an alternative. All three methods estimate an unknown set of underlying signals (components or factors) and an unknown mixing (weight) matrix, but differ in the assumptions of these quantities.

Defined above, an FA model is a probabilistic graphical model that parses the data into factors (which are hidden variables) premultiplied by a mixing matrix (factor-loading matrix) plus an additive noise term. Thus, FA is a general case, of which PCA and ICA are special cases requiring square mixing matrices and zero-noise. ICA further requires factors to be stationary, independent, and non-Gaussian.

PCA separates the underlying signals by linear transformation so that they are orthogonal to each other and are ordered in terms of explained variance of the data. An eigenvalue decomposition or singular value decomposition (SVD) is used to compute these principal components. Assumptions of this method are that the data has zero mean and that the meaningful underlying signals are actually orthogonal to each other (which is not usually true for brain activity, but can sometimes be a reasonable approximation). The data y (of dimension $K \times N$ with K measurement sensors and N time points) is decomposed by SVD into $W\Sigma V^T$, with Σ a diagonal matrix of the eigenvalues and ΣV^T the principal components. A plot of the ordered eigenvalues usually depicts a transition point where the

first few eigenvalues are stronger than the remaining ones such that they can be assumed to represent "signal" and the remaining eigenvalues represent "noise." These "noise" eigenvalues can be set to zero and the reverse transform can be applied to obtain a denoised version of the sensor data. This method is used for the eigen-subspace beamformer (Sekihara et al., 2002b). A potential problem of this method is that a strong but non-brain signal, such as eyeblinks, cannot be distinguished from true signals of interest.

ICA separates signals that are assumed non-Gaussian and independent of each other, but are allowed to be correlated (i.e. non-orthogonal). The generative model for data y is $y = As$ where A is the square mixing matrix and s are the non-Gaussian independent components. The independence implies that knowledge of one signal does not help to better predict another signal. Many methods for actually computing the unknown mixing matrix and components exist. In general, they can be classified as methods that do or do not take higher than second-order statistics into account. For more details on ICA and the methods for computing the components, see Hyvarinen et al. (2001) and Cardoso (1998). There is not an ordering of components like for PCA; commonly the components are displayed and visually selected for which appear to be signal and the weighting of the remaining components can be set to zero. This method can work well for removing sources of interference such as eyeblinks which have a characteristic time course.

Neither PCA or ICA allow for noise at the sensor level. Thus, any sensor

noise is effectively projected to the components. Furthermore, both PCA and ICA involve computation/estimation of a square weight (mixing) matrix, which can be difficult and/or unnecessary when the true signal space is much less than the number of sensors and a large number of sensors are present. Factor analysis is a dimension reduction method which explicitly takes into account both sensor noise and the dimension reduction via a non-square mixing matrix. The data y is assumed to come from the model $y = Ax + v$, where A is in general non-square (size $K \times L$, where $L \leq K$), the dimension of x is then $L \times N$, and v accounts for sensor noise (dimension $K \times N$). The signal factors x and noise factors v are assumed Gaussian, with x drawn from zero-mean, identity precision distribution and v drawn from a zero-mean with diagonal precision λ distribution. The following is assumed for $n = 1 : N$:

$$p(x) = \prod_n p(x_n); \quad p(x_n) = \mathcal{N}(x_n|0, I), \quad (2.1)$$

$$p(v) = \prod_n p(v_n); \quad p(v_n) = \mathcal{N}(v_n|0, \lambda), \quad p(\lambda) = \text{const.} \quad (2.2)$$

$$p(y|x, A, \lambda) = \prod_n p(y_n|x_n, A, \lambda); \quad p(y_n|x_n, A, \lambda) = \mathcal{N}(y_n|Ax_n, \lambda) \quad (2.3)$$

Thus the data y is assumed Gaussian with mean Ax and precision λ . The maximum likelihood approach could be taken to solve for the unknown factors x by maximizing the data with respect to (w.r.t.) A and λ ; however this produces a set of nonlinear equations which are difficult to solve. Instead, a method called Expectation-Maximization is used.

2.2.2 Expectation-Maximization

Expectation-Maximization (EM) is an iterative algorithm for computing underlying unknown (hidden/missing/latent) quantities. In the E-step, the posterior mean of the unknown factors is computed for a given value of parameters (either an initial estimate or the value from the previous M-step). In the M-step, the values of the parameters (such as A and λ) that maximize the averaged complete data likelihood are found. The likelihood is guaranteed not to decrease with each EM step and the algorithm can iterate until the likelihood appears to converge.

The data likelihood is defined as the logarithm of the observed data: $L = \log p(y|\Theta)$. The best solution for the parameters would be the one that maximized the joint likelihood of the data and hidden variables (if indeed they were known). This joint likelihood is termed the complete data likelihood, since it involves all variables, both observed and hidden, written as $l = \log p(y, x|\Theta)$, where Θ is the set of all model parameters, in this case $\Theta = \{A, \lambda\}$. However, the factors are unknown, but, given the data, a posterior estimate can be computed. Either just the maximum *a posteriori* estimate can be used or an average over all values of x weighted by their probability, hence creating the averaged complete data likelihood, given by $\bar{l}(y|\Theta) = E_{p(x|y)}l(y, x|\Theta) = \int dx p(x|y)l(y, x|\Theta)$.

2.2.3 Factor Analysis estimates

Assuming that the factors are independent and identically distributed (i.i.d.) allows the posterior estimate of the factors to factorize over time (indicated by subscript n). Since the distribution of the posterior of the factors is Gaussian, the easiest way to compute the mean and precision is to take the derivative of the log posterior with respect to the factor. The slope is the precision, and the value of the factor that makes the derivative zero is the mean. Specifically,

$$\begin{aligned}\log p(x_n|y_n) &= \log p(y_n|x_n) + \log p(x_n) \\ &= -\frac{1}{2}(y_n - Ax_n)^T \lambda (y_n - Ax_n) - \frac{1}{2}x_n^T x_n\end{aligned}\tag{2.4}$$

This gives the following for the mean and precision:

$$\begin{aligned}p(x_n|y_n) &= \mathcal{N}(x_n|\bar{x}_n, \Gamma) \\ \bar{x}_n &= \Gamma^{-1}A^T \lambda y \\ \Gamma &= A^T \lambda A + I\end{aligned}\tag{2.5}$$

For all variables throughout the rest of this dissertation, a bar over the variable indicates its posterior mean.

Given these posterior estimates, the averaged complete data likelihood can be maximized for the optimum parameter values. Taking the derivative of $\bar{l}(y|\Theta) = E_{p(x|y)}[l(y, x|\Theta)]$ w.r.t. A and λ gives

$$\begin{aligned}A &= R_{yx}R_{xx}^{-1} \\ \lambda &= \frac{1}{N}\text{diag}(R_{yy} - AR_{xy})\end{aligned}\tag{2.6}$$

where N is the total number of time points. The likelihood function using the updated values for parameters and factors, which is guaranteed not to decrease with each EM step, is given by:

$$\mathcal{L} = \sum_{n=1}^N \log p(y_n) = \frac{N}{2} \log |\Sigma^{-1}| - \frac{1}{2} \sum_{n=1}^N y_n^T \Sigma^{-1} y_n \quad (2.7)$$

with the covariance of $p(y_n)$ given by $\Sigma = AA^T + \lambda^{-1}$.

2.2.4 Variational Bayesian Expectation-Maximization

For simple models, the posterior distribution can often be computed directly. However, for more interesting and realistic models, the posterior is often computationally intractable. In these cases, some approximation must be made. Two main branches are sampling methods and variational methods. Sampling methods, such as Markov Chain Monte Carlo methods (Jun et al., 2005; Gelman and Rubin, 1996), extensively estimate points in the distribution. However, this is dependent on the sampled points and can be quite computationally costly. Variational methods instead choose to factorize the joint distribution over factors and parameters assuming conditional independence of the factors and parameters, also termed the mean field approximation.

The energy of a global configuration x, y is the negative of the complete data likelihood. Instead of computing the averaged (expected) energy w.r.t. to the true posterior distribution $p(x, A|y)$ in order to get the averaged complete data likelihood defined earlier, one can instead compute this expected energy w.r.t.

any arbitrary “free” distribution $q(x, A|y)$. The negative of a quantity termed “free energy” (from statistical physics) is defined as this expected energy w.r.t. q plus the entropy of q :

$$\mathcal{F} = E_{q(x, A|y, \Theta)} [l(y, x, A|\Theta)] + H_q \quad (2.8)$$

This can be related to the data likelihood as follows:

$$\begin{aligned} \mathcal{F}(q, \Theta) &= E_{q(x, A|y, \Theta)} [l(y, x, A|\Theta)] + H_q \\ &= \int dx dA q(x, A|y, \Theta) [\log p(y, x, A|\Theta) - \log q(x, A|y, \Theta)] \\ &= \int dx dA q(x, A|y, \Theta) \left[\log \frac{p(y, x, A|\Theta)}{q(x, A|y, \Theta)} \right] \\ &= \int dx dA q(x, A|y, \Theta) \left[\log \frac{p(x, A|y, \Theta)}{q(x, A|y, \Theta)} + \log p(y|\Theta) \right] \\ &= \log p(y|\Theta) - \text{KL}[q(x, A|y, \Theta) || p(x, A|y, \Theta)] \\ &= \mathcal{L}(\Theta) - \text{KL}[q(x, A|y, \Theta) || p(x, A|y, \Theta)] \end{aligned} \quad (2.9)$$

where Θ is the set of all model parameters and H_q is the entropy of $q(u, B|y, \Theta)$. The Kullback-Leibler divergence between distributions q and p , defined $\text{KL}(q||p) = \int dx q(x) [\log q(x) - \log p(x)]$, is always nonnegative and can equal zero if the two distributions are equal; thus, \mathcal{F} is a lower bound to the log likelihood function \mathcal{L} . Note further that, if q equals p , then:

$$\mathcal{F} = \mathcal{L} = \bar{l} + H_p \quad (2.10)$$

Instead of maximizing the likelihood function itself, this new function \mathcal{F} is parameterized by the arbitrary distribution q . Now, finding the q that maximizes \mathcal{F} means that q will be as close to the true posterior distribution as possible, and that \mathcal{F} will be as close to the true likelihood \mathcal{L} .

A full Bayesian EM algorithm iterates between maximizing \mathcal{F} w.r.t. q to find the best full posterior distribution (not just the MAP estimate) and the maximizing \mathcal{F} w.r.t. the parameters. When it is mathematically tractable to compute this exact q , each E-step brings \mathcal{F} equal to \mathcal{L} and each M-step increases \mathcal{L} .

The parameters can be treated equally with the unknown factors, so that a full distribution of the parameters can be found as well (rather than just the MAP estimate). However, this now means the parameters and hidden variables are correlated. Exact inference on this model is intractable using the true joint posterior over the background factors u and mixing matrix B . The marginal conditional posterior estimates of the background factors and mixing matrix is desired, given the data. We could choose to estimate just the maximum *a posteriori* (MAP) estimates, but this ignores the full distribution including its precision.

$$p(x, A|y) \approx q(x, A|y) = q(x|y)q(A|y) \quad (2.11)$$

This variational approximation leads to the final version of maximizing the likelihood: Variational Bayesian Expectation Maximization (VB-EM). This approximation restricts the posterior to a product of factor distributions, but allows the solution to be computed analytically. \mathcal{F} is equal to $\log p(y|\Theta)$ when the approximation in Eq. (2.11) is true, thus making the KL-divergence zero. The VB-EM algorithm (Attias, 1999; Ghahramani and Beal, 2001; Beal, 2003) iteratively maximizes \mathcal{F} w.r.t. each factorized distribution to, at least, a local maximum of \mathcal{F} , al-

ternating w.r.t. the posteriors $q(x|y)$ and $q(A|y)$. In the VBE-step, \mathcal{F} is maximized w.r.t. $q(x|y)$, keeping $q(A|y)$ constant, and the sufficient statistics of the hidden variables are computed. In the VBM-step, \mathcal{F} is maximized w.r.t. $q(A|y)$, keeping $q(x|y)$ constant, to solve for the full distribution of A and the MAP estimate of any additional parameters or hyperparameters are computed as they would be in a traditional non-VB M-step. The $q(x, A|y) = q(x|y)q(A|y)$ that maximizes \mathcal{F} is the distribution that minimizes the KL-divergence, thus making the mean field approximation as close to equivalence as possible under the constraint implied by Eq. (2.11). A derivation of a similar model is described in the appendix of Nagarajan et al. (2006) and the variational approximation has been shown to be a powerful and flexible tool (Jordan et al., 1999). It can further be shown that, given our previous assumptions of Gaussian prior distributions, the form of the posterior $q(x|y)$ must also be Gaussian, by showing that the function $\log q$ is quadratic in x .

2.2.5 Variational Bayesian Factor Analysis (VBFA)

Even though a prior distribution of the factors x was assumed, a full Bayesian approach requires prior distributions on the parameters as well so that full posterior distributions can be computed. The prior terms usually involve hyperparameters, which here are updated in the VBM-step. In addition to the prior distributions on x and v mentioned in section 2.2.1, a conjugate prior is used for the

mixing matrix A , as follows,

$$p(A) = \prod_{kl} p(A_{kl}); \quad p(A_{kl}) = \mathcal{N}(A_{kl}|0, \lambda_k \alpha_l) \quad (2.12)$$

where α_l is a hyperparameter over the l th column of A and λ_k is the precision of the k th sensor. The matrix $\alpha = \text{diag}(\alpha_1, \dots, \alpha_L)$ provides a robust mechanism for automatic model order selection, so that the optimal size of A is inferred from the data through α (for general reference, see Neal (1996)). Examples of the effectiveness of this automatic dimension selection in the VBFA framework are shown in Nagarajan et al. (2007); further examples of its use in the source reconstruction algorithms in this chapter are shown in the next chapter.

The first-stage E-step maximizes \mathcal{F} w.r.t. $q(x|y)$, while holding $q(A|y)$ fixed, using calculus of variations, yielding:

$$\log q(x|y) = E_{q(A|y)} (\log p(y, x, A|\Theta)) \quad (2.13)$$

This $q(x|y)$ is the best function to approximate $p(x|y)$. It can be shown that this q has a Gaussian form given the previous Gaussian assumptions. To find the mean and precision of $q(x|y)$, take the derivative w.r.t. x as in FA, and use equations A.6

and A.9.

$$\begin{aligned}
\frac{\partial}{\partial x_n} \log q(x|y) &= \frac{\partial}{\partial x_n} E_{q(A|y)} [\log p(y, x, A|\Theta)] = 0 \\
\frac{\partial}{\partial x_n} E_{q(A|y)} [\log p(y|x, A, \Theta) + \log p(x|\Theta)] &= 0 \\
\frac{\partial}{\partial x_n} E_{q(A|y)} \left[-\frac{1}{2} \sum_{n=1}^N ((y_n - Ax_n)^T \lambda (y_n - Ax_n) - x_n^T x_n) \right] &= 0 \\
E_{q(A|y)} [A^T \lambda y_n - A^T \lambda Ax_n] &= 0 \\
A^T \lambda y_n &= A^T \lambda Ax_n
\end{aligned} \tag{2.14}$$

Using Eq. (A.13), the following posterior estimates are obtained for the factors:

$$\begin{aligned}
q(x|y) &= \prod_n q(x_n|y_n); \quad q(x_n|y_n) = \mathcal{N}(x_n|\bar{x}_n, \gamma) \\
\bar{x}_n &= \gamma^{-1} \bar{A}^T \lambda y_n; \quad \gamma = \bar{A}^T \lambda \bar{A} + K \psi_A^{-1} + I
\end{aligned} \tag{2.15}$$

The last line for γ is the final version using the estimate determined below in Eq. (2.18) that the precision of A is $\lambda \psi_A$. In the VBM-step, the full posterior distribution of the mixing matrix A is computed, including its precision matrix ψ_A , and the MAP estimates of the noise precision λ and the hyperparameter α . The VBM step finds the distribution that best approximates $p(A|y)$:

$$\log q(A|y) = E_{q(x|y)} (\log p(y, x, A|\Theta)) \tag{2.16}$$

Similar to the VBE-step, this q also has a Gaussian form. Its mean and precision

are found by taking the derivative w.r.t. A :

$$\begin{aligned}
\frac{\partial}{\partial A_k} \log q(A|y) &= \frac{\partial}{\partial A_k} E_{q(x|y)} [\log p(y, x, A|\Theta)] = 0 \\
\frac{\partial}{\partial A_k} E_{q(x|y)} [\log p(y|x, A, \Theta) + \log p(A|\Theta)] &= 0 \\
\frac{\partial}{\partial A_k} E_{q(x|y)} \left[-\frac{1}{2} \sum_{n=1}^N (y_n - Ax_n)^T \lambda (y_n - Ax_n) - \frac{1}{2} \sum_{k=1}^K A_k \lambda_k \alpha A_k^T \right] &= 0 \\
E_{q(x|y)} \left[\sum_{n=1}^N (\lambda_k y_n x_n^T - \lambda_k A_k x_n x_n^T) - A_k \lambda_k \alpha \right] &= 0 \\
\lambda_k R_{yx}^k - \lambda_k A_k R_{xx} - A_k \lambda_k \alpha &= 0 \\
\lambda_k R_{yx}^k &= A_k (\lambda_k R_{xx} + \lambda_k \alpha)
\end{aligned} \tag{2.17}$$

Thus, the posterior distribution for A is:

$$\begin{aligned}
q(A|y) &= \mathcal{N}(A|\bar{A}, \lambda\psi_A) \\
\bar{A} &= R_{yx} \psi_A^{-1}; \quad \psi_A = (R_{xx} + \alpha)
\end{aligned} \tag{2.18}$$

The matrices, such as R_{yx} , represent the posterior covariance between the two subscripts.

$$R_{yy} = \sum_{n=1}^N y_n y_n^T, \quad R_{yx} = \sum_{n=1}^N y_n \bar{x}_n^T \tag{2.19}$$

To find the MAP estimates in the M-step, take the derivative of \mathcal{F} w.r.t. the

parameter. To obtain the MAP estimate of λ :

$$\begin{aligned}
\frac{\partial \mathcal{F}}{\partial \lambda_k} &= \frac{\partial}{\partial \lambda_k} E_{q(x|y)} E_{q(A|y)} [\log p(y|x, A, \lambda) + \log p(A|\lambda)] = 0 \\
\frac{\partial}{\partial \lambda_k} E_{q(x|y)} E_{q(A|y)} &\left[\frac{N}{2} \log |\lambda| - \frac{1}{2} \sum_{n=1}^N (y_n - Ax_n)^T \lambda (y_n - Ax_n) + \frac{1}{2} \log |\lambda| \right. \\
&\quad \left. - \frac{1}{2} \sum_{k=1}^K A_k \lambda_k \alpha A_k^T \right] = 0 \\
E_{q(x|y)} E_{q(A|y)} &\left[\frac{N+1}{2} \lambda_k^{-1} - \frac{1}{2} \sum_{n=1}^N (y_{nk} y_{nk} - 2A_k x_n y_{nk} + A_k x_n x_n^T A^T) \right. \\
&\quad \left. - \frac{1}{2} A_k \alpha A_k^T \right] = 0 \\
\frac{N+1}{2} \lambda_k^{-1} &= \frac{1}{2} R_{yy} - \bar{A} R_{xy}^k + \frac{1}{2} A_k \bar{R}_{xx} A_k^T + \frac{1}{2} A_k \bar{\alpha} A_k^T \\
(N+1) \lambda_k^{-1} &= \\
R_{yy} - 2\bar{A} R_{xy}^k + \bar{A}_k R_{xx} \bar{A}_k^T + \text{Tr}(R_{xx} \lambda^{-1} \psi_A^{-1}) + \bar{A}_k \alpha \bar{A}_k^T + \text{Tr}(\alpha \lambda^{-1} \psi_A^{-1}) \\
(N+1) \lambda_k^{-1} &= R_{yy} - 2\bar{A} R_{xy}^k + \bar{A}_k \psi_A \bar{A}_k^T + \lambda_k^{-1} \\
\lambda_k^{-1} &= \frac{1}{N} (R_{yy} - 2\bar{A}_k R_{xy}^k + \bar{A}_k \psi_A \bar{A}_k^T) \\
\lambda_k^{-1} &= \frac{1}{N} (R_{yy}^k - \bar{A}_k R_{xy}^k) \lambda = N [\text{diag}(R_{y'y'} - \bar{A} R_{xy'})]^{-1} \quad (2.20)
\end{aligned}$$

The MAP estimate of α is similarly found by maximizing \mathcal{F} :

$$\begin{aligned}
\frac{\partial \mathcal{F}}{\partial \alpha} &= \frac{\partial}{\partial \alpha} E_{q(x|y)} E_{q(A|y)} [\log p(A)] = 0 \\
\frac{\partial}{\partial \alpha} E_{q(x|y)} E_{q(A|y)} &\left[\frac{K}{2} \log |\alpha| - \frac{1}{2} \sum_{k=1}^K A_k \lambda_k \alpha A_k^T \right] = 0 \\
\frac{K}{2} \alpha^{-1} &= \frac{1}{2} A \bar{\lambda} A^T \alpha^{-1} = \text{diag}\left(\frac{1}{K} \bar{A}^T \lambda \bar{A} + \psi_A\right); \quad (2.21)
\end{aligned}$$

The final maximized likelihood is computed and can be plotted to ensure it

does not decrease with any E- or M-step due to implementation errors.

$$\mathcal{L} = \frac{N}{2} \log \frac{|\lambda|}{|\gamma|} + \frac{K}{2} \log \frac{|\alpha|}{|\psi_A|} - \frac{1}{2} \sum_{n=1}^N (y_n^T \lambda y_n + x_n^T \gamma x_n) \quad (2.22)$$

2.3 Stimulus Evoked Factor Analysis (SEFA)

SEFA (Nagarajan et al., 2007) uses the computational framework of VBFA but includes the additional concept of how MEG/EEG data are collected in order to further separate out noise components. In the stimulus evoked paradigm, some baseline control data is collected for several hundred milliseconds termed the “pre-stimulus period”, then a stimulus occurs, evoking a neural response in the “post-stimulus period.” The pre-stimulus period is typically used to determine the mean baseline to subtract from the post-stimulus period as well as to measure the baseline variance in the data so that the post-stimulus can be compared to see if it indeed is above baseline.

The key idea of SEFA is that background activity such as ongoing brain activity unrelated to the stimulus, other biological noise such as eyeblinks and heart-beat, other room noise, and sensor noise will be present in both pre-stimulus and post-stimulus periods. However, only the evoked brain sources of interest will be present in post-stimulus period alone and *not* the pre-stimulus period. Note this assumption is valid for the evoked response paradigm, but not the event-related synchronization/desynchronization analysis, which is not the focus of this thesis but is discussed in the last chapter. The VBFA model is extended to include the

background sources:

$$\begin{aligned} y_n &= Bu_n + v_n & n &= -N_{pre}, \dots, -1 \\ y_n &= Ax_n + Bu_n + v_n & n &= 0, \dots, N_{post} - 1 \end{aligned} \quad (2.23)$$

Time ranges from $-N_{pre} : 0 : N_{post} - 1$ where N_{pre} (N_{post}) indicates the number of time samples in the pre-(post-)stimulus period. The $K \times M$ matrix B and the $M \times 1$ vector u_n represent the background mixing matrix and background factors, respectively. The quantities A and x are as defined above in VBFA. Note that the size of A and B do not need to be the same; in fact, L can be roughly the size of the number of expected evoked sources (typically less than 10) whereas M is usually larger, but should scale in relation to the total number of sensors. Furthermore, N_{pre} and N_{post} do not need to be the same size, but if N_{pre} is much less than N_{post} there will be identifiability issues between A and B . The distribution of the data now is:

$$\begin{aligned} p(y|x, u, A, B, \lambda) &= \prod_n p(y_n|x_n, u_n, A, B, \lambda) \\ p(y_n|x_n, u_n, A, B, \lambda) &= \begin{cases} \mathcal{N}(y_n|Bu_n, \lambda) & n = -N_{pre}, \dots, -1 \\ \mathcal{N}(y_n|Ax_n + Bu_n, \lambda) & n = 0, \dots, N_{post} - 1 \end{cases} \end{aligned} \quad (2.24)$$

To describe the full model in the Bayesian framework, prior probability dis-

tributions are given to these quantities:

$$p(x) = \prod_n p(x_n); \quad p(x_n) = \mathcal{N}(x_n|0, I), \quad (2.25)$$

$$p(u) = \prod_n p(u_n); \quad p(u_n) = \mathcal{N}(u_n|0, I), \quad (2.26)$$

$$p(v) = \prod_n p(v_n); \quad p(v_n) = \mathcal{N}(v_n|0, \lambda), \quad p(\lambda) = \text{const.} \quad (2.27)$$

and, similar to VBFA, hyperparameters are used for the mixing matrices:

$$p(A) = \prod_{kl} p(A_{kl}); \quad p(A_{kl}) = \mathcal{N}(A_{kl}|0, \lambda_k \alpha_l) \quad (2.28)$$

$$p(B) = \prod_{km} p(B_{km}); \quad p(B_{km}) = \mathcal{N}(B_{km}|0, \lambda_k \beta_m) \quad (2.29)$$

There are two slightly different ways in which all the posterior estimates of the factors and parameters can be computed. One would be to use all the data at once to compute x , u , A and B . However, there is greater chance of identifiability problems between A and B , especially if the amount of pre-stimulus data is small relative to post-stimulus data. Furthermore, a better initial estimate of A can be made if B is already known, rather than learning both A and B from less reliable initial estimates. Thus, SEFA is best computed by using a two-stage procedure. First, B , u and λ are learned from just the pre-stimulus data alone (essentially just doing regular VBFA on the pre-stimulus data). Then, B and λ are held fixed and A , x and u are computed using just the post-stimulus data. Note that u needs to be recomputed for the post-stimulus period, since the projection from the data to the noisy source space is defined by B which is fixed, but the actual realization of noise strength in the post-stimulus changes from time point to time point.

Nagarajan et al. (2007) describe the VBEM updates for SEFA using the first method (learning A and B at once) although in practice the two-step method is more commonly used. Here, the explicit update rules for the two-stage procedure are given. The first are identical to VBFA applied to pre-stimulus data: The following posterior estimates are obtained for the factors in the first-stage VBE-step:

$$q(u|y) = \prod_n q(u_n|y_n); \quad q(u_n|y_n) = \mathcal{N}(u_n|\bar{u}_n, \gamma)$$

$$\bar{u}_n = \gamma^{-1} \bar{B}^T \lambda y_n; \quad \gamma = \bar{B}^T \lambda \bar{B} + K \psi_B^{-1} + I \quad (2.30)$$

In the first-stage VBM-step, the full posterior distribution of the background mixing matrix B is computed, including its precision matrix ψ_B , and the MAP estimates of the noise precision λ and the hyperparameter β .

$$q(B|y) = \mathcal{N}(B|\bar{B}, \psi_B)$$

$$\bar{B} = R_{yu} \psi_B; \quad \psi_B = (R_{uu} + \beta)^{-1}$$

$$\beta^{-1} = \text{diag}\left(\frac{1}{K} \bar{B}^T \lambda \bar{B} + \psi_B\right); \quad \lambda^{-1} = \frac{1}{N} \text{diag}(R_{yy} - \bar{B} R_{yu}^T) \quad (2.31)$$

Now that B and λ have been learned from the data, the statistics of these noise sources are assumed not to change. Unless there is a systematic bias in a noise source introduced by the experiment itself such that some noise is only present in post-stimulus time period, this assumption is reasonably valid.

The free energy function that is maximized in the second step of SEFA is:

$$\begin{aligned} \mathcal{F}(q, \Theta) &= \int dx du dA q(x, u, A|y, \Theta) [\log p(y, x, u, A|\Theta) - \log q(x, u, A|y, \Theta)] \\ &= \log p(y|\Theta) - \text{KL}[q(x, u, A|y, \Theta)||p(x, u, A|y, \Theta)] \end{aligned} \quad (2.32)$$

In the second-stage VBE-step, the posterior distribution of the background factors given the data is computed by maximizing the above expression for \mathcal{F} .

First, define

$$\kappa_n = \begin{pmatrix} x_n \\ u_n \end{pmatrix}; \quad \Omega = \begin{pmatrix} A & \bar{B} \end{pmatrix}; \quad \bar{\Omega} = \begin{pmatrix} \bar{A} & \bar{B} \end{pmatrix}; \quad (2.33)$$

Then, taking the derivative of optimal posterior $q(\kappa|y)$ gives:

$$\begin{aligned} \frac{\partial}{\partial \kappa_n} \log q(\kappa|y) &= \frac{\partial}{\partial \kappa_n} E_{q(A|y)} [\log p(y, \kappa, A|\Theta)] = 0 \\ \frac{\partial}{\partial \kappa_n} E_{q(A|y)} [\log p(y|\kappa, \Omega, \Theta) + \log p(\kappa|\Theta)] &= 0 \\ \frac{\partial}{\partial \kappa_n} E_{q(A|y)} \left[-\frac{1}{2} \sum_{n=1}^N ((y_n - \Omega \kappa_n)^T \lambda (y_n - \Omega \kappa_n) - \kappa_n^T \kappa_n) \right] &= 0 \\ E_{q(A|y)} [\Omega^T \lambda y_n - \Omega^T \lambda \Omega \kappa_n - \kappa_n] &= 0 \\ \bar{\Omega}^T \lambda y_n &= (\bar{\Omega}^T \lambda \bar{\Omega} + I) \kappa_n \end{aligned} \quad (2.34)$$

The second-stage VBE-step results as:

$$q(\kappa_n|y_n) = \mathcal{N}(\kappa_n|\bar{\kappa}_n, \Gamma);$$

$$\bar{\kappa}_n = \Gamma^{-1}\bar{\Omega}^T \lambda y_n;$$

$$\Gamma = \overline{\Omega^T \lambda \Omega} + I = \bar{\Omega}^T \lambda \bar{\Omega} + K\Psi^{-1} + I \quad (2.35)$$

$$= \begin{pmatrix} \bar{A}^T \\ \bar{B}^T \end{pmatrix} \lambda \begin{pmatrix} \bar{A} & \bar{B} \end{pmatrix} + K \begin{pmatrix} 0 & \Psi_A^{-1} \\ 0 & 0 \end{pmatrix} + I \quad (2.36)$$

In the second-stage VBM-step, the posterior distribution of the interference mixing matrix A is updated, including its precision Ψ_A . Note that \bar{B} and the sensor noise precision λ were learned and fixed from the pre-stimulus data. The MAP value of the hyperparameter α is learned here from the post-stimulus data.

$$\begin{aligned} \frac{\partial}{\partial A_k} \log q(A|y) &= \frac{\partial}{\partial A_k} E_{q(\kappa|y)} [\log p(y, \kappa, A|\Theta)] = 0 \\ \frac{\partial}{\partial A_k} E_{q(\kappa|y)} [\log p(y|\kappa, A, \Theta) + \log p(A|\Theta)] &= 0 \\ \frac{\partial}{\partial A_k} E_{q(\kappa|y)} \left[-\frac{1}{2} \sum_{n=1}^N (y_n - Ax_n - \bar{B}u_n)^T \lambda (y_n - Ax_n - \bar{B}u_n) - \frac{1}{2} \sum_{k=1}^K A_k \lambda_k \alpha A_k^T \right] &= 0 \\ E_{q(\kappa|y)} \left[\sum_{n=1}^N (\lambda_k y_n \kappa_n x_n^T - \lambda_k A_k x_n x_n^T) + \lambda_k \bar{B}_k u_n x_n^T - A_k \lambda_k \alpha \right] &= 0 \\ \lambda_k R_{yx}^k - \lambda_k A_k R_{xx} - \lambda_k \bar{B}_k^T R_{ux} - A_k \lambda_k \alpha &= 0 \\ \lambda_k (R_{yx}^k - \bar{B}_k^T R_{ux}) &= A_k (\lambda_k R_{xx} + \lambda_k \alpha) \end{aligned} \quad (2.37)$$

Thus, the posterior distribution for A is:

$$q(A|y) = \mathcal{N}(A|\bar{A}, \lambda\Psi_A)$$

$$\bar{A} = (R_{yx} - \bar{B}R_{ux})\Psi_A^{-1}; \quad \psi_A = (R_{xx} + \alpha) \quad (2.38)$$

The matrices, such as R_{ux} , represent the posterior covariance between the two subscripts

$$R_{xu} = \sum_{n=1}^N \bar{x}_n \bar{r}_n^T + N\Sigma_{xu} = R_{ux}^T; \quad R_{xx} = \sum_{n=1}^N \bar{x}_n \bar{x}_n^T + N\Sigma_{xx} \quad (2.39)$$

where $\Sigma = \Gamma^{-1}$ is specified as:

$$\Sigma = \begin{pmatrix} \Sigma_{xx} & \Sigma_{xu} \\ \Sigma_{ux} & \Sigma_{uu} \end{pmatrix} \quad (2.40)$$

The update for α is the same as in VBFA. The likelihood function that is guaranteed to converge to at least a local minimum is:

$$\mathcal{L}^r = \frac{N}{2} \log \frac{|\lambda|}{|\Gamma|} - \frac{1}{2} \sum_{n=1}^N \left(y_n^T \lambda y_n - \bar{x}_n'^T \Gamma \bar{x}_n' \right) + \frac{K}{2} \log |\alpha| |\Psi_A| \quad (2.41)$$

Performance of SEFA to denoise the sensor data, compute a regularized data covariance matrix, and estimate the source factors are described further in Nagarajan et al. (2007).

Like PCA, VBFA (or SEFA) can be applied to the post-stimulus (or total) data window of interest to obtain a denoised estimate of the data covariance, which can be used as an input to some source localization method such as beamforming. The factors can also be further used as temporal basis functions such that neural activity in source space is a linear combination of the factors. The next two sections focus on source reconstruction methods that use the probabilistic graphical model framework introduced with the denoising methods like SEFA. These methods learn hidden variables representing localized source activity while accounting for noise confounds.

2.4 SAKETINI

SAKETINI is a new probabilistic modeling framework for MEG/EEG source localization which estimates Source Activity using Knowledge of Event Timing and Independence from Noise and Interference. SAKETINI builds on the ideas of SEFA for denoising but additionally computes source localized estimates of neural activity. The framework uses a probabilistic hidden variable model that describes the observed sensor data in terms of activity from unobserved brain and interference sources. The unobserved source activities and model parameters are inferred from the data by a Variational-Bayesian Expectation-Maximization algorithm. The algorithm then creates a spatiotemporal image of brain activity by scanning the brain, inferring the model parameters and variables from sensor data, and using them to compute the likelihood of a dipole at each grid location in the brain.

The data is assumed collected in the stimulus evoked paradigm in the same way as described above for SEFA, such that the stimulus onset or some other experimental marker indicated the “zero” time point. Ongoing brain activity, biological noise, background environmental noise, and sensor noise are present in both pre-stimulus and post-stimulus periods; however, the evoked neural sources of interest are only present in the post-stimulus time period. The sensor data can be described as coming from four types of sources: (I) evoked source at a particular voxel (grid point), (II) all other evoked sources not at that voxel, (III) all

background noise sources with spatial covariance at the sensors (including brain, biological, or environmental sources), and (IV) sensor noise.

Section 2.4.1 describes the full SAKETINI model. The use of VBFA (see section 2.2.5) to infer the pre-stimulus model describing source types (III) and (IV) is described in section 2.4.2. Then certain quantities are fixed and the full model describing the remaining source types (I) and (II) is inferred from the post-stimulus data, which is described in section 2.4.3. After inference of the model, a map of the source activity is created, as well as a map of the likelihood of activity across voxels. In section 2.4.4, the existing MVAB method is shown to be an approximation of SAKETINI. Performance of SAKETINI compared to existing methods is shown in the next chapter.

2.4.1 SAKETINI: Full model overview

Let y_n denote the $K \times 1$ vector of data from sensor data for time point n , where K is the number of sensors (typically ~ 200). Time ranges from $-N_{pre} : 0 : N_{post} - 1$ where N_{pre} (N_{post}) indicates the number of time samples in the pre-(post-)stimulus period. The generative model for data y_n is

$$y_n = \begin{cases} Bu_n + v_n & n = -N_{pre}, \dots, -1 \\ F^r s_n^r + A^r x_n^r + Bu_n + v_n & n = 0, \dots, N_{post} - 1 \end{cases} \quad (2.42)$$

The $K \times Q$ forward lead field matrix F^r represents the physical (and linear) relationship between a dipole source at voxel r for each dipole orientation q along

a coordinate basis and its influence on sensor $k = 1 : K$ (see section 1.4 for more details). As discussed in section 1.4, the single-shell model describing the head will be used throughout the methods and results described here thus $Q = 2$. This forward field (here and in the next method NSEFALoc) is actually normalized by its column norm as was done in the MVAB case to avoid artificially high weighting towards the center of the head.

The source activity s_n^r is a $Q \times 1$ vector of dipole strength at time n for the voxel r . The $K \times L$ matrix $A^{\setminus r}$ and the $L \times 1$ vector $x_n^{\setminus r}$ represent the post-stimulus mixing matrix and evoked non-localized factors, respectively, corresponding to source type (II) discussed above, where the $\setminus r$ superscript indicates for all voxels *not* at voxel r . We leave these superscripts off in the rest of this chapter for clarity. The $K \times M$ matrix B and the $M \times 1$ vector u_n represent the background mixing matrix and background factors, respectively, corresponding to source type (III) discussed above. The $K \times 1$ vector v_n represents the sensor-level noise, which is assumed to be drawn from a Gaussian distribution with zero-mean and precision (inverse covariance) defined by the diagonal matrix λ . All quantities depend on r in the post-stimulus period except for B, u_n and λ , which will be learned from the pre-stimulus data, and fixed as the other quantities are learned for each voxel. Note, however, that the posterior update for u_n (represented by \bar{u}_n) does depend on the voxel r . The graphical model is shown in Fig. 2.1. This generative model becomes a probabilistic model when we specify prior distributions, as described

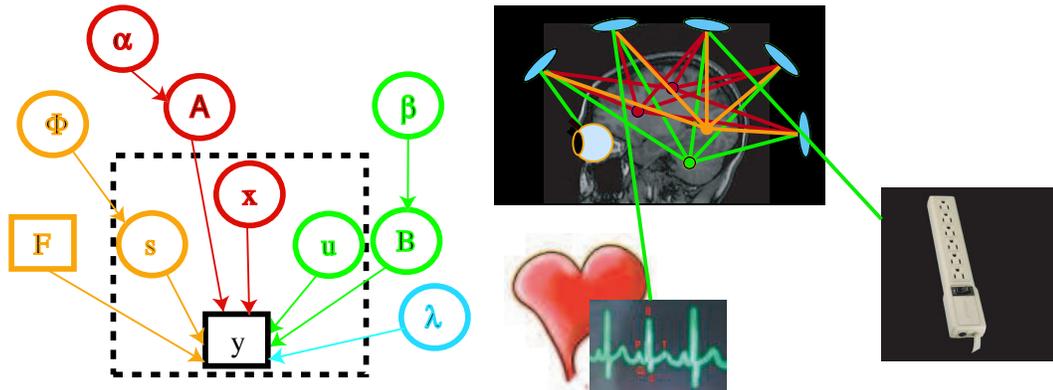


Figure 2.1 (Left) SAKETINI Graphical Model. Variables dependent on time are inside dotted box; parameters independent of time outside dotted box. Values in circles are unknown and learned from the model, and values in squares are known. After applying a simpler model to the pre-stimulus data to learn the sensor noise precision (λ) and interference mixing matrix (B), these quantities are then known and placed in squares. The source activity (s) at a particular voxel specified by the forward field (F) is then estimated, while the influence of evoked activity at other locations is kept separate (unknown factors x projected by unknown matrix A). Hyperparameters α and β help learn the correct dimension of A and B . (Right) Representation of factors influencing the data recorded at the sensors. In orange, a post-stimulus source at the voxel of interest, focused on by the lead field F . In red, other post-stimulus sources not at that particular voxel. In green, all background sources, including on-going brain activity, eyeblinks, heartbeat, and electrical noise. In blue, thermal noise present in each sensor.

in the next two sections. (See Jordan (1999) and Beal et al. (2003) for additional information on related graphical models.)

