Revealing Hidden Neural Processes
Signal Processing with MEG in the Human Brain

MERIT 2008

Kevin Kahn
Graduate Student: Nai Ding
Mentor: Dr. Jonathan Simon

Abstract

When analyzing the neural responses to auditory stimuli using magnetoencephalography (MEG), two peaks with high power not associated with the evoked response were seen. Traditional analysis could not be used to examine these peaks because of the induced nature of this activity. New techniques using phase difference between channels within each trial and their coherences were developed and verified the existence of a dipole after alterations of coherence and power thresholds. After removing the evoked response, the neural source was localized to the left frontal region, implicating higher order processing.

Introduction

Positron emission tomography (PET) and magnetic resonance imaging (MRI) have been widely used to measure brain activity in a clinical setting. However, both these technologies offer poor temporal resolution and measure neural activity in a roundabout way. PET and MRI actually only measure neuron cell metabolism through glucose usage and hemoglobin-oxygen association respectively. Higher neural activity is then inferred from higher cell metabolism. These methods offer clear 3D images, but don't give a clear idea of what is happening in the brain when looking at smaller time scales.

Electroencephalography (EEG) and magnetoencephalography (MEG) are powerful alternatives in measuring brain activity. By measuring electric potentials, and the resulting magnetic fields from those potentials, EEG and MEG can be used to analyze neural activity at a much higher temporal resolution. MEG is advantageous since the magnetic fields are not distorted by human tissue like electric fields in EEG are[2]. This study focuses only on the use of MEG in detecting neural activity.

MEG is particularly useful in detecting auditory responses because the orientation of the neural sources in auditory cortex is roughly parallel instead of perpendicular to the scalp. This orientation results in a magnetic dipole across the surface of a subject's head. In contrast, neural current running perpendicular to the subject's head would not show a dipole[2]. Thus, magnetic dipoles in MEG studies are important and imply a localizable neural source. Figure 1 shows an example of the dipoles seen during the M100 peak (peak found through MEG approximately 100 ms after an auditory stimulus[4]) of the auditory response. Two magnetic dipoles are visible here, one in the left hemisphere and one in the right hemisphere. The neural source of these dipoles localize to their respective left or right auditory cortex.
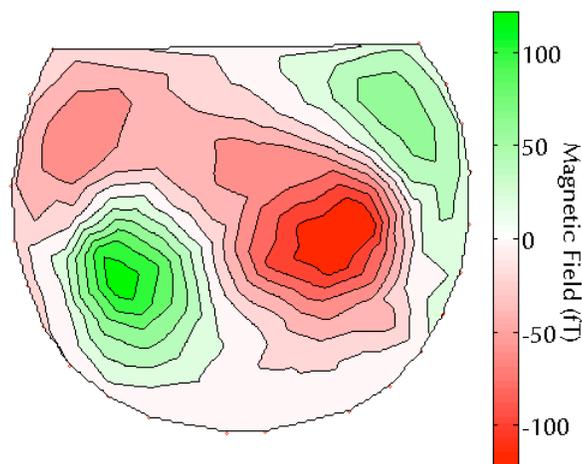
Fig 1. Head map of magnetic field of evoked M100 response

MEG has been used extensively in previous studies to analyze auditory responses because of the auditory cortex's orientation. These studies include analyzing latency in the M100 peak from averaged responses in the time domain in response to differing stimuli[4], phase coherence in the theta band (4-8 Hz) of the frequency domain as a means to determine sentence intelligibility[3], and induced versus evoked activity in the gamma band (30-50 Hz) of the frequency domain for visual object recognition[5].

Evoked activity is defined as any activity that is phase locked with the onset of a stimulus while induced activity is defined as anything that is not evoked[5]. The M100 peak in the first study is phase locked since the peak occurs consistently at some time around 100 ms after the stimulus onset. Activity shown through phase coherence in the second study is also phase locked, though the activity is not necessarily visible through just averaging over trials[3]. The first two studies listed above are therefore both studies of evoked responses.

In contrast, this study focuses on the detection and localization of induced activity using both the time and frequency domain. Induced activity can result from an increase in power in response to a stimulus, similar to what is found in evoked activity but with a variable time of response[5]. This variation in timing results in this response not being visible when data is averaged over time. Traditional techniques can therefore not be used since this type of activity will disappear. New techniques had to be developed when analyzing the induced activity from the auditory response.