2.4.2 SAKETINI: Separation of background sources learned from pre-stimulus data

We learn the background mixing matrix B and sensor noise precision $\lambda = \text{diag}(\lambda_1, \dots, \lambda_K)$ from the pre-stimulus data using a variational-Bayesian factor analysis model. All variables, parameters and hyperparameters are hidden and

are learned from the pre-stimulus data. We assume Gaussian prior distributions on the background factors and sensor noise; we assume a flat prior on the sensor noise precision. We further assume the background factors are independent and identically distributed (i.i.d.) across time. The following is assumed for $n = -N_{pre} : -1$:

$$p(u) = \prod_n p(u_n); \quad p(u_n) = \mathcal{N}(u_n|0, I), \quad (2.43)$$

$$p(v) = \prod_n p(v_n); \quad p(v_n) = \mathcal{N}(v_n|0, \lambda), \quad p(\lambda) = \text{const.} \quad (2.44)$$

$$p(y|u, B, \lambda) = \prod_n p(y_n|u_n, B, \lambda); \quad p(y_n|u_n, B, \lambda) = \mathcal{N}(y_n|Bu_n, \lambda) \quad (2.45)$$

To complete the specification of this model, we need prior distributions on the model parameters. In the same way as for VBFA, a conjugate prior for the background mixing matrix B is used, as follows:

$$p(B) = \prod_{km} p(B_{km}); \quad p(B_{km}) = \mathcal{N}(B_{km}|0, \lambda_k \beta_m) \quad (2.46)$$

where β_m is a hyperparameter over the m th column of B and λ_k is the precision of the k th sensor, where again the hyperparameter β helps in model order (dimension) selection. The variational approximation is again used so that an analytic solution exists for the posterior updates.

$$p(u, B|y) \approx q(u, B|y) = q(u|y)q(B|y) \quad (2.47)$$

In the E-step, maximizing \mathcal{F} w.r.t. $q(u|y)$ holding $q(B|y)$ fixed gives the following posterior estimate update for the factors.

$$q(u|y) = \prod_n q(u_n|y_n); \quad q(u_n|y_n) = \mathcal{N}(u_n|\bar{u}_n, \gamma)$$

$$\bar{u}_n = \gamma^{-1} \bar{B}^T \lambda y_n; \quad \gamma = \bar{B}^T \lambda \bar{B} + K \psi_B^{-1} + I \quad (2.48)$$

In the M-step, we compute the full posterior distribution of the background mixing matrix B , including its precision matrix ψ_B , and the MAP estimates of the noise precision λ and the hyperparameter β .

$$q(B|y) = \mathcal{N}(B|\bar{B}, \psi_B)$$

$$\bar{B} = R_{yu} \psi; \quad \psi = (R_{uu} + \beta)^{-1}$$

$$\beta^{-1} = \text{diag}\left(\frac{1}{K} \bar{B}^T \lambda \bar{B} + \psi_B\right); \quad \lambda^{-1} = \frac{1}{N} \text{diag}(R_{yy} - \bar{B} R_{yu}^T) \quad (2.49)$$

2.4.3 SAKETINI: Localizing evoked factors learned from post-stimulus data

In the stimulus-evoked paradigm, the source strength at each voxel is learned from the post-stimulus data. The background mixing matrix B and sensor noise precision λ are fixed, after having been learned from the pre-stimulus data, as described in the previous section. Those quantities are assumed to remain constant through the post-stimulus period and are independent of source location. We assume Gaussian prior distributions on the source factors and interference factors. We further make the assumption that both of these factors are i.i.d. across time.

The source factors have prior precision given by the $Q \times Q$ matrix Φ , which relates to the strength of the dipole in each of Q directions. We have used the superscript r here as a reminder that s is for only one voxel at a time, although in general we will leave off this superscript.

$$p(s^r) = \prod_n p(s_n^r); \quad p(s_n^r) = \mathcal{N}(s_n^r|0, \Phi^r) \quad (2.50)$$

The interference and background factors are assumed to have identity precision (uninformed prior).

$$\begin{aligned} p(x) &= \prod_n p(x_n); \quad p(x_n) = \mathcal{N}(x_n|0, I); \\ p(u) &= \prod_n p(u_n); \quad p(u_n) = \mathcal{N}(u_n|0, I) \end{aligned} \quad (2.51)$$

We also use a conjugate prior for the interference mixing matrix A , where $\alpha = \text{diag}(\alpha_1, \dots, \alpha_L)$ is a hyperparameter, similar to the expression for the background mixing matrix, which helps in model order selection for determining the size of A from the data.

$$p(A) = \prod_{kl} p(A_{kl}); \quad p(A_{kl}) = \mathcal{N}(A_{kl}|0, \lambda_k \alpha_l) \quad (2.52)$$

We now specify the full model of the data:

$$\begin{aligned} p(y|s, x, u, A, B, \lambda) &= \prod_n p(y_n|s_n, x_n, u_n, A, B, \lambda); \\ p(y_n|s_n, x_n, u_n, A, B, \lambda) &= \mathcal{N}(y_n|F s_n + A x_n + B u_n, \lambda) \end{aligned} \quad (2.53)$$

Exact inference on this model is intractable using the joint posterior over the evoked non-localized factors x and mixing matrix A ; thus a variational-Bayesian

approximation for the posteriors is again used.

$$p(s, x, u, A|y) \approx q(s, x, u, A|y) = q(s, x, u|y)q(A|y) \quad (2.54)$$

All variables, parameters and hyperparameters are unknown and are learned from the post-stimulus data, iterating through each voxel across the brain, using the VB-EM algorithm. The equation for the free energy \mathcal{F} that is maximized in this problem is:

$$\begin{aligned} \mathcal{F}(q, \Theta) &= \int dx ds dA q(s, x, u, A|y, \Theta) [\log p(y, s, x, u, A|\Theta) - \log q(s, x, u, A|y, \Theta)] \\ &= \log p(y|\Theta) - \text{KL}[q(s, x, u, A|y, \Theta)||p(s, x, u, A|y, \Theta)] \end{aligned} \quad (2.55)$$

The derivations details are similar to that for SEFA (section 2.3). In the E-step, the posterior distribution of the background factors given the data is computed:

$$\begin{aligned} \Omega &= \begin{pmatrix} F & A & \bar{B} \end{pmatrix}; \quad \bar{\Omega} = \begin{pmatrix} F & \bar{A} & \bar{B} \end{pmatrix}; \\ \kappa_n &= \begin{pmatrix} s_n \\ x_n \\ u_n \end{pmatrix}; \quad \underline{\Phi} = \begin{pmatrix} \Phi & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{pmatrix}; \quad \Psi = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Psi_{AA} & 0 \\ 0 & 0 & 0 \end{pmatrix} \end{aligned} \quad (2.56)$$

Then, taking the derivative of optimal posterior $q(\kappa|y)$ gives:

$$\begin{aligned}
\frac{\partial}{\partial \kappa_n} \log q(\kappa|y) &= \frac{\partial}{\partial \kappa_n} E_{q(A|y)} [\log p(y, \kappa, A|\Theta)] = 0 \\
\frac{\partial}{\partial \kappa_n} E_{q(A|y)} [\log p(y|\kappa, A, \Theta) + \log p(\kappa|\Theta)] &= 0 \\
\frac{\partial}{\partial \kappa_n} E_{q(A|y)} \left[-\frac{1}{2} \sum_{n=1}^N ((y_n - \Omega \kappa_n)^T \lambda (y_n - \Omega \kappa_n) - \kappa_n^T \Phi \kappa_n) \right] &= 0 \\
E_{q(A|y)} [\Omega^T \lambda y_n - (\Omega^T \lambda \Omega + \Phi) \kappa_n] &= 0 \\
\overline{\Omega}^T \lambda y_n &= (\overline{\Omega}^T \lambda \Omega + \underline{\Phi}) \kappa_n
\end{aligned} \tag{2.57}$$

The SAKETINI VBE-step results as:

$$\begin{aligned}
q(\kappa_n|y_n) &= \mathcal{N}(\kappa_n | \bar{\kappa}_n, \Gamma); \\
\bar{\kappa}_n &= \Gamma^{-1} \overline{\Omega}^T \lambda y_n; \\
\Gamma &= \overline{\Omega}^T \lambda \overline{\Omega} + \underline{\Phi} = \overline{\Omega}^T \lambda \overline{\Omega} + K \Psi^{-1} + \underline{\Phi} \\
&= \begin{pmatrix} F^T \\ \bar{A}^T \\ \bar{B}^T \end{pmatrix} \lambda \begin{pmatrix} F & \bar{A} & \bar{B} \end{pmatrix} + K \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Psi_{AA} & 0 \\ 0 & 0 & 0 \end{pmatrix} + \begin{pmatrix} \Phi & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{pmatrix}
\end{aligned} \tag{2.58}$$

In the SAKETINI VBM-step, we update the posterior distribution of the interference mixing matrix A including its precision Ψ_{AA} . Note that the lead field F is fixed and known based on the geometry of the sensors relative to the head (same quantity assumed known by beamforming and other approaches) and \bar{B} was learned and fixed from the pre-stimulus data. The sensor noise precision λ is

also kept fixed from the pre-stimulus period.

$$\begin{aligned}
\frac{\partial}{\partial A_k} \log q(A|y) &= \frac{\partial}{\partial A_k} E_{q(\kappa|y)} [\log p(y, \kappa, A|\Theta)] = 0 \\
\frac{\partial}{\partial A_k} E_{q(\kappa|y)} [\log p(y|\kappa, A, \Theta) + \log p(A|\Theta)] &= 0 \\
\frac{\partial}{\partial A_k} E_{q(\kappa|y)} \left[-\frac{1}{2} \sum_{n=1}^N (y_n - F s_n - A x_n - \bar{B} u_n)^T \lambda (y_n - F s_n - A x_n - \bar{B} u_n) \right. \\
&\quad \left. - \frac{1}{2} \sum_{k=1}^K A_k \lambda_k \alpha A_k^T \right] = 0 \\
E_{q(\kappa|y)} \left[\sum_{n=1}^N (\lambda_k y_n x_n^T - \lambda_k F_k s_n x_n^T - \lambda_k A_k x_n x_n^T) + \lambda_k \bar{B}_k u_n x_n^T - A_k \lambda_k \alpha \right] &= 0 \\
\lambda_k R_{yx}^k - \lambda_k F_k R_{sx} - \lambda_k A_k R_{xx} - \lambda_k \bar{B}_k^T R_{ux} - A_k \lambda_k \alpha &= 0 \\
\lambda_k (R_{yx}^k - F_k R_{sx} - \bar{B}_k^T R_{ux}) &= A_k (\lambda_k R_{xx} + \lambda_k \alpha) \tag{2.59}
\end{aligned}$$

Thus, the posterior distribution for A is:

$$\begin{aligned}
q(A|y) &= \mathcal{N}(A|\bar{A}, \lambda \Psi_A) \\
\bar{A} &= (R_{yx} - F R_{sx} - \bar{B} R_{ux}) \Psi_A^{-1}; \quad \psi_A = (R_{xx} + \alpha) \tag{2.60}
\end{aligned}$$

The MAP values of the hyperparameter α and source factor precision Φ are learned here from the post-stimulus data.

$$\Phi^{-1} = \frac{1}{N} R_{ss}; \quad \alpha^{-1} = \text{diag}\left(\frac{1}{K} \bar{A}^T \lambda \bar{A} + \Psi_{AA}\right) \tag{2.61}$$

Most of the expressions for posterior covariance were defined previously in sec-

tion 2.3, so the remaining are defined here:

$$R_{yu} = \sum_{n=1}^N y_n \bar{u}_n^T, \quad R_{yx} = \sum_{n=1}^N y_n \bar{s}_n^T \quad (2.62)$$

$$R_{ss} = \sum_{n=1}^N \bar{s}_n \bar{s}_n^T + N\Sigma_{ss}, \quad R_{sx} = \sum_{n=1}^N \bar{s}_n \bar{x}_n^T + N\Sigma_{sx} \quad (2.63)$$

$$R_{su} = \sum_{n=1}^N \bar{s}_n \bar{u}_n^T + N\Sigma_{su} \quad (2.64)$$

where $\Sigma = \Gamma^{-1}$ is specified as:

$$\Sigma = \begin{pmatrix} \Sigma_{ss} & \Sigma_{sx} & \Sigma_{su} \\ \Sigma_{sx}^T & \Sigma_{xx} & \Sigma_{xu} \\ \Sigma_{su}^T & \Sigma_{xu}^T & \Sigma_{uu} \end{pmatrix} \quad (2.65)$$

In each iteration of VB-EM, the marginal likelihood is increased. The variational likelihood function (the lower bound on the exact marginal likelihood) is given as follows:

$$\mathcal{L}^r = \frac{N}{2} \log \frac{|\lambda| |\Phi^r|}{|\Gamma^r|} - \frac{1}{2} \sum_{n=1}^N \left(y_n^T \lambda y_n - \bar{x}_n'^r T \Gamma^r \bar{x}_n'^r \right) + \frac{K}{2} \log |\alpha^r| |\Psi^r| \quad (2.66)$$

This likelihood function is dependent on the source voxel r , since the variables are dependent on r , and thus a map of the likelihood across the brain can be displayed. Furthermore, an image of the source power estimates and the time course of activity at each voxel could be plotted.

The computational complexity of SAKETINI is on the order $O(KLNS)$, roughly equivalent to a single dipole scan, which is of order $O(N(K^2 + S))$. These are much smaller than the complexity of a multi-dipole scan which is order

$O(NS^P)$ where P is the number of dipoles, and if S represents roughly several thousand voxels. We further note that the number of hidden variables to be estimated is less than the number of data points observed, thus not posing problems for estimation accuracy. See Fig. 3.6 for performance as a function of data points.

2.4.4 Relationship of SAKETINI to MVAB

Minimum variance adaptive beamforming (MVAB) is one of the best-performing source localization techniques (Sekihara et al., 2001, 2005, 2002a; Küçükaltun-Yildirim et al., 2006). MVAB estimates the dipole source time series by $\hat{s}_n = W_{MVAB}y_n$, where

$$W_{MVAB} = (F^T R_{yy}^{-1} F)^{-1} F^T R_{yy}^{-1} \quad (2.67)$$

Thus, MVAB also has computational complexity equivalent to a single-dipole scan, on the order $O(N(K^2 + S))$. MVAB is known to fail when two sources are highly correlated in time (Sekihara et al., 2002a), although methods exist to correct this issue (Dalal et al., 2006; Brookes et al., 2007). In this section, we derive that MVAB is an approximation to SAKETINI.

We start by rewriting Eq. (2.42) as $y_n = F s_n + z_n$, where z_n is termed the total noise and is given by $z_n = A x_n + B u_n + v_n$. It has mean zero and precision matrix $\Upsilon = (A A^T + B B^T + \lambda^{-1})^{-1}$. Assuming we have estimated the model parameters

A , B , and λ , the MAP estimate of the dipole source time series is $\bar{s}_n = W y_n$, where

$$W = \Omega^{-1} F^T \Upsilon; \quad \Omega = F^T \Upsilon F + \Phi \quad (2.68)$$

The Φ term exists from the prior distribution of s_n . It can be shown that Eq. (2.68) is equivalent to Eq. (2.67).

In the infinite data limit, the data covariance satisfies $R_{yy} = F \Phi^{-1} F^T + \Upsilon^{-1}$. Its inverse is found, using the matrix inversion lemma (see Eq. (A.11)), to be $R_{yy}^{-1} = \Upsilon - \Upsilon F \Omega^{-1} F^T \Upsilon$. Hence, we obtain

$$F^T R_{yy}^{-1} = (I - F^T \Upsilon F \Omega^{-1}) F^T \Upsilon = \Phi \Omega^{-1} F^T \Upsilon \quad (2.69)$$

where the last step used the expression for Ω . Next, we approximate $\Omega \approx F^T \Upsilon F$, which is equivalent to the prior term on s_n having zero precision. We then obtain

$$\begin{aligned} W &\approx (F^T \Upsilon F)^{-1} F^T \Upsilon = (F^T \Upsilon F)^{-1} \Omega \Phi^{-1} \cdot \Phi \Omega^{-1} F^T \Upsilon \\ &= (F^T R_{yy}^{-1} F)^{-1} F^T R_{yy}^{-1} = W_{MVA B} \end{aligned} \quad (2.70)$$

where the last step uses Eq. (2.69).

2.5 NSEFALoc

The second main type of source reconstruction method developed using probabilistic graphical models is described here. Neurodynamic Stimulus Evoked Factor Analysis Localization (NSEFALoc) uses SEFA to find a set of temporal basis functions of which all stimulus evoked activity is assumed to be a linear mixture and then finds the best mixture (spatial weighting) of these basis functions

at each source voxel. Similar to SAKETINI, NSEFALoc uses a probabilistic hidden variable model that describes the observed sensor data in terms of activity from unobserved brain and interference sources. The unobserved source activities and model parameters are inferred from the data by a Variational Bayesian Expectation-Maximization algorithm. These techniques create an image of brain activity by scanning the brain, inferring the models from sensor data, and using them to compute the maximized likelihood of the data with the best set of parameters at each voxel, creating a spatial map to indicate the most likely locations of sources.

Like SAKETINI, the probabilistic models for the two versions of NSEFALoc are based on a physical description of neural activity, in which brain sources are modeled by current dipoles. For a given volume conductor model, the $K \times Q$ forward lead field matrix F^r represents the physical relationship between a dipole at voxel r and its influence on sensor $k = 1 : K$ (Sarvas, 1987). Unlike SAKETINI, both versions of NSEFALoc assume the source activity is a linear combination of $L \times N$ temporal basis functions ϕ computed from the data using SEFA (section 2.3), spatially weighted at each voxel by a $Q \times L$ dipole mixing matrix G^r .

Figure 2.2 depicts the generative model and the graphical models for the processing steps of NSEFALoc. Evoked brain activity, biological noise, other room interference and sensor noise all contribute to the measured sensor data (top row). The first step is to apply SEFA and, within the NSEFALoc framework, the evoked

factors are termed ϕ (formally x in the SEFA section 2.3). The second row of Fig. 2.2 depicts the graphical model for SEFA and the temporal basis functions ϕ learned from SEFA: the post-stimulus evoked factors separated from other ongoing background noise sources. Next, the graphical models for NSEFALoc1 and NSEFALoc2 are shown; either can be run on the data to obtain a likelihood map indicating the location of sources and the source time course estimates. Each method is now described in detail in the next two subsections.

2.5.1 NSEFALoc1: Spatial weighting of temporal basis functions

The NSEFALoc1 model and its solution are related to that proposed by Dogandzic and Nehorai (2000) and Baryshnikov et al. (2004). NSEFALoc1 differs from their work by precomputing the basis functions from the data (see section 2.3) and by placing a prior Wishart distribution on the full-rank precision matrix. The NSEFALoc1 generative model for the $K \times 1$ sensor data y_n is:

$$y_n = F^r G^r \phi_n + w_n^r \quad (2.71)$$

The noise w_n^r is modeled by zero-mean Gaussian with a general $K \times K$ precision matrix Λ . Both the parameters G and Λ are unknown. For a large number of sensors K , the precision matrix becomes quite large and difficult to infer accurately from the data. It may also become ill-conditioned. Hence, a prior probability using a Wishart distribution is used for Λ :

$$p(\Lambda) = \mathcal{W}(\Lambda|\nu, \Sigma_0) \propto |\Lambda|^{\nu/2} e^{-\frac{1}{2}Tr\Sigma_0\Lambda} \quad (2.72)$$

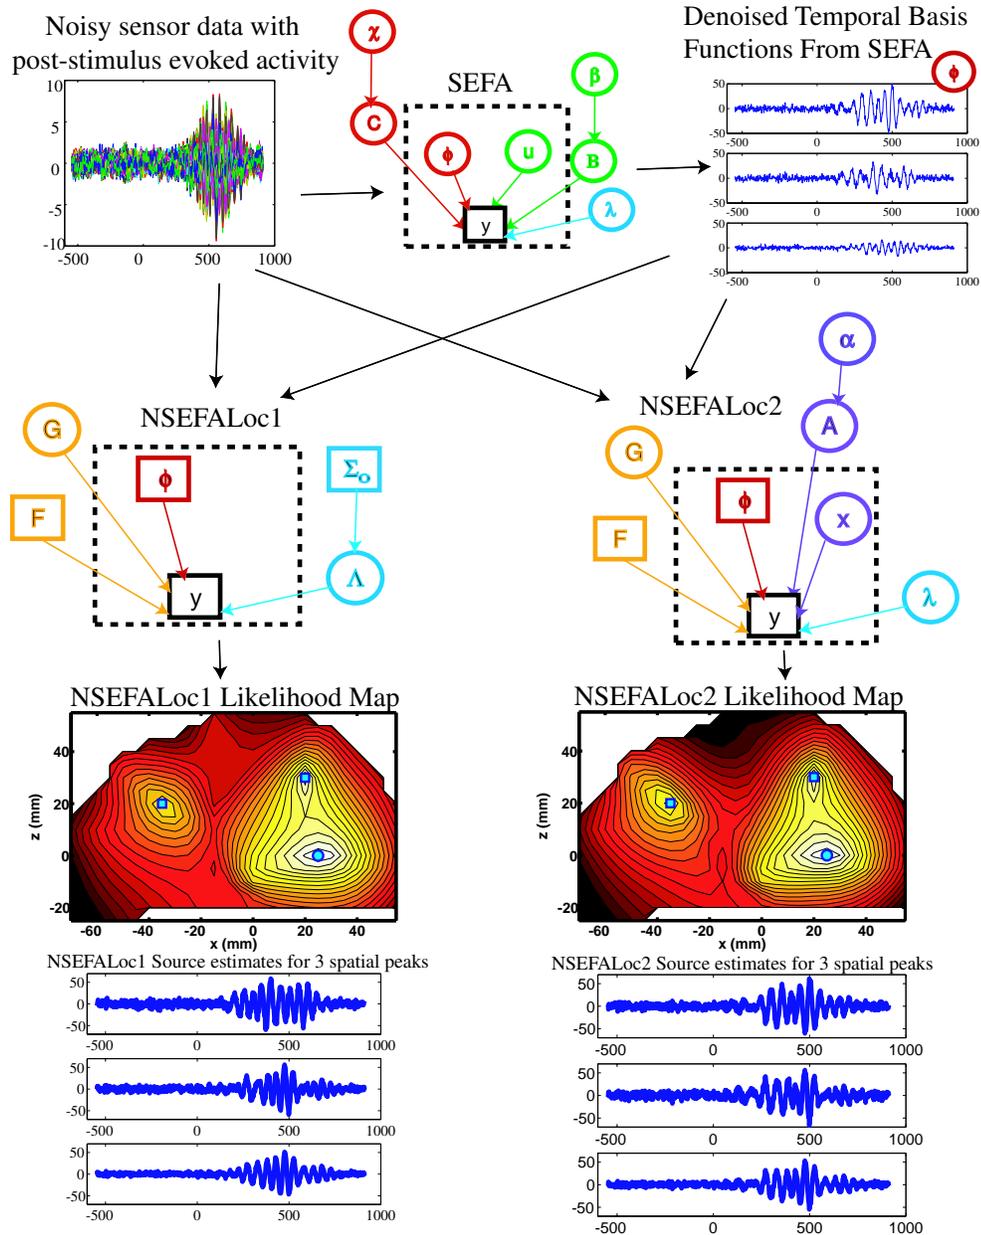


Figure 2.2 Graphical models for NSEFALoc1 and NSEFALoc2. Noisy sensor data is first processed by SEFA to determine the non-localized evoked factors, of which a linear mixture can produce any localized evoked source. These can be used as temporal basis functions ϕ to input as fixed bases to both NSEFALoc1 and NSEFALoc2, which estimate the spatial weighting G of these temporal basis functions for each voxel. The likelihood map can be displayed, and the source estimate at its spatial peaks can be plotted.

where Σ_0 and ν are hyperparameters. A Wishart distribution is related to a multivariate Γ distribution.

We choose $\nu = K + 2$ for the distribution to be normalizable. Whereas Σ_0 could be inferred by directly measuring the sample covariance, instead Variational Bayesian Factor Analysis (VBFA) is used on the pre-stimulus data (section 2.3). From VBFA, λ_0 is the diagonal sensor precision and B_0 is the interference mixing matrix, so $\Sigma_0 = (B_0 B_0^T + \lambda_0^{-1})^{-1}$.

For each scanned voxel, we consider the likelihood function:

$$\mathcal{L}^r = \log p(y|\Lambda) + \log p(\Lambda) \quad (2.73)$$

$$\log p(y|\Lambda) = \sum_{n=1}^N \log p(y_n|\Lambda)$$

The model parameters G^r and Λ^r are inferred from the data by maximizing \mathcal{L}^r . Again the superscript r indicating the dependence on voxel r is omitted for clarity.

$$p(y_n|\Lambda) = \mathcal{N}(y_n|FG\phi_n, \Lambda) \quad (2.74)$$

The estimates $\hat{\Lambda}$ and \hat{G} are the value of each that maximizes the likelihood \mathcal{L} .

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \Lambda} = \frac{\partial}{\partial \Lambda} \left(\frac{N}{2} \log |\Lambda| - \frac{1}{2} \sum_{n=1}^N (y_n - FG\phi_n)^T \Lambda (y_n - FG\phi_n) \right. \\ \left. + \frac{\nu}{2} \log |\Lambda| - \frac{1}{2} \text{Tr}(\Sigma_0 \Lambda) \right) = 0 \end{aligned} \quad (2.75)$$

Here we use equation A.5 so that the second term of the r.h.s. equals:

$$-\frac{1}{2} \sum_{n=1}^N \text{Tr} (\Lambda (y_n - FG\phi_n)(y_n - FG\phi_n)^T) \quad (2.76)$$

The above was rewritten so that equation A.2 could be used to simplify the derivative:

$$\frac{\partial \mathcal{L}}{\partial \Lambda} = \frac{N}{2} \Lambda^{-1} - \frac{1}{2} \sum_{n=1}^N (y_n - FG\phi_n)(y_n - FG\phi_n)^T + \frac{\nu}{2} \Lambda^{-1} - \frac{1}{2} \Sigma_0 = 0 \quad (2.77)$$

Solving for Λ^{-1} :

$$\hat{\Lambda}^{-1} = \frac{1}{N + \nu} (R_{yy} - FGR_{\phi y} - R_{y\phi}G^T F^T + FGR_{\phi\phi}G^T F^T + \Sigma_0) \quad (2.78)$$

This expression can be simplified by using the property from equation A.12 so that

$$\hat{\Lambda}^{-1} = \frac{1}{N + \nu} \sum_{n=1}^N (y_n - FG\phi_n)(y_n - FG\phi_n)^T + \Sigma_0 \quad (2.79)$$

$$= \frac{1}{N + \nu} \sum_{n=1}^N (y_n - FG\phi_n)y_n^T + \Sigma_0 \quad (2.80)$$

$$= \frac{1}{N + \nu} (R_{yy} - FGR_{\phi y} + \Sigma_0) \quad (2.81)$$

Now Λ^{-1} is a function of G (since G was assumed known when taking the derivative above); an expression for Λ^{-1} not dependent on G is needed. To solve for G (assuming Λ^{-1} is known):

$$\hat{G} = \operatorname{argmax}(\mathcal{L}) = \operatorname{argmax}(\log p(y|\Lambda)) \quad (2.82)$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial G} &= \frac{\partial}{\partial G} \left(\frac{N}{2} \log |\Lambda| - \frac{1}{2} \sum_{n=1}^N (y - FG\phi_n)^T \Lambda (y - FG\phi_n) \right) = 0 \quad (2.83) \\ &= \frac{\partial}{\partial G} \left(\sum_{n=1}^N (\phi_n^T G^T F^T \Lambda y + y^T \Lambda F G \phi_n - \phi_n^T G^T F^T \Lambda F G \phi_n) \right) = 0 \end{aligned}$$

Taking the derivative and using equations A.7, A.8, and A.9:

$$2F^T \Lambda R_{y\phi} - 2F^T \Lambda F G R_{\phi\phi} = 0 \quad (2.84)$$

$$F^T \Lambda (R_{y\phi} - F G R_{\phi\phi}) = 0 \quad (2.85)$$

$$G = (F^T \Lambda F)^{-1} F^T \Lambda R_{y\phi} R_{\phi\phi}^{-1} \quad (2.86)$$

Now G is a function of Λ . With two equations and two unknowns, it should be possible to find an expression for G independent of Λ . $F^T \Lambda$ appears in the expression for G , so pre-multiply $F^T \Lambda$ onto the expression for Λ^{-1} :

$$F^T = \frac{1}{N + \nu} F^T \Lambda (R_{yy} - F G R_{\phi\phi} + \Sigma_0) \quad (2.87)$$

Post-multiply equation 2.85 by $R_{\phi\phi}^{-1} R_{\phi y}$ and get

$$F^T = \frac{1}{N + \nu} F^T \Lambda (R_{yy} - R_{y\phi} R_{\phi\phi}^{-1} R_{\phi y} + \Sigma_0) \quad (2.88)$$

$$F^T S^{-1} = F^T \Lambda \quad (2.89)$$

where

$$S = \frac{1}{N + \nu} (R_{yy} - R_{y\phi} R_{\phi\phi}^{-1} R_{\phi y} + \Sigma_0) \quad (2.90)$$

Now, G can be rewritten:

$$G^r = (F^T S^{-1} F)^{-1} F^T S^{-1} R_{y\phi} R_{\phi\phi}^{-1} \quad (2.91)$$

Since G is now known, this value can be plugged into equation 2.78 to find Λ . The maximized likelihood is then:

$$\mathcal{L}^r = \frac{N + \nu}{2} \log |\Lambda^r| + \text{const.} \quad (2.92)$$

whose spatial peaks correspond to the most likely source locations.

2.5.2 NSEFALoc2: Spatial weighting of temporal basis functions

In contrast to NSEFALoc1, the contributions to sensor measurements not arising from a dipole source at the voxel r are now more explicitly modeled. The $L \times 1$ unknown interference factors $x_n^{\setminus r}$ correspond to activity in all voxels excluding r , A is a $K \times L$ unknown mixing matrix, and the sensor noise has unknown diagonal precision λ . The corresponding generative model for the sensor data is:

$$y_n = F^r G^r \phi_n + A^{\setminus r} x_n^{\setminus r} + v_n^r \quad (2.93)$$

The following conditional probabilities complete specification of the model:

$$p(y_n | x_n, A, \lambda) = \mathcal{N}(y_n | FG\phi_n + Ax_n, \lambda) \quad (2.94)$$

$$p(x_n) = \mathcal{N}(x_n | 0, I), \quad p(v_n) = \mathcal{N}(v_n | 0, \lambda) \quad (2.95)$$

Notice that in place of the $(K^2 + K)/2$ elements of the precision matrix Λ in NSEFALoc1, now just the $KL + K$ elements of A and λ need to be inferred from the data. Since typically $L \ll K$ ($L \lesssim 10$ and at UCSF $K = 275$), NSEFALoc2 has significantly less parameters and can thus be inferred more accurately.

The VB-EM algorithm is again used to infer the unknown quantities from the data. The E-step computes the posterior distribution $p(x_n | y_n)$ over the hidden variables x_n and the M-step updates the parameters G, A, λ . Specifically, in the E-step, $p(x_n | y_n)$ is found by taking the derivative of the free energy, where here

$\Theta = \{G, \lambda, \alpha\}$:

$$\mathcal{F}(q, \Theta) = \int dx dA q(x|y, \Theta)q(A|y, \Theta) [\log p(y|x, A, \Theta) + \quad (2.96)$$

$$\log p(x) + \log p(A|\lambda, \alpha) - \log q(x|y, \Theta) - \log q(A|y, \Theta)] \quad (2.97)$$

Recall that taking the derivative of \mathcal{F} w.r.t. $q(x|y)$ yields:

$$\log q(x|y) = E_{q(A|y)} (\log p(y, x, A|\Theta)) \quad (2.98)$$

and that $q(x|y)$ was already shown to be a Gaussian distribution under similar conditions. Furthermore, the mean of a Gaussian is the value that makes the derivative zero and the variance of the Gaussian is the slope of gradient:

$$\begin{aligned} \frac{\partial}{\partial x_n} E_{q(A|y)} [\log p(y|x, A, \Theta) + \log p(x)] &= 0 \\ \frac{\partial}{\partial x_n} E_{q(A|y)} \left[-\frac{1}{2} \sum_{n=1}^N (y_n - FG\phi_n - Ax_n)^T \lambda (y_n - FG\phi_n - Ax_n) \right. \\ &\quad \left. - \frac{1}{2} \sum_{n=1}^N x_n^T x_n \right] = 0 \end{aligned} \quad (2.99)$$

Moving the derivative inside the expectation and using Eqs. A.9, A.6 and A.13,

$$\begin{aligned} E_{q(A|y)} [A^T \lambda y_n - A^T \lambda FG\phi_n - A^T \lambda Ax_n - x_n] &= 0 \\ \bar{A} \lambda y_n - \bar{A}^T \lambda FG\phi_n &= (A^T \bar{\lambda} A + I) x_n \end{aligned} \quad (2.100)$$

Finally,

$$\begin{aligned} \bar{x}_n &= \Gamma^{-1} \bar{A}^T \lambda (y_n - FG\phi_n) \\ \Gamma &= A^T \lambda A + K\Psi + I \\ p(x_n|y_n) &= \mathcal{N}(x_n|\bar{x}_n, \Gamma) \end{aligned} \quad (2.101)$$

In the M-step, the full posterior over A is found by finding the $q(A|y)$ that best approximates $p(A|y)$ and the MAP estimates of the parameters G and λ and hyperparameter α are found. Thus, solving for A is a VBM-step and solving for G , λ , and α are M-steps.

$$\begin{aligned} \log q(A|y) &= E_{q(x|y)} (\log p(y, x, A|\Theta)) \\ \frac{\partial}{\partial A_k} E_{q(x|y)} [\log p(y|x, A, \Theta) + \log p(A)] &= 0 \\ \frac{\partial}{\partial A_k} E_{q(x|y)} \left[-\frac{1}{2} \sum_{n=1}^N (y_n - FG\phi_n - Ax_n)^T \lambda (y_n - FG\phi_n - Ax_n) \right. \\ &\quad \left. - \frac{1}{2} \sum_{k=1}^K A_k \lambda_k \alpha A_k^T \right] = 0 \end{aligned} \quad (2.102)$$

$$E_{q(x|y)} \left[\sum_{n=1}^N (\lambda_k y_n x_n^T - \lambda_k F_k G \phi_n x_n^T - \lambda_k A_k x_n x_n^T) - A_k \lambda_k \alpha \right] = 0 \quad (2.103)$$

$$\lambda_k R_{yx}^k - \lambda_k F_k G R_{\phi x} = A_k (\lambda_k R_{xx} + \lambda_k \alpha) \quad (2.104)$$

The posterior distribution of A is thus:

$$p(A|y) = \mathcal{N}(A|\bar{A}, \lambda\Psi)$$

$$\bar{A} = (Ryx - FGR_{\phi x} \Psi^{-1})$$

$$\Psi = R_{xx} + \alpha \quad (2.105)$$

To find the MAP estimate of G , the derivative of the free energy is taken:

$$\begin{aligned}
\frac{\partial \mathcal{F}}{\partial G} &= \frac{\partial}{\partial G} E_{q(x|y)} E_{q(A|y)} [\log p(y|x, A, G)] = 0 \\
\frac{\partial}{\partial G} E_{q(x|y)} E_{q(A|y)} &\left[-\frac{1}{2} \sum_{n=1}^N (y_n - FG\phi_n - Ax_n)^T \lambda (y_n - FG\phi_n - Ax_n) \right] = 0 \\
E_{q(x|y)} E_{q(A|y)} &\left[\sum_{n=1}^N (F^T \lambda y_n \phi_n^T - F^T \lambda FG\phi_n \phi_n^T - F^T \lambda Ax_n \phi_n^T) \right] = 0 \\
F^T \lambda R_{y\phi} - F^T \lambda FG R_{\phi\phi} - F^T \lambda \bar{A} R_{x\phi} &= 0 \\
\bar{G} &= (F^T \lambda F)^{-1} F^T \lambda (R_{y\phi} - \bar{A} R_{x\phi}) R_{\phi\phi}^{-1} \tag{2.106}
\end{aligned}$$

Or, one may plug in the value for \bar{A} to obtain:

$$\begin{aligned}
\bar{G} &= (F^T \lambda F)^{-1} F^T \lambda (R_{y\phi} - R_{yx} \Psi^{-1} R_{x\phi}) \\
&\quad (R_{\phi\phi} - R_{\phi x} \Psi^{-1} R_{x\phi})^{-1} \tag{2.107}
\end{aligned}$$

Solving for λ in NSEFALoc2 is very similar to solving for λ in VBFA, by letting:

$$y'_n = y_n - FG\phi_n \tag{2.108}$$

Then, take the derivative of \mathcal{F} w.r.t. λ to obtain a similar solution:

$$\lambda = N [\text{diag}(R_{y'y'} - \bar{A} R_{xy'})]^{-1} \tag{2.109}$$

The MAP estimate of α in NSEFALoc2 is identical to that for VBFA since it is only dependent on the prior term for A .