Materials and Methods

All data was collected prior to my participation in the project. Data consisted of the magnetic field measured over 157 channels on the surface of the head during auditory experiments. Data was collected with two subjects, ten different stimuli, and seventy trials for each stimulus. The subject was asked to listen to numerous stimuli and press a button when a target stimulus was heard. The stimuli are listed in Table 1 (target stimulus not included since data was not analyzed)[7]:

Types of Auditory Stimuli
White noise amplitude modulated at 3.5 Hz
Pink noise ripple, modulated at 3.5 Hz, 0.2 cycles/octave
Pink noise
Chimera: Pink noise envelope with sentence fine structure
Chimera: Pink noise fine structure with sentence envelope
Sentence 1
Sentence 2
Pure tone
Chimera: Pure tone fine structure with sentence envelope
Pure tone modulated at 3.5 Hz

Table 1. List of auditory stimuli used for MEG data collection

Data Analysis

   The motivation for this neural signal processing research starts from the analysis of power (magnitude squared) in the frequency domain averaged over the seventy trials. The average power (P_{av}) of each channel at 3.5 Hz is plotted with contours across the head in Figure 2. This head map is interesting for two reasons. The first reason is that the levels of activity in the left and right hemisphere are drastically different from each other. The second reason, which is important in terms of this research, is that there are two strong peaks in the left hemisphere of the frontal lobe. This raises the possibility of the presence of a magnetic dipole in this area, which would in turn imply a compact neural source. This power head map in itself is not enough to confirm a dipole however, since the two peaks could have arisen from two sources or two sinks.

(Eq. 1)

$$P_{av} = \frac{\sum_{n=1}^{N} \overline{A}_n^2}{N}$$
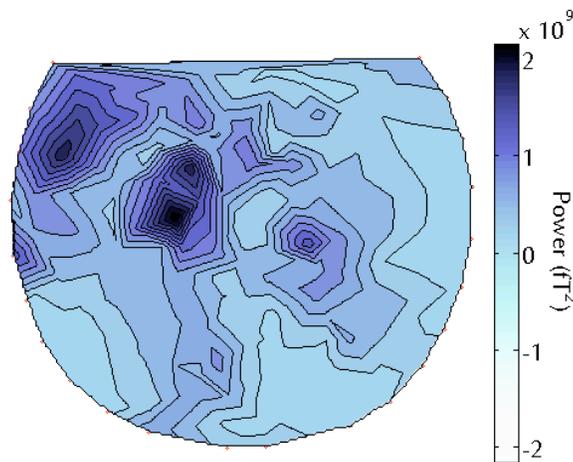
Fig 2. Head map of power at 3.5 Hz

   The following analysis techniques will only be explained in reference to the first subject's response to white noise amplitude modulated at 3.5 Hz for the moment. The results for each stimulus and subject were all qualitatively similar. The goal of this paper is to show the development of a set successful techniques used to analyze this newfound neural process and not to focus on the small differences found between subject and stimuli.

   The next step is to try to analyze this potential dipole using traditional methods. These methods include examining the signal in the frequency domain averaged over trials and the coherence (C_{phase}) of the phase in the frequency domain between trials. Averaging over trials is frequently used to increase signal to noise ratio for evoked

responses, but in this case, the two peaks vanish completely as shown in Figure 3. This head map resembles the evoked transient auditory response.

$$C_{phase} = \left(\frac{\sum_{n=1}^{N} \cos(\theta_n)}{N} \right)^2 + \left(\frac{\sum_{n=1}^{N} \sin(\theta_n)}{N} \right)^2$$
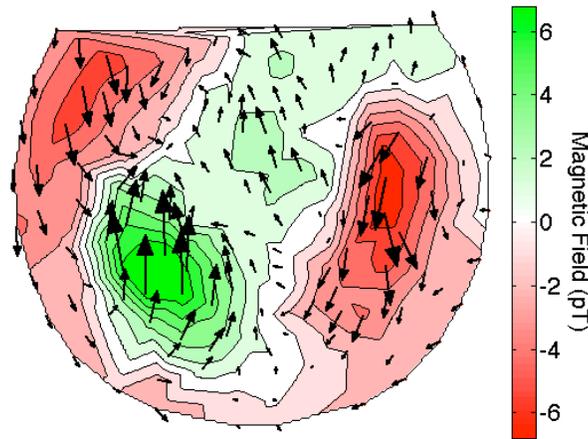
Fig 3. Magnetic field amplitude and phase at 3.5 Hz. Direction of arrow represents phase (Cartesian, right=0 degrees)

The phase coherence map in Figure 4 show that the two regions where the peaks were found are not coherent with respect to the stimulus onset. Phase coherence would not give proof of a dipole regardless of the head map, but it provides guidance and intuition in further data analysis. The observations from these two figures show the activity is induced and this poses a problem. Traditional methods require averaging over trials in order to remove noise from the signal. However, since the activity is induced, it is not phase locked and therefore will disappear with averaging, so these techniques will not work.
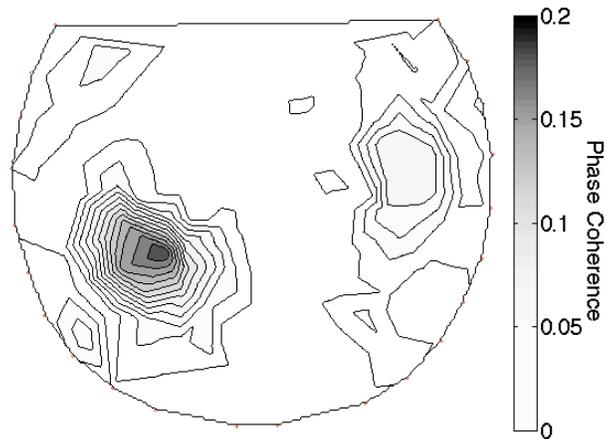
Fig 4. Head map of phase coherence at 3.5 Hz