The maximized likelihood function for NSEFALoc2 is the following, where the dependency on voxel location is made explicit:

$$\mathcal{L}^r = \frac{N}{2} \log \frac{|\lambda/2\pi|}{|\Gamma|} + \frac{K}{2} \log |\alpha \Psi| - \frac{1}{2} \sum_{n=1}^N ((y_n - F\bar{G}\phi_n)^T \lambda (y_n - F\bar{G}\phi_n - \bar{x}_n^T \Gamma \bar{x}_n)) \tag{2.110}$$

2.6 Conclusion

The proposed methods have been described in a graphical model framework, which is a powerful and flexible technique for describing probabilistic dependencies between observed and unobserved quantities. We have chosen a scanning-based method to formulate the problem, rather than to solve the full tomographic problem which is very ill-parameterized. The variational-Bayesian Expectation-Maximization algorithm is used to solve for the values of the unknown quantities that maximize the marginal log likelihood. A map of the maximized likelihood is used to depict localized sources and estimates of their time course are extracted from those spatial peaks.

Acknowledgments

I am extremely grateful to Hagai Attias and Sri Nagarajan for the theoretical background and development of methods presented in this chapter.

Chapter 3

Probabilistic Methods for Neural Source Reconstruction: Performance in Simulations and Real Data ¹

3.1 Overview

From the introduction, it is clear that improved performance for noisy data with correlated sources is a desirable trait for a new source reconstruction method. The simulations and real data tested here were chosen to illustrate these issues and demonstrate improved performance of SAKETINI and NSEFALoc over existing methods. In simulations, several parameters such as location of sources,

¹Portions of this chapter originally appeared as an article in *NeuroImage* (Zumer et al., 2007b), NIPS conference paper (Zumer et al., 2007a), and ISBI conference paper (Zumer et al., 2006)

rotating or fixed dipoles, SNR, and type of background noise were varied. The effect of number of sensors and timepoints (total data available) was also tested. Finally, robustness to choice of number of basis functions or factors by the user is shown. Furthermore, performance of all methods is compared using some real-data examples from an auditory evoked MEG dataset and a low-SNR somatosensory MEG dataset.

Both localization ability and source time course estimation were used as the two main criteria for measuring performance in simulations. However, performance quantification is still an open question in the MEG literature. The localization error of a single simulated source can be calculated, but this becomes more difficult when multiple simulated sources are used, since the distinction between a mislocalized source versus a false positive is blurred. The existence of false positives or true negatives is dependent on the method of thresholding, which is not clearcut. Thresholds determined by bootstrapping, using subsets of trials and/or comparing to a control period, has been employed for some methods of source reconstruction (Sekihara et al., 2004; Darvas et al., 2005). Since a rigorous thresholding method has not yet been developed for these new probabilistic source reconstruction methods, care was taken to not let choice of threshold affect the results. Individual examples of performance are illustrated with the threshold set to zero, with qualitative conclusions drawn. Then, localization performance over the variety of simulation parameters mentioned above was quantified at multiple

thresholds and averaged.

Simulations and real data were analyzed using NUTMEG (Neurodynamic Utility Toolbox for MEG) (Dalal et al., 2004), a toolbox developed using MATLAB (MathWorks, Natick, MA, USA), obtainable from <http://bil.ucsf.edu>. NUTMEG is useful for coregistration of fiducial points to a structural MRI, selection of volume-of-interest, computation of forward field, filtering and other denoising preprocessing methods, as well as a variety of source reconstruction methods, including MVAB, sLORETA and now SAKETINI and NSEFALoc.

3.2 MEG Simulations

Simulations were created using a variety of realistic source configurations reconstructed on a 5mm voxel grid. A single-shell spherical volume conductor model was used to calculate the forward lead field (Sarvas, 1987). Simulated datasets were constructed by placing Gaussian-damped sinusoidal time courses at specific locations inside a voxel grid based on realistic head geometry. Sources were set to be active only during a post-stimulus period, which always comprised 62.5% of the total data available, while the remaining 37.5% was pre-stimulus data. Typically 700 total datapoints were used, unless specified otherwise.

In previous work testing algorithms, simulations are often designed with only simulated sources and Gaussian sensor noise. This clearly does not reflect true data which has interference sources somewhere in source space contribut-

ing to covariance across sensors. However, to test performance on this type of data and for better comparison with previous studies, this type of simulation was tested, termed “sensor noise only”.

In another set of simulations, termed “simulated interference” cases, background activity was drawn from the Gaussian distributions assumed by the model. The simulated sources mentioned above were projected to the sensors using the computed forward field. Noise drawn from a Gaussian distribution was added to each sensor across all time points. Furthermore, to simulate “ongoing brain activity,” time courses were drawn from a Gaussian noise distribution and were placed in 30 random locations throughout the brain voxel grid, occurring in both pre- and post-stimulus periods. Their activity was projected onto the sensors and added to both Gaussian sensor noise and source activity. These simulated background brain sources add noise to the sensors in a spatially-correlated manner.

In order to test simulation performance using data with more realistic (and unknown) statistical distributions, a second set of simulations was created termed “real brain noise.” Real MEG sensor data was collected from a CTF MEG System with 275 axial gradiometers while a human subject was alert but not performing tasks or receiving stimuli. This real background data thus includes real sensor noise plus real ongoing brain activity that could interfere with evoked sources and adds spatial correlation to the sensor data. Since throughout this work averaged

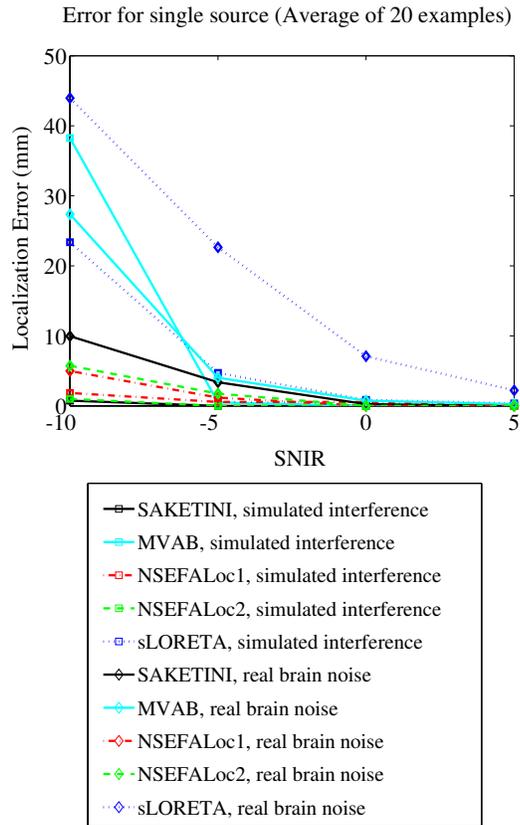


Figure 3.1 Localization error of all methods for a single source placed randomly in a realistic brain volume as a function of SNIR and type of background noise.

data is used, this real data was binned into 100 trials of 700 data points each and averaged. The Signal to Noise Ratio (SNR) and the corresponding Signal to Noise-plus-Interference Ratio (SNIR) were varied. SNIR is calculated from the ratio of the sensor data resulting from sources only to the sensor data from noise plus interference.

3.2.1 Localization accuracy of a single source

The localization error for a single active source was first examined. The single source was placed randomly within the voxel grid space. Fig. 3.1 shows localization error of SAKETINI and NSEFALoc relative to MVAB and sLORETA. The largest peak in the SAKETINI and NSEFALoc likelihood maps were found and the three-dimensional Euclidean distance from this point to the true source was recorded. Each datapoint is an average of 20 realizations of the source configuration and source time course. This simulation was performed for a variety of SNIRs and for both cases of simulated and real brain noise. This figure clearly shows that the error in localization is smaller for SAKETINI (black), NSEFALoc1 (red), and NSEFALoc2 (green) than for MVAB (cyan) and sLORETA (blue), especially in the highest noise setting. Errors for the “sensor noise only” case were not shown since they were zero or essentially zero for all methods for all values of SNR. Standard error for each data point was usually around 1mm, not worse than 4mm in the lowest SNIR. Most methods had higher localization error in the real brain noise case rather than simulated interference, however the lowest SNIR case for MVAB shows worse error for simulated interference. While the exact reason for this is unknown, it is possible that some of the simulated interference might have randomly had a high temporal correlation with the true source while the real brain noise did not.

3.2.2 Multiple sources, including correlated sources

Next, performance of the proposed models was tested for three simultaneously active sources. Fig. 3.2 shows performance in two examples with three uncorrelated sources each fixed in orientation across time, with real brain noise added with SNIR of 5dB. The differences between the two examples are the random realization of the source time course. Two examples are shown to illustrate the changes in results as the source time course changes. For each example, the likelihood (or power) map is on the left and the estimated time courses (cyan) of the three sources overlaid onto the true time courses (black) are shown on the right.

These examples in Fig. 3.2 show that all methods localize all three sources either perfectly or near perfectly. The spatial peaks for MVAB are so focal they are hidden by the blue dots. The sLORETA power map shows some difficulty in finding the lower left source perfectly, but does show a peak nearby. The SAKETINI likelihood map localizes all three sources perfectly, though in the left example, there is a slight secondary peak around $(x = -25, z = -15)$ that could show up as a false positive depending on the threshold. The NSEFALoc1 likelihood map finds all three sources perfectly, though in Example 2, there is a possible false positive around $(x = -25, z = 5)$. The NSEFALoc2 likelihood map also finds all three sources perfectly, but also with a possible false positive around $(x = -30, z = 35)$.

Fig. 3.2 also shows all methods' ability to estimate the source time course.

These examples show that SAKETINI and sLORETA estimate the shape and amplitude very well for all sources. NSEFALoc2 also estimates the time courses well, although there is some cross talk in Example 1, source R. NSEFALoc1 shows more severe crosstalk errors. MVAB estimates the time courses reasonably well, although a slight mis-estimation of amplitude is seen.

The performance of all methods was further tested when two of the three sources are highly correlated in time. Fig. 3.3 shows the results from two different simulation examples where two sources are highly correlated in time ($\rho = .95$). The same real brain noise was added at SNIR=5dB and all other aspects of this simulation were the same as the uncorrelated case above, except that the P'-source time course was adjusted to correlate strongly with the P-source time course. In all cases, the estimated time course plotted is the one extracted from the true location, regardless of the localization map peak locations.

These examples illustrate the failure of MVAB for correlated sources. The MVAB power map in Example 1 finds the uncorrelated Q-source, but largely mislocalizes the P-source and only weakly finds the P'-source. The reduction in power is seen in the P time course plot. The MVAB power map for Example 2 localizes all three sources within a reasonable error; however the amplitude of the peak location the two correlated sources (P and P') is much reduced and might not be detected depending on the threshold, which is also indicated by the large reduction in time series amplitude. On the other hand, sLORETA is not in theory

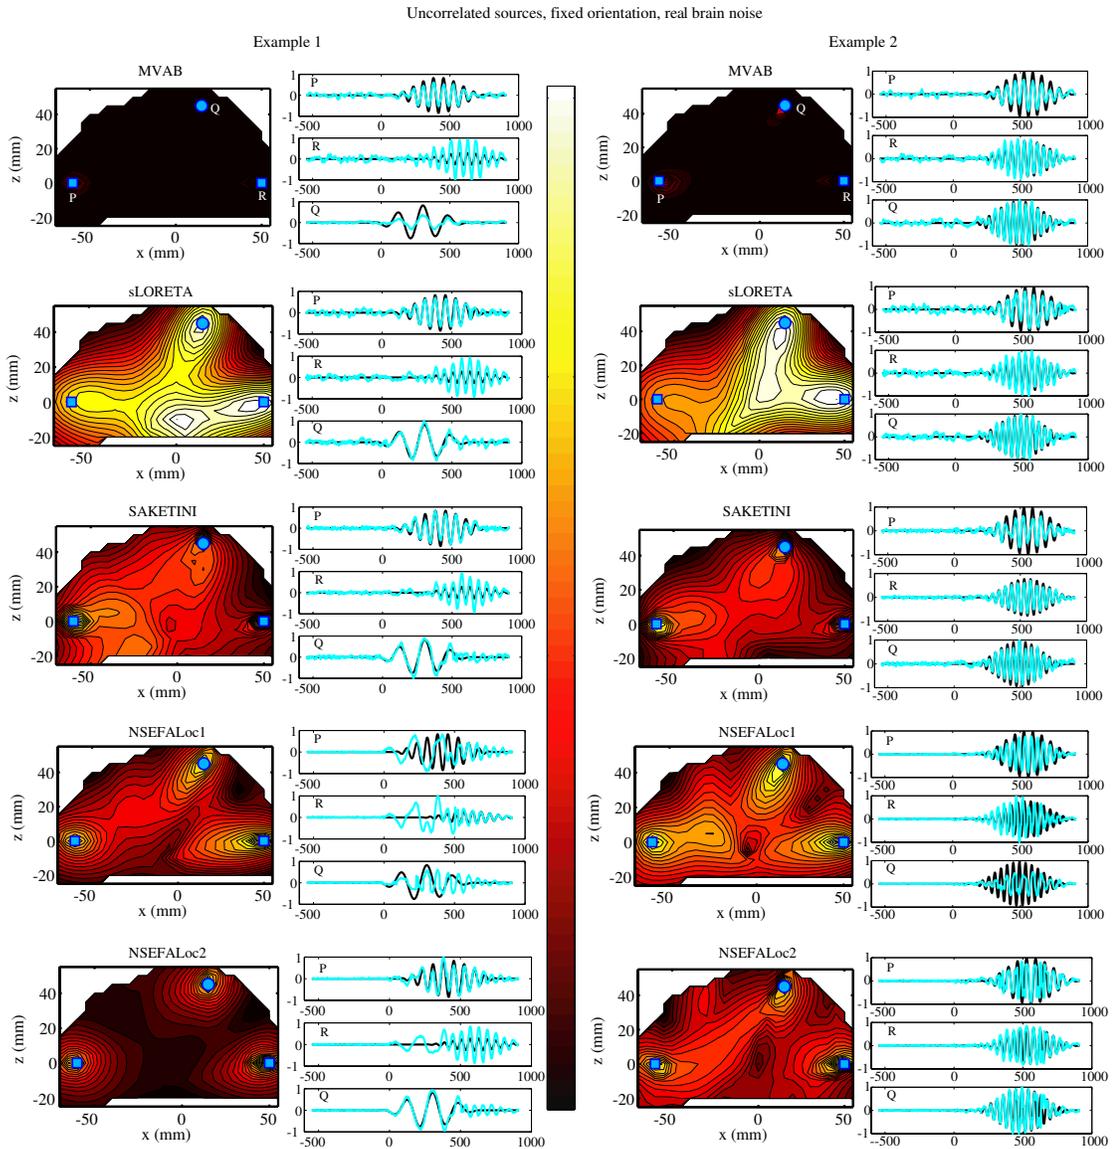


Figure 3.2 Performance of all methods for three uncorrelated sources in two different examples of source time course, with real brain noise added. (Left) Blue circles and squares indicate location of three true sources, labeled in the MVAB map as P, Q and R. Intensity of map corresponds to normalized likelihood map for SAKETINI, NSEFALoc1 and NSEFALoc2 and normalized power map for MVAB and sLORETA. (Right) Black lines indicate simulated time series for each of the three source locations, with the labels P, Q and R indicating the correspondence of time plot to location on the map. Cyan lines indicate each method's estimates of the source time series at those three locations.

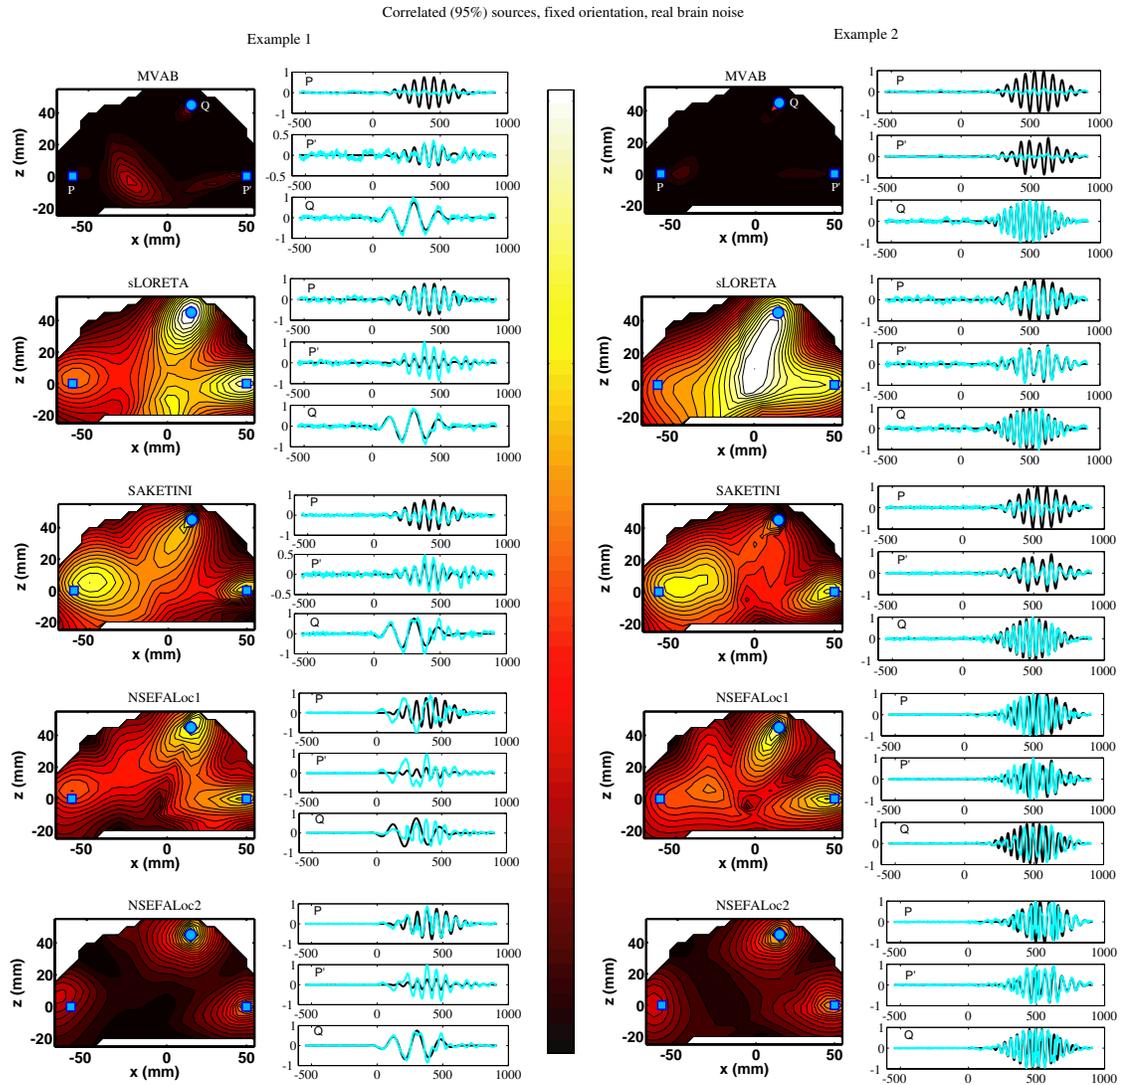


Figure 3.3 Performance of all methods when two of three sources are highly correlated in time. On the left for each example, blue squares, labeled P and P', indicate location of true sources highly correlated with each other, while the circle source (labeled Q) is uncorrelated with the other two. Intensity of map corresponds to normalized likelihood map (SAKETINI, NSEFALoc) or normalized power map (MVAB, sLORETA). In the right for each example, black lines indicate simulated time series for each of the three source locations; cyan lines indicate estimates of the source time series at those three locations. The labels P, P' and Q are included in each time series plot to indicate correspondance with the location on the map.

supposed to be sensitive to correlated sources; sLORETA finds all three sources in Example 1 (though one is weak and a center-of-the-head false positive is of larger amplitude) but, in Example 2, fails to show distinct peaks for the two correlated source locations P and P' . Despite these localization issues, the time course estimation of sLORETA in both shape and amplitude were very accurate.

Overall, the probabilistic methods SAKETINI and NSEFALoc2 localize the sources and estimate source time course better than MVAB and sLORETA in these examples of correlated sources in Fig. 3.3, while NSEFALoc1 performance is in-between. The SAKETINI likelihood map localizes all three sources with high amplitude in both examples, although the P -source is mislocalized slightly. The shape of the SAKETINI time course was very well estimated; the P -source time course is of lower amplitude due to the slight mislocalization of that source. NSEFALoc1 localizes the sources well in Example 1, but fails to find the P -source in Example 2, with a distant false positive instead. The time course estimates by NSEFALoc1 suffer from quite a bit of crosstalk. Finally, the NSEFALoc2 likelihood maps localize all three sources clearly with only a slight localization error in the P -source. Furthermore, the NSEFALoc2 time course estimation is quite good in both shape and amplitude, with much less crosstalk errors than NSEFALoc1 even though the same set of temporal basis functions were used.

3.2.3 Simulation: average performance

In addition to the above illustrative examples, average results over many realizations of simulations were examined. Several simulation parameters were varied across different simulations. (i) Two different source configurations were used: one with three sources more near the surface as depicted in Fig. 3.2 and the other configuration with three deeper sources. (ii) The orientation of the source was fixed in half the simulations and allowed to rotate over time in the other half. (iii) The correlation of two of the three sources with each other was set to be $\rho = 0$, $\rho = 0.95$, or $\rho = 1$; the third source was always uncorrelated with the other two sources. (iv) Each combination of parameters were tested for 10 different randomly generated source time courses and source orientations. (v) In addition to the true source contribution to the sensor data, the three cases of sensor noise only, simulated Gaussian interference or real brain noise were tested. (vi) SNIR was set at 5dB, 0dB, or -5dB, with corresponding SNR for each case of 10dB, 5dB, and 0dB. Thus, a total of 1080 simulations were run using all combinations of simulation parameters.

Performance was measured in two ways: localization ability and estimation of time course. To assess localization ability, it is important to take into account source strength, source localization error, and presence of false positives. Thus the ROC (receiver-operator characteristic) method was modified for brain imaging results as suggested by Darvas et al. (2004), which is a measure of hit rate versus

false positive rate. The free-response ROC (FROC) curve in particular allows for multiple “hits” per image (Bunch et al., 1978).

A “local peak” is defined here as a voxel that is greater in value than its 26 three-dimensional neighbors. A “hit” is defined as a local peak that is within a specified distance of the true location *and* above a certain threshold. A “miss” is defined as a true source location that has no “hit” within the specified distance. A “false positive” is a local peak above a certain threshold but further than the specified distance from a true source location. A “true negative” is any voxel that is none of the above.

FROC curves are generated by varying the threshold and allowable distance error, thus varying the tradeoff of sensitivity and specificity. The following distances were used as allowable localization error of a local peak to a true location in order to be counted as a hit: $5 * \sqrt{3}\text{mm}$, $10 * \sqrt{3}\text{mm}$, or $15 * \sqrt{3}\text{mm}$. The threshold was varied to be 30%, 50%, 70% or 90% of the maximum value in the whole image. Thus, a hit rate (HR) and false positive rate (FR) was recorded for each of 12 combinations of threshold/error for each of the 1080 simulations.

Since these HR vs FR points do not increase monotonically, as they would if threshold were the only criteria varied, we chose to use the measure of A' . A' is a way to approximate the area under the ROC curve for one HR/FR point (Snodgrass and Corwin, 1988). The larger the area under the ROC, the better the method is performing, since this means a higher HR relative to FR for specified

thresholds/localization errors. For each simulation, the twelve computed A' values were averaged to give one A' value per simulation. The SAKETINI likelihood map was used as the spatial map to test localization. For MVAB and sLORETA, the power map (sum of squares of post-stimulus time points) was used.

The performance of the proposed methods is now shown according to the metrics of A' (area under ROC curve) and time course estimation. Fig. 3.4(a) plots A' for each method, for each value of source correlation and SNIR, and for all types of interference. For uncorrelated sources in all cases of noise type, SAKETINI localizes sources the best according to A' . NSEFALoc1 and NSEFALoc2 both also show very high A' , close to SAKETINI, and higher than MVAB and sLORETA. For the perfectly correlated source cases, NSEFALoc2 localizes sources best for all noise types.

The other main test of performance was ability to estimate the source time course. The estimated time courses for all three methods were obtained from the true source locations, regardless of whether their respective localization maps found that source as a "hit." The correlation of the true time course with the estimated time course was computed for each simulation and the averages are plotted in Fig. 3.4(b). Similar to the A' results, for uncorrelated sources, SAKETINI estimates time courses better than other methods for all types of noise/interference. In the sensor-noise only case, sLORETA (blue) estimates the time course better than other methods regardless of source correlation, as previously understood for

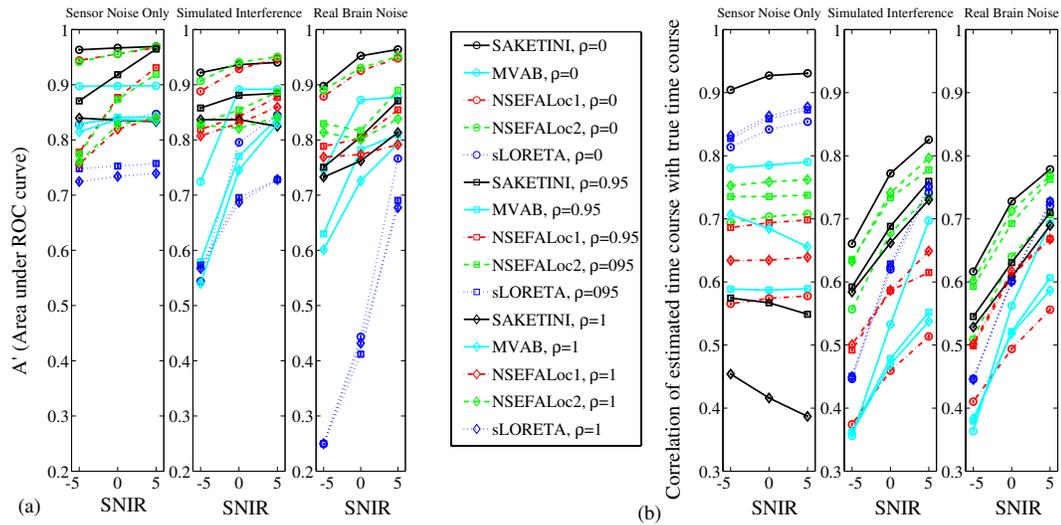


Figure 3.4 Performance of SAKETINI relative to MVAB and sLORETA for variety of simulated datasets. Each datapoint is an average of 40 simulations, consisting of two different source locations and either a fixed or rotating source orientation. Standard errors were less than 0.05 for all points (not shown). (a) A measure of area under ROC curve A' is plotted for the three methods as a function of SNIR for both simulated and real brain interference. See text for discussion of the A' metric. (b) The correlation of the estimated with the true time course is plotted for each method.

this method. However, when other source-space interference or real brain noise is added, this advantage of sLORETA is lost. Instead, similar to the A' results, NSEFALoc2 estimates the source time course the best when sources are perfectly correlated in simulated interference and real brain noise cases.

3.2.4 Simulation: effects of number of sensors and time points

Previous studies have shown advantage of sensor arrays with larger number of channels (Hamalainen et al., 1993). Likewise, increased amount of data points across time usually lead to improved estimation of unknown quantities. There-

fore, the next set of simulations sought to determine how few sensors and how few time points were needed to preserve the results found above.

To test the effect of the number of sensors, simulations were created similarly to those discussed above with three uncorrelated sources. Two values of SNIR were created using real brain noise: 0dB and -10 dB. Ten different realizations of source time course and orientation were tested for each case. All simulations discussed previously were created using the full 275 channel array from the CTF system. Here, only a random subset of sensors were selected, using 150, 74, 37, 25 or 15 sensors. The numbers 74 and 37 were specifically chosen to correspond to the BTi commercial MEG system previously installed in the UCSF lab until 2004 and from which some data in the next chapter was collected.

Fig. 3.5 shows simulation performance resulting from reduced number of sensors. Fig. 3.5(a) shows the A' metric as described above; Fig. 3.5(b) shows estimation accuracy of the source time courses. The top row in both (a) and (b) is for SNIR=0dB and the bottom row is for SNIR= -10 dB. SAKETINI, NSEFALoc1 and NSEFALoc2 did not show any major degradation in performance for either A' or time course estimation in the moderate SNIR value of 0dB; only minor decline in performance is seen for 25 sensors or less. For the very noisy case of SNIR= -10 dB, A' begins to decline more with only 37 sensors or less; time course estimation for the noisy SNIR= -10 dB is poor for all number of sensors. In contrast to the probabilistic methods, both MVAB and sLORETA show decline in performance for

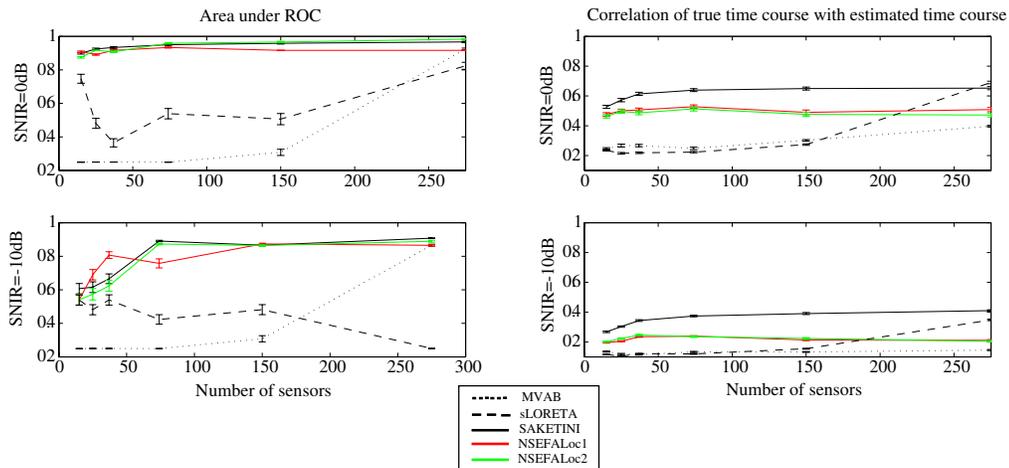


Figure 3.5 (a) A' and (b) time course estimation as a function of the number of MEG sensors for simulated data with 3 uncorrelated sources. The top row shows $\text{SNIR} = 0\text{dB}$ and bottom row shows $\text{SNIR} = -10\text{dB}$ using real brain noise. Error bars represent standard error.

both measures in the reduction from 275 to 150 sensors, but then plateaus for fewer sensors. In all cases of 150 sensors or fewer, all 3 probabilistic methods outperform MVAB and sLORETA for both metrics. Interestingly, the MVAB shows an improvement in source localization A' as the number of sensors becomes very small, probably due to less error in inverting the data covariance matrix.

To test the effect of the number of data points, the full set of 275 channels were used, but the available amount of data points was reduced. All previous simulations have used 700 total data points, where 62.5% were in the post-stimulus period. The ratio of data points in the post-stimulus period was kept the same, but the total number was reduced to 300, 200, 150, 100, or 50 time points.

Fig. 3.6 shows the performance results of all methods with decreased number of time points available. Fig. 3.6(a) shows the A' localization accuracy metric;

Fig. 3.6(b) shows the time course estimation. The top row in both (a) and (b) is for SNIR=0dB and the bottom row is for SNIR=-10dB. The A' results show that all three probabilistic methods outperform MVAB and sLORETA for all numbers of total data points; A' performance begins to decline for 150 or fewer data points. sLORETA is a non-data-dependent method thus the inverse weight is not affected by number of time points available. The MVAB is dependent on the data to provide an estimate of the data covariance matrix. Since the simulations in both top and bottom rows are with relatively high noise (SNIR = 0dB and -10dB, respectively), the data covariance estimate might not change much with decreased data, since it is already noisy (note the time course correlation does not reach above 0.5 for any number of data points tested at SNIR=-10dB).

On the other hand, the time course estimation results show that NSEFALoc1, NSEFALoc2 and sLORETA do not show a decline in performance with fewer data points and that all three methods generally perform equally well and better than MVAB or SAKETINI. This is most likely due to NSEFALoc methods not requiring many data points in the first step of using SEFA to find temporal basis functions; once the temporal bases have been found, less data is then needed for localization of these bases. SAKETINI is the only method to show a decline in time course estimation with fewer data points (less than 150 total). SAKETINI relies on an initial estimate of interference sources from the pre-stimulus data; perhaps this estimate is not as robust when only 30 or less pre-stimulus data points are present

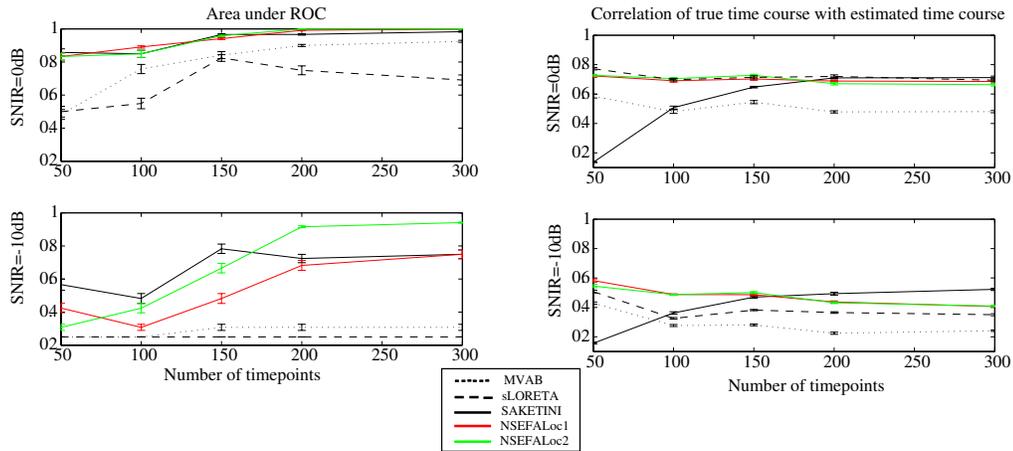


Figure 3.6 (a) A' and (b) time course estimation as a function of the number of total data points for simulated data with 3 uncorrelated sources. The top row of each shows SNIR = 0dB and bottom row shows SNIR = -10dB using real brain noise. Error bars represent standard error.

to estimate the 30 interference sources projected to 275 sensors, allowing for some cross-talk of interference with “evoked” time estimates.

3.2.5 Model Order and Initialization

In general, it is recommended to use as many sensors and time points as are available, and some care in data collection can contribute to this. However, the experimenter is still faced with choosing the number of factors and/or temporal basis functions for SAKETINI and NSEFALoc, even though the hyperparameters will zero-out extra dimensions to some extent.

Fig. 3.7 shows how the hyperparameters reduce the influence of extra columns of the mixing matrix when more than necessary are included. In this example, three sources were placed in a brain volume, uncorrelated with each

other and each orientation (of two) was uncorrelated (rotating dipole); thus six independent time courses were projected to the sensors, plus real brain noise added. SEFA was applied to this simulated sensor data with either $L = 4, 5, 6,$ or 7 columns for the evoked mixing matrix A . When more than enough evoked factors are used (seven in the the lower right panel), it is clear the extra factor's influence is much reduced (less than 5% of the largest). Furthermore, by including this extra component, the amplitude of α^{-1} does not change for the other components. Thus, this example indicates the usefulness of the hyperparameters in determining the number of factors to explain the evoked activity.

In practice, SEFA seems to be a very good way to obtain TBFs of the denoised evoked activity from real data when the true time course is not known. However, their use as in input to the NSEFALoc class of models should be tested. Fig. 3.8 compares NSEFALoc1 and NSEFALoc2 with three types of temporal basis functions: the true source time sources, those obtained from SEFA (as the models were intended) or from PCA. The simulations were the same as in Fig. 3.4 with three sources placed and three levels of correlation between two of the three sources, and either rotating or fixed dipole orientation. In all cases, the number of temporal basis functions was held the same (fewer were used when sources were known to be correlated for all three types of basis functions). In real brain noise and simulated interference, using PCA to obtain TBFs results in the worst performance for both metrics of A' and time course estimation. In comparing SEFA with the

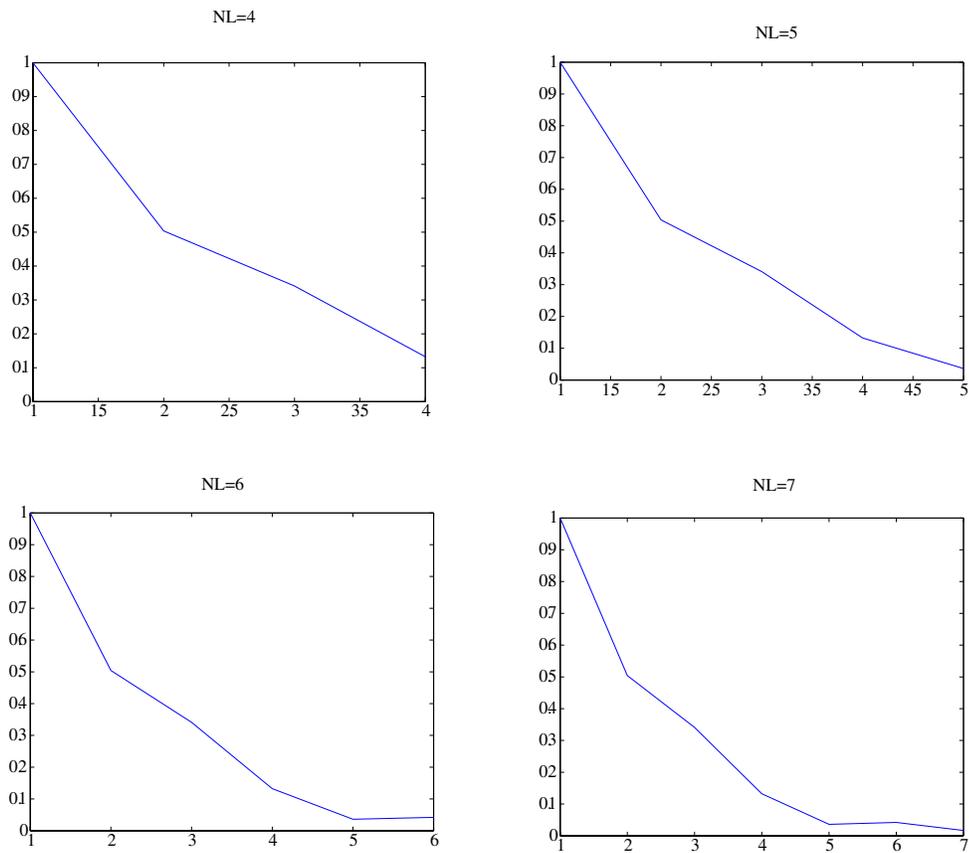


Figure 3.7 Number of temporal basis functions selected for NSEFALoc: Plot of normalized inverse of α hyperparameter from SEFA

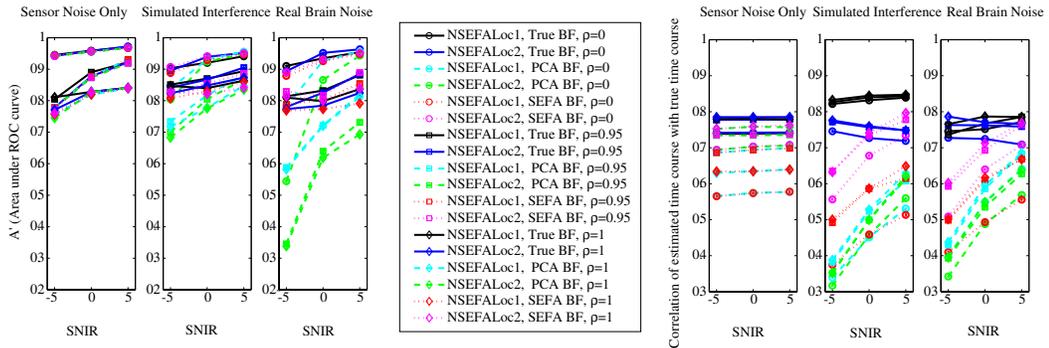


Figure 3.8 Performance of NSEFALoc1 and NSEFALoc2 as a function of three types of temporal basis function used: the true source time sources, those obtained from SEFA (as the models were intended) or from PCA. The plots on the left are of A' metric for localization ability and the plots on the right are for time course estimation accuracy (similar to Fig. 3.4).

true time courses, A' is not affected, but time course estimation is worse when using SEFA compared to true; however, NSEFALoc2 with SEFA TBFs performs reasonably close to the true TBFs (while NSEFALoc1 is considerably worse).

Fig. 3.9 shows examples of SAKETINI and NSEFALoc2 performance in the same simulated dataset while the number of non-localized evoked factors are varied. The main plots are of α^{-1} (which controls number of evoked factors not at the voxel of interest) normalized to the first hyperparameter (which is omitted from the plot). Each line within all plots is the value for each of 1760 voxels analyzed.

In the top row of Fig. 3.9, NSEFALoc2 was run with either 25 (upper left plot) or 10 (upper right plot) dimensions for A . From using 25 dimensions, it seems that the (inverse) hyperparameters get close to zero after about 10 and that the extra dimensions are not contributing much. By using only 10, the values stay roughly the same but do change somewhat. Furthermore, the localization

results (inset in each plot) are roughly the same yet some differences exist. The left inset (25 dimension result) shows all three sources with strong likelihood, but the lower source is blurred with the source above it, whereas the right inset (10 dimension result) shows all three sources as distinct peaks although the lower source is weaker, and also has a possible false positive near it.

In the bottom row of Fig. 3.9, SAKETINI was run with either 10 (lower left plot) or 5 (lower right plot) dimensions for A (the evoked non-localizing mixing matrix). It appears that as many as 10 might be needed in this case, since they don't tend to reduce to very low values until about the tenth one. The localization though, does not change; in both cases, the third lower source is not found as a separate peak.

Finally, Fig. 3.10 shows performance (through the A' metric and time course estimation) of all the methods as the number of dimensions was varied, averaged over many simulations. For MVAB, the x-axis represents the number of eigenvalues, for SAKETINI it is the dimension of the evoked non-localized factors x , for NSEFALoc1 and NSEFALoc2 it is the number of temporal basis functions, and it is meaningless for sLORETA whose performance does not depend on such a parameter. In all simulations, three sources are present with fixed dipole orientation; the bottom row two of the three sources are perfectly correlated; thus three dimensions is needed for the top row while only two is needed for the bottom row.

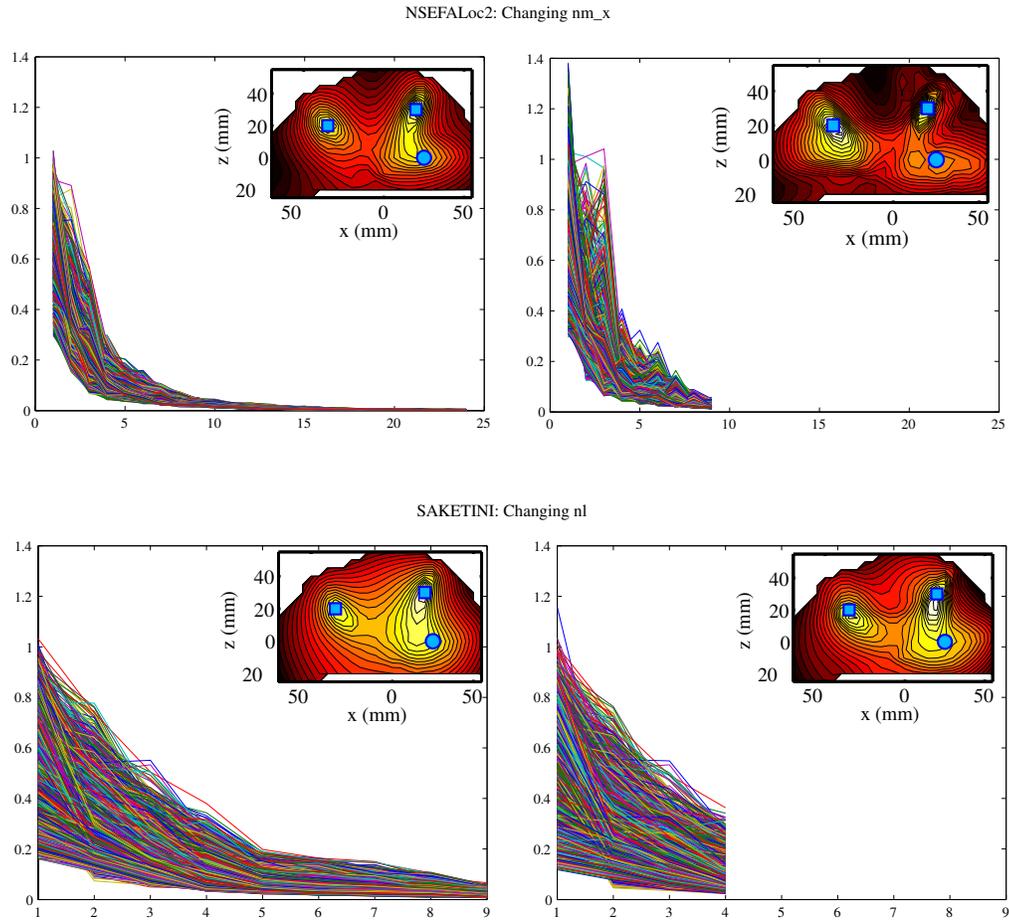


Figure 3.9 Plots of α^{-1} hyperparameter for NSEFALoc2 (top) and SAKETINI (bottom). Each line within all plots is the value for each of 1760 voxels analyzed. The first hyperparameter is normalized to one but is not shown. Inset in each plot is the localization result with the given number of dimension chosen, and blue symbols indicating correct location. All three sources were rotating orientation and uncorrelated with each other.

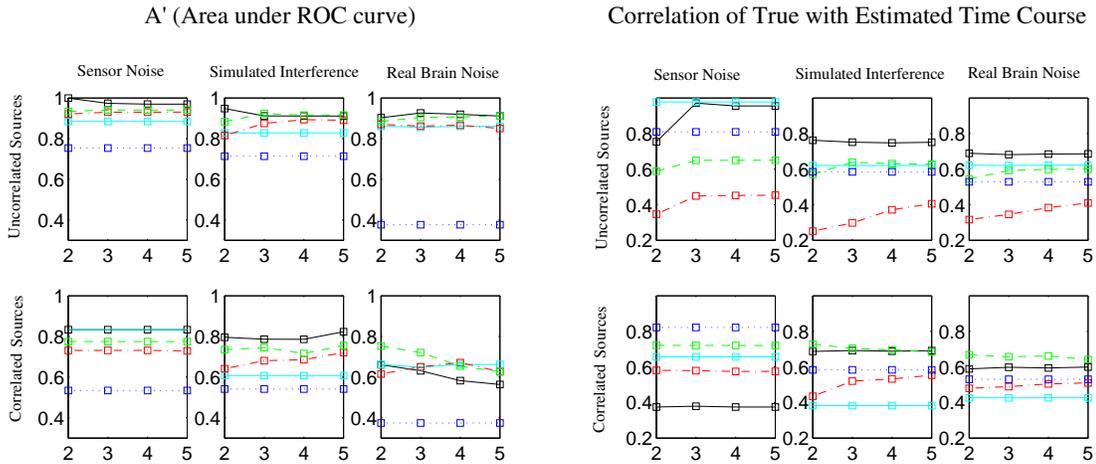


Figure 3.10 Performance as a function of number of temporal basis functions or evoked factors. Each line color is same as in Fig. 3.4 (black = SAKETINI, red = NSEFALoc1, green = NSEFALoc2, cyan = MVAB, blue = sLORETA).

Overall, the lines for all methods are relatively flat, indicating not too large of a dependence on number of dimension reduction. The time course estimation in sensor-noise only shows the clear improvement in using at least three dimensions. The correlated-source-case (bottom row) in interference or real brain noise show the clear advantage of both NSEFALoc2 and SAKETINI over MVAB for time course estimation. Interestingly, the A' metric gets worse for SAKETINI and NSEFALoc2 in correlated sources in real brain noise as the number of factors/TBFs increase; this is probably due to incorrectly trying to fit the extra components to the wrong location confused by the correlated sources.

3.3 Real MEG Data

Several real datasets were analyzed with the proposed method and compared to existing methods. For all data, the 275-channel CTF MEG System in a magnetically shielded room was used. All subjects gave written, informed consent to participate in each study, according to UCSF institutional review board approval.

3.3.1 Bilateral auditory data

The first dataset was obtained by presenting 120 repetitions of a 1kHz tone binaurally to a healthy subject, at an intertrial interval of 1.4s. The trials were averaged locked to stimulus onset. This auditory stimulus is known to invoke bilateral auditory cortex to be active simultaneously, causing problems for the MVAB's ability to localize the auditory sources.

The well-characterized auditory evoked M100 peak was observed and the sensor map for this latency is shown in Fig. 3.11(a). Bilateral auditory cortex is clearly activated. As mentioned above, MVAB performance deteriorates when neural sources are highly correlated in time and, in general, bilateral auditory cortex is known to be correlated (Dalal et al., 2006). Fig. 3.11(b) shows the MVAB performance on this dataset: the MVAB fails to find either left or right auditory cortex and instead finds a false source in the center. However, the other four methods are able to find bilateral auditory cortex (Fig. 3.11(c) - (f)). In this example, sLORETA and SAKETINI found both sides equally well. NSEFALoc1 and

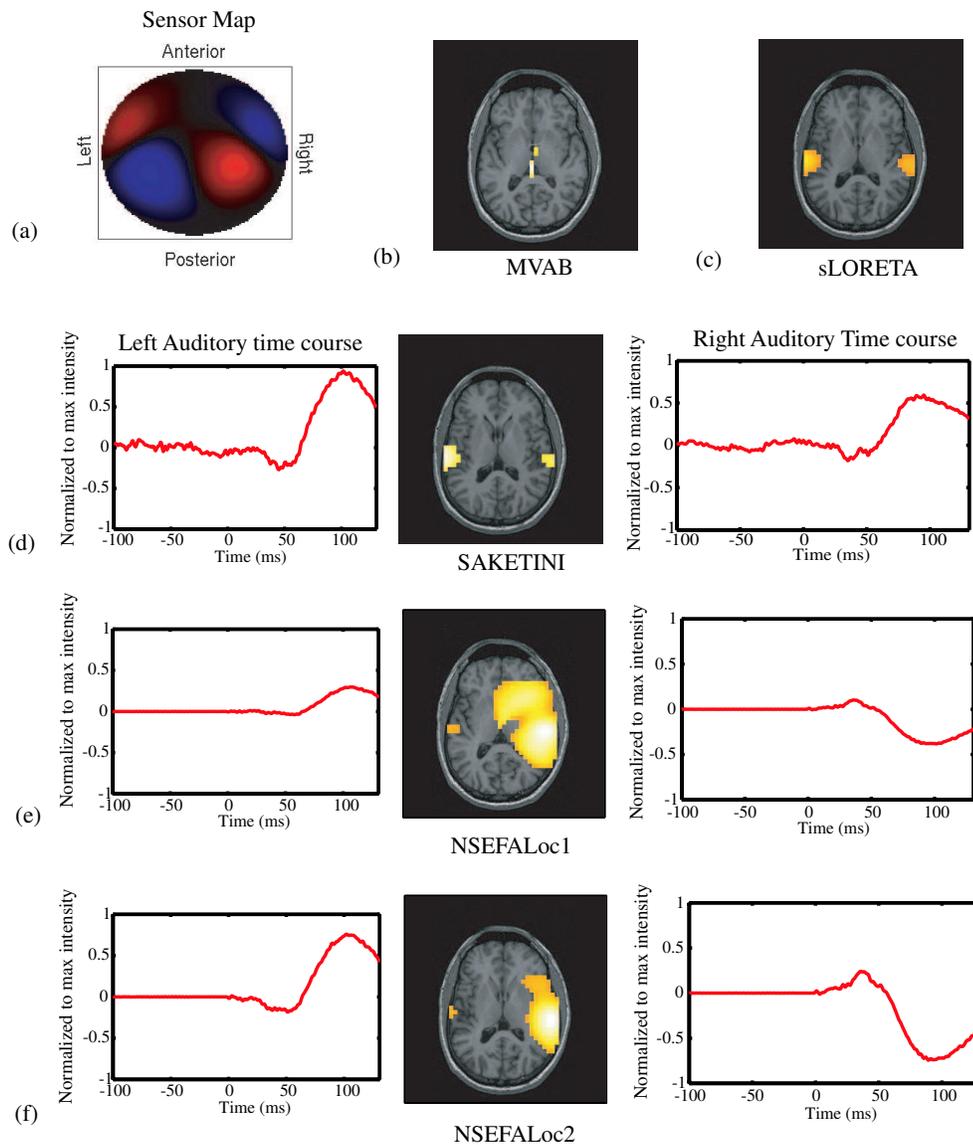


Figure 3.11 Performance of methods using a real auditory evoked MEG dataset. (a) Sensor map showing bilateral activation at M100 latency. (b) MVAB power map, showing failure to localize sources to auditory cortex. (c) sLORETA power map localizes bilateral auditory cortex. (d) SAKETINI likelihood map localizes activity in bilateral auditory cortex. Time series from left and right peaks are also shown, where intensity is normalized to maximum value across all voxels and timepoints. (e) and (f) NSFEALoc1 and NSFEALoc2 likelihood maps localize activity in bilateral auditory cortex, although the threshold had to be lowered to see the left peak. Time series from left and right peaks are also shown, where intensity is normalized to maximum value across all voxels and timepoints.

NSFEALoc2 found the right side stronger than the left, perhaps because the TBFs estimated from SEFA slightly favored the actual time course on the right over the left. Additionally shown in Fig. 3.11(d), (e), and (f) are the time courses of the SAKETINI, NSEFALoc1, and NSEAFLoc2 source estimates extracted from the peak voxel on each side. In the SAKETINI case, the correlation of the left with the right source estimates was found to be 0.93.

3.3.2 High noise somatosensory data

Next, a somatosensory dataset is used since the localization of primary somatosensory cortex is relatively easy for all methods when many trials are available to average. A small diaphragm was placed on the subject's right index finger and was driven by compressed air. The stimulus was given 256 times every 500ms. However, if the available data is restricted to only a small subset of trials, the lower SNR can become limiting for all methods. First, all methods were applied to the average of all 256 trials to assess performance for the standard (high) SNR case. Then all methods were applied to the average of only the first 5 trials. To further test if the performance was consistent across other sets of just 5-trial averages, all methods were also applied to the 5-trial average of trials 6-10, 11-15, and 16-20. The average was taken of these four different results. Any location found consistently will show up in the average.

The left panel of Fig. 3.12(a) shows typical somatosensory evoked MEG data

with the largest peak at 50ms, expected to be coming from primary somatosensory cortex in the posterior wall of the central sulcus. The next five panels of Fig. 3.12(a) show localization performance of all five methods, which all accurately localize activity to the contralateral primary somatosensory cortex. However, performance changes when only 5 trials are used in the average. The left panel of Fig. 3.12(b) shows the sensor data averaged over trials 1-5 of the same somatosensory dataset. The next five panels of Fig. 3.12(b) show errors in localization in all methods. However, SAKETINI, NSEFALoc1 and NSEFALoc2 all show the less error, relative to the peak location found using all 256 trials, than MVAB and sLORETA. Other averages of 5 trials showed varied performance, but, when averaging four different sets of 5-trial averages together, SAKETINI, NSEAFLoc1 and NSEFALoc2 showed localization fairly close to primary somatosensory cortex, as shown in Fig. 3.12(c), whereas MVAB and sLORETA mislocalize this source towards the center of the head.

3.4 Discussion

As seen in simulations and real data, the proposed models are not as sensitive to temporally correlated sources as the standard formulation of MVAB. However, it is possible to reduce the MVAB's dependence on correlated sources through a modified weight matrix computed subject to additional constraints, if a rough idea of the location of sources is known (Dalal et al., 2006).

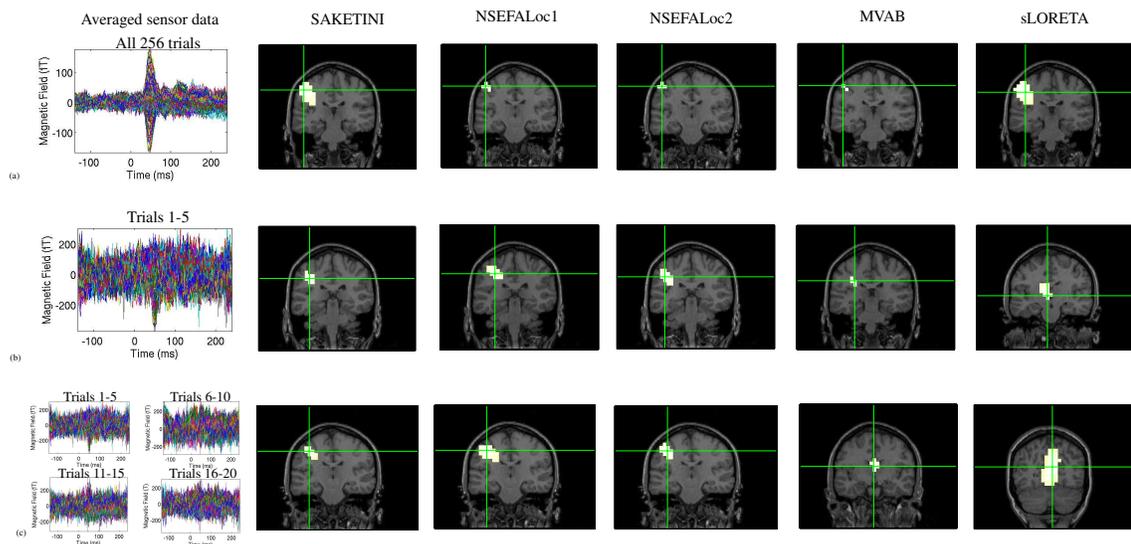


Figure 3.12 Performance of methods using real somatosensory data as a function of the number of trials. Left column shows sensor data averaged over varied number of trials, while remaining columns show localization performance of the five methods. Row (a) shows performance of the three methods applied to the average of all 256 trials. Row (b) shows the localization performance to the average of only the first 5 trials. In order to show performance over other subsets of 5-trial averages, the spatial maps in row (c) are spatial averages of the localization of 4 different 5-trial averages. See Methods for details. Crosshairs in localization maps show peak location within “active” voxels at the slice of peak location, where the threshold for “active” was defined at 90% of the maximum for all maps.

The models take advantage of knowledge of when sources of interest are not occurring (such as in the pre-stimulus period of a evoked response paradigm). They learn the statistical structure of the interference sources from the pre-stimulus period and then can suppress these signals in the post-stimulus period. Set in a probabilistic framework, the use of prior probabilities and hyperparameters assist in computation of unknown quantities. The flexibility of the models allows for many possible extensions.

A consideration of algorithms that iteratively update values is the longer computation time. Both SAKETINI and NSFEEALoc algorithms infer the model values for each voxel and convergence typically occurs after about 10-20 VB-EM iterations, which takes roughly one second per voxel on a standard Linux personal computer with 2.0GHz processor. Changing the number of channels or timepoints also affects computation time.

All methods which do not have a closed-form solution require initialization of the values to be iteratively updated. It was found that choice of initialization can change the final results somewhat but not largely, and these effects were not extensively examined. After finding one method of initialization that worked well in a few test simulations, that set was used for all results shown. Of note, in testing data with only sensor noise (no simulated or real brain interference at all), a greater dependence on initialization values was seen. This is worth mentioning since many source reconstruction algorithms are only tested on simulated sources

plus sensor noise, thus comparison with other methods on sensor-noise-only data should only be done with this in mind. However, since this type of data does not follow the assumptions of the models that the pre-stimulus data has spatially-covarying noise, and is therefore not realistic, analysis was not focused on sensor-noise only data.

Acknowledgments

I thank Sri Nagarajan, Hagai Attias, Kenneth Hild, Kensuke Sekihara and Ben Inglis for helpful discussions, Sarang Dalal for help with NUTMEG programming, Mary Mantle, Anne Findlay, and Susanne Honma for help with data collection.

Chapter 4

Comparison of Methods and Consistency of Results

4.1 Overview

The previous chapters have shown the mathematical detail and performance of new probabilistic methods for MEG source reconstruction. Performance was characterized in many simulations where the true source location was known, and SAKETINI and NSEFALoc were shown to be an improvement over MVAB and sLORETA, especially in data with high noise and correlated sources. In some simulations, real brain noise was added to make the simulations as realistic as possible. A few examples of performance in real data were shown. However, a more exhaustive comparison of the methods' performance in real data should

be performed, where assumptions of Gaussianity and stationarity could cause errors. Datasets where simple stimuli evoked responses in well-characterized areas, such as digit and lip somatosensory and auditory, were collected for this purpose. The methods' performances were characterized in several ways, where only some metrics are reliant on an estimate of the true source location.

The underlying neural causes of MEG data are generally agreed to be from the currents flowing in the apical dendrites of large pyramidal neurons (Hamalainen et al., 1993). However, the exact link between local neural activity as measured with classical electrophysiological methods and noninvasive neurodynamic data is not clear. An initial study by Teale et al. (1994) has examined the auditory evoked fields recorded with MEG in a nonhuman primate and localized the responses to auditory cortex. However, this animal model needs to be extended to other brain areas as well as validation of sources from invasive methods in the same animals.

An on-going project with Dr. Nagarajan, Dr. Disbrow and others is to develop a non-human primate model for examining the underlying neurophysiology of signals recorded in a wide variety of MEG experiments. An initial feasibility study has been performed testing whether spatially separate regions of finger, lip and auditory areas could be distinguished using MEG data recorded from anesthetized monkeys given tactile or auditory stimulation. Detailed analysis of the results of the MEG data, both for species comparison and the separability of

the brain regions, are presented in the next chapter, along with use of these MEG reconstructions as a seed location for diffusion-tensor MRI tractography.

The focus of this chapter is the comparison of the methods developed in this dissertation against existing methods in these additional real-data examples, both in humans and in monkeys. The comparison is performed on two forms of the averaged data: broad/unfocused (all sensors onto all brain voxels) and focused for optimal detection of single primary cortical field (select sensors onto focused brain regions). Since the true neural location is not obtainable for the healthy human subjects and the ground-truth electrophysiology locations from the monkeys are not yet available, the goal for this analysis is to test which methods give the most reasonable results, roughly in the correct region and with minimal extra regions. The monkey data includes more interference sources than the human data, thus testing the noise suppressing properties of the new proposed methods.

4.2 Somatosensory and Auditory Evoked MEG data in Human and Nonhuman Primates

4.2.1 Data acquisition and processing

Four healthy human volunteers (three male and one female, ages 23 – 28) participated in both MEG and MRI experiments which were performed on separate days. Four adult male macaque monkeys (9 – 15 kg) were scanned using

both MEG and MRI on separate days. Localization of primary somatosensory and auditory cortex obtained using MEG was coregistered to MR data. All studies were performed with approval of the UCSF Committee on Animal Research or the Committee on Human Research.

Stimuli

Tactile and auditory stimuli were used in MEG experiments to localize the corresponding primary sensory cortices. Tactile stimulation was delivered by compressed air-driven balloon diaphragms clipped onto either the hands or lips. In the human, two clips were placed on the right lateral upper lip and two clips on the right lateral lower lip; all four clips were activated simultaneously. Taps for both hand and lip (duration 140ms, 25PSI for monkeys and 18PSI for humans) were presented in blocks of 256 trials with stimulus-onset-asynchrony (SOA) 1000ms \pm 50ms for monkeys and SOA 500ms \pm 50ms for humans. White noise at low levels was presented continuously to the human subjects during tactile stimulation via binaural earplugs to mask the sound made by the tappers; monkeys wore earplugs for this purpose. Auditory stimulation consisted of bursts of white noise (350Hz - 17.6kHz, 400ms duration) delivered via binaural earplugs, with 128 epochs per block and an SOA of 1500ms \pm 50ms.

For hand stimulation in the monkey, two clips delivering simultaneous taps were placed on the tip of the thumb (digit one, D1) and the thenar area; three monkeys (M2, M3 and M4) were stimulated on the right hand and the other monkey

(M1) on the left hand. In the humans, four clips delivering simultaneous stimulation were placed on the subject's right hand: one each on the distal and middle segments of D2 and D3. For lip stimulation in the monkey, two clips were placed on lateral upper and lower lips, on the same side as the hand that was stimulated.

MEG acquisition: Humans

All human subjects were scanned using the same 275-channel CTF Omega 2000 system with 3rd-order gradient correction (VSM MedTech, Coquitlam, B.C., Canada) used for data discussed in Chapter 3. The localizing coils were taped on the subjects head prior to placement in the whole-head helmet sensor array. One coil was placed at the nasion, and two additional coils were placed 1cm away from the left and right pre-auricular points in the direction of the nasion. The position of these coils was measured relative to the sensor array before and after every block of trials. In all blocks, the subjects moved less than the criterion of 5mm. These points were later coregistered with the high-resolution structural MR image. On separate days, a structural MRI and diffusion tensor image were collected (for details, see Section 5.4.2).

Somatosensory data were collected in 400ms epochs, with a 150ms pre-stimulus period, sampled at 1200Hz. Auditory data were collected in 600ms epochs with 100ms pre-stimulus period, and were sampled at 1200Hz. Each human subject was scanned only once.

MEG acquisition: Monkeys

Macaque monkeys were initially anesthetized using intra-muscular (I.M.) ketamine hydrochloride (10mg/kg) and anesthesia was maintained with boluses of ketamine hydrochloride given intra-venously (I.V.) or I.M. every 15 – 30 minutes (3–5mg/kg/hr) and optionally boluses of I.V. midazolam (0.05–0.36mg/kg/hr). Heart rate, respiratory rate, temperature and SpO₂ were documented every 30–60 minutes. Heated water bottles and blankets were used to keep the animal warm. An experimenter stayed in the shielded room to monitor the monkey at all times. Any large movements that might cause a large change in head position relative to the sensors, due for example, to a cough, yawn or global arm movements, were noted and that data block was not analyzed. Sessions lasted 4-6 hours including preparation time, and then monkeys were returned to their home cage for recovery.

In a shielded room, a 37-channel magnetometer (BTi, San Diego, CA) was used for MEG recordings of two monkeys (M1 and M2). The monkey's head geometry of the bony portions of the forehead and longitudinal fissure was recorded using a digital sensor position indicator (FASTRAK, by Polhemus, Colchester, VT). Central forehead and left and right pre-auricular points were landmarked to define the MEG spatial reference frame for later coregistration with the structural MR image. The sensor array was centered over the central sulcus contralateral to the site of stimulation. Somatosensory data were collected in 300ms epochs, with

150ms pre-stimulus period, sampled at 300Hz. Auditory data were collected in 600ms epochs, with 100ms pre-stimulus, sampled at 1kHz. M1 was scanned on two separate days while M2 was scanned only one day.

The same 275-channel CTF Omega 2000 system used for humans was also used to scan monkeys M3 and M4. Three localizing coils were glued and taped onto the central forehead and 1cm from the left and right pre-auricular points in the direction of the central coil. Digitized headshape was not necessary in addition to the three localizing coils. At the start and end of each trial, current ran through the coils to permit precise localization of those three points relative to the sensor array. These points were later coregistered with the high-resolution structural MR image, which included MR-visible fiducials at the same locations. Somatosensory data were collected in 400ms epochs, with a 150ms pre-stimulus period, sampled at 1200Hz. Auditory data were collected in 600ms epochs with 100ms pre-stimulus period, and were sampled at 1200Hz. For all monkeys, sometimes the same stimulation was repeated on the same day or on separate days.

On separate days, a structural MRI and diffusion tensor image were collected on these same monkeys (for details, see Section 5.4.2).

Pre-processing of human data and forward field computation

The Neurodynamic Utility Toolbox for MEG (NUTMEG; (Dalal et al., 2004)) was used for all data pre-processing. The three fiducial points were coregistered with the structural MRI. The spherical homogeneous single-shell model was fit-

ted to the scalp derived from the structural MRI. Individual trials with eyeblinks or other movement artifacts were discarded prior to averaging. For the remaining blocks, after removal of trials with large artifacts, the raw sensor data was averaged and the mean of the pre-stimulus period was subtracted from the whole waveform.

Comparisons across all five reconstruction methods were made on the full dataset (all sensors included) and reconstructed over the whole brain at a 5mm-spaced grid. Roughly 250 post-stimulation time points were used. This human data was not very noisy, but using all sensors could lead to problems in the reconstruction by MVAB due to correlated sources (especially in auditory data, but also between primary and secondary somatosensory cortex). At least two brain areas are expected active in each dataset.

To focus the data on just the contralateral primary sensory cortex, sources were estimated over a 2mm-spaced 3D grid over a focused volume of interest (VOI). (2mm was also the resolution chosen to match up with creating a seed for the DTI tractography described in the next chapter.) Due to computational limits for MVAB, the VOI was roughly one tenth of the total brain volume, but large enough to encompass the primary somatosensory and auditory areas and surrounding regions. All sensors were used to reconstruct somatosensory activations, whereas only sensors on the hemisphere of interest were used for auditory data due to the correlated source problem for MVAB. All five reconstruction

methods were tested and compared on this reduced/focused dataset.

The individual subject's structural MRI was normalized to MNI space using SPM2 (<http://www.fil.ion.ucl.ac.uk/spm>). The same spatial warping was then applied to the MEG reconstruction (as implemented in NUTMEG) and the MNI coordinate of the peak MEG voxel was noted as well as the peak latency.

Pre-processing of monkey data and forward field computation

NUTMEG was also used for monkey data pre-processing, which involved the same steps as for human data. Individual trials with eyeblinks or other movement artifacts were discarded prior to averaging. Any stimulation block during which large movements were observed was not analyzed further. For the remaining blocks, after removal of trials with large artifacts, the raw sensor data was averaged and the mean of the pre-stimulus period was subtracted from the whole waveform.

The fiducial points, and digitized headshape if available, were coregistered with the structural MRI. A VOI was defined to include the whole cerebrum. A grid size of 2mm was chosen since this is within (or less than) the normal error of MEG reconstruction, 2mm roughly corresponded to the voxel size of the diffusion tensor images discussed in the next chapter, and any smaller grid would have put computational limits on the total volume reconstructed using the MVAB approach.

The forward lead field was calculated based on a single-shell spherical vol-

ume conductor model, of which the sphere center must be defined. In humans, the center of the head and the center of the sensor array roughly coincide and either one can be used. Due to the much smaller head size of monkeys, it was found that a sphere center roughly half way between that defined by the head and the sensor curvature gave the most consistent results. The same coregistration, head model and computed forward field were used for all five methods (MVAB, SAKETINI, NSEFALoc1 and 2, and sLORETA).

The localization of the MEG data on both structural MR images was performed for the two monkeys that it was available (M3 and M4). Since the coregistration of the monkey MEG data to the structural MRI for purpose of computing the lead field is more difficult and thus more variable, it was important to include localizations from both MR images since one was not any *a priori* better than the other. Thus, the total number of monkey MEG localizations examined in the “focused” case was 13 digit sets, 17 lip sets, and 6 auditory sets.

The time window using roughly the first 50 – 70 time points post-stimulation was used to calculate the data covariance for MVAB analysis; this same window was also used at the post-stimulus period for SAKETINI and NSEFALoc, which used roughly 30 pre-stimulus time points.

Since the monkey’s head is smaller than a human’s head, most of the sensors were not near the head and were therefore just recording room noise rather than brain signals, as well as including very large heartbeat artifact. Therefore, using

all sensors is a good test of the denoising prospects of SAKETINI and NSEFALoc.

The second comparison was to focus only on reconstruction of the contralateral primary sensory area. For M3 and M4 scanned in the 275-channel MEG system, only the 37 channels closest to the cortical region of interest were chosen for three reasons. The sensors in the 275-channel system are spaced about as far apart as the 37-channel system, so there was not a difference in the tradeoff of spatial coverage versus spatial sampling; thus it is fair to include data from both MEG machines. Heartbeat artifacts still affected these 37 sensors in all four monkeys, so this data also tested the ability of the new algorithms to remove this noise source.

In human brain imaging studies, the individual subject's data is often warped to a standard normalized MNI/Talairach space. However, since such an MRI template does not yet exist for the macaque monkey brain, each animal's structural MRI was registered to the "AC-PC coordinate" frame using just translations and rotations (no warping). This coordinate system is defined with the anterior commissure (AC) as the origin, the y-axis running through the posterior commissure (PC), and the z-axis running through longitudinal cerebral fissure.

4.2.2 Sensor data results

Clear response peaks were seen in all three types of stimulation in all four human subjects. The sensor data appeared consistent with previous studies of somatosensory and auditory stimulation in humans. A discussion of latencies

and separability of areas is in the next chapter. The averaged sensor data from one human subject (H4) are shown in Fig. 4.1 as an example. The left column shows a two-dimensional representation of the sensor map at the peak time latency; the right column shows the time series from all sensors used for localization in the “processed” data case.

Response peaks were also consistently observed in the monkey averaged MEG sensor data. A clear response in the sensor data was observed for digit stimulation in 10/10 blocks total from all animals. Similarly, a clear response from lip stimulation was observed in 12/13 blocks. Clear auditory responses were seen in 4/4 runs from monkeys M1 and M3. However, only 1/4 auditory blocks showed a clear peak in M2 and 0/5 auditory blocks in M4 showed a distinguishable peak.

Fig. 4.2 shows an example set of digit, lip and auditory responses in sensor data from M1 in order to introduce the qualitative similarities and differences from human data. All clear responses from other monkeys were generally similar to this example. The circle plots on the left of the figure show a two-dimensional representation of the sensor array and its spatial distribution of the evoked magnetic fields at the latency of the largest peak (indicated by red vertical bar on time-series plot, right). These plots show the topography of MEG data, with positive and negative adjacent regions, indicating that the neural source is below the zero-crossing between these regions. The right time-series plot of Fig. 4.2 shows the average temporal dynamics of each sensor overlaid on a common baseline. The

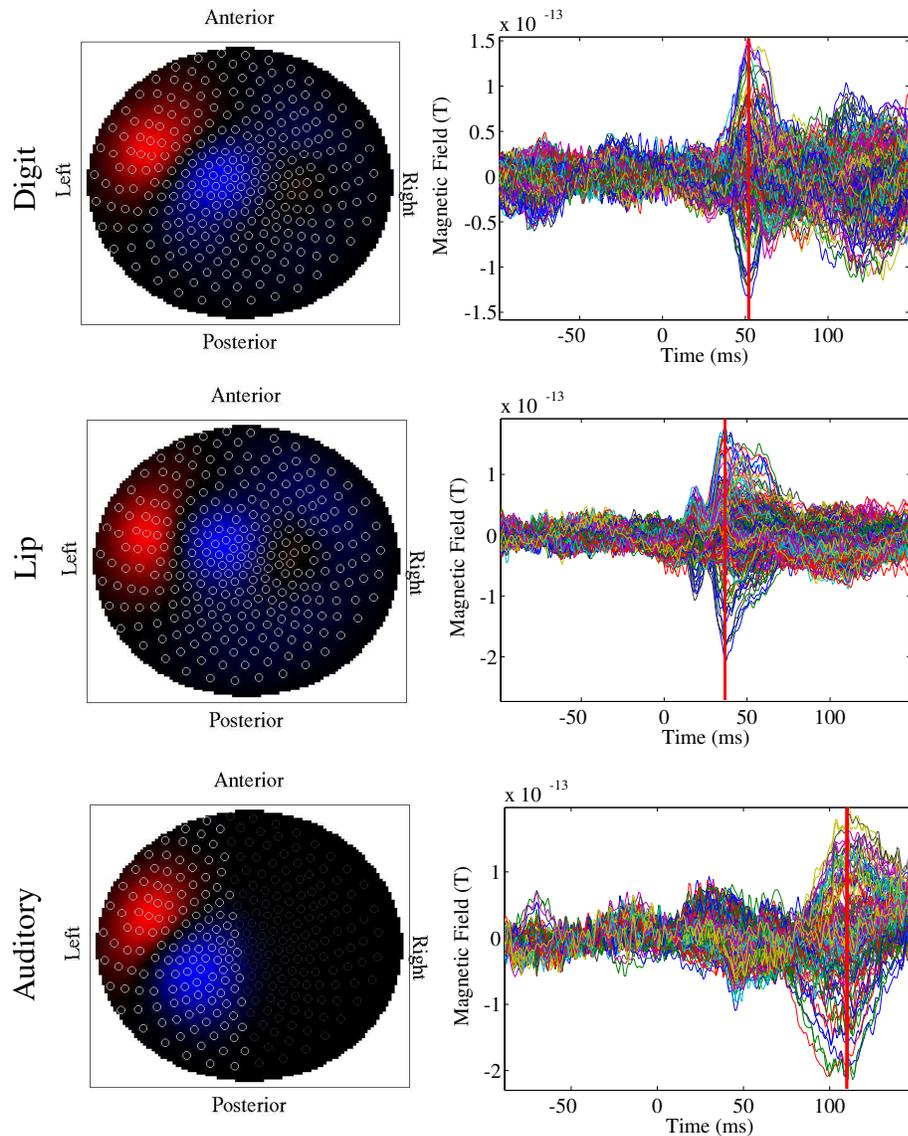


Figure 4.1 Averaged MEG sensor data from human subject H4. Top row: hand stimulation; Middle row: Lip stimulation; Bottom row: Auditory stimulation. Left column: Two-dimensional representation of sensor array, with red indicating magnetic fields going into the head and blue coming out, occurring at the latency of the red vertical bar in the adjacent right plot. Each white circle represents a sensor. Right column: Averaged time-series for all sensors overlaid on each other, with stimulus onset occurring at time=0ms. Digit and lip plots show overlay of all 275 sensors used in beamforming analysis; auditory plot shows only the left 132 sensors used in beamforming analysis.

first peak latencies in these examples are at 13ms for digit stimulation (top row), 7ms for lip stimulation (middle row), and 62ms for auditory stimulation (bottom row).