Therefore, a different approach must be developed analyzing this potential dipole. The challenge is to develop a method that does not remove induced activity but still removes noise. A simple analysis technique is to multiply the fields of channels from different peaks of the possible dipole and to examine the distribution of numbers through a histogram. If the two channels are from a sink and a source from a dipole, then the product of the magnetic field should be negative. This simple technique was performed on the data after band passing it between 3 and 7 Hz. The histogram in Figure 5 shows there is a larger amount of data in the negative region as well as a larger magnitude in the actual negative numbers. The two regions are therefore out of phase and this gives some evidence that the peaks do relate to a dipole. However, this result is still not particularly conclusive since the two regions are not necessarily 180 degrees out of phase. Channels from opposing sides of the dipole ideally will be 180 degrees out of phase. Since the product of the channels is not always negative, either the peaks are not a dipole or there is some additive noise interfering with the channels. Unrelated neural activity might also offset the phase difference from exactly 180 degrees.
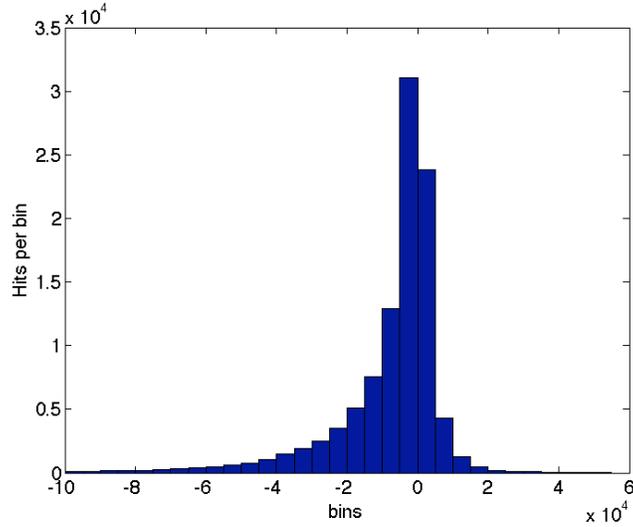
Fig 5. Histogram of product of channels in possible dipole.
Negative numbers show channels are out of phase.

Another, more robust approach, would be to specifically search for the presence of approximately 180 degree phase shift. To accomplish this, we can examine the phase differences between channels within each trial and the coherences of those phase differences. Analyzing patterns within each trial is critical here since induced responses can only be observed in single trials. The phase differences at 3.5 Hz are first calculated for each trial and then the differences themselves are averaged (circular average, θ_{avg}) over the seventy trials.

(Eq. 3)

$$\theta_{avg} = \tan^{-1}\left(\frac{\sum_{n=1}^{N}\sin(\theta_n)}{\sum_{n=1}^{N}\cos(\theta_n)}\right)$$

The following few figures connect channels that are consistently out of phase (within 10 degrees of 180) and consistently in phase (with 10 degrees of 0). In Figure 6, "consistently" is defined as phase differences that have a coherence of at least 0.2. This is initially a reasonable threshold since 0.2 seems to be a relatively high coherence looking back at the coherences shown in Figure 4. However, there are so many connections it is hard to distinguish anything intelligible from this figure. Any connections that might be occurring in the middle are obscured by the large number of lines running across the entire head map.
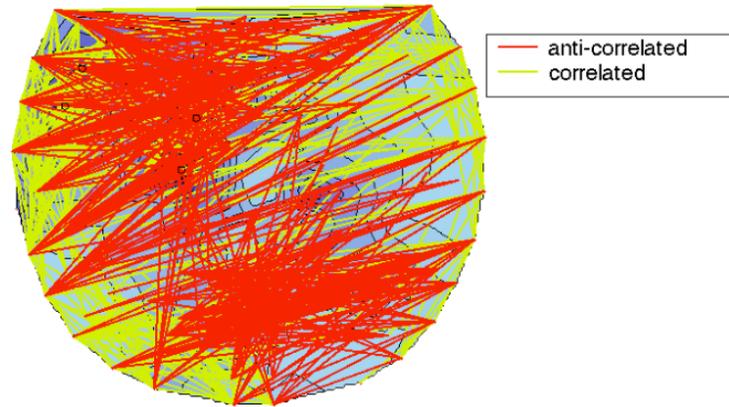
Fig 6. Channel correlations imposed on power head map. Phase differences must have been within 10 degrees of 0 or 180 degrees with a coherence of at least 0.2 in order to be connected.