The overall signal-to-noise ratio was not as high in the monkey data as is normally seen in human MEG data, but was sufficiently high to detect peaks and perform source localization. Thus, anesthetic effects may have reduced the monkeys' responses somewhat, but not prohibitively. The latencies of the responses in monkeys were clearly shorter than in humans to identical stimuli as expected due to shorter transit distances.

4.3 Comparison of Methods on Full Data

4.3.1 Analysis of method comparison on full data

The comparison of methods on the data described above was performed in two main ways. As mentioned above, the first type of analysis ("unprocessed") for method comparison on the above data used all sensors available and reconstructed over all voxels in the brain. Here, only data from M3 and M4 are analyzed since the difference between this analysis and the forthcoming analysis on M3/M4 on reduced sensors was to use all 275 sensors rather than the 37 closest to the source. The analysis on M1/M2 did not change (using all 37 sensors) so the M1/M2 results are described thoroughly in the next subsection on focused data.

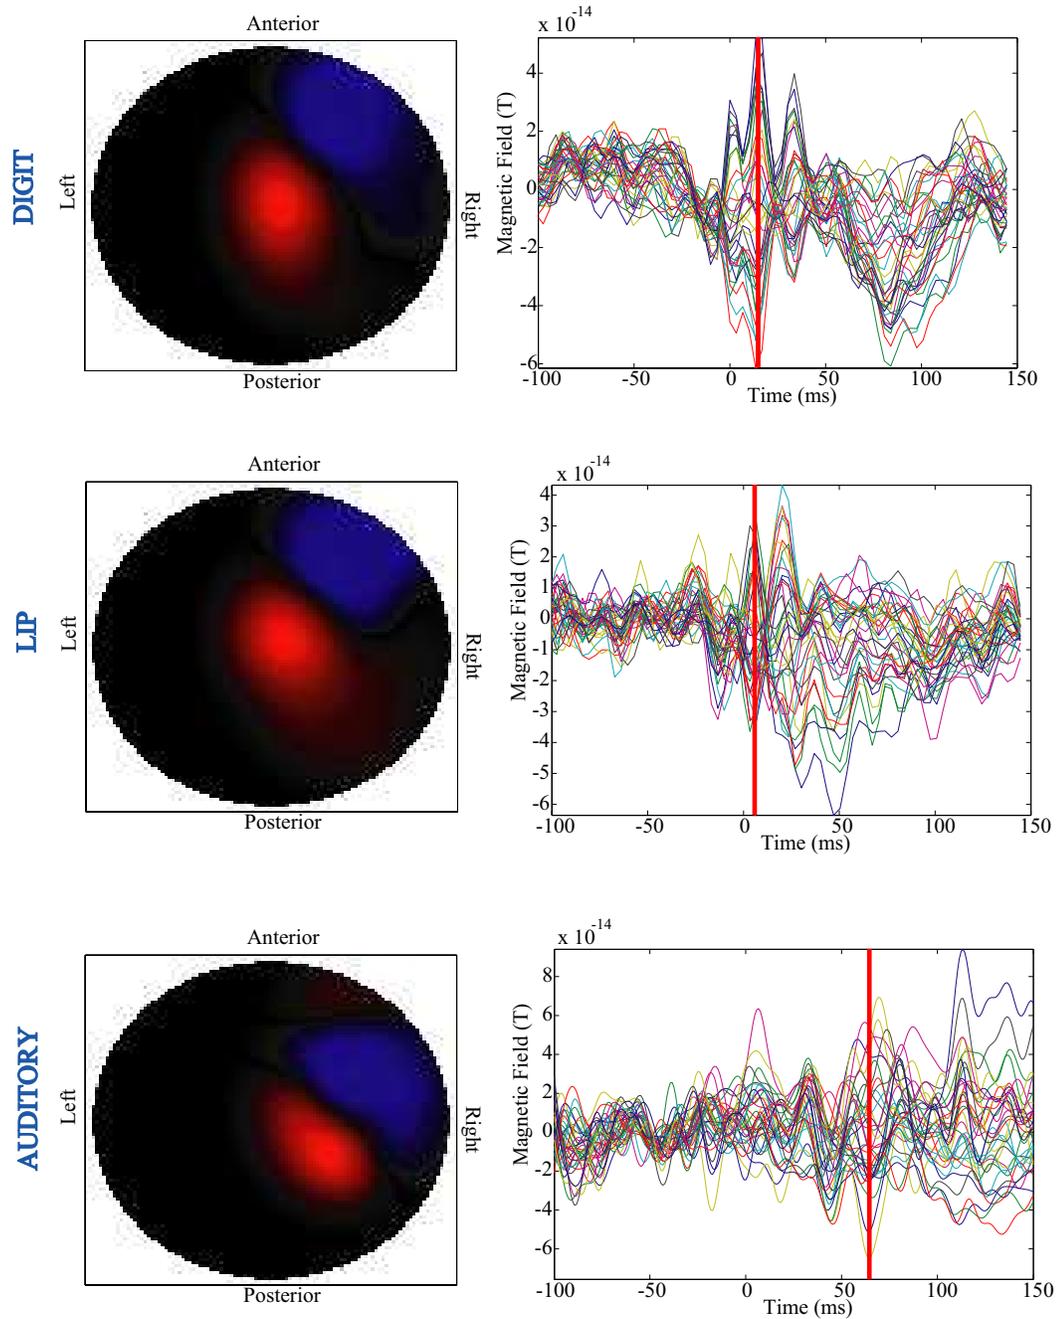


Figure 4.2 Averaged MEG sensor data from monkey M1. Top row: hand stimulation; Middle row: Lip stimulation; Bottom row: Auditory stimulation. Left column: Two-dimensional representation of sensor array, with red indicating magnetic fields going into the head and blue coming out, occurring at the latency of the red vertical bar in the adjacent right plot. Right column: Averaged time-series for all 37 sensors overlaid on each other, with stimulus onset occurring at time=0ms.

sLORETA requires a regularization term which was empirically determined to give the best results on other similar human MEG data and was left fixed for all reconstructions reported here. All methods except sLORETA have a choice in number of dimension reduction related to signal space. Here, "dimension reduction" refers to the number of eigenvalues selected for the eigenspace MVAB, the number of factors chosen for SAKETINI or the number of temporal basis functions (TBFs) for NSEFALoc1/2. Values for dimension reduction 1-4 were tested for each method. For all methods, the value of dimension reduction that gave the most realistic result (best matching estimated true locations) was used.

All results were studied for multiple spatial peaks, not just the maximum peak. The A' metric was introduced in the last chapter as a single number which quantifies the tradeoffs of threshold on hits versus rejections and of allowed localization error on hits versus false positives. However, A' requires knowledge of the true source location. Since there is no exact answer with real data as to where the spatial peak should be, these results are relative to best approximations rather than ground truth. While electrocorticography data can be very useful for obtaining this information, this is only possible in patients requiring neurosurgery and thus is not applicable for routine neuroscience MEG studies in healthy control subjects. Thus, to obtain the "ground-truth" coordinates in the healthy human subjects studied, several researchers very familiar with the localization of primary digit, lip and auditory areas were polled to mark on the structural MRI where they

thought the true location was in each of the 4 subjects. Their judgements were guided by anatomical landmarks, including the central sulcus and the “Omega” shaped knob of the pre-central gyrus at the hand area. The MNI coordinates of these opinions were averaged to form the “ground-truth” location. Additionally, location of secondary somatosensory cortex for digit and lip was found using the MNI database of coordinates built-in to NUTMEG, adopted from code by John Ashburner (http://www.sph.umich.edu/~nichols/JG2/get_orig_coord2.m). All coordinates labeled “Brodmann area 43” were averaged to form the digit subregion and the lip subregion was lateral to the digit.

The primary reason for collecting and analyzing MEG data in primates is to ultimately compare the MEG localizations with the location giving the strongest neural response to the same stimulation as recorded by standard, acute electrophysiological procedures. The analysis to obtain the ground truth locations of activation will be performed by collaborators and unfortunately will not be completed in time for this dissertation. Thus, for now, the same way to obtain true source locations was used, of polling researchers familiar with monkey neuroanatomy for their best estimate of the true source location, again guided by anatomical landmarks including the central sulcus.

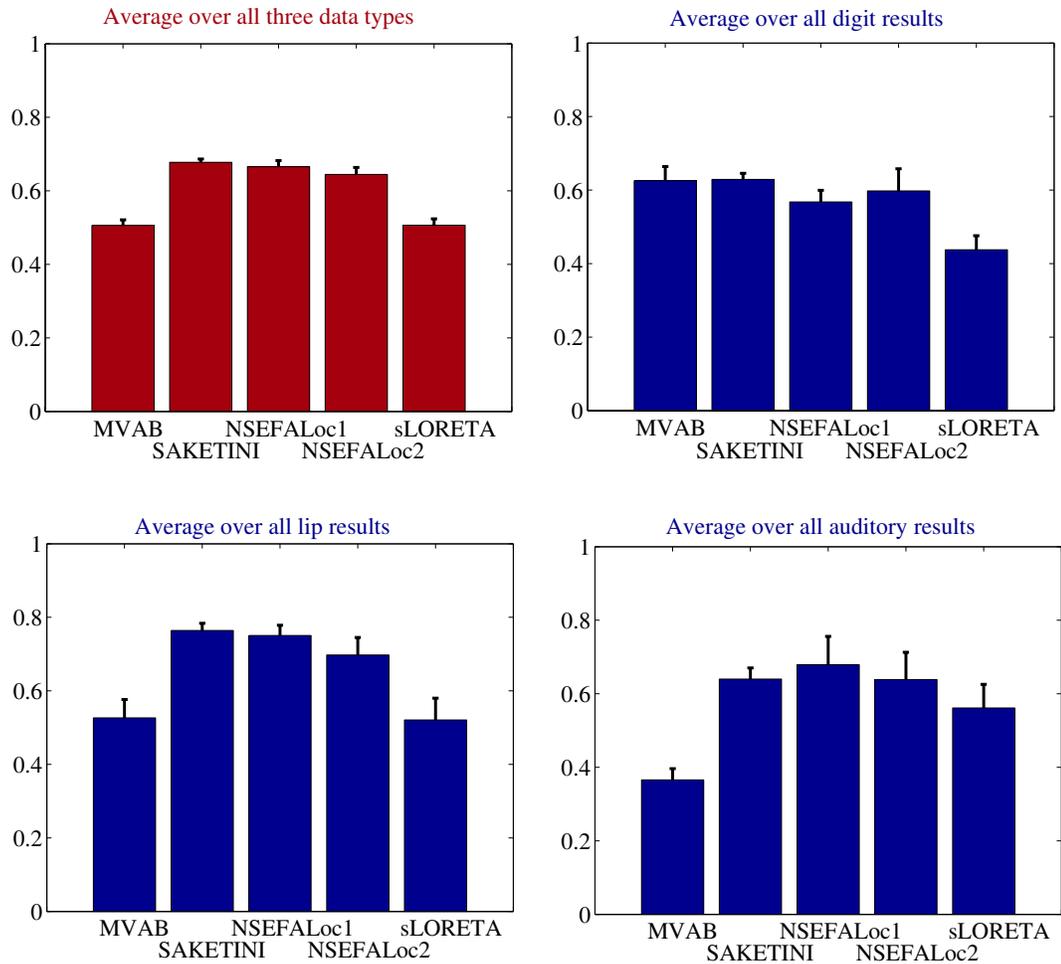
In both cases of human and monkey estimates, there is some error expected in this best guess of ground truth estimate. However, it is not expected that the error would be biased in any particular way. Furthermore, the estimates were

made without knowledge of specific results of any of the reconstruction methods, thus should not be biased towards any one method over the others. Finally, the A' metric is computed at several allowable errors of the reconstruction from the true location ($[5\ 10\ 15] * \text{sqrt}(3)\text{mm}$); so the largest error allowed should help reduce any final differences.

4.3.2 Results and discussion of method comparison on full data

Fig. 4.3 shows that, averaging over all four subjects and three stimulation types, all three new methods (SAKETINI, NSEFALoc1 and NSEFALoc2) have a higher A' value than MVAB and sLORETA. The results differ slightly when focusing on just one type of stimulation. In digit stimulation results, A' is equally high for SAKETINI, NSEFALoc2 and MVAB, while NSEFALoc1 and sLORETA have a slightly lower value. For digit, all four areas of contra- and ipsi-lateral primary and secondary somatosensory cortices were used as true locations. In lip stimulation results, A' followed the same pattern as the average over all stimulation types. For lip same as for digit, the four areas of contra- and ipsi-lateral primary and secondary somatosensory cortices were also used as true locations, but recall that primary and secondary lip areas are much closer to each other than digit. That SAKETINI, NSEFALoc1 and NSEFALoc2 perform well in this case means this close distance (of roughly 25mm) was not a problem to resolve. Finally, the auditory A' results again show SAKETINI, NSEFALoc1 and NSEFALoc2 with

Figure 4.3 A' value for human MEG data under digit, lip and auditory stimulation



the highest values, sLORETA only slightly less, but MVAB much lower. Here, only left and right primary auditory cortices were used as the two true locations, but their activity is expected to be highly correlated in time, causing well-known problems for MVAB.

From a clinical standpoint, individual results matter more for a patient's diagnosis than averages across subjects. For example, it could be the case that cS2 is never found when cS1 is found, and the previously discussed Fig. 4.3 would not

make this distinction. Fig. 4.4 shows the reconstruction of all methods on subject H4's digit MEG data. The left column shows a maximum intensity projection (MIP) onto a coronal view, while the right column shows a MIP onto an axial view. Blue squares indicate locations found as a "hit" (above given threshold and distance error from true location) for each reconstruction method. In this example, SAKETINI fits three hits (cS1, cS2, and iS1) while NSEFALoc1 and NSEFALoc2 find two hits (cS1 and iS1), and MVAB and sLORETA only find cS1. All methods show a false positive near the center of the head, but at varying intensities relative to the true locations.

For the monkey data, Fig. 4.5 shows that, averaging over all runs and stimulation types, all three new methods (SAKETINI, NSEFALoc1 and NSEFALoc2) have a higher A' value than MVAB and sLORETA, just like in the human data, but with NSEFALoc2 the highest by a slight amount, then NSEFALoc1, SAKETINI, MVAB and then sLORETA. These results can be expected to be slightly different from the human data for many reasons, but mainly since only one hemisphere was close enough to the sensors to be detected, hence fewer regions are possible to be active (only contralateral S1 and S2 for digit and lip, and contralateral A1 for auditory). Since only one auditory side is expected, the correlated source problem for MVAB does not occur here.

When focusing just on digit results, NSEFALoc2, SAKETINI and NSEFALoc1 had a slightly higher A' than MVAB, and all were much better than sLORETA.

Figure 4.4 Digit localization results from all methods on subject H4. The left column shows a maximum intensity projection (MIP) onto a coronal view, while the right column shows a MIP onto an axial view. Blue squares indicate locations found as a “hit” (above given threshold and distance error from true location) for each reconstruction method.

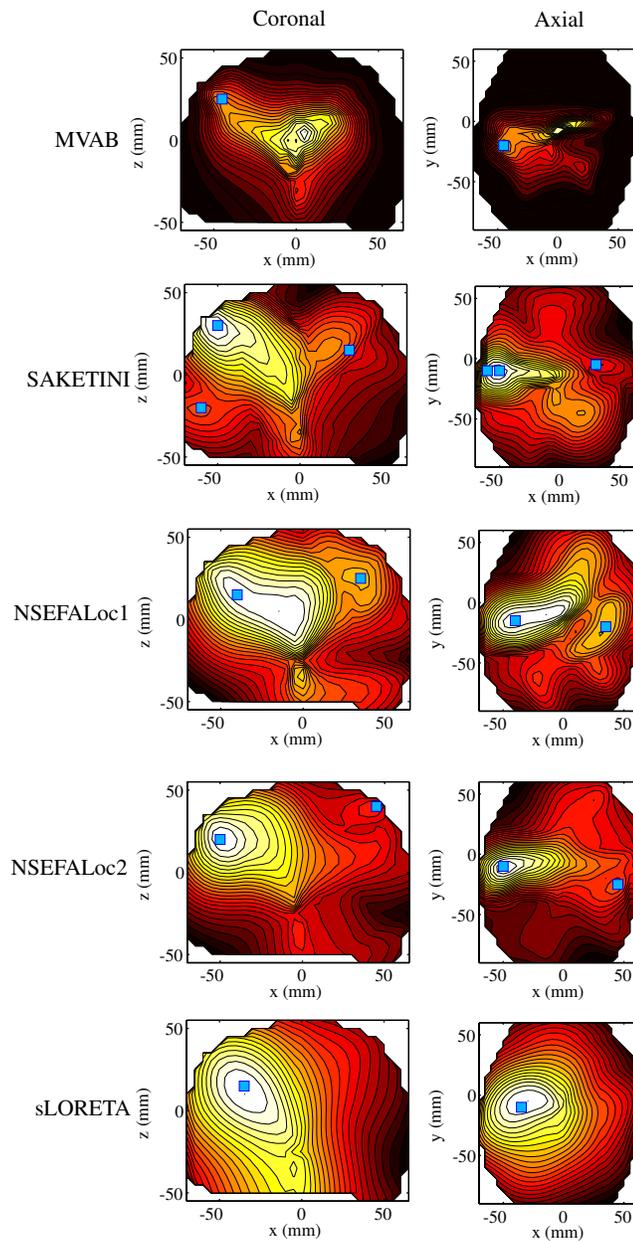
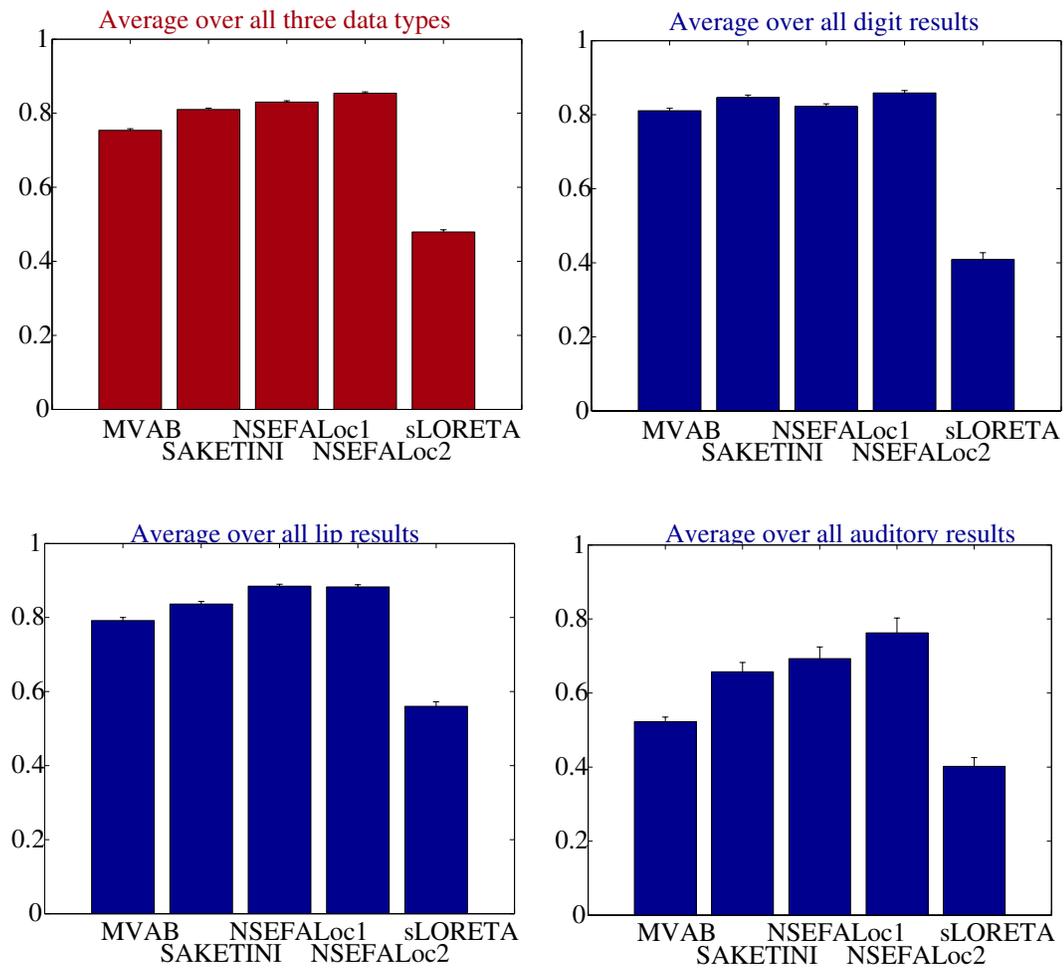


Figure 4.5 A' value for monkey MEG data under digit, lip and auditory stimulation



Often the activated area extended from cS1 down to cS2 and so was not counted as a separate spatial peak. In just lip results, recall that primary and secondary lip areas are much closer to each other than digit. NSEFALoc1 and NSEFALoc2 have the highest A' , then followed by SAKETINI, MVAB and then sLORETA. The same pattern is also seen for auditory-only stimulation average.

4.4 Comparison of Methods on Focused Data

4.4.1 Analysis of method comparison on focused data

In the “focused” case, the data had been processed to reduce noise and correlated source problems, to focus on just a single primary sensory response. The data reduction included removing extra sensors in the CTF monkey data so that only the 37 sensors closest to the brain area of interest were included, since the rest were recording noise and large heartbeat artifact. In the human analysis, all sensors were used for digit and lip stimulation, but only left hemisphere sensors were used for auditory due to the correlated source problem. Furthermore, the VOI used in the human analysis was not the whole brain, but a smaller region including the expected region discretized at 2mm-sided voxels; this region extended as far as possible while still limiting the total number of voxels to roughly 15,000 for computational-limit reasons. The result of this smaller VOI was that opposite hemisphere voxels were not included; in fact, roughly only a tenth of

the total brain volume was used.

The A' metric was again used to analyze the reconstructions from these focused datasets. However, the emphasis here is more on improved accuracy of the single primary peak, not on finding multiple peaks. Yet this is still not entirely clear-cut when the true location is not known with certainty.

In most cases, the largest spatial peak was the one in the area near the expected location. However in a few instances, there was a clear spatial peak near the expected region but it was not the largest. This was noted as a binary result (1 if absolute maximum, 0 if not) and averaged over all results, and referred to as “maximum peak index”.

Furthermore, while the coordinate of just the peak voxel was recorded (and not coordinates of nearby activated voxels), the spatial extent of the activation near the peak is also of interest for several reasons. All methods examined here are preferable over traditional single dipole fits in that they attempt to measure the spatial extent of the activation; the spatial precision of a localization method matters, not just accuracy.

Two metrics were used to quantify the spatial extent. First, the threshold was set to obtain either a 21-voxel cluster in monkeys or 100-voxel cluster in humans. The percentage of the peak voxel’s strength for that cluster that the threshold was set to, was termed “percentage threshold.” Using the 2mm voxel resolution, these cluster sizes correspond to roughly the same anatomical size across the two

species, and to the same size used as seeds for DTI tractography (discussed in the next chapter). A low percentage threshold here indicates a small cluster and/or high spatial precision (corresponding to a small full-width half-maximum, another common measure of spatial precision).

The second metric used to quantify spatial resolution, termed “single cluster index”, was related to quantifying whether, for the given threshold above, other voxels/clusters besides the anticipated one were above this threshold. This was quantified as the total number of voxels above that threshold regardless of which cluster they belonged to, divided by the expected number of active voxels if no other distant clusters were active. For example, in a localization of monkey data, the threshold was set to produce a cluster of 21 voxels near the expected region. If no other voxels were above this threshold, the single cluster index equals 1. However, if another distant cluster of 30 voxels was above this threshold, then the index would equal $51/21 = 2.43$.

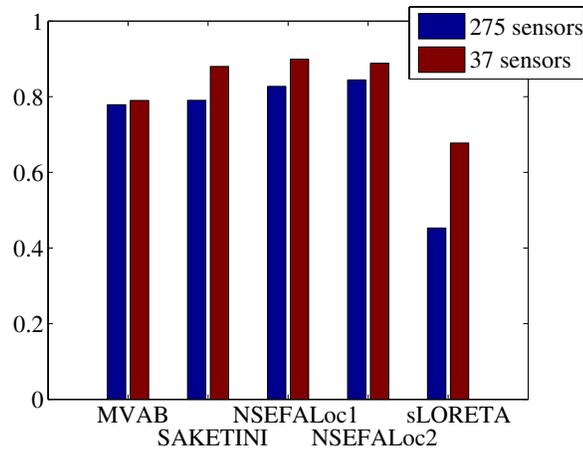
Finally, the actual peak coordinates of all five methods on all datasets are examined. Each individual localization was compared with the “ground-truth” estimates discussed previously and the three-dimensional Euclidean distance was measured. This error for each of 13 digit, 17 lip and 6 auditory results in monkeys, and 4 digit, 4 lip, and 4 auditory results in humans was averaged. This average over individual errors is termed the “individual error” in the results plot. Additionally, the error between the average localization result and “ground-truth”

estimate for each of digit, lip and auditory regions was measured, and termed “average error” in the results below. These two different metrics focus on variation across subjects and overall bias, respectively. For example, each individual result could have large error about some mean estimate, leading to large individual error, but if the mean estimate was exactly equal to the “ground-truth” estimate, then the average error would be zero. Alternatively, there could be small individual variation about some biased mean, leading to an individual error equal to the average error.

4.4.2 Results and discussion of method comparison on focused data

Overall, all five methods perform well on the processed monkey and human MEG data, with some differences between methods. All methods’ results were dependent on dimension reduction to some extent. The dimension that gave the best result for each method was used.

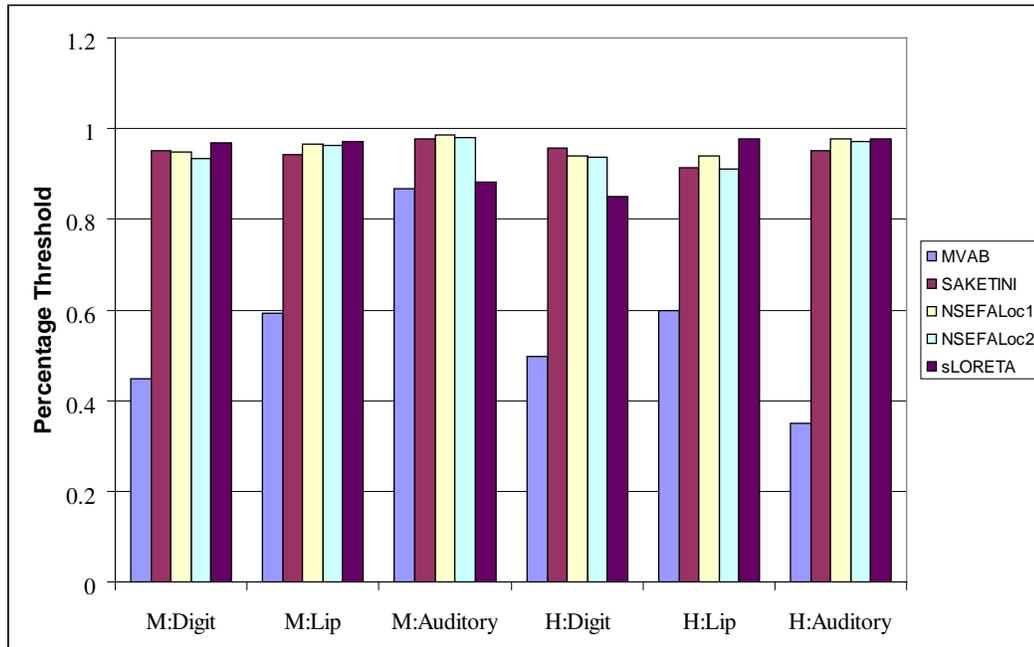
Fig. 4.6 shows the same A' metric discussed previously, here averaged only over M3 and M4 data which was collected in the 275-channel MEG system. Since the data from M1 and M2 was not changed here from the previous section, it is not included in this plot. This figure shows that, for all methods, using just the 37 channels closest to the source, rather than all 275 channels, did improve the reconstruction according to the A' metric. Surprisingly, it would have been

Figure 4.6 Comparing A' for all 275 versus reduced 37 sensors in two monkeys

expected that this improvement be strongest for MVAB and sLORETA and modest for SAKETINI and NSEFALoc1/2 since they should be more immune to the noisy sensors. Instead, it was strong for sLORETA, moderate for the three probabilistic methods, and slight for MVAB. However, since A' is a measure of both accuracy and minimal false positives, further investigation is needed to tease out these components, as discussed in the next few figures. The following analysis and figures are based on averages over all four monkeys (37 channels for all) and the focused human data as already defined.

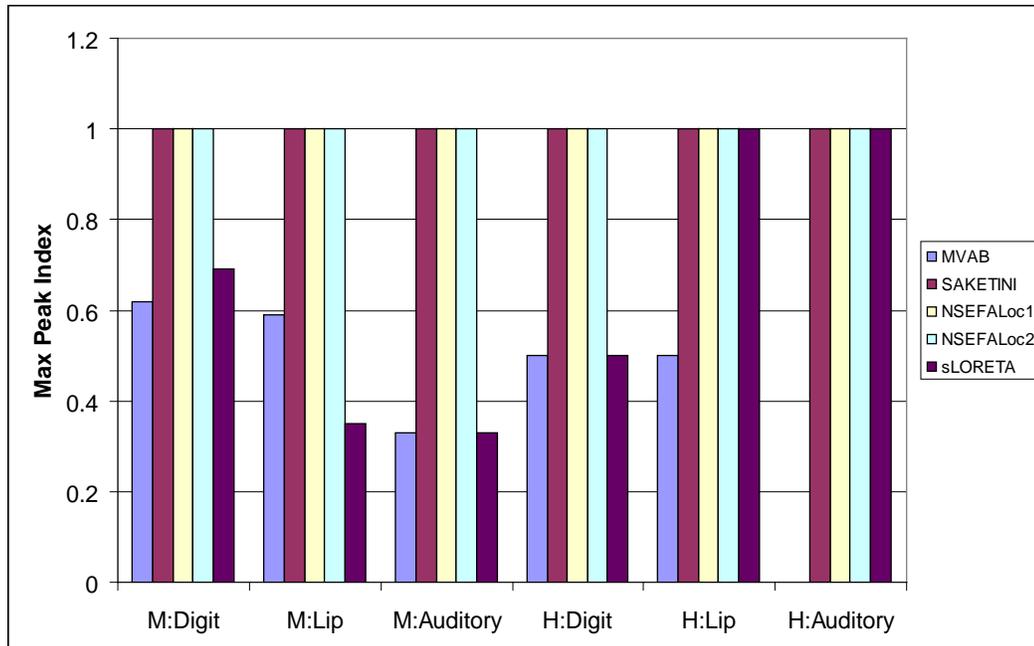
The average “percentage thresholds” (threshold level needed to create the cluster of either 21 or 100 voxels for the monkey or human data, respectively), are shown in Fig. 4.7. MVAB clearly has the tightest spatial resolution. It is well known that sLORETA has less spatial precision than MVAB (Sekihara et al., 2005). The spatial resolutions from the likelihood maps from the new probabilistic algorithms are even less than sLORETA. It is possible that some manipulation of the

Figure 4.7 Threshold comparison: The threshold was set to obtain either a 21-voxel cluster in monkeys or 100-voxel cluster in humans. The percentage of the peak voxel’s strength for that cluster is the “percentage threshold.”



likelihood (such as a generalized likelihood ratio test, Bayesian Information Criterion or Akaike Information Criterion) would produce a map with tighter spatial resolution. Furthermore, plotting the spatial map of SAKETINI or NSEFALoc *source power* for a given time point (\bar{s}_n) displays a spatial resolution similar to MVAB. However, these maps are not used as the primary method for localization since they are not as robust as the likelihood maps for spatial localization of peaks.

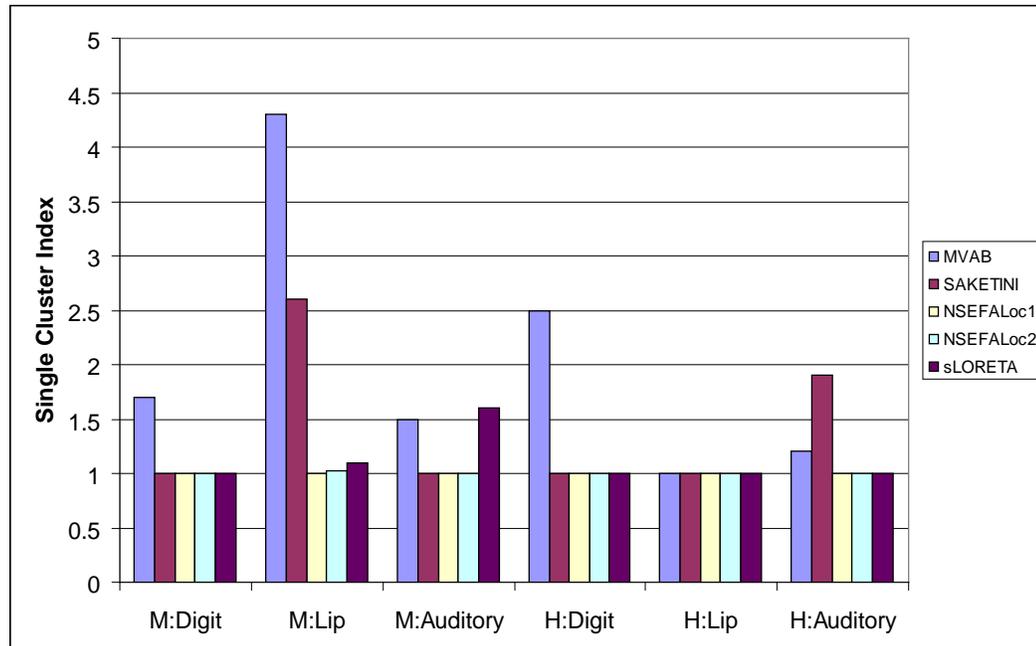
All three new probabilistic methods outperform MVAB and sLORETA in indicating the maximum peak at or near the expected location (see Fig. 4.8). Even if the source power at the early primary peak may not be as strong as later secondary peaks, the likelihood maps of SAKETINI and NSEFALoc are able to indi-

Figure 4.8 Peak maximum comparison

cate a strong probability of activity at those locations nevertheless.

The average single cluster indices are shown in Fig. 4.9. In general for all methods, the other active voxels causing this index to exceed 1 were not in any consistent location or in a location that would make sense based on the stimulation (such as S2 or posterior temporal cortex); thus they appeared to be false positives. NSEFALoc1 and NSEFALoc2 performed best in this metric of specificity, followed by sLORETA, then SAKETINI and MVAB.

Finally, the differences (errors) between the peak localization of each method relative to a “ground-truth” estimate (described above) are shown in Fig. 4.10. For the monkey data, MVAB and NSEFALoc1 show the smallest individual and average error, with SAKETINI showing comparable results, but NSEFALoc2 and

Figure 4.9 Single Cluster Index comparison

sLORETA showing higher errors. For the human data, MVAB and SAKETINI show the lowest errors for both individual and average, with NSEFALoc1, NSEFALoc2 and sLORETA showing higher errors. Overall, the differences between methods of the errors were only 2-3mm, within the expected error of MEG in general, and certainly within a range expected due to the method of selection of ground-truth coordinates. Thus, MVAB and SAKETINI have slightly smaller overall errors than the other methods, but this difference is small relative to the total errors.

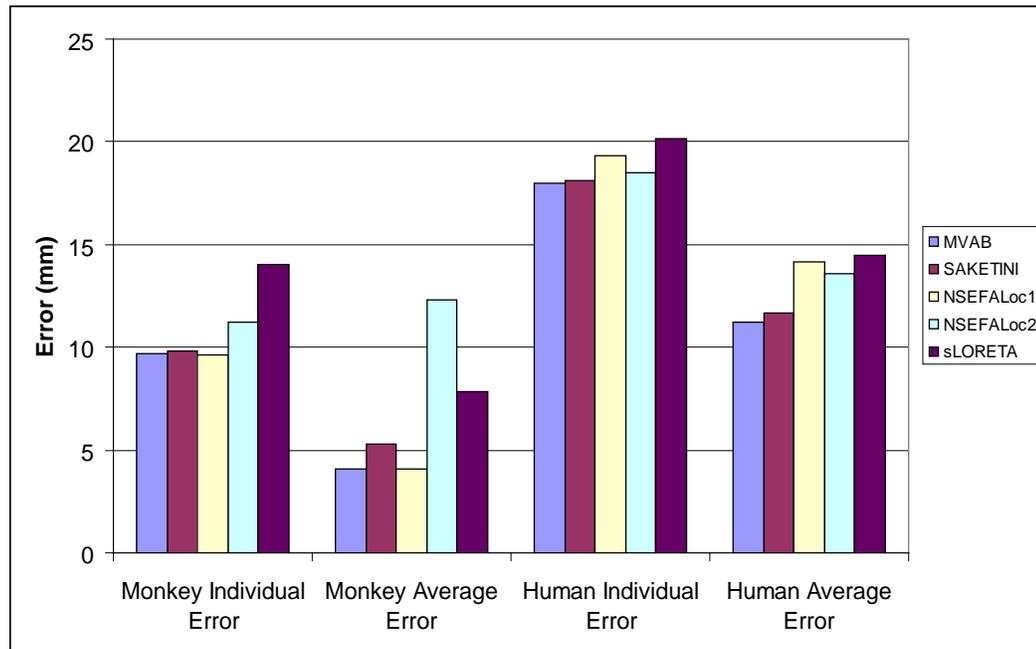


Figure 4.10 Average localization errors of all methods in both monkeys and humans, for digit, lip and auditory stimulation.