In an effort to visually clean up this head map, the coherence threshold was raised to 0.35 for Figure 7. In this case, the out of phase connections are mainly between the leftmost and the rightmost channels. Therefore, there seems to be a magnetic dipole across the entire head. It is unlikely that the source of this is neural the large dipole size suggests a distant source. The differences in the trial concatenated waveforms between the leftmost and rightmost channels were low passed at 20 Hz and averaged to give Figure 8. The partially periodic spiking is occurring at roughly 1-2 Hz and resembles the QRS complex, which is the stereotypical spike from an electrocardiogram (ECG or EKG). Therefore, these high correlation connections across the head map are a cardiac artifact from the subject's heartbeat.
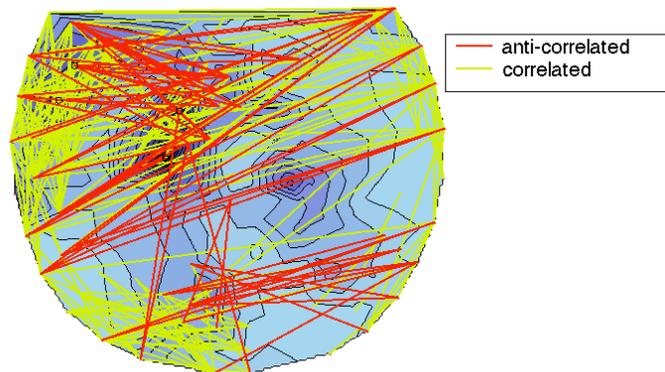
Fig 7. Channel correlations with a higher coherence threshold (0.35)
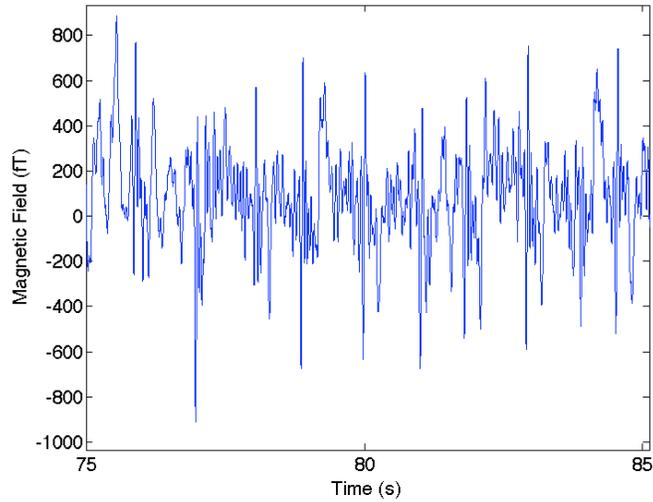
Fig 8. Average difference between leftmost and rightmost channels. The large peaks correspond to QRS complex of the heartbeat.

Independent component analysis (ICA)[6] was then used to separate out the heartbeat. The heart component was manually identified as the only signal with the signature peaks from the QRS complex and is shown in Figure 9. The original waveforms for each channel were then projected onto this heart signal and those projections were then removed from each channel signal. The magnitude of the projections for each channel is shown in a head map in Figure 10. We should expect a dipole across the entire head but obviously this is not the case. There seems to be some neural activity removed along with the heartbeat.
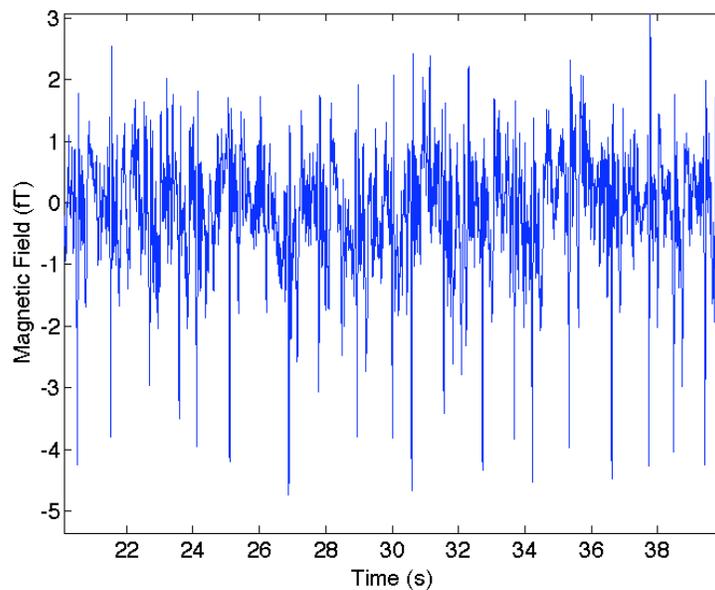
Fig 9. Heartbeat signal extracted through independent component analysis
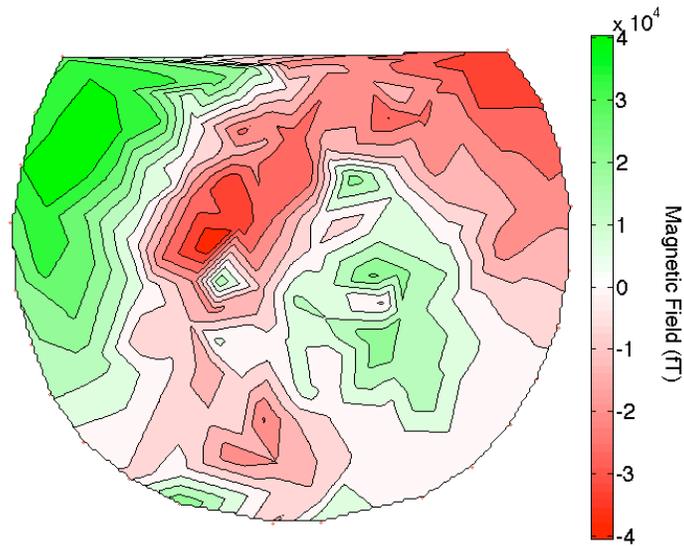
Fig 10. Magnitude of projection each channel signal onto
the heart beat signal