4.5 Conclusion

The MEG derived source localization results from both human and non-human primates showed peaks in expected primary sensory areas, with the digit representation superior to lip representation along the central sulcus and the auditory representations most inferior in the superior temporal lobe. The feasibility of using a nonhuman primate model for validating MEG source localization methods has been shown. The comparison across methods show that when all sensors (with artifacts) are used to reconstruct over the whole brain (with potentially correlated sources), SAKETINI and NSEFALoc2 outperform the other methods on average using the A' metric. When the data is focused on a single brain region by

reducing number of sensors and/or total voxels reconstructed over, MVAB has roughly equal localization errors but more false positives than SAKETINI and NSEFALoc.

Acknowledgments

I thank Sri Nagarajan for discussion and comments on this analysis, Liz Disbrow, Zhao Zhu, Marianne Lowenthal and Leah Krubitzer for help with animal experiment preparation, Susanne Honma, Anne Findlay, Evelyn Proctor, Kavita Vashi, Leighton Hinkley and Geoff Rau for assistance with data acquisition, Ron Baireuther, Robert English-Fernandez, Patricia Ramsey and Grace Kim for assistance with anesthesia, and Sarang Dalal for comments on MEG processing. I also thank the NIH for funding R01DC4855 and R01NS44590.

Chapter 5

Separability of functionally-defined regions and networks through analysis of MEG data and MEG-seeded DTI tractography

5.1 Overview

The brain is partitioned into discrete functional areas with specific purposes, which is why only certain areas are active for a given task and can be distinguished from other brain areas using MEG (and fMRI) data. However, each location certainly does not act in isolation, but is highly interconnected in organized

networks. These networks need to be examined both anatomically (by tracing axonal connections) and functionally (by looking for causal influences of the activity of one region on another). Dissociation between networks based on functional and anatomical considerations is needed. Functional (and effective) connectivity has been studied by noninvasive measures in humans (e.g. Biswal et al. (1995); Friston et al. (2003); Sun et al. (2005); Gross et al. (2001)) though further validation, such as from electrocorticographical measures, is still needed.

Most of what is known about the functional organization of the cortex, and the connections of individual cortical fields that compose the processing networks of the brain, comes from invasive electrophysiological recording and neuroanatomical studies in macaque monkeys and other mammals. The results from these studies have, to a large extent, guided psychological and brain imaging experiments in humans. However, this extrapolation is limited due to the increased complexity of human neuroanatomy and the indirect nature of imaging measurements. Experiments to validate the many noninvasive imaging methods that have been developed in recent years lag behind their use to study human brain function. While a macaque monkey model for MEG would help bridge the gap between a wealth of pre-existing electrophysiological and neuroanatomical data from the monkey and the less well studied human brain, data on the feasibility of MEG in the monkey is sparse.

Anatomical connectivity has been extensively studied, especially in macaque

monkeys using invasive neuroanatomical techniques. However, measuring cortical and subcortical connections in humans had been limited to postmortem studies. The advent of white-matter tractography based on diffusion tensor imaging (DTI) has allowed noninvasive measurements of white-matter tracts in both healthy and patient populations. One can trace the connections starting from a seed region in the brain. This procedure is termed tractography. Seed regions are frequently chosen anatomically, but functionally defining seed regions, as in the work described here, is becoming increasingly popular.

In this chapter, the use of MEG to determine separate brain regions is demonstrated, based on differences of latency, spatial reconstruction localization, and pattern of tracts anatomically connected to a functional region as determined by MEG-seeded DTI tractography. Diffusion tensor images were acquired from the same human subjects and same macaque monkeys from which MEG data was collected to test the feasibility of MEG-seeded DTI tractography in both human and non-human primates.

5.2 MEG latencies across stimulation type and species

The previous chapter illustrated the feasibility of reconstructing meaningful results from MEG data recorded from macaque monkeys. A comparison of the

monkey and human peak latencies can further assist the usefulness of a primate model for MEG. In the sensor data, multiple peaks were seen in the data from each monkey, with some variability in peak magnitude and latency. Table 5.1 shows the peak latencies seen in each human and monkey and the peak with the strongest amplitude is marked with an asterisk.

5.2.1 MEG latencies in humans

As mentioned previously, the human data was consistent with previous studies of digit, lip and auditory stimulation. A normal amount of variability across subjects was also seen, regarding peak amplitude and latency (Table 5.1). For digit stimulation, the peak averaging 44ms was largest for all subjects. An early peak was sometimes seen around 22ms (2/4 subjects). Later peaks were seen in all subjects, some clustering around 85ms (3/4 subjects), some around 120ms (3/4 subjects), and some around 140ms (2/4 subjects). These later peaks around 120ms probably correspond to the late response seen in somatosensory responses elsewhere, usually originating from S2 (Hari et al., 1983).

There was more variability in the human lip data regarding peak latency than for digit stimulation. All (4/4) subjects show an early 19ms peak, which was largest in 2/4 subjects. All four subjects also showed a peak around 40ms, which was largest for the two remaining subjects. Later peaks were also seen around 55ms (2/4 subjects) and 105ms (3/4 subjects).

There also existed variability in the human auditory data. All subjects showed a peak around 106ms, consistent with the M100 commonly reported (Pantev et al., 1990) as well as a peak around 165ms. The 106ms peak was largest in 2/4 subjects; the other two subjects showed the largest peak at 165ms. Two subjects showed a peak around 55ms, consistent with the auditory M50 reported elsewhere (Reite et al., 1988). All four subjects also showed a peak around 200ms, consistent with the auditory M200 (Jacobson et al., 1991).

5.2.2 MEG latencies in monkeys

The responses seen in monkeys were generally similar to those seen in humans, only with shorter latencies. The most consistent largest peak from digit stimulation in monkeys was seen around 15ms (range 11-27ms, 9/10 datasets) and this peak was the one focused on for source localization of primary somatosensory cortex for creating seeds for DTI tractography. An early peak was sometimes seen at 3ms (7/10); in two of these cases this early peak was larger than 15ms peak. Later peaks were also seen around 35ms (range 30-40ms, 7/10) and around 60ms (range 47-82ms, 8/10).

Like in the human lip data, the monkey lip-stimulated responses showed some consistent peaks, but there was more variability relative to digit stimulation regarding which peak was largest. An early peak was seen in 11/12 datasets around 10ms (range 3-15ms) and was the largest in 7/12 datasets. A second peak

Table 5.1 Latencies of peaks observed across all sessions of MEG recordings from monkeys and humans. Peaks with largest amplitude are marked with an asterisk.

Monkey	Digit				Lip			Auditory		
M1	3	13*	33 (50%)	64-82	3-12*	20-23	45-54	43	64*	
M2	3	18* (67%)	33	47-53	7-10	27*	47-57	50*	110	
M3		11*	30	63	13		43*	43*	77	
M4		25*	40 (50%)		15* (67%)	27	48-55			
Human	Digit				Lip			Auditory		
H1	22	46*	85	122	19*	42	108	102*	175	200
H2		50*	60	116 140	19	37*	55	55	112*	165 208
H3	23	39*	96	134	19*	33	55 108	105	155*	195
H4		41*		121	19	36*	101	61	103	153* 200

around 25ms was seen in 10/12 datasets (range 23-30ms) and was largest in 3/12 datasets. Finally, a later peak around 50ms (range 42-57ms, 12/12) and was the largest in 2/12 datasets.

The auditory responses also showed some consistency and some variability in peak latencies. A 45ms peak was seen in 4/5 datasets (range 43-45ms) and was the largest in 2/5 datasets. A slightly later peak was also seen (range 50-65ms) in 4/5 datasets and was the largest in 3/5 datasets. Finally, a late peak was sometimes seen (77ms in one case and 110ms in a second case).

5.2.3 Discussion of MEG latencies

Comparing human and monkey digit responses

The latencies of responses recorded using MEG in monkeys generally showed a pattern similar to human data, though they were shorter in non-human

than in human primates, most likely due to longer conduction distances. Variability in latency and morphology of response peaks also occurred in both species and again the pattern of variation was similar across species.

This shortening of latency in nonhuman primates relative to humans is in agreement with previous studies. In particular, McCarthy et al. (1991) studied the subdural recordings obtained in macaque monkeys from median nerve stimulation and compared them to the human results from the studies by Allison et al. (1989a) and Wood et al. (1988). They found a N10-P20 wave generated from area 3b and a P12-N25 wave from area 1 in monkeys, which they proposed to be the correlates of the human N20-P30 and P25-N35.

Tactile stimulation was chosen over electrical median nerve stimulation in this study because electrical pulses to the median nerve non-selectively stimulate nerves of various receptor types and receptive fields on the hand. Traditionally, median nerve stimulation has been employed to provide strong stimulation for mapping, especially during surgery.

Similarly, Gardner et al. (1984) presented airpuff stimuli to the hand or forearm of alert macaque monkeys while recording epidural SEPs. They found P15 and P25 peaks followed by a large N43 and P70. These latencies are consistent with the MEG data reported here and are slightly delayed relative to the median-nerve stimulated latencies of P12-P20 reported by McCarthy et al. (1991). They also suggest that the latencies do not change substantially as a function of anesthe-

sia. However, the amplitude of later peaks might be suppressed more with anesthetic level, suggested by the data presented here and by Allison et al. (1989b).

The human MEG latency results presented here are in agreement with Forss et al. (1994), which studied the difference of SEF data from airpuff to the digits versus electric stimuli to the median nerve. They found that the responses to airpuffs were about 5ms later than median nerve stimulation and that the later peak at 65ms was larger than the early 25ms peak, in direct contrast to median nerve stimulation.

A reasonable speculation is that the 14ms peak seen in the monkey MEG digit data corresponds to the M40 wave in humans, that the occasional earlier peak at 7ms in monkeys corresponds to the occasional M20 in humans, and that the later 30ms/40ms peak in the monkey responses might correspond to the human M80 (see Table 5.1).

Comparing human and monkey lip responses

The data presented here are in general agreement with previous studies of evoked potentials/fields in monkeys and humans, although fewer studies exist for lip than for digit responses. In monkeys, McCarthy and Allison (1995) used electrical stimulation applied to median nerve, lip, tongue and palate and found that the lip response was just lateral to the hand response along the postcentral gyrus and that the latencies were slightly shorter for lip relative to hand (N9-P14 for lip and N10-P20 for hand). The evoked fields resulting from lip stimulation

in humans have been studied by Disbrow et al. (2003); Hoshiyama et al. (1996); Nakahara et al. (2004). Similar to the human results shown here, they all found a clear 20ms peak as well as a later 35ms peak. The human data presented here are in agreement with these results, although neither study report that the early 20ms peak might be larger than the 35ms peak. A reasonable speculation is that the monkey 10ms and 25ms peaks seen here correspond with the human 20ms and 35ms peaks. Ipsilateral responses in monkeys could not be measured since the sensor array could only be placed by one hemisphere at a time; thus further examination of possible ipsilateral responses in humans was not pursued.

Comparing human and monkey auditory responses

In both humans and monkeys, the localization of the MEG data from auditory stimulation to the superior temporal lobe agrees with the known location of primary auditory cortex (Brugge and Merzenich, 1973; Celesia, 1976). Secondary auditory areas adjacent to and surrounding A1, such as the belt and parabelt, could also be activated by the stimuli presented in this study. We did not expect to distinguish primary auditory cortex from adjoining belt regions using MEG source localization, but to localize the region of A1 and surrounding auditory areas to seed the DTI tractography.

Evidence exists that the human M100 in particular seems to arise from either A1 or the planum temporale (Godey et al., 2001; Liegeois-Chauvel et al., 1994). In monkeys, Arezzo et al. (1975) show early (12-22ms) responses in the monkey

originating from primary A1 while later responses (N40 and N60) arose from areas on the superior temporal plane more anterior to A1 and the N70 and N100 arose from more posterior areas. They propose that the monkey N70 corresponds to the human N100, which they suggest have dual origin in A1 and posterior regions in both species. Furthermore, different components might be distinguished based on refractory and/or habituation properties, as examined by both Budd et al. (1998) and Hari et al. (1982), with the component with shorter recovery period coming from the superior temporal plane and the component with longer recovery period more diffuse and frontal. This proposal is consistent with the data in this study, with peaks in the range from 60ms to 70ms.

5.3 Spatial Separability of MEG MVAB Results

5.3.1 Spatial Separability of human MEG MVAB results

The spatial distribution of the localized sources in the human brain from tactile and auditory stimulation generally agreed with known neuroanatomy. Fig. 5.1 shows the estimated beamformer source localization overlaid on the structural MRI in the left column and the corresponding estimated temporal activity at the peak voxel in the right column (reconstructed from the same H4 dataset as in Fig. 4.1). Both digit and lip localizations are entirely in the post-central gyrus; the auditory localization in Fig. 5.1 is entirely in the superior temporal gyrus.

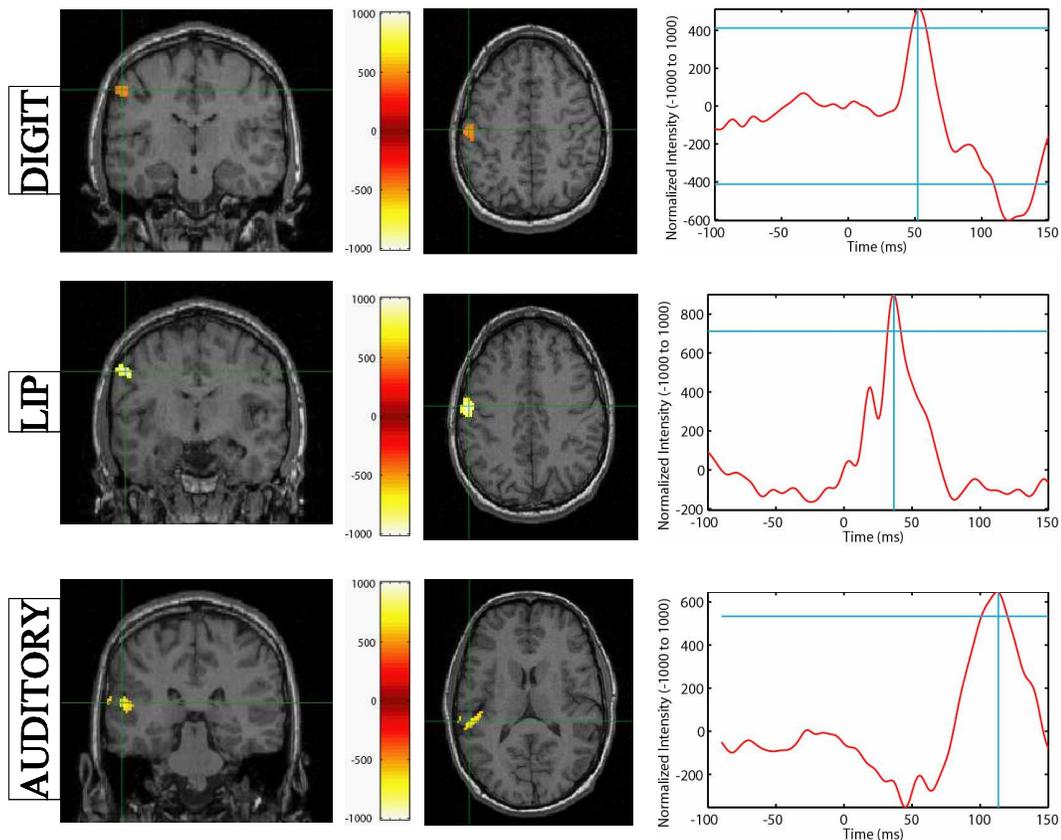


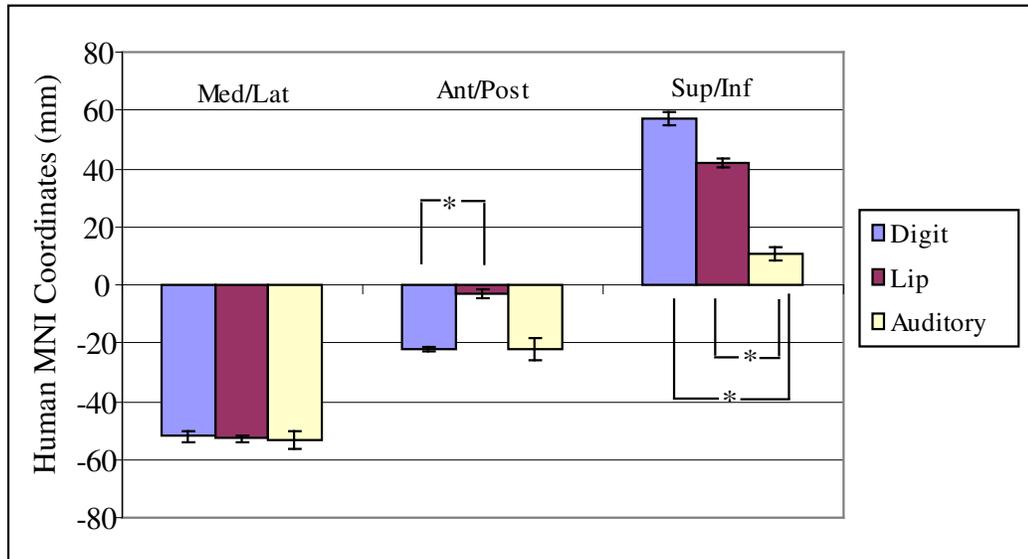
Figure 5.1 Beamformer source localization of human MEG dataset. Top row: digit localization; Middle row: lip localization; Bottom row: auditory localization. Left and middle column showing MEG source activity overlaid on coronal and axial slices of structural MRI, occurring at the peak latency indicated by blue vertical bar in adjacent time-series. The 100 most active voxels at the peak latency are shown, which were subsequently used to seed the DTI tractography (see Methods and Fig. 3). The amplitude threshold used to produce these 100 voxels is indicated by the horizontal blue lines. Crosshairs on MRI indicate voxel with largest magnitude activity. Right column indicates source time-series estimated by beamforming at the peak voxel in crosshairs.

Fig. 5.1 shows the top 100 voxels at the peak latency of the first major neural response (~ 45 ms for digit, ~ 35 ms for lip, and ~ 110 ms for auditory) that were used to create the seed ROI for DTI tracking. This ROI was selected out of over 20,000 total voxels reconstructed in the 2mm grid, leading to an active region of about 0.8cc out of ~ 175 cc examined. The voxel coordinates were then transformed from the original MRI to the space of the DTI images and was used as the starting point for tractography.

The group localization data from the human subjects are reported in Fig. 5.2. The auditory coordinate was found to be significantly inferior to both digit and lip. Though the digit representation tended to be superior to that of the lip, the standard deviation was large. The location of the digit representation was significantly posterior to the lip localization, following the path of the central sulcus. No other comparisons were significant.

5.3.2 Spatial separability of monkey MEG MVAB results

The results from the previous chapter show that when extra care is taken with removing noisy sensors and using the best set of eigenvalues for the MVAB, results can be obtained equivalent to using the more sophisticated methods with built-in denoising. The spatial distribution of the localized sources in the brain for tactile and auditory stimulation generally agreed with known neuroanatomy. Fig. 5.3 shows the MVAB source reconstruction for the MEG dataset from monkey

Figure 5.2 Average MVAB localization results in four humans

M1. The data from M1 demonstrate that the digit localization was more medial and superior than the lip localization, corresponding to the known somatotopic organization in primary somatosensory cortex. The localization of auditory cortex was inferior to that of tactile stimulation, on the opposite side of the lateral sulcus. The spatial map of activation overlaid on the structural MRI is shown for the latency with largest amplitude (largest peak latency). The reconstructed time series (on the left) is shown for the voxel with maximal activation (in green crosshairs on MRI). The largest peak latencies in this example are at 13ms for digit stimulation, 7ms for lip stimulation, and 62ms for auditory stimulation. The 21 most-active voxels are displayed, which are the same used to create the seed ROI for DTI tracking.

MVAB localization results from all four monkeys showed peaks in roughly

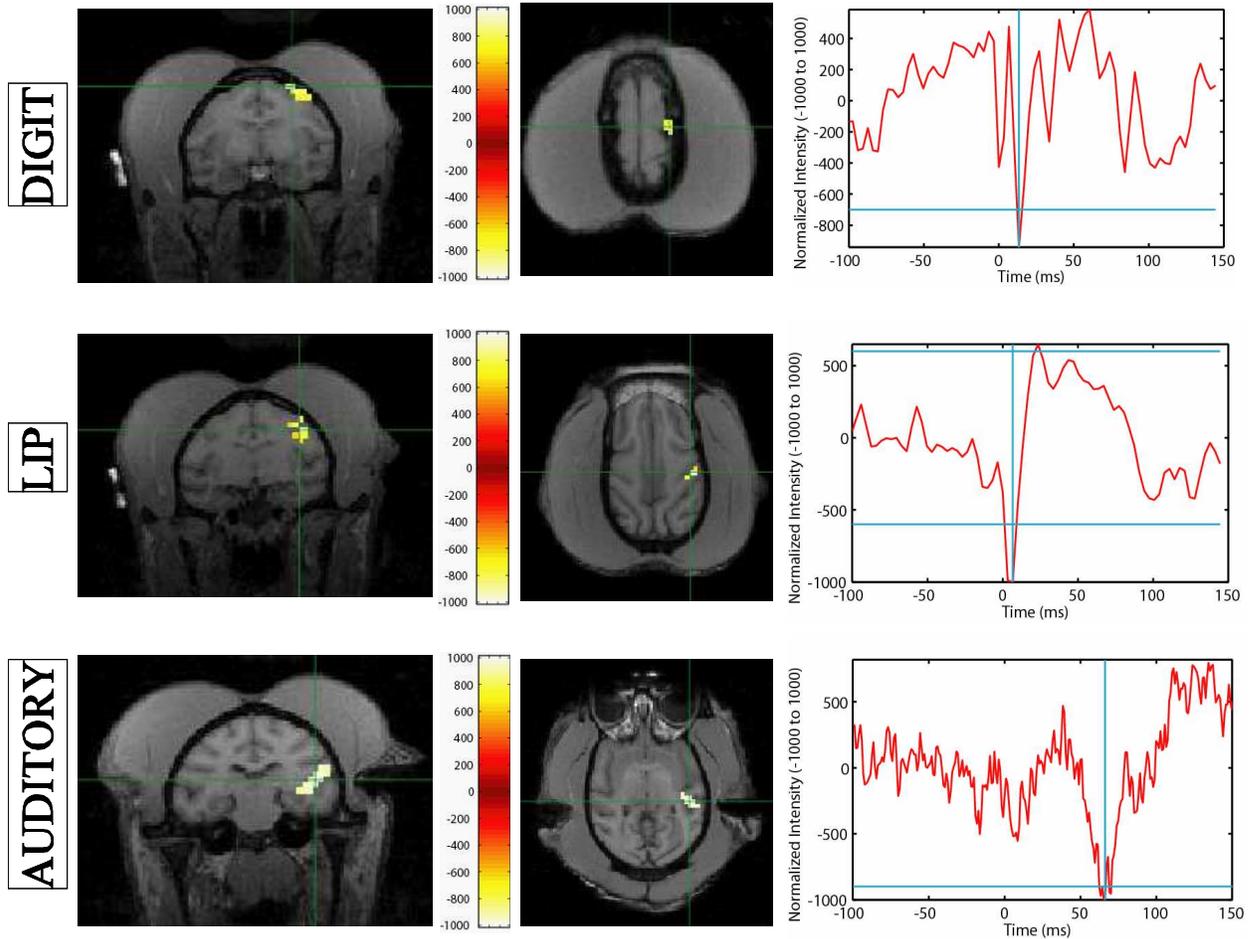


Figure 5.3 Beamformer source localization of MEG dataset from monkey M1. Top row: digit localization; Middle row: lip localization; Bottom row: auditory localization. Left and middle column showing MEG source activity overlaid on coronal and axial slices of structural MRI, occurring at the peak latency indicated by blue vertical bar in adjacent time-series. The 21 most active voxels at the peak latency are shown, which were subsequently used to seed the DTI tractography (see section 5.4 and Fig. 5.6) The amplitude threshold used to produce these 21 voxels is indicated by the horizontal blue lines. Crosshairs on MRI indicate voxel with greatest magnitude of signal. Right column indicates source time-series estimated by beamforming at the peak voxel in crosshairs.

the anticipated latency (based on sensor data) and location (based on known neuroanatomy), similar to Fig. 5.3. However, some variability existed as well. For instance, the digit localization in Fig. 5.3 does extend into pre-central gyrus, although no motor activity is expected in this task. Only the runs with clear responses in the sensor data (described in the previous chapter) were processed for source localization and further quantification of latency and localization. In 7/10 digit datasets, the beamformer reconstruction produced a spatial peak in the expected location and latency (at least one dataset from each animal). In the remaining three datasets, the source estimates were (1) more medial and inferior (far into white matter), (2) more medial and posterior (in white matter and too posterior from the central sulcus), and (3) near the central sulcus but much lower in power relative to an activation on the opposite hemisphere at a later latency. Likewise, 7/12 lip datasets showed activation in the expected region (at least one dataset from each animal). The source estimates in the remaining five datasets were (1) too posterior, (2) in the center of the head, (3 and 4) very inferior into temporal lobe, and (5) near central sulcus but much weaker than other sources in frontal areas. Finally, 4/5 auditory data sets were localized to the expected temporal lobe area (in 3/4 monkeys). The remaining localization result was in the most inferior portion of the temporal lobe.

The AC-PC coordinates of the voxel with largest MEG MVAB activation across space for the time peak of interest ($\sim 12\text{ms}$ for digit, $\sim 8\text{ms}$ for lip, and

~ 50 ms for auditory) was noted as well as the actual peak latency. For comparing the mediolateral variation across stimulation types and animals, the absolute value of the mediolateral coordinate was used since M1 was stimulated on the left while the other three animals were stimulated on the right. The active region (ROI) used to seed the DTI tracts were chosen to be the top 21 most active voxels (about 0.17cc) at the afore-mentioned time latency, out of an average of 12,582 voxels (range 9295-18003) (~ 100 cc) total in the chosen VOI covering the whole cerebrum. The voxel coordinates were then transformed to the space of the DTI images. Coregistration was focused on aligning the central sulcus from the structural MRI to the non-diffusion-weighted ($b=0 \text{ s} \cdot \text{mm}^{-2}$) echo-planar image, rather than applying a global coregistration, warping or distortion-reduction algorithm.

In some reconstructions, a few extra voxels in other regions further away sometimes also showed similar strength of activity. However, since these other activations were not consistent across runs or animals, and were not in plausible locations, they were not quantified further or used in calculation of the ROIs used to seed the DTI tracking. The successful MEG MVAB reconstructions which showed peaks in roughly the expected areas were the ones used for DTI tracking. In the cases of the two monkeys that had two different DTI scans (M3 and M4), all MVAB localizations were roughly similar on both. This finding supports that the coregistration from the MEG fiducials to the MR structural image to the DT

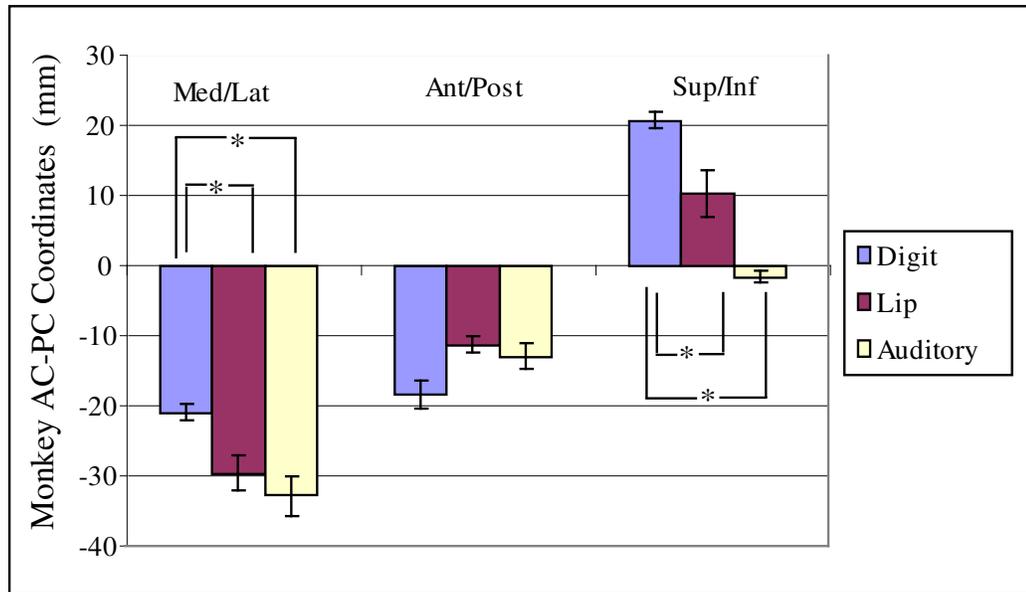
image was not different between the two scans for either monkey.

Average latencies and peak source locations using MVAB across multiple datasets from the four monkeys are listed in Fig. 5.4, with the localizations defined in AC-PC coordinate system (see Section 4.2.1). Note the standard error of the localization was under 3mm in all cases, well within expected errors of MEG source localization, and indicating overall agreement across animals.

Comparing the MVAB localizations across stimulation types, the digit localizations were found to be significantly different in the superior-inferior dimension and in the medial-lateral dimensions from both the lip and auditory localizations. No other significant localization differences were found. While lip and auditory could not be found significantly different from each other, in most cases lip localizations were clearly above the lateral sulcus while auditory localizations were below the lateral sulcus. Significance of differences between the coordinates of the hand, face and auditory representations for both monkeys and humans were tested using two-tailed pairwise Students t-tests corrected for family-wise errors (thus for three tests $p < 0.016$ is significant).

5.3.3 Discussion of MEG MVAB results and neuroanatomical variation

The somatosensory MEG localization results are in agreement with the known somatotopic representation of the body along the post-central gyrus in

Figure 5.4 Average MVAB localization results in four monkeys

both humans and monkeys with the foot area most medial and superior, progressing to the hand, then face area at the most lateral and inferior position (Kaas et al., 1979; Penfield and Boldrey, 1937; Wood et al., 1988). Some variability in estimated source location was seen around the central sulcus. From these and other studies including Zhu et al. (2007), it is assumed that the tactile stimulus is activating cutaneous receptors in both areas 3b and 1, and possibly area 2. Due to the spatial resolution limitations of MEG, discrimination between areas 3b, 1 and 2 from the MEG data based on localization is not expected, as these areas are only several millimeters apart.

These MEG localization results support the assertion that the posterior bank of the central sulcus is the cortical origin of the large early peak seen in the MEG data. Since the spatial resolution of MEG is not precise enough to distinguish

between areas 3b and 1, other data acquired invasively in humans for clinical purposes can be examined in conjunction with the noninvasive MEG results to help give support to a more precise statement about the underlying cortical origin. Previous studies (Wood et al., 1988; Allison et al., 1989b) have examined the cortical origin of median-nerve electrically-stimulated evoked responses through use of subdurally-placed electrodes in human patients requiring neurosurgery. They found the strongest peak in area 3b active first (N20-P30), with area 1 peaking 5ms later (P25-N35). A later wave occurred as N45-P80-N180 from area 3b and P50-N90-P190 from area 1 (Allison et al., 1989b).

Since neurons in area 3b are generally located in the posterior bank of the central sulcus, their tangential orientation relative to the surface of the head make them most likely to be detected with MEG, while the radially-oriented sources of area 1 on the crown of the post-central gyrus are least likely to be detected with MEG (Hamalainen et al., 1993). Therefore it is most likely that area 3b is contributing primarily to the MEG data shown here, although it has been shown that the location of the maps can be shifted dramatically relative to the sulcus, with 3a on the anterior bank, in the depth of the sulcus or even on the posterior bank of the sulcus (Huffman and Krubitzer, 2001). Inter-subject differences in cortical gyral anatomy and the relative location of areas 3b and 1, as well as head shape, suggest that area 1 may contribute more in some subjects than others. Thus while some of the variation observed is undoubtedly due to technical considerations, it

is likely that it reflects anatomical variation as well.

While Sutherling et al. (1988) have shown good correspondence of median-nerve stimulated responses measured by both MEG and ECoG, validation of localization of a more wide variety of evoked responses is still needed (Halgren, 2004; Dalal et al., 2007; Hoogenboom et al., 2006). Furthermore, the relationship between MEG time series data and the underlying neural activity is not precisely known. It is hoped that by comparing this MEG data to local field potential recordings obtained in the same animals to the same stimuli, a better understanding may be gained of the relative contribution of the somatosensory areas to the signals recorded in MEG and the factors affecting these contributions.

5.4 DTI Tractography using MEG MVAB Localization

5.4.1 Tractography overview

DTI is a specific type of MRI acquisition that is sensitive to water diffusion in several arbitrary directions. In gray matter or cerebral spinal fluid (CSF), water is free to diffuse in any direction, resulting in a low anisotropy (directionality) within a voxel. In voxels containing coherently arranged white matter tracts, water is restricted to moving in the direction of these tracts, thereby resulting in a high degree of anisotropic diffusivity in the voxel. The direction with greatest diffusivity is parallel to the axonal bundles and is the primary eigenvector. Using

the primary eigenvector of diffusion in each voxel, one can trace the connections starting from a seed region in the brain, termed tractography.

Studies of the validity and reliability of DTI are sparse. Faithful reproductions of the major known fasciculi in human cortex have been shown (Catani et al., 2002). These results were obtained by starting from seed regions in white matter, marked on the fractional anisotropy image using neuroanatomical references to place two regions of interest in white matter that are both known to participate in the same tract. However, common tractography methods are limited in accuracy due to crossing fibers and partial volume effects. The accuracy with which such algorithms specify axonally-connected regions is unknown.

Several groups have demonstrated the feasibility of DTI in macaque monkeys and showed tracts started from anatomically-defined seed regions in white matter as a first step towards validation of these data (Croxson et al., 2005; Parker et al., 2002; Ramnani et al., 2005; Tuch et al., 2005). Several regions have been examined, such as tracts to prefrontal cortex, the corticospinal tract and optic radiations, cerebral peduncle, and the dorsolateral convexity, pontine decussation, and pulvinar and temporal subcortical white matter, respectively. Dauguet et al. (2007) have performed DTI tractography in macaque monkeys from anatomically defined somatosensory and motor areas and further have compared these tracts with WGA-HRP injected tracts in the same animals. They showed a generally good agreement between the methods but also demonstrated limitations of DTI

far from the initial seed region. However, the ability to determine the validity of connections seeded from known, functionally-defined gray-matter regions-of-interest (ROIs), is of major interest to neuroscientists and neurosurgeons.

This MEG-seeded DTI experiment needed to use a well-established MEG source reconstruction method with minimal bias to obtain the most reliable source estimates for seed locations. MVAB has been used for over five years and has been shown to have minimal bias relative to sLORETA even in noisy data (Sekihara et al., 2005). Even though the new probabilistic models proposed in Chapter 2 have been shown to improve performance over MVAB especially in noisy data (Chapters 3 and 4), it has also been shown in Chapter 4 that, when the data is focused on a particular brain region and sensors with large artifacts are removed, MVAB performs just as well as the new probabilistic methods. Thus, it was favored to use MVAB as the more established method, and DTI tracts have been seeded only from MVAB results.

It is beyond the scope of this dissertation to examine the effects of slightly different seed locations (based on different MEG localizations) on DTI tractography results. It is likely that differences in tract patterns are most due to different seed areas (i.e. digit versus lip) than shifts of a given seed location within that area. The MVAB-based MEG-seeded DTI results shown here are mainly to demonstrate the feasibility and significance of this type of multimodal data integration. Hence, this section focuses on the additional information of separability

obtained from distinct anatomical connectivity patterns resulting from different functionally-defined seed locations.

5.4.2 DTI acquisition and processing

The same four healthy human volunteers (three male and one female, ages 23-28) used for MEG experiments in the previous chapter also participated in the DT-MRI experiments. Additionally, the same four adult male macaque monkeys (9-15 kg) were scanned using DT-MRI at the same time as their structural MRI described in the previous chapter.

Localization of primary somatosensory and auditory cortex obtained using MVAB MEG was coregistered to MR data (described in previous chapter) and used to seed white-matter tracking from diffusion tensor images collected on the same subjects. All studies were performed with approval of the UCSF Committee on Animal Research or the Committee on Human Research. The acquisition parameter setting and analysis of DTI data were selected and performed by Drs. Henry and Berman.

MR and DTI acquisition: Humans

All MR scans were performed on a GE Signa 1.5T scanner (GE, Milwaukee, WI). A standard GE birdcage volume coil was used for all subjects and scans. Donut-shaped MR-contrast fiducials were placed on the subjects head in the same

positions as the localizing coils in the MEG session to assist in coregistration of MEG and MR datasets. After a sagittal localizing scan, a 3D-SPGR image was acquired (flip angle= 40 deg, TR= 27ms, TE=6ms, FOV= 240 × 240mm, 1.5mm slice thickness, 256 × 256 × 124 pixels, in-plane resolution 0.94 × 0.94mm). Spin-echo echo-planar diffusion-tensor images were subsequently acquired with the following parameters: TR= 6.2s, TE= 105.3ms, NEX= 9, six (plus null) diffusion-sensitizing directions, $b = 1000 \text{ s} \cdot \text{mm}^{-2}$, in-plane resolution 1.72 × 1.72mm, and 2.1mm thick slices. The DTI volume spanned from the middle of the pons to the vertex of the brain.

MR and DTI acquisition: Monkeys

Macaque monkeys were initially anesthetized using I.M. ketamine hydrochloride (10mg/kg). Atropine (0.04mg/kg) was also given I.M. to help reduce secretions. An I.V. catheter was placed and animals received lactated Ringers solution (10mL/kg/hr). Animals were intubated and anesthesia was maintained using 1-2% isoflurane. Heart rate and SpO₂ were monitored at all times with an MR-compatible pulse-oximeter. Heart rate, respiratory rate, temperature and SpO₂ were documented every 30-60 minutes. Heated water bottles and blankets were used to keep the animal warm. An experimenter stayed in the MR room at all times.

All MR scans were performed on the same GE Signa 1.5T scanner as used for humans. The monkeys were placed in the prone position and their heads

were secured in an MR-compatible stereotaxic frame. A five inch diameter surface coil was placed on the superior part of the head. Donut shaped MR contrast fiducials were glued and taped onto the same three positions as the localizing coils in the MEG session to assist in coregistration. After a sagittal localizing scan, a 3D-SPGR image was acquired (flip angle= 40 deg, TR= 27ms, TE= 7ms, FOV= 190 × 190mm, 1.0mm slice thickness, 256 × 256 × 124 pixels, in-plane resolution 0.74 × 0.74mm). Monkeys M1, M3 and M4 all received two scans on different days; monkey M2 received only one scan.

Spin-echo echo-planar diffusion weighted images were subsequently acquired with six diffusion sensitizing directions at $b=1000\text{s}\cdot\text{mm}^{-2}$ in addition to an image volume with no diffusion weighting. Slight alterations in scan parameters were made across subjects to optimize scan quality by minimizing distortion. The range of parameters used was: 5-7.9s TR, 70-105ms TE, 1.4-1.72mm isotropic voxel size, and 12-24 NEX. Also, 1-shot EPI was used for 2 animals (M1 and M2) and 2-shot EPI was used for the other two animals (M3 and M4). The number of slices was varied to ensure whole brain coverage, ranging from 23-39 slices depending on the orientation of the brain in the scanner.

DTI processing: humans and monkeys

The diffusion tensor and associated diffusion eigenvectors (Basser et al., 1994) were calculated at each voxel using in-house software developed by Dr. Berman written in the C programming language. The principal eigenvector of each voxel

was assumed to reflect the orientation of axons in that voxel. Fiber tracking was performed using software written in Interactive Data Language (ITT Visual Information Solutions, Boulder, CO) by Dr. Berman and based on the Fiber Assignment by Continuous Tracking method (Mori et al., 1999). Starting regions for fiber tracking were defined by MEG and contained 21 voxels for monkeys and 100 voxels for humans. In each starting voxel, 27 densely seeded starting points arranged in a $3 \times 3 \times 3$ grid were used to launch fiber tracks. Fiber trajectories were then created by following the primary eigenvector from voxel to voxel in 3D. The tract continued until it exited the brain or the angle between the eigenvectors of two successive voxels was above 50-80 degrees, indicating unrealistic white matter pathways.

5.4.3 DTI results

DTI results: Humans

The tracts reconstructed from the digit, lip and auditory seeds in the human DTI data show qualitative differences. Fig. 5.5 shows the complete set of digit, lip and auditory tracts from subject H1, as well as examples from subjects H2 and H3.

Digit seeded data showed local connections in adjacent gyri (4/4), tracts to S2 (2/4) and tracts descending the thalamocortical pathway (2/4) (Fig. 5.5b). Short tracts to nearby posterior parietal and motor areas were also consistently seen.

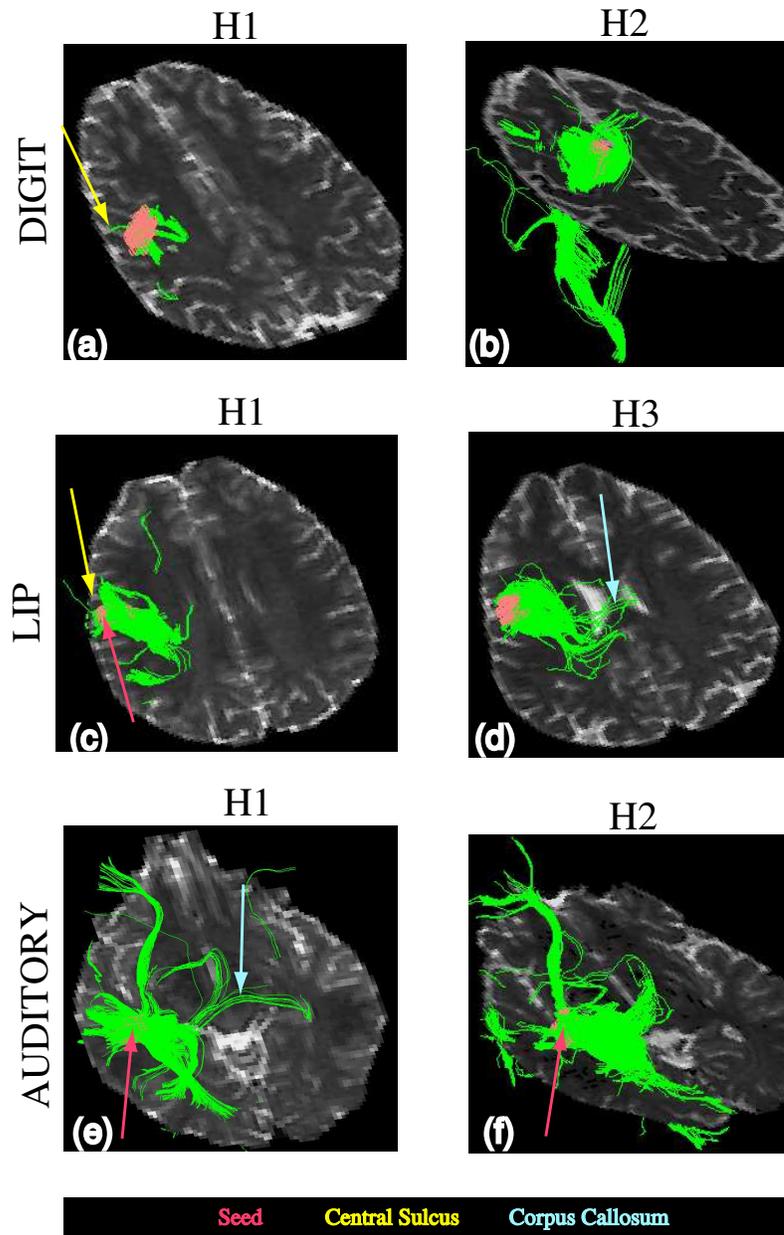


Figure 5.5 DTI tractography results in humans. Top row: digit-seeded tracts; Middle row: lip-seeded tracts; bottom row: auditory-seeded tracts. The green lines indicate the tracts that met the anisotropy and angle criteria (see Methods), as seeded from the ROI created from the MEG source localization (pink), which is obscured by tracks and/or slice in some images. The MRI slice shown is a few slices inferior of each ROI and is from the non-diffusion-weighted echo-planar image.

In general, the tracts seeded from the lip representation were similar to digit-seeded tracts showing local connections (4/4), tracts to the thalamus (3/4) and presumed connections with S2. Unlike the pattern seen for the digit representation, connections also crossed the midline at the corpus callosum (3/4) (Fig. 5.5d) consistent with the crossing location identified by Hofer and Frahm (2006). An additional tract connecting the frontal lobe is seen in subject H1 (Fig. 5.5c), possibly due to the ROI containing some temporal lobe voxels.

The auditory-seeded tracts in humans show local temporal lobe connections (4/4, Figs. 5.5 e and f), tracts crossing the midline in the corpus callosum (2/4, Fig. 5.5e), tracts to the frontal lobe (3/4, Figs. 5.5e and f), and tracts to posterior temporal cortex (3/4). Thus, the auditory-seeded tracts show connection patterns that are distinctly different from the digit- and lip-seeded tracts, as well as showing general agreement with known connections.

DTI results: Monkeys

The hand, face and auditory ROIs derived from MEG MVAB localization in monkeys also resulted in qualitatively different tracts obtained from the same diffusion tensor image. Fig. 5.6 shows a full set of examples from M1 and M2, as well as hand and lip from M4 and auditory from M3 for additional examples.

The digit representation seed showed connections locally (5/9) and with posterior parietal area (2/9), as well as with the presumed location of S2 (5/9, example Fig. 5.6b) and the thalamus (1/9). Nine total digit-seeded results were

obtained from five digit ROIs seeded onto one DTI scan each (from M1 and M2) and two digit ROIs were seeded onto two DTI scans each (two DTI scans from M3 and two DTI scans from M4). Previous neuroanatomical studies in monkeys have shown that the hand area in S1 is connected to ipsilateral S2, posterior parietal cortex, and the thalamus (Krubitzer and Kaas, 1990). Thus, the tracts seen are in general agreement with known connections of the digit portion of area 3b. However, there was variability across the tractography results. In 4/9 cases, projections to frontal motor area were seen (example Fig. 5.6a): the ROIs in 3/4 of these cases included voxels extending back to the intraparietal sulcus. In only one case expected connections were absent, with only local connections (example Fig 5.6c). In one other case, a tract to the prefrontal region was seen.

In general, the tracts seeded from the lip representation were similar to digit-seeded tracts, showing local connections (11/11), tracts to the thalamus (5/11) and presumed connections with S2. Unlike the pattern seen for the digit representation, connections also crossed the midline at the corpus callosum (6/11) (Figs. 5.6d, e and f) consistent with the crossing location identified by Hofer and Frahm (2006). From previous neuroanatomical studies, the face representation in S1 has also been shown to be connected to ipsilateral S2, posterior parietal cortex, and thalamus, but, unlike the hand representation, has additional connections to S1 in the opposite hemisphere (Killackey et al., 1983). Similar to digit-seeded tractography, the lip-seeded tractography is therefore consistent with known neu-

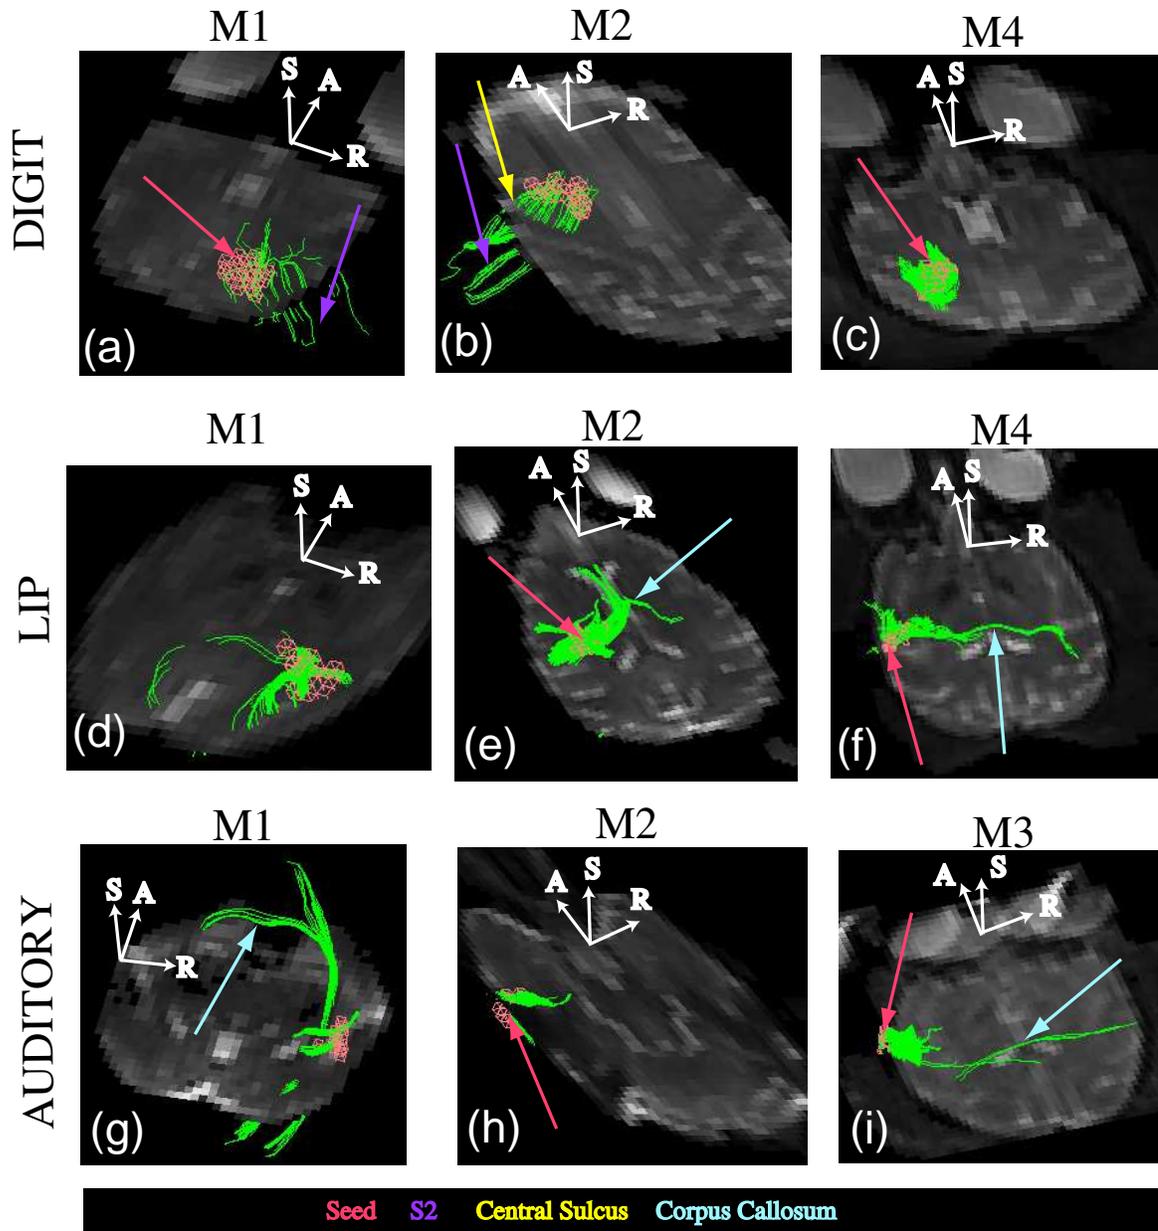


Figure 5.6 DTI tractography results in monkeys. Top row: digit seed, middle row: lip seed, bottom row: auditory seed. 21 voxels for each ROI in pink, tracts in green, overlaid on slice of non-diffusion-weighted echo-planar image chosen to highlight location of tracts relative to main features such as the corpus callosum. In some cases, the pink ROI is covered by the green tracts and/or MRI slice, but the pink arrows indicate their location.

roanatomical connections but also shows variability like the digit data. The tractography result in monkey M2 (Fig 5.6e) shows an additional tract in a medial area coursing anteriorly, most likely due to crossing fibers. Three additional data sets showed connections crossing the midline in locations that appear to be either anterior or posterior to position of the callosum. This may be due to some voxels in the lip ROI extending into areas of S2 or temporal lobe.

The auditory-seeded tracts in monkeys show local temporal lobe connections (5/5, Fig. 5.6h), tracts crossing the midline in the corpus callosum (3/5, Figs. 5.6g and i), and tracts to the frontal lobe (1/5, Fig. 5.6g). One case shows a connection to thalamus; two cases show connections crossing the midline in an area not consistent with the location of the corpus callosum. Both auditory seeds in monkey M1 also produced posterior-going tracts within temporal lobe (Fig. 5.6g). Kaas and Hackett (1998) have demonstrated primary auditory cortex (A1) to have dense local connections with the belt and parabelt regions as well as being connected to contralateral A1. Furthermore, they showed frontal connections to the orbito-frontal region and to frontal eye fields. Thus, the auditory-seeded tracts show connection patterns that are distinctly different from the digit- and lip-seeded tracts, as well as showing general agreement with known connections.

Comparing and averaging across monkey and human DTI results

Showing consistency between monkey and human DTI tractography results helps for two reasons. First, evidence of similar results in monkeys compared

with humans lends support that issues such as greater partial voluming and distortion in the monkey data do not hamper the monkey data significantly for the type of questions asked here. Second, after monkey DTI tractography has been compared with the tracer histology in the same animals, the amount of confidence given to each DTI tract based on this validation could be more convincingly extrapolated to the human tracts if the human and monkey tracts are similar. Figure 5.7 shows the percentage of scans showing a given tract, averaged over all results.

Fig. 5.7(a) shows the percentage of tracts seen averaged over all seed types, comparing monkey with human results. This figure shows that connections to local regions, thalamus, corpus callosum crossings, and S2 are consistent across both species, while connections to posterior parietal, motor and other regions are more variable across species. This may be due to U-fibers connecting primary somatosensory one gyrus forward (motor) or back (posterior parietal) were too difficult to detect in the relatively larger voxels in monkey DTI.

Fig. 5.7(b) shows the percentage of tracts per seed types, comparing monkey with human results. These results support the following four claims that 1) lip and auditory regions have connections across the corpus callosum (roughly half the scans) while digit area does not, 2) somatosensory areas have connections with posterior parietal and motor areas while auditory does not, 3) auditory has connections with frontal areas while somatosensory does not, and 4) digit S1 has connections to S2 while auditory does not and lip S1 to S2 connections cannot be

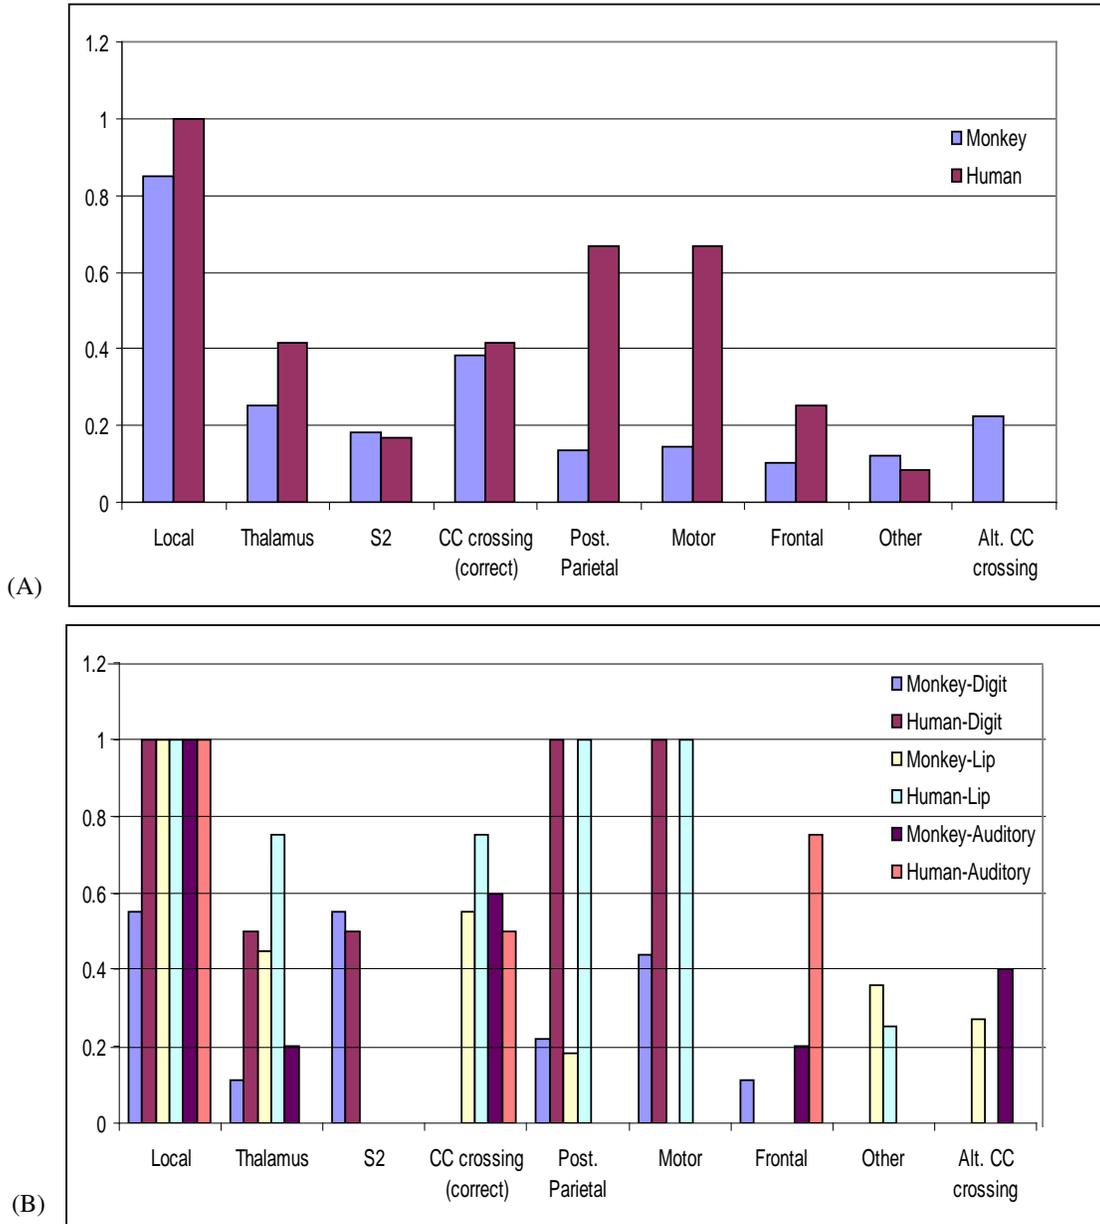


Figure 5.7 DTI tractography summary in monkeys and humans. (Average across 9 Digit seeds, 11 Lip seeds, and 5 Auditory seeds on 6, 6, and 4 monkey DTI scans, respectively, and 4 Digit, 4 Lip and 4 Auditory seeds on 4 DTI scans in humans)

determined at this resolution. It can be seen that claims 1 and 4 are supported by both monkey and human data independently, while claims 2 and 3 are more strongly supported by the human data.

5.4.4 Discussion of MEG-seeded tractography

Expected connections

Previous neuroanatomical tracer studies in monkeys have shown that the hand representation in 3b is connected to ipsilateral secondary somatosensory cortex (S2), posterior parietal cortex, as well as thalamus (Burton et al., 1995; Darian-Smith and Darian-Smith, 1993; Darian-Smith et al., 1993; Jones and Powell, 1969a,b; Vogt and Pandya, 1978). The face representation in 3b is similarly connected but also has connections to contralateral S1 via the corpus callosum (Killackey et al., 1983). Since the lip representation in S2 is adjacent to the lip representation in S1 (e.g. Krubitzer et al. (1995)), the connection between these two regions is difficult to distinguish from local connections at this level of resolution. Thus, the hand and face representations each have a unique pattern of connections that we were able to distinguish in the data reported here. Finally, primary auditory cortex (A1) has been shown to have dense connections with the belt and parabelt regions as well as being connected to the contralateral A1 via the corpus callosum, to the orbito-frontal region via arcuate fasciculus and to frontal eye fields (Kaas and Hackett, 1998; Morel et al., 1993; Luethke et al., 1988).

Thus, it was possible to identify a unique pattern of connections based on the MEG-derived auditory seed that was similar to pre-existing non-human primate data.

Hofer and Frahm (2006) have used DTI to illustrate the partitioning of the corpus callosum in humans, by defining anatomical-based ROIs near specific areas of cortex. They found an anterior-to-posterior parcellation of the corpus callosum connecting prefrontal, premotor, motor, primary sensory, parietal, occipital and temporal cortex. The connections seen in the MEG-based lip-seeded data are consistent with their findings of a crossing roughly in the midsection of the corpus callosum. However, the auditory crossings seen in the data here are further anterior than seen in Hofer and Frahm (2006). This difference could be due to technical constraints (see below) or anatomical factors.

While the gross pattern shown in the DTI data was similar to pre-existing anatomical data, the rich detail obtained using neuroanatomical tracing techniques was absent from the imaging data sets. A promising direction for validation and study of non-invasive tractography has been demonstrated by Saleem et al. (2002), who injected a Mn_2^+ tracer along with WGA-HRP. The two tracers followed very similar and highly specific paths, thus validating that Mn_2^+ (an MRI-visible contrast agent) could be used for development of DTI tractography methods. In addition, collaborators in Drs. Disbrow and Krubitzer's labs have gathered neuroanatomical data from the monkeys used in this study in order to

perform a comparison of invasive and non-invasive techniques.

Functional based seed regions for DTI tractography

Previous DTI studies have demonstrated feasibility of collecting DTI data from monkeys and performing tractography (Crosson et al., 2005; Parker et al., 2002; Ramnani et al., 2005; Tuch et al., 2005; Dauguet et al., 2007). However, in these studies, the seed region was chosen based on anatomical landmarks, not obtained from any functional specification. Furthermore, their seed regions all began in a user-defined white-matter area. Previous studies have shown the importance of choosing a seed region based on a functionally defined seed, for example using fMRI. Due to anatomical variation of functional localization in different individuals, an anatomical-based ROI could misrepresent the connections studied from that region. Guye et al. (2003) have shown tractography results from a seed in the motor region defined by functional MRI in humans. Kim et al. (2003) demonstrated feasibility of using an fMRI-derived seed in cat visual cortex for DTI tractography. Alternatively, in this work, the feasibility of creating seed regions for tractography from MEG source localization estimates was demonstrated. The use of both fMRI and MEG for functionally defining a seed region for DTI tractography most likely will increase in coming years, both for neuroscience questions and presurgical planning. Functional mapping and DTI are complementary techniques since MEG can locate functionally specific cortices while DTI tractography can delineate the connecting white matter.

One potential concern for using a functionally-defined ROI is that it mostly will exist in gray matter with low diffusion anisotropy. In considering the number of voxels to use to define the ROI, fewer voxels would more precisely define the region, but would also more likely be exclusively in gray-matter. DTI fiber tracking from gray matter structures requires that track trajectories traverse the gray-white matter boundary or that the starting region includes small adjacent regions of white matter. Increasing the number of voxels too much would lead to imprecisely defined regions. It was found that using the top 21 voxels (0.17cc) from monkey MEG activation and the top 100 voxels (0.8cc) from human activation was a good compromise between these considerations. Alternatively, one could modify some aspect of the DTI acquisition in order to assist seeding in/near gray matter. Ronen et al. (2003) has shown that if a lower b-value is used, the slower diffusion in gray matter can be measured and tracts initiated in white matter can continue further into gray matter than if $b = 1000 \text{ s} \cdot \text{mm}^{-2}$ was used, although this would require additional scan time.

Another potential concern of using MEG derived locations as fiber tracking seed regions involves the inherent spatial uncertainty due to the ill-posed MEG inverse problem. The main drawback of the MVAB method used here is when two sources are correlated in time, such as can sometimes be the case for bilateral auditory cortex, the beamformer is known to produce source power that is artificially reduced, sometimes to the point of not being detected. Analysis of auditory

data was restricted to sensors covering only one hemisphere at a time, which by-passes this potential problem. While later components of somatosensory responses can include concurrent activation of S1 and S2, the early components are primarily dominated by S1 (Hari et al., 1984). Thus the coherent source problem is minimal. The localization results from the new probabilistic methods discussed in the previous chapter show similar results when specific sensors were used, so choice of MVAB over these other methods should not result in any largely different tractography results (i.e. the majority of variation is in seed selection and/or DTI methodology, not source localization).

Effects of DTI acquisition parameters and tractography algorithms

Some variability in the tracts destination was seen between individuals in both the monkey data and human data. Several technical issues may contribute to some of this variability. First, the human data was collected at a greater anatomical resolution than the monkey data based on relative brain structure size. Even though the voxel size was approximately the same for both species, the monkey data will exhibit more partial-volume overlap of cytoarchitectonically different regions within a given voxel. In addition, several groups have suggested that a more robust estimation of primary diffusion direction can be made by acquiring more than the six necessary directions (at least 20 but up to 60) rather than repeating acquisition of the 6 directions (Jones, 2004; Poonawalla and Zhou, 2004; Tuch et al., 2002). High angular resolution diffusion imaging, including q-ball

techniques, have successfully identified crossing fiber population within cortical gray matter (Tuch, 2004; Tuch et al., 2005; Wedeen et al., 2005). Hopefully future studies in primates will incorporate these technological improvements in DTI acquisition.

Ciccarelli et al. (2003) showed the amount of variability occurring across normal human subjects and that this variability should be taken into account before comparing with clinical populations. They emphasize that different tracts have different amounts of potential tracking problems from crossing fibers or abrupt turns, and that the variability for a given tract should be considered. Between session reproducibility and between subject variability of diffusion MR and tractography measures were examined by Heiervang et al. (2006). They found that fractional anisotropy was more reproducible than total tract volume, that tracts defined by two ROIs (such as start and end points) are more reproducible than just one seed ROI, and that measuring more than 12 diffusion directions did not increase reproducibility but that less than 12 did. These are further points to incorporate for future studies.

In this work, a common, robust tractography algorithm was used (Mori et al., 1999); many other tractography algorithms have been proposed. Some probabilistic techniques report the uncertainty or probability of a track (Behrens et al., 2003; Hagmann et al., 2003; Parker et al., 2003). Still, there is no consensus in the field of how to report whether a particular track is significant or not. Hopefully, even-

tually comparing the DTI data obtained in this study with the neuroanatomical tracer results obtained in the same monkeys will validate and provide a gold-standard dataset for comparing various tractography methods.

5.5 Conclusion

MEG can be used to distinguish primary somatosensory from auditory cortex and to identify specific body part representations within somatosensory cortex both in macaque monkeys and in humans. These data, obtained using non-invasive techniques, were in general agreement with previous electrophysiological recording studies performed invasively in macaque monkeys and humans (Penfield and Boldrey, 1937; Wood et al., 1988). Using these functionally defined somatosensory and auditory regions of cortex, distinct patterns of connections are seen using DTI tractography for somatosensory versus auditory cortex and for the hand and face representations within the somatosensory strip. These data, obtained using non-invasive techniques, were in general agreement with previous neuroanatomical studies performed invasively in macaque monkeys (Morel et al., 1993).

Acknowledgments

I thank Liz Disbrow and Sri Nagarajan for supervision of this project, Liz Disbrow, Zhao Zhu, Marianne Lowenthal and Leah Krubitzer for help with animal

preparation, Jeff Berman and Roland Henry for DTI acquisition parameters and tractography processing, Susanne Honma, Anne Findlay, Evelyn Proctor, Kavita Vashi, Leighton Hinkley and Geoff Rau for assistance with data acquisition, Ron Baireuther, Robert English-Fernandez, Patricia Ramsey and Grace Kim for assistance with anesthesia, and Sarang Dalal for comments on MEG processing. I also thank the NIH for funding R01DC4855 and R01NS44590.

Chapter 6

Conclusion

6.1 Overview

New probabilistic methods for estimating neural activity from noisy MEG data have been developed and described here. The mathematical framework was presented in detail in Chapter 2, the performance of these methods on simulated and example real data was shown in Chapter 3, a more thorough comparison of the methods on many real datasets is shown in Chapter 4, and steps towards validation and multi-modal data integration were described in Chapter 5. The basic idea is that if the noise and interference are explicitly accounted for in the data model, their effect on source localization is reduced.

6.2 Future Directions

6.2.1 Direct extensions of models

The graphical models proposed here are just three variations of a class of models that could be used to perform source reconstruction for MEG and EEG data. Certainly many aspects of these models can be modified in order to relax assumptions made here. One assumption is that the background interference sources are stationary across the pre-stimulus and post-stimulus periods. If background is non-stationary, the model could be modified to include a term similar to Φ which would influence the posterior update for \bar{u} . This would occur, for instance, if the power of ongoing brain activity decreased as a result of the stimulus/event onset (such as in the ERS/ERD findings).

One could further relax assumptions about the source and interference factors. They could be modeled by non-Gaussian distributions, although care should be taken in choice of distribution so that they could still be solved for analytically. The temporal progression/smoothness of the factors could be explicitly modeled, such as using an auto-regressive model or Kalman filter.

Since Section 2.4.4 showed how MVAB is an approximation of SAKETINI under certain assumptions, it might also be possible to use advantages learned from SAKETINI to help MVAB. Specifically, in observing some individual results from SAKETINI in the data in Chapter 4, the source estimates from SAKETINI

very closely resembled MVAB source estimates. Both might be localized to a reasonably-accurate location or both might be quite mis-localized. In the cases of mis-localization of the source power estimates, the SAKETINI likelihood map localized correctly. Perhaps some equivalent of a likelihood map for MVAB could be derived which might be used as an improved estimate of source localization.

6.2.2 Improved model order selection

The models themselves have two built-in ways of assisting in the dimension choice. First, the hyperparameters over the columns of the mixing matrices A and B are updated in the M-step; thus their final values can be plotted (such as in Fig. 3.7) and used as indication of which columns (and therefore which factors) contribute more, similar to a plot of eigenvalues for PCA. As long as the number of dimensions chosen is larger than the true dimension, the extra dimensions can be “zero-ed out” by the model as those corresponding hyperparameters decrease.

A second way to examine model order is by learning the model for several values of model order, and then comparing the final maximized likelihood value. In general, the greater the model order, the easier to explain the data, thus the likelihood will increase for greater model order. However, there is usually an inflection point in the plot of likelihood versus model order, again similar to a plot of eigenvalues for PCA, around which the ideal model order resides.

The hyperparameters help decide the true number of factors/dimensions

needed to describe the data, both for pre-stimulus interference factors and for evoked non-localized factors. The figures at the end of Chapter 3 show their utility to some extent, in the sense that as long as the user selects a higher number of factor than needed, the hyperparameters will remove the extra factors that are not needed (by setting their influence to zero). However, their influence is usually not set to exactly zero, but something rather small relative to the other needed factors. The principle of Occam's Razor is paraphrased in several ways and basically states that a model should be as simple as possible but still complex enough to explain the data properly. Applied to probabilistic modeling, the maximized data likelihood will keep increasing as the model order/dimension increases since the data can be better explained. However, this more complex model might not fit the next dataset as well, or be "over-fitting" of certain irregular features of the current dataset. Certain metrics such as Bayesian Information Criteria and Akaike Information Criteria are a function of the data likelihood that penalize model complexity. These metrics might be of use in dimension selection for SEFA, SAKETINI and NSEFALoc in addition to relying on hyperparameters.

6.2.3 Time-frequency analysis

All of the variations of MVAB rely on some estimate of the data covariance matrix in calculating the weight matrix (inverse solution). Instead of averaging the data across trials first, and then computing the data covariance matrix em-

phasizing the stimulus-phase-locked evoked activity, the data covariance could be computed on each single trial first, and then the data covariance matrices averaged. This allows for signals which are not exactly time-locked to the stimulus onset to affect the result. Additionally, the data may be filtered into specific frequency bands. Robinson and Vrba (1999) first proposed this modification to the beamformer to include a frequency-based weight matrix, which then can be used to compare increases or decreases of power in a specific frequency band relative to a control time window. This is the basis of SAM (synthetic aperture magnetometry) included in the software package for CTF MEG systems. Most analyses using SAM report a few static images showing increases/decreases across the brain of the few frequency bands (alpha, beta, and gamma).

It would be of more interest to see how a given location varies over time across frequencies. Gaetz and Cheyne (2003) and others use a broadband data covariance to compute the weight, apply it to broadband data (such as 1 – 70Hz), then in source space at spatial peaks of interest, compute a time-frequency spectrogram. However, this broadband weight might miss out on high-frequency (low-amplitude) signal since it is biased towards low-frequency (high-amplitude) signals. Dalal et al. (2007) proposed to compute a separate weight for each time-frequency window of interest, which allows each weight to be tuned to that segment, allowing for reconstructions that are tighter in spatial, temporal and frequency resolution with less cross-talk. These “active” time-frequency reconstruc-

tions are compared relative to a frequency-specific control time-window.

Thus, in considering SAKETINI and NSEFALoc, if the background is stationary, but the source activations vary drastically over the post-stimulus period, it might help to break up the post-stimulus data into partitions of relatively similar activity. Specifically, the data could be decomposed into time-frequency windows where inference of this model on each window would result in a separate likelihood map. The temporal-frequency dependence of the localizations could then be displayed using the likelihood map rather than source activity map, resulting in different locations being active (above a threshold) at different times/frequencies. Working out exactly how to learn the unknown mixing matrices, such as on concatenated single trials, needs to be determined.

6.2.4 Improved head modeling and forward field errors

In all results shown, we used a single-shell spherical conductor model for generation of the forward field F . Certainly more sophisticated forward fields could be computed using multisphere methods (such as in CTF software), elliptical-shell models (Dassios and Kariotou, 2004) or BEM models (van't Ent et al., 2001), which could further improve performance for all methods. Distributed sources could be modeled explicitly using a forward field computed from patch bases (Limpiti et al., 2006).

However, errors in coregistration of MEG or EEG data to the structural MRI

can affect any computation of forward field. If one wishes to restrict the source activity to one orientation based on cortical surface, errors in this segmentation can lead to errors in orientation estimation. Localization of EEG data has larger errors since tissue conductance values cause uncertainty in the forward field calculation. With MVAB, the array mismatch (use of incorrect forward field) can be partly reduced by adding noise to the data covariance diagonal (Tikhonov regularization), but with the tradeoff of reduced spatial resolution. The eigenspace MVAB reduces this loss of spatial resolution.

Another advantage of the probabilistic framework of the proposed models could be to include a hyperparameter over the forward field F that would enable the algorithm to make small changes to the computed F , using the data to “learn” an improved estimate of the forward field. Since the forward field in the simplest single-shell model is really only a function of the three fiducial points plus sphere center, the full composite lead field (number of sensors by number of voxels) could be allowed to vary slightly only as a function of these 4 transformation coordinates (rather than letting each column vary independently).

All of these tweaks of the inverse method to overcome incorrect forward fields could be examined by purposely added noise to the forward field or shifting the grid that the forward field is based on.

6.2.5 Spatial smoothness and tomographic reconstruction

The scanning framework of solving for one voxel's activity at a time can be relaxed. Spatial smoothness or spatial priors from other modalities, such as structural or functional MRI, could be incorporated. Furthermore, one is not limited to s_n in a single voxel; the above formulation holds for any P arbitrarily chosen dipole components, no matter which voxels they belong to, and for any value of P . Of course, as P increases the inferred value of Φ becomes less accurate, and one might choose to restrict it to a diagonal or block-diagonal form.

6.2.6 Combine MEG data with other modalities

Many debate the benefits or drawbacks of collecting MEG versus EEG data. If possible, collecting both simultaneously would be optimal. Sources weakly picked up by one modality can often be picked up better by the other (except deep sources). One drawback to using simultaneous MEG/EEG is that extra electronic and biological noise enters the MEG sensors from the EEG wires. However, the methods developed in this dissertation are specifically designed to remove such interference, thus should be tested in combined MEG/EEG data.

As already discussed in Chapter 5, MEG localizations can be used to seed white-matter tractography using DTI data. Another possible direction is to use diffusion images to help compute a voxel-by-voxel estimate of tissue conductivity, to be used as an input to finite-element-modeling for EEG forward field (Tuch

et al., 1999).

Finally, MEG and EEG data could be combined with fMRI data to hopefully improve the spatial certainty of MEG/EEG and the temporal shortcomings of fMRI. Many methods for such combination have already been proposed (Liu et al., 1998; Baillet and Garnero, 1997; Ahlfors and Simpson, 2004; Schulz et al., 2004; Mattout et al., 2003; Babiloni et al., 2004; Sato et al., 2004; Auranen et al., 2005; Im and Lee, 2006; Nangini et al., 2007) which mainly use fMRI as a spatial prior into the MEG/EEG inverse solution. These ideas could easily be extended to use of fMRI for altering the source power in SAKETINI and NSEFALoc (such as a prior on s in SAKETINI and a prior on G in NSEFALoc).