Afterwards, the connections were then redrawn for the signal with the heart beat removed in Figure 11. A lot of the connections crossing the entire head map were removed. However, it is likely that some neural connections were removed since Figure 9 looks like a mix of the expected heart dipole and some neural activity. Therefore, it was possible to remove much of the heartbeat, but not without removing some neural activity also.
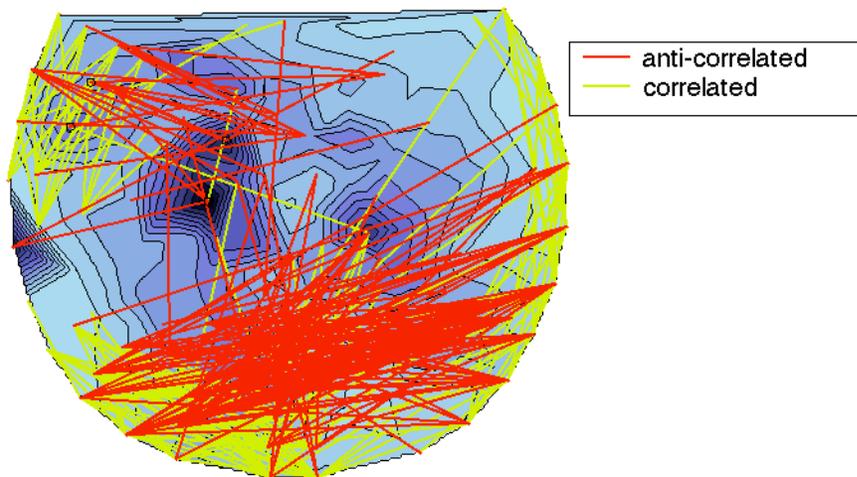
Fig 12. Channel correlations with a power threshold (1.5
times the mean of channel powers)

Another approach for the cleaning of the connections is to use of a power threshold. Connections for Figure 12 were only drawn if the power of the two channels being connected exceeded a certain power threshold (1.5 times the mean channel power in this case). Including a power threshold makes sure that the connection is based on data with a high signal to noise ratio. Finally, enough of the miscellaneous connections are removed to see the pattern relating to the original two peaks in power in Figure 2. Anti-correlation connections are shown between the two peaks while correlation connections are shown within the two peaks (among other areas). This exactly what is expected from a noisy dipole, so now it is finally safe to say that the two peaks in power are actually from a dipole. In the end, the heart beat removal was not necessary at all. The channels showing heartbeat were very coherent, but had a low power. Therefore, these connections were automatically disregarded with a power threshold.
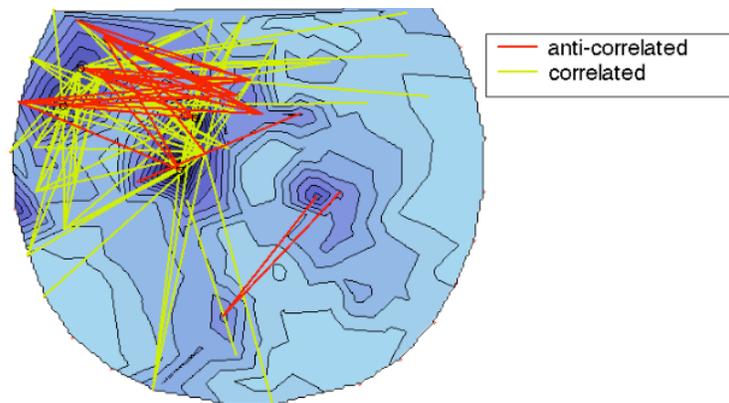
Fig 12. Channel correlations with a power threshold (1.5 times the mean of channel powers)