As simultaneously collecting EEG data with fMRI becomes feasible with appropriate ballistocardiogram artifact removal techniques, it is becoming more popular to use EEG power time-series (from electrodes, not source estimates) to convolve with the hemodynamic (impulse) response function to create a new regressor to look for significant activations in the fMRI related to the EEG. Enough evidence exists that there are some common underlying factors affecting both MEG/EEG and fMRI (e.g. Logothetis et al. (2001)). Braitenberg and Schuz (1991) show that there are roughly six excitatory synapses for every one inhibitory synapse; thus the metabolic demand for all neural activity is mostly created by signals that are also measurable by MEG/EEG. There seems to be a specific trend that fMRI is more related to increases in gamma frequency and decreases in alpha

frequency (Foucher et al., 2003; Laufs et al., 2003) and a model to account for this (Kilner et al., 2005), although the relationship is not entirely clear-cut (Winterer et al., 2006).

Ultimately, a more integrated graphical model that includes hidden nodes representing different aspects of neural activity which would drive post-synaptic potentials and metabolic demand in possibly different ways so that MEG/EEG and fMRI measurements could be allowed to represent these different aspects. Such an idea has been proposed (Babajani et al., 2005; Babajani and Soltanian-Zadeh, 2006) but further work is still needed. Measuring intermediate biological factors such as blood flow (via arterial spin labeling MRI) simultaneously could also help in inferring a large integrated model.

6.2.7 Networks and causal models

A large branch of systems-level neuroscience and brain imaging are focusing on neural networks, based on the fact that activity in one brain region influences other brain region's activity. While the initial focus of fMRI studies was to find "the area of the brain that does X", more recent studies have focused on causal influences of one region onto another (e.g. Friston et al. (2003)). Each neuron connects with roughly one thousand other neurons; it is widely hypothesized that it is the strength of connections across many brain regions that matters for learning and memory and other higher cognitive functions, not simply activity in

an isolated region.

The source reconstruction methods discussed in this dissertation have not included any prior information of causal influences amongst brain regions, though they certainly could. To start with, the temporal smoothness of evoked potentials should be used simply to provide an auto-regressive form to the temporal statistics of source estimates. Secondly, source activity at a distant region with known important causal influences onto the region of interest should be accounted for in the model to help learn the source activity there.

6.2.8 Relating scales of measurement

For now, invasive studies in animals (and in humans requiring surgery) are still needed to help validate these noninvasive imaging measures of brain function and explain how they relate to local field potential and spike recordings. Interesting work is being pursued on this topic (e.g. Moran et al. (2007); Sotero et al. (2007); Robinson et al. (2005)), though there still is a large discord between the large-scale noninvasive measures of activity (on the order of several cubic millimeters for both MEG and fMRI) and the individual neuron level/spike-rate coding. Even if the inverse solution to MEG and EEG was easy and exact, we still are missing out on the relationship to individual neurons and whether this missing information matters for a given scientific question of interest. New hardware advances for brand new ways to non-invasively measure brain activity are

sure to come (both whole-head MEG and fMRI are not even twenty years old at this point). Advances in modeling local/individual neural output to larger-scale measures are strongly needed to reduce, and maybe one day eliminate, invasive recordings.

Appendix A

Appendix of Useful Matrix Relationships

$$\text{Tr}(AB) = \text{Tr}(BA) \quad (\text{A.1})$$

$$\frac{\partial \text{Tr}(AB)}{\partial A} = B^T \quad (\text{A.2})$$

$$\frac{\partial \text{Tr}(A^T BC)}{\partial B} = AC^T \quad (\text{A.3})$$

$$\frac{\partial \text{Tr}(ABA^T)}{\partial A} = 2AB \quad (\text{A.4})$$

$$A^T BA = \text{Tr}(BAA^T) \quad (\text{A.5})$$

$$\frac{\partial x^T b}{\partial x} = \frac{\partial b^T x}{\partial x} = b \quad (\text{A.6})$$

$$\frac{\partial a^T X b}{\partial X} = ab^T \quad (\text{A.7})$$

$$\frac{\partial a^T X^T b}{\partial X} = ba^T \quad (\text{A.8})$$

$$\frac{\partial a^T X a}{\partial a} = (X + X^T)a \quad (\text{A.9})$$

$$\frac{\partial b^T X^T D X c}{\partial X} = D^T X b c^T + D X c b^T \quad (\text{A.10})$$

Matrix Inversion Lemma:

$$\begin{aligned} (A^{-1} + CD^{-1}C^H)^{-1} &= A - AC(D + C^HAC)^{-1}C^H A \\ (A^{-1} - BD^{-1}C)^{-1} &= A + AB(D - CAB)^{-1}CA \end{aligned} \quad (\text{A.11})$$

If μ is the mean of y over all n , then:

$$\sum_{n=1}^N (y_n - \mu)(y_n - \mu)^T = \sum_{n=1}^N (y_n - \mu)y_n^T \quad (\text{A.12})$$

For expectations over n-th order terms,

$$E_x(x^T Ax) = \bar{x}^T A \bar{x} + \text{Tr}(AR_{xx}) \quad (\text{A.13})$$

Bibliography

- Ahlfors, S. P. and Simpson, G. V. (2004). Geometrical interpretation of fMRI-guided MEG/EEG inverse estimates. *Neuroimage*, 22:323–332.
- Allison, T., McCarthy, G., Wood, C. C., Darcey, T. M., Spencer, D. D., and Williamson, P. D. (1989a). Human cortical potentials evoked by stimulation of the median nerve. i. cytoarchitectonic areas generating short-latency activity. *J Neurophysiol*, 62(3):694–710.
- Allison, T., McCarthy, G., Wood, C. C., Williamson, P. D., and Spencer, D. D. (1989b). Human cortical potentials evoked by stimulation of the median nerve. ii. cytoarchitectonic areas generating long-latency activity. *J Neurophysiol*, 62(3):711–22.
- Arezzo, J., Pickoff, A., and Vaughan, H. G., J. (1975). The sources and intracerebral distribution of auditory evoked potentials in the alert rhesus monkey. *Brain Res*, 90(1):57–73.
- Attias, H. (1999). Inferring parameters and structure of latent variable models by variational bayes. In *Proc. 15th Conf. Uncert. Art. Intell.*, pages 21–30.
- Auranen, T., Nummenmaa, A., Hamalainen, M. S., Jaaskelainen, I. P., Lampinen, J., Vehtari, A., and Sams, M. (2005). Bayesian analysis of the neuromagnetic inverse problem with $l(p)$ -norm priors. *Neuroimage*, 26:870–884.
- Babajani, A., Nekooei, M.-H., and Soltanian-Zadeh, H. (2005). Integrated MEG and fMRI model: Synthesis and analysis. *Brain Topogr*, 18:101–113.
- Babajani, A. and Soltanian-Zadeh, H. (2006). Integrated MEG/EEG and fMRI model based on neural masses. *IEEE Trans Biomed Eng*, 53:1794–1801.
- Babiloni, F., Mattia, D., Babiloni, C., Astolfi, L., Salinari, S., Basilisco, A., Rossini, P. M., Marciani, M. G., and Cincotti, F. (2004). Multimodal integration of EEG, MEG and fMRI data for the solution of the neuroimage puzzle. *Magn Reson Imaging*, 22:1471–1476.

- Baillet, S. and Garnero, L. (1997). A Bayesian approach to introducing anatomic-functional priors in the EEG/MEG inverse problem. *IEEE Trans Biomed Eng*, 44:374–385.
- Baryshnikov, B. V., Van Veen, B. D., and Wakai, R. T. (2004). Maximum likelihood dipole fitting in spatially colored noise. *Neurol. Clin. Neurophysiol.*, 2004:53–53.
- Basser, P. J., Mattiello, J., and LeBihan, D. (1994). Estimation of the effective self-diffusion tensor from the nmr spin echo. *J Magn Reson B*, 103(3):247–54.
- Beal, M. J. (2003). *Variational Algorithms for Approximate Bayesian Inference*. PhD thesis, University of London.
- Beal, M. J., Jojic, N., and Attias, H. (2003). A graphical model for audiovisual object tracking. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25:828–836.
- Behrens, T. E., Woolrich, M. W., Jenkinson, M., Johansen-Berg, H., Nunes, R. G., Clare, S., Matthews, P. M., Brady, J. M., and Smith, S. M. (2003). Characterization and propagation of uncertainty in diffusion-weighted mr imaging. *Magn Reson Med*, 50(5):1077–88.
- Bijma, F., de Munck, J. C., and Heethaar, R. M. (2005). The spatiotemporal MEG covariance matrix modeled as a sum of Kronecker products. *NeuroImage*, 27:402–415.
- Biswal, B., Yetkin, F. Z., Haughton, V. M., and Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med*, 34:537–541.
- Blagoev, K. B., Mihaila, B., Travis, B. J., Alexandrov, L. B., Bishop, A. R., Ranken, D., Posse, S., Gasparovic, C., Mayer, A., Aine, C. J., Ulbert, I., Morita, M., Müller, W., Connor, J., and Halgren, E. (2007). Modelling the magnetic signature of neuronal tissue. *NeuroImage*, 37:137–148.
- Braitenberg, V. and Schuz, A., editors (1991). *Anatomy of the cortex: Statistics and Geometry*. Springer-Verlag, Berlin.
- Brookes, M. J., Stevenson, C. M., Barnes, G. R., Hillebrand, A., Simpson, M. I. G., Francis, S. T., and Morris, P. G. (2007). Beamformer reconstruction of correlated sources using a modified source model. *NeuroImage*, 34:1454–1465.
- Brugge, J. F. and Merzenich, M. M. (1973). Responses of neurons in auditory cortex of the macaque monkey to monaural and binaural stimulation. *J Neurophysiol*, 36(6):1138–58.

- Budd, T. W., Barry, R. J., Gordon, E., Rennie, C., and Michie, P. T. (1998). Decrement of the n1 auditory event-related potential with stimulus repetition: habituation vs. refractoriness. *Int J Psychophysiol*, 31(1):51–68.
- Bunch, P., Hamilton, J., Sanderson, G., and Simmons, A. (1978). A free response approach to the measurement and characterization of radiographic-observer performance. *J. App. Photo. Eng.*, 4:166–172.
- Burton, H., Fabri, M., and Alloway, K. (1995). Cortical areas within the lateral sulcus connected to cutaneous representations in areas 3b and 1: a revised interpretation of the second somatosensory area in macaque monkeys. *J Comp Neurol*, 355(4):539–62.
- Cardoso, J.-F. (1998). Blind signal separation: Statistical principles. In *Proc. IEEE*, volume 86, pages 2009–2025.
- Catani, M., Howard, R. J., Pajevic, S., and Jones, D. K. (2002). Virtual in vivo interactive dissection of white matter fasciculi in the human brain. *Neuroimage*, 17(1):77–94.
- Celesia, G. G. (1976). Organization of auditory cortical areas in man. *Brain*, 99(3):403–14. 0006-8950 (Print) Journal Article.
- Ciccarelli, O., Parker, G. J., Toosy, A. T., Wheeler-Kingshott, C. A., Barker, G. J., Boulby, P. A., Miller, D. H., and Thompson, A. J. (2003). From diffusion tractography to quantitative white matter tract measures: a reproducibility study. *Neuroimage*, 18(2):348–59.
- Crone, N. E., Sinai, A., and Korzeniewska, A. (2006). High-frequency gamma oscillations and human brain mapping with electrocorticography. *Prog Brain Res*, 159:275–295.
- Croxson, P. L., Johansen-Berg, H., Behrens, T. E., Robson, M. D., Pinsk, M. A., Gross, C. G., Richter, W., Richter, M. C., Kastner, S., and Rushworth, M. F. (2005). Quantitative investigation of connections of the prefrontal cortex in the human and macaque using probabilistic diffusion tractography. *J Neurosci*, 25(39):8854–66.
- Dalal, S. S., Guggisberg, A. G., Edwards, E., Sekihara, K., Findlay, A. M., Canolty, R. T., Berger, M. S., Knight, R. T., Barbaro, N. M., Kirsch, H. E., and Nagarajan, S. S. (2007). Five-dimensional neuroimaging: Localization of the time-frequency dynamics of cortical activity. *submitted*.
- Dalal, S. S., Sekihara, K., and Nagarajan, S. S. (2006). Modified beamformers for coherent source region suppression. *IEEE Trans. Biomed. Eng.*, 53:1357–1363.

- Dalal, S. S., Zumer, J. M., Agrawal, V., Hild, K. E., Sekihara, K., and Nagarajan, S. S. (2004). NUTMEG: A neuromagnetic source reconstruction toolbox. *Neurol. Clin. Neurophysiol.*, page 52.
- Dale, A. M., Liu, A. K., Fischl, B. R., Buckner, R. L., Belliveau, J. W., Lewine, J. D., and Halgren, E. (2000). Dynamic statistical parametric mapping: Combining fMRI and MEG for high-resolution imaging of cortical activity. *Neuron*, 26:55–67.
- Darian-Smith, C. and Darian-Smith, I. (1993). Thalamic projections to areas 3a, 3b, and 4 in the sensorimotor cortex of the mature and infant macaque monkey. *J Comp Neurol*, 335(2):173–99.
- Darian-Smith, C., Darian-Smith, I., Burman, K., and Ratcliffe, N. (1993). Ipsilateral cortical projections to areas 3a, 3b, and 4 in the macaque monkey. *J Comp Neurol*, 335(2):200–13.
- Darvas, F., Pantazis, D., Kucukaltun-Yildirim, E., and Leahy, R. M. (2004). Mapping human brain function with MEG and EEG: methods and validation. *NeuroImage*, 23 Suppl 1:S289–S299.
- Darvas, F., Rautiainen, M., Pantazis, D., Baillet, S., Benali, H., Mosher, J. C., Garnero, L., and Leahy, R. M. (2005). Investigations of dipole localization accuracy in MEG using the bootstrap. *Neuroimage*, 25:355–368.
- Dassios, G. and Kariotou, F. (2004). On the exterior magnetic field and silent sources in magnetoencephalography. *Abstract and Applied Analysis*, 2004(4):307–314. doi:10.1155/S1085337504306032.
- Dauguet, J., Peled, S., Berezovskii, V., Delzescaux, T., Warfield, S. K., Born, R., and Westin, C. F. (2007). Comparison of fiber tracts derived from in-vivo dti tractography with 3d histological neural tract tracer reconstruction on a macaque brain. *Neuroimage*.
- de Munck, J. C., Bijma, F., Gaura, P., Sieluzycski, C. A., Branco, M. I., and 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.
- Disbrow, E. A., Hinkley, L. B., and Roberts, T. P. (2003). Ipsilateral representation of oral structures in human anterior parietal somatosensory cortex and integration of inputs across the midline. *J Comp Neurol*, 467(4):487–95.
- Dogandzic, A. and Nehorai, A. (2000). Estimating evoked dipole responses in unknown spatially correlated noise with EEG/MEG arrays. *IEEE Trans. Sig. Proc.*, pages 13–25.

- Forss, N., Salmelin, R., and Hari, R. (1994). Comparison of somatosensory evoked fields to airpuff and electric stimuli. *Electroencephalogr Clin Neurophysiol*, 92(6):510–7.
- Foucher, J. R., Otzenberger, H., and Gounot, D. (2003). The BOLD response and the gamma oscillations respond differently than evoked potentials: an interleaved EEG-fMRI study. *BMC Neurosci*, 4:22–22.
- Friston, K. J., Harrison, L., and Penny, W. (2003). Dynamic causal modelling. *Neuroimage*, 19:1273–1302.
- Gaetz, W. C. and Cheyne, D. O. (2003). Localization of human somatosensory cortex using spatially filtered magnetoencephalography. *Neurosci Lett*, 340:161–164.
- Gardner, E. P., Hamalainen, H. A., Warren, S., Davis, J., and Young, W. (1984). Somatosensory evoked potentials (seps) and cortical single unit responses elicited by mechanical tactile stimuli in awake monkeys. *Electroencephalogr Clin Neurophysiol*, 58(6):537–52.
- Gelman, A. and Rubin, D. B. (1996). Markov chain Monte Carlo methods in biostatistics. *Stat. Meth. Med. Res.*, 5:339–355.
- Ghahramani, Z. and Beal, M. (2001). Graphical models and variational methods. In Oppen, M. and Saad, D., editors, *Advanced Mean Field Methods — Theory and Practice*. MIT Press.
- Godey, B., Schwartz, D., de Graaf, J. B., Chauvel, P., and Liegeois-Chauvel, C. (2001). Neuromagnetic source localization of auditory evoked fields and intracerebral evoked potentials: a comparison of data in the same patients. *Clin Neurophysiol*, 112(10):1850–9.
- Gross, J., Kujala, J., Hamalainen, M., Timmermann, L., Schnitzler, A., and 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.
- Guye, M., Parker, G. J., Symms, M., Boulby, P., Wheeler-Kingshott, C. A., Salek-Haddadi, A., Barker, G. J., and Duncan, J. S. (2003). Combined functional mri and tractography to demonstrate the connectivity of the human primary motor cortex in vivo. *Neuroimage*, 19(4):1349–60.
- Hagmann, P., Thiran, J. P., Jonasson, L., Vandergheynst, P., Clarke, S., Maeder, P., and Meuli, R. (2003). Dti mapping of human brain connectivity: statistical fibre tracking and virtual dissection. *Neuroimage*, 19(3):545–54.
- Halgren, E. (2004). How can intracranial recordings assist meg source localization? *Neurol Clin Neurophysiol*, 2004:86.

- Hamalainen, M., Hari, R., Ilmoniemi, R. J., Knuutila, J., and Lounasmaa, O. V. (1993). Magnetoencephalography-theory, instrumentation, and applications to noninvasive studies of the working human brain. *Rev. Mod. Phys.*, 65:413–497.
- Hari, R., Hamalainen, M., Kaukoranta, E., Reinikainen, K., and Teszner, D. (1983). Neuromagnetic responses from the second somatosensory cortex in man. *Acta Neurol Scand*, 68(4):207–12.
- Hari, R., Kaila, K., Katila, T., Tuomisto, T., and Varpula, T. (1982). Interstimulus interval dependence of the auditory vertex response and its magnetic counterpart: implications for their neural generation. *Electroencephalogr Clin Neurophysiol*, 54(5):561–9.
- Hari, R., Reinikainen, K., Kaukoranta, E., Hamalainen, M., Ilmoniemi, R., Penttinen, A., Salminen, J., and Teszner, D. (1984). Somatosensory evoked cerebral magnetic fields from si and sii in man. *Electroencephalogr Clin Neurophysiol*, 57(3):254–63.
- Heiervang, E., Behrens, T. E., Mackay, C. E., Robson, M. D., and Johansen-Berg, H. (2006). Between session reproducibility and between subject variability of diffusion mr and tractography measures. *Neuroimage*, 33(3):867–77.
- Hofer, S. and Frahm, J. (2006). Topography of the human corpus callosum revisited—comprehensive fiber tractography using diffusion tensor magnetic resonance imaging. *Neuroimage*, 32(3):989–94.
- Hoogenboom, N., Schoffelen, J.-M., Oostenveld, R., Parkes, L. M., and Fries, P. (2006). Localizing human visual gamma-band activity in frequency, time and space. *NeuroImage*, 29:764–773.
- Hoshiyama, M., Kakigi, R., Koyama, S., Kitamura, Y., Shimojo, M., and Watanabe, S. (1996). Somatosensory evoked magnetic fields following stimulation of the lip in humans. *Electroencephalogr Clin Neurophysiol*, 100(2):96–104.
- Huffman, K. J. and Krubitzer, L. (2001). Area 3a: topographic organization and cortical connections in marmoset monkeys. *Cereb Cortex*, 11(9):849–67.
- Huizenga, H. M., de Munck, J. C., Waldorp, L. J., and 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.
- Hyvarinen, A., Karhunen, J., and Oja, E., editors (2001). *Independent Component Analysis*. John Wiley and Sons, Inc., New York, NY.
- Im, C.-H. and Lee, S. Y. (2006). A technique to consider mismatches between fMRI and EEG/MEG sources for fMRI-constrained EEG/MEG source imaging: a preliminary simulation study. *Phys Med Biol*, 51:6005–6021.

- Jacobson, G. P., Ahmad, B. K., Moran, J., Newman, C. W., Tepley, N., and Wharton, J. (1991). Auditory evoked cortical magnetic field (m100-m200) measurements in tinnitus and normal groups. *Hear Res*, 56(1-2):44–52.
- Jerbi, K., Mosher, J. C., Baillet, S., and Leahy, R. M. (2002). On MEG forward modelling using multipolar expansions. *Phys. Med. Biol.*, 47:523–555.
- Jones, D. K. (2004). The effect of gradient sampling schemes on measures derived from diffusion tensor mri: a monte carlo study. *Magn Reson Med*, 51(4):807–15.
- Jones, E. G. and Powell, T. P. (1969a). Connexions of the somatic sensory cortex of the rhesus monkey. i. ipsilateral cortical connexions. *Brain*, 92(3):477–502.
- Jones, E. G. and Powell, T. P. (1969b). Connexions of the somatic sensory cortex of the rhesus monkey. ii. contralateral cortical connexions. *Brain*, 92(4):717–30.
- Jordan, M., Ghaharamani, Z., Jaakkola, T., and Saul, L. (1999). An introduction to variational methods for graphical models. In Jordan, M., editor, *Learning in Graphical Models*. MIT Press, Cambridge, MA, USA.
- Jordan, M. I., editor (1999). *Learning in Graphical Models*. MIT Press, Cambridge, MA, USA.
- Jun, S. C., George, J. S., Paré-Blagoev, J., Plis, S. M., Ranken, D. M., Schmidt, D. M., and 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., and Wood, C. C. (2006). Improving source detection and separation in a spatiotemporal Bayesian inference dipole analysis. *Phys. Med. Biol.*, 51:2395–2414.
- Jung, T. P., Makeig, S., Humphries, C., Lee, T. W., McKeown, M. J., Iragui, V., and Sejnowski, T. J. (2000). Removing electroencephalographic artifacts by blind source separation. *Psychophysiology*, 37:163–178.
- Kaas, J. H. and Hackett, T. A. (1998). Subdivisions of auditory cortex and levels of processing in primates. *Audiol Neurootol*, 3(2-3):73–85.
- Kaas, J. H., Nelson, R. J., Sur, M., Lin, C. S., and Merzenich, M. M. (1979). Multiple representations of the body within the primary somatosensory cortex of primates. *Science*, 204(4392):521–3.
- Killackey, H. P., Gould, H. J., r, Cusick, C. G., Pons, T. P., and Kaas, J. H. (1983). The relation of corpus callosum connections to architectonic fields and body surface maps in sensorimotor cortex of new and old world monkeys. *J Comp Neurol*, 219(4):384–419.

- Kilner, J. M., Mattout, J., Henson, R., and Friston, K. J. (2005). Hemodynamic correlates of EEG: a heuristic. *Neuroimage*, 28:280–286.
- Kim, D. S., Kim, M., Ronen, I., Formisano, E., Kim, K. H., Ugurbil, K., Mori, S., and Goebel, R. (2003). In vivo mapping of functional domains and axonal connectivity in cat visual cortex using magnetic resonance imaging. *Magn Reson Imaging*, 21(10):1131–40.
- Krubitzer, L., Clarey, J., Tweedale, R., Elston, G., and Calford, M. (1995). A redefinition of somatosensory areas in the lateral sulcus of macaque monkeys. *J Neurosci*, 15(5 Pt 2):3821–39.
- Krubitzer, L. A. and Kaas, J. H. (1990). The organization and connections of somatosensory cortex in marmosets. *J Neurosci*, 10(3):952–74.
- Küçükaltun-Yildirim, E., Pantazis, D., and Leahy, R. M. (2006). Task-based comparison of inverse methods in magnetoencephalography. *IEEE Trans. Biomed. Eng.*, 53:1783–1793.
- Laufs, H., Kleinschmidt, A., Beyerle, A., Eger, E., Salek-Haddadi, A., Preibisch, C., and Krakow, K. (2003). EEG-correlated fMRI of human alpha activity. *Neuroimage*, 19:1463–1476.
- Leahy, R. M., Mosher, J. C., Spencer, M. E., Huang, M. X., and Lewine, J. D. (1998). A study of dipole localization accuracy for MEG and EEG using a human skull phantom. *Electroencephalogr Clin Neurophysiol*, 107:159–173.
- Liegeois-Chauvel, C., Musolino, A., Badier, J. M., Marquis, P., and Chauvel, P. (1994). Evoked potentials recorded from the auditory cortex in man: evaluation and topography of the middle latency components. *Electroencephalogr Clin Neurophysiol*, 92(3):204–14.
- Limpiti, T., Van Veen, B. D., and Wakai, R. T. (2006). Cortical patch basis model for spatially extended neural activity. *IEEE Trans. Biomed. Eng.*, 53:1740–1754.
- Liu, A. K., Belliveau, J. W., and Dale, A. M. (1998). Spatiotemporal imaging of human brain activity using functional MRI constrained magnetoencephalography data: Monte Carlo simulations. *Proc Natl Acad Sci U S A*, 95:8945–8950.
- Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., and Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 412:150–157.
- Luethke, L. E., Krubitzer, L. A., and Kaas, J. H. (1988). Cortical connections of electrophysiologically and architectonically defined subdivisions of auditory cortex in squirrels. *J Comp Neurol*, 268(2):181–203.

- Mattout, J., Péligrini-Issac, M., Bellio, A., Daunizeau, J., and Benali, H. (2003). Localization Estimation Algorithm (LEA): a supervised prior-based approach for solving the EEG/MEG inverse problem. *Inf. Process. Med. Imaging*, 18:536–547.
- Mattout, J., Phillips, C., Penny, W. D., Rugg, M. D., and Friston, K. J. (2006). MEG source localization under multiple constraints: an extended Bayesian framework. *NeuroImage*, 30:753–767.
- McCarthy, G. and Allison, T. (1995). Trigeminal evoked potentials in somatosensory cortex of the macaca mulatta. *J Neurosurg*, 82(6):1015–20.
- McCarthy, G., Wood, C. C., and Allison, T. (1991). Cortical somatosensory evoked potentials. i. recordings in the monkey macaca fascicularis. *J Neurophysiol*, 66(1):53–63.
- Moran, R. J., Kiebel, S. J., Stephan, K. E., Reilly, R. B., Daunizeau, J., and Friston, K. J. (2007). A neural mass model of spectral responses in electrophysiology. *NeuroImage*, 37:706–720.
- Morel, A., Garraghty, P. E., and Kaas, J. H. (1993). Tonotopic organization, architectonic fields, and connections of auditory cortex in macaque monkeys. *J Comp Neurol*, 335(3):437–59.
- Mori, S., Crain, B. J., Chacko, V. P., and van Zijl, P. C. (1999). Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Ann Neurol*, 45(2):265–9. 0364-5134 (Print) Journal Article Research Support, Non-U.S. Gov't.
- Mosher, J. C. and Leahy, R. M. (1998). Recursive MUSIC: a framework for EEG and MEG source localization. *IEEE Trans. Biomed. Eng.*, 45:1342–1354.
- Nagarajan, S., Attias, H., Hild, K., and Sekihara, K. (2005). Stimulus evoked independent factor analysis of MEG data with large background activity. In *Adv. Neur. Info. Proc. Sys.* 18.
- Nagarajan, S. S., Attias, H. T., Hild, K. E., and Sekihara, K. (2006). A graphical model for estimating stimulus-evoked brain responses from magnetoencephalography data with large background brain activity. *NeuroImage*, 30:400–416.
- Nagarajan, S. S., Attias, H. T., Hild II, K. E., and Sekihara, K. (2007). A probabilistic algorithm for robust interference suppression in bioelectromagnetic sensor data. *Stat. in Med.*, *accepted*.

- Nakahara, H., Nakasato, N., Kanno, A., Murayama, S., Hatanaka, K., Itoh, H., and Yoshimoto, T. (2004). Somatosensory-evoked fields for gingiva, lip, and tongue. *J Dent Res*, 83(4):307–11.
- Nangini, C., Tam, F., and Graham, S. J. (2007). A novel method for integrating MEG and BOLD fMRI signals with the linear convolution model in human primary somatosensory cortex. *Hum Brain Mapp*.
- Neal, R. M., editor (1996). *Bayesian Learning for Neural Networks*. Springer, New York.
- Okada, Y., Lähteenmäki, A., and Xu, C. (1999). Comparison of MEG and EEG on the basis of somatic evoked responses elicited by stimulation of the snout in the juvenile swine. *Clin Neurophysiol*, 110:214–229.
- Pantev, C., Hoke, M., Lehnertz, K., Lutkenhoner, B., Fahrendorf, G., and Stoher, U. (1990). Identification of sources of brain neuronal activity with high spatiotemporal resolution through combination of neuromagnetic source localization (nmsl) and magnetic resonance imaging (mri). *Electroencephalogr Clin Neurophysiol*, 75(3):173–84.
- Parker, G. J., Haroon, H. A., and Wheeler-Kingshott, C. A. (2003). A framework for a streamline-based probabilistic index of connectivity (pico) using a structural interpretation of mri diffusion measurements. *J Magn Reson Imaging*, 18(2):242–54.
- Parker, G. J., Stephan, K. E., Barker, G. J., Rowe, J. B., MacManus, D. G., Wheeler-Kingshott, C. A., Ciccarelli, O., Passingham, R. E., Spinks, R. L., Lemon, R. N., and Turner, R. (2002). Initial demonstration of in vivo tracing of axonal projections in the macaque brain and comparison with the human brain using diffusion tensor imaging and fast marching tractography. *Neuroimage*, 15(4):797–809.
- Pascual-Marqui, R. D. (2002). Standardized low-resolution brain electromagnetic tomography (sLORETA): technical details. *Meth. Find. Exp. Clin. Pharmacol.*, 24 Suppl D:5–12.
- Penfield, W. and Boldrey, E. (1937). Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain*, 60:389–443.
- Phillips, C., Mattout, J., Rugg, M. D., Maquet, P., and Friston, K. J. (2005). An empirical Bayesian solution to the source reconstruction problem in EEG. *NeuroImage*, 24:997–991011.
- Poonawalla, A. H. and Zhou, X. J. (2004). Analytical error propagation in diffusion anisotropy calculations. *J Magn Reson Imaging*, 19(4):489–98.

- Pulvermüller, F., Lutzenberger, W., Preissl, H., and Birbaumer, N. (1995). Spectral responses in the gamma-band: Physiological signs of higher cognitive processes? *Neuroreport*, 6:2059–2064.
- Ramnani, N., Behrens, T. E., Johansen-Berg, H., Richter, M. C., Pinski, M. A., Andersson, J. L., Rudebeck, P., Ciccarelli, O., Richter, W., Thompson, A. J., Gross, C. G., Robson, M. D., Kastner, S., and Matthews, P. M. (2005). The evolution of prefrontal inputs to the cortico-pontine system: Diffusion imaging evidence from macaque monkeys and humans. *Cereb Cortex*.
- Reite, M., Teale, P., Zimmerman, J., Davis, K., and Whalen, J. (1988). Source location of a 50 msec latency auditory evoked field component. *Electroencephalogr Clin Neurophysiol*, 70(6):490–8.
- Robinson, P. A., Rennie, C. J., Rowe, D. L., O'Connor, S. C., and Gordon, E. (2005). Multiscale brain modelling. *Philos Trans R Soc Lond B Biol Sci*, 360:1043–1050.
- Robinson, S. and Vrba, J. (1999). Functional neuroimaging by synthetic aperture magnetometry (SAM). *Recent Advances in Biomagnetism*.
- Ronen, I., Kim, K. H., Garwood, M., Ugurbil, K., and Kim, D. S. (2003). Conventional dti vs. slow and fast diffusion tensors in cat visual cortex. *Magn Reson Med*, 49(5):785–90.
- Saleem, K. S., Pauls, J. M., Augath, M., Trinath, T., Prause, B. A., Hashikawa, T., and Logothetis, N. K. (2002). Magnetic resonance imaging of neuronal connections in the macaque monkey. *Neuron*, 34(5):685–700.
- Sarvas, J. (1987). Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem. *Phys. Med. Biol.*, 32:11–22.
- Sato, M.-a., Yoshioka, T., Kajihara, S., Toyama, K., Goda, N., Doya, K., and Kawato, M. (2004). Hierarchical Bayesian estimation for MEG inverse problem. *NeuroImage*, 23:806–826.
- Schmidt, D. M., George, J. S., and Wood, C. C. (1999). Bayesian inference applied to the electromagnetic inverse problem. *Human Brain Mapping*, 7:195–212.
- Schulz, M., Chau, W., Graham, S. J., McIntosh, A. R., Ross, B., Ishii, R., and Pantev, C. (2004). An integrative MEG-fMRI study of the primary somatosensory cortex using cross-modal correspondence analysis. *Neuroimage*, 22:120–133.
- Sekihara, K., Nagarajan, S., Poeppel, D., and Marantz, A. (2002a). Performance of an MEG adaptive-beamformer technique in the presence of correlated neural activities: Effects on signal intensity and time-course estimates. *IEEE Trans. Biomed. Eng.*, 49:1534–1546.

- Sekihara, K., Nagarajan, S. S., Poeppel, D., Marantz, A., and Miyashita, Y. (2001). Reconstructing spatio-temporal activities of neural sources using an MEG vector beamformer technique. *IEEE Trans. Biomed. Eng.*, 48:760–771.
- Sekihara, K., Nagarajan, S. S., Poeppel, D., Marantz, A., and Miyashita, Y. (2002b). Application of an MEG eigenspace beamformer to reconstructing spatio-temporal activities of neural sources. *Human Brain Mapping*, 15:199–215.
- Sekihara, K., Sahani, M., and Nagarajan, S. (2005). Localization bias and spatial resolution of adaptive and non-adaptive spatial filters for MEG source reconstruction. *NeuroImage*, 25:1056–1067.
- Sekihara, K., Sahani, M., and Nagarajan, S. S. (2004). Bootstrap-based statistical thresholding for MEG source reconstruction images. *Conf Proc IEEE Eng Med Biol Soc*, 2:1018–1021.
- Snodgrass, J. G. and Corwin, J. (1988). Pragmatics of measuring recognition memory: applications to dementia and amnesia. *J. Exp. Psych.: General*, 117:34–50.
- Sotero, R. C., Trujillo-Barreto, N. J., Iturria-Medina, Y., Carbonell, F., and Jimenez, J. C. (2007). Realistically coupled neural mass models can generate EEG rhythms. *Neural Comput*, 19:478–512.
- Sun, F. T., Miller, L. M., and D’Esposito, M. (2005). Measuring temporal dynamics of functional networks using phase spectrum of fMRI data. *Neuroimage*, 28:227–237.
- Sutherling, W. W., Crandall, P. H., Darcey, T. M., Becker, D. P., Levesque, M. F., and Barth, D. S. (1988). The magnetic and electric fields agree with intracranial localizations of somatosensory cortex. *Neurology*, 38(11):1705–14.
- Teale, P., Delmore, J., Simon, J., and Reite, M. (1994). Magnetic auditory source imaging in macaque monkey. *Brain Res Bull*, 33(5):615–20.
- Tuch, D. S. (2004). Q-ball imaging. *Magn Reson Med*, 52(6):1358–72.
- Tuch, D. S., Reese, T. G., Wiegell, M. R., Makris, N., Belliveau, J. W., and Wedeen, V. J. (2002). High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magn Reson Med*, 48(4):577–82.
- Tuch, D. S., Wedeen, V. J., Dale, A. M., George, J. S., and Belliveau, J. W. (1999). Conductivity mapping of biological tissue using diffusion MRI. *Ann N Y Acad Sci*, 888:314–316.
- Tuch, D. S., Wisco, J. J., Khachaturian, M. H., Ekstrom, L. B., Kotter, R., and Vanduffel, W. (2005). Q-ball imaging of macaque white matter architecture. *Philos Trans R Soc Lond B Biol Sci*, 360(1457):869–79.

- van't Ent, D., de Munck, J. C., and Kaas, A. L. (2001). A fast method to derive realistic BEM models for E/MEG source reconstruction. *IEEE Trans Biomed Eng*, 48:1434–1443.
- Vogt, B. A. and Pandya, D. N. (1978). Cortico-cortical connections of somatic sensory cortex (areas 3, 1 and 2) in the rhesus monkey. *J Comp Neurol*, 177(2):179–91.
- Waldorp, L. J., Huizenga, H. M., Grasman, R. P. P. P., Böcker, K. B. E., de Munck, J. C., and Molenaar, P. C. M. (2002). Model selection in electromagnetic source analysis with an application to VEFs. *IEEE Trans. Biomed. Eng.*, 49:1121–1129.
- Wedeen, V. J., Hagmann, P., Tseng, W. Y., Reese, T. G., and Weisskoff, R. M. (2005). Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magn Reson Med*, 54(6):1377–86.
- Winterer, G., Carver, F. W., Musso, F., Mattay, V., Weinberger, D. R., and Coppola, R. (2006). Complex relationship between BOLD signal and synchronization/desynchronization of human brain MEG oscillations. *Hum Brain Mapp*.
- Wipf, D., Ramirez, R., Palmer, J., Makeig, S., and Rao, B. (2007). Analysis of empirical Bayesian methods for neuroelectromagnetic source localization. In Schölkopf, B., Platt, J., and Hoffman, T., editors, *Adv. Neur. Info. Proc. Sys. 19*. MIT Press, Cambridge, MA.
- Wood, C. C., Spencer, D. D., Allison, T., McCarthy, G., Williamson, P. D., and Goff, W. R. (1988). Localization of human sensorimotor cortex during surgery by cortical surface recording of somatosensory evoked potentials. *J Neurosurg*, 68(1):99–111.
- Yetik, I. S., Nehorai, A., Muravchik, C. H., and Haueisen, J. (2005). Line-source modeling and estimation with magnetoencephalography. *IEEE Trans. Biomed. Eng.*, 52:839–851.
- Zhu, Z., Zumer, J. M., Lowenthal, M. E., Padberg, J., Recanzone, G. H., Krubitzer, L. A., Nagarajan, S. S., and Disbrow, E. A. (2007). The relationship between magnetic and electrophysiological signals from monkey somatosensory cortex. *Submitted*.
- Zumer, J., Attias, H., Sekihara, K., and Nagarajan, S. (2006). Two probabilistic algorithms for MEG/EEG source reconstruction. *IEEE ISBI: From Nano to Macro*.
- Zumer, J., Attias, H., Sekihara, K., and Nagarajan, S. (2007a). A probabilistic algorithm integrating source localization and noise suppression for MEG and EEG data. In B. Schölkopf and J. Platt and T. Hoffman, editor, *Adv. Neur. Info. Proc. Sys. 19*. MIT Press, Cambridge, MA.

Zumer, J. M., Attias, H. T., Sekihara, K., and Nagarajan, S. S. (2007b). A probabilistic algorithm integrating source localization and noise suppression for MEG and EEG data. *NeuroImage*, 37:102–115.

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.

Johannes Zimmer
Author Signature

11/07/07
Date