Now, the entire time series of data must be compressed into a single magnetic field head map for the localization algorithm to find a compact neural source. Numerous different methods of compression were tested. The data was first band passed between 3 and 7 Hz. The first method was to use principal component analysis (PCA) to decompose the 157 channels of data into 157 orthogonal waveforms. The coefficients of the most powerful waveforms were used to generate the compressed head map. The second method and third methods are essentially weighted averages. After band-passing the signal, the orientation of the dipole is seen switching constantly between source-sink and sink-source for a movie of head maps over a time span. This result is to be expected because of the removal of the dc component and leads to the second and third methods of compression.

The second method weights the head map at each time instance by either +1, or -1 depending on the orientation of the dipole. If the dipole is present and oriented in one direction, a weight of +1 is used, while if the orientation is in the opposite direction, a weight of -1 is used. If there is no dipole evident, then that head map is not counted as part of the average. The presence and orientation of the dipole are determined by the

signs of the four most powerful channels of the dipole from Figure 2. If the signs follow what is expected of a dipole the corresponding nonzero weight is selected. If the signs don't follow a dipole, then the head map isn't included in the average at all. The last method is very similar, except with a magnetic field threshold. In order for a nonzero weight to be assigned, the four channels must have the correct orientation along with sufficient magnitude in the field. This threshold eliminates head maps that don't really contribute toward the dipole.

The next several methods focused on removing the evoked response from the signal so only the induced activity can be seen. The first method found the evoked response with the magnetic fields averaged over the trials. The second method used de-noising source separation (DSS) to separate out the evoked response[1]. With these two methods, the projection of the original data onto the evoked response was then removed to leave a hopefully cleaner induced response. The last method used tried to avoid the evoked response completely. The evoked response to our example stimulus should be largest at 3.5 Hz since that is the modulation rate of the stimulus. Since the two peaks in power are still visible (though slightly weaker) at higher frequencies until about 9 Hz, it might be possible to avoid the evoked response if the signal is narrowly band passed between 4.5 and 5.5 Hz. The region around 7 Hz is avoided because of harmonics of 3.5 Hz. After each of these three methods were used, the respectively "cleaned" signals were run through the weighted average with magnetic field threshold technique to produce a final head map for each method. Figure 13 shows the final head map computed by each of the six different methods. Most of the head maps are qualitatively similar, but some localize better than others as will be seen.
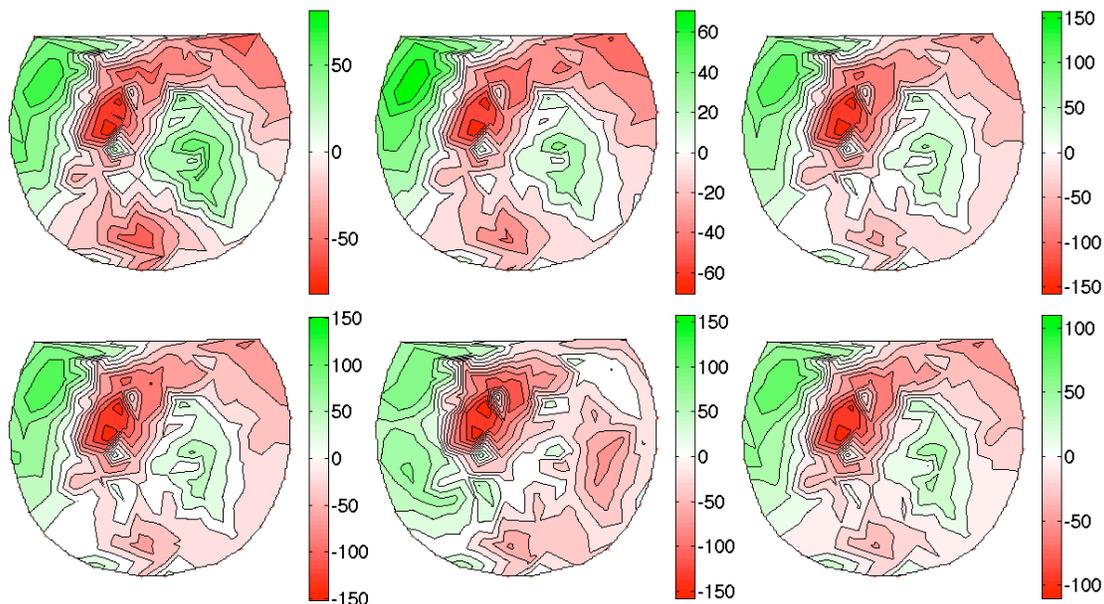
Fig 13. Head maps of magnetic field (fT) from different signal compression methods. Different methods shown in top row, from left to right: principle component analysis, weighted average without a power threshold, weighted average with a power threshold. Different methods of removing evoked response shown in bottom row, from left to right: removing average, removing evoked calculated by DSS, narrow band passing between 4.5 and 5.5 Hz.

The final step is to run the localization algorithm on the different compressed head maps. When running the algorithm, the channels used for localization must be selected. In this case, different channel combinations were tested. The different combinations are: all 157 channels, left hemisphere channels only, left hemisphere channels without channel 73 (the anomalous channel seen in all the head maps), and subjective selection based on dipole areas. The success of the localization is based on the correlation coefficient between the original head map and the theoretical head map generated by the localized dipole. The correlation coefficients for each compression method and group of channels are shown in Table 2.

Correlation Coefficients (%) for different channels and methods				
	Entire Head	Left Hem.	Left Hem. (-chan 73)	Dipole Only (subjective)
Principle component	53.15	72.96	74.82	83.47
Weighted average	-	-	-	-
Weight w/ threshold	69.51	83.6	85.13	88.79
Removing average	68.49	82.79	84.31	88.48
Removing DSS	77.43	84.84	86.49	90.15
Narrow band passed	68.28	84.95	86.52	88.95

Table 2. Correlation coefficients in percentages showing levels of success for localization using different channels and compression methods. A dash (-) means the source was localized outside of the head.

The correlation coefficients shown in Table 2 show that the dipole fits are generally good, but not in general reaching the desirable 90% level. This means that there is definitely a localizable neural source, but the head map needs to be further cleaned of noise (sensor, environmental, and biological) in order to get a precise localization. Using only the weighted average without a power threshold did not work at all. Even for the cases with passable correlation coefficients, the source localized was outside of the head. The principle component analysis did not work as well as the other methods in any of the channel selections. Removing the evoked response identified by DSS seemed to work the best. It gave the best localization when using the entire head map and was the only method that gave a correlation coefficient above 90% (though based on a subjective selection of channels).

Excluding the head map yielded from principle component analysis and weighted averaging without a threshold, all the other head maps gave qualitatively similar localizations. The neural source was found in the left hemisphere of the temporal lobe as shown in Figure 14. Figure 15 shows the theoretical head map generated by that dipole

as well as the residuals. The residuals on the left hemisphere of the head map appear noisy and random showing a good overall fit. However, the right hemisphere still shows what might be some unrelated neural activity so further cleaning of the signal will be necessary for a better localization.
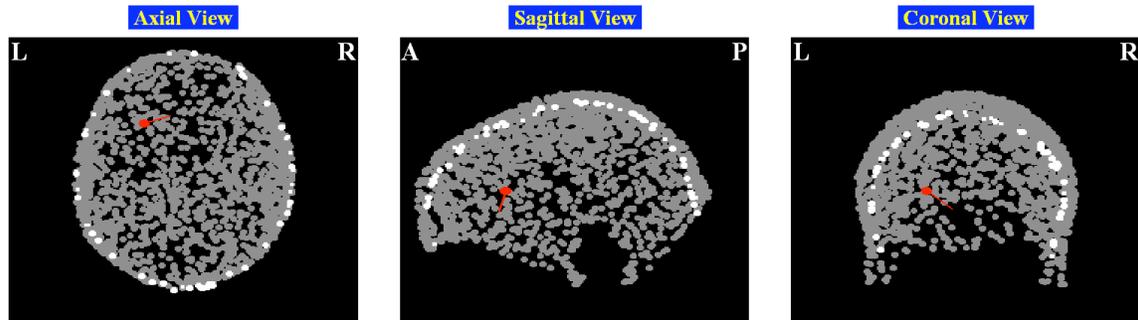
Fig 14. Localization of neural source. Used data with removal evoked response identified by DSS and localization from the left hemisphere without channel 73.
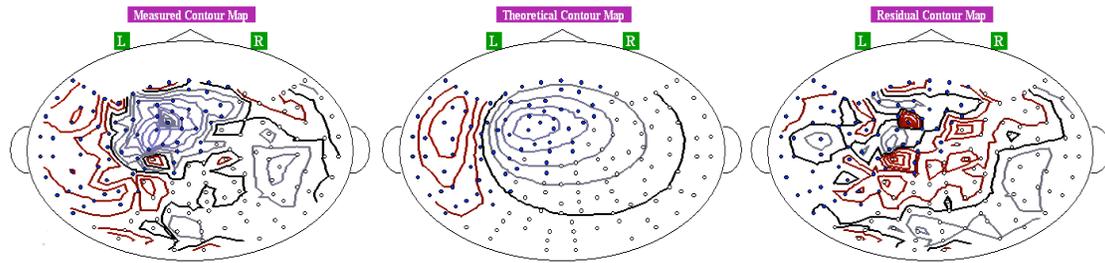
Fig 15. Head maps of magnetic field given from localization. From left to right: original head map of measured data, theoretical head map reconstructed from localized dipole, head map of residuals between measured and theoretical.

Conclusions

Through this research, the two peaks hidden in the head map of power were identified successfully as a dipole and a rough localization to the left hemisphere of the frontal lobe was made. Though the localization was successful, the ideal 90% correlation coefficient was not reached through a robust technique. The signal needs to be further cleaned of noise, especially biological, before a more precise localization can be possible. In each of the signal compression head maps, there still seems to be some neural activity present. Ideally, the head map would include the dipole and everything else would be noise.

Though the spatial source of the neural activity is known, it is unknown what this means physiologically. The source probably has something to do with higher order processing because it was found in the frontal lobe. Though the data analyzed was an auditory response, the presence of the dipole seemed to be unaffected by the type of stimulus or even the presence of the stimulus. Though not shown here, it is still present in the silence following the stimulus, though this might just be explained by a slow dying auditory response. One possibility is that the dipole might be related to the task-

dependent thinking and categorizing of the stimuli as a target or non-target. This idea is merely a conjecture of a curious mind and should not be taken as a serious deduction.

In addition to the actual findings of the study, the research also found new techniques and insight for analysis of induced activity within the brain. The most important technique developed was the examination of phase differences within each trial. One source of induced activity is the increase of power at a certain frequency after a stimulus without any change in phase. Another is the increase of power at a variable time after the stimulus. In both cases however, when there is a large amount of neural activity, the dipoles generated will have peaks 180 degrees out of phase with each other within the trial. It is important to only examine phase differences that have a high coherence and are therefore consistent throughout the trials. The inclusion of the coherence of these phase differences removes the necessity of averaging data over trials in order to remove noise. As seen above in Figure 12, it is also important to focus on connections between channels with the largest magnetic field. However, this would remove any dipoles from weaker neural sources. It would be interesting in the future to connect channels based on similarity of magnetic field instead of just having a large one. It is less likely that an extremely powerful sink is related of an extremely weak source. With a large enough data set, it might even be possible to reveal magnetic dipoles that are very low in power but maintain a high coherence over time.

Analyzing induced activity through this research has shown that a deeper classification might be necessary over the common evoked vs. induced response dichotomy. There will be a necessity in the future to distinguish an actual induced response from general related but non-stimulus locked activity: Induced activity (by the current definition) is simply anything that is not evoked or phase locked. An induced response, however should be an actual response locked to the stimulus in some manner other than in phase. The two possible sources explained above are from induced responses specifically since they are actual responses to stimuli. In contrast, induced activity also currently includes activity completely unrelated to the stimulus. One (trivial) example would be the heartbeat. These techniques revealed a heartbeat simply because it is not phase locked to the stimulus, but other neural responses might be equally unrelated to the specific stimulus preceding them. The neural process studied is most likely itself not even an induced response since it seems to be unchanging in time and independent of the stimulus. This shows the utility of the techniques developed here. It is possible to distinguish neural activity that is completely unrelated to the stimulus. This technique would be able to assist with identifying unwanted biological noise as well as analyzing induced activity and responses.

References

[1]   Cheveigne, A., and Simon, J.Z. Denoising based on spatial filtering. Journal of
        Neuroscience Methods. 331-339. (2008)

[2]   Hamalainen, M., Hari, R., et. al. Magnetoencephalography – theory,
        instrumentation, and applications to noninvasive studies of the working human
        brain. Reviews of Modern Physics, Vol. 65, No.2 (1993).

[3]   Luo, H., and Poeppel, D. Phase Patterns of Neuronal Responses Reliably
        Discriminate Speech in Human Auditory Cortex. Neuron 54, 1001-10 (2007).

[4]   Roberts, T.P.L., Ferrari, P., Stufflebeam S.M., and Poeppel D. Latency of the
        Auditory Evoked Neuromagnetic Field Components: Stimulus Dependence and
        Insights Toward Perception. J. Clin. Neurophsiol. Vol. 17, No.2, (2000).

[5]   Tallon-Baudry, C., and Bertrand, O. Oscillatory gamma activity in humans and
        its role in object representation. Trends in Cognitive Sciences. Vol. 3, No. 4.
        (1999)

[6]   Helsinki University of Technology – Adaptive Informatics Research Center. The
        FastICA package for MATLAB (2007). Retrieved from:
        http://www.cis.hut.fi/projects/ica/fastica/

[7]   Xiang, J., Y. Wang and J. Z. Simon (2005) *Magnetoencephalography (MEG)
        Response to Speech and Speech-like Modulations*, Society for Neuroscience
        abstracts.