UCLA
UCLA Electronic Theses and Dissertations

Title
Improvements to Simultaneous Electroencephalography – functional Magnetic Resonance Imaging and Electroencephalographic Source Localization

Permalink
https://escholarship.org/uc/item/3gg3z2q6

Author
Rodriguez, Cameron

Publication Date
2016

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

UNIVERSITY OF CALIFORNIA
Los Angeles

IMPROVEMENTS TO SIMULTANEOUS ELECTROENCEPHALOGRAPHY – FUNCTIONAL MAGNETIC RESONANCE IMAGING AND ELECTROENCEPHALOGRAPHIC SOURCE LOCALIZATION

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy
in Bioengineering

by

Cameron Rodriguez

2016

© Copyright by

Cameron Rodriguez

2016

ABSTRACT OF THESIS

Improvements to

Simultaneous Electroencephalography –

functional Magnetic Resonance Imaging and

Electroencephalographic Source Localization

By

Cameron Rodriguez

Doctor of Philosophy in Bioengineering

University of California, Los Angeles, 2016

Professor Mark S. Cohen, Chair

Both the method of simultaneous electroencephalography – functional magnetic resonance imaging (EEG-fMRI) and the method of electroencephalographic source localization are opening up new understandings in neuroscience, but both are also technically very challenging. Simultaneous EEG-fMRI combines the strengths of each modality: the temporal resolution of electroencephalography (EEG) with the spatial resolution of functional magnetic resonance imaging (fMRI). EEG source localization provides a non-invasive means of locating in the brain the sources of electrical signals measured on the surface of the head. The potential for these methods is not limited to research neuroscience; each has much to offer in fields like clinical neurology and psychiatry. In clinical neurology, for example, both methods have been used in the location of epileptic foci. Simultaneous EEG-fMRI is technically very challenging as

each modality serves as a source of artifact and safety concerns for the other. Two contaminants prominent in EEG acquired in the MRI environment are the gradient artifact (GA) and the ballistocardiogram artifact (BCG). The GA occurs only during an active MRI acquisition, while the BCG is always present, as its origins arise from the cardiovascular system of the subject. The current method to remove the GA from the EEG recording is called moving windowed average template subtraction (MWATS), the windows of which are time-locked to the repetition time of the MR acquisition. The problem with MWATS is that there are a number of options on how to implement it and little consensus or consistency in the field on how it is done. Current methods of removing the BCG artifact from the EEG recording all rely on establishing its timing based upon a simultaneously acquired electrocardiogram (ECG) signal. One issue with using the ECG signal is that the differential method for acquiring it implemented by the manufactures of the MR compatible EEG systems has a high rate of failure.

Source localization is technically very challenging in part because it relies on multiple measurements for its solution. Two of the measurements required for accurate source location are the measurement of the EEG electrode positions with respect to the underlying anatomy and the measurement of the underlying anatomy itself. At present, these measurements are made separately: the underlying anatomy is measured with an anatomical MRI scan, while the electrode positions are measured with either an electromagnetic digitization device or a photogrammetry system, neither of which is MR compatible. As such the two measurements must be registered to one another, typically through the location of fiducial markers on the surface of the head in both of the measurements. By aligning the markers, the two measurements can be brought into

register. The issues with this technique are two fold. First, aligning the separate measurements introduces a potential source of error. Second, both of the current methods for measuring the electrode positions are very time consuming. This thesis focuses in robust solutions to these issues.

Specifically what follows is: 1) an analysis of the consequences of all the principal options for implementing the moving windowed average template subtraction; 2) a method to eliminate the need of the separately recorded ECG signal for establishing the timing of the BCG artifact; and 3) a method for automatically locating and identifying EEG electrodes in an anatomical MRI scan. The analysis of the MWATS method permits more informed decisions to be made in its implementation. The method to eliminate the need for the ECG signal is to both simplify the experimental setup and to make the EEG-fMRI recording more robust. Automatically locating and identifying EEG electrodes in an anatomical MRI scan eliminates both a measurement step and a registration step. The elimination of the registration step eliminates a potential source of error. In addition, the combination of the elimination of the measurement step and the automatic way in which the algorithm progresses greatly reduces the time required for this process.

The thesis of Cameron Rodriguez is approved.

Daniel B. Ennis

Martin M. Monti

Daniel Jiong Jion Wang

Mark S. Cohen, Chair

University of California, Los Angeles

2016

DEDICATION

There are several people to whom I wish to dedicate this work: my father and mother, my brother, my wife, and my three sons. Their love, support, and the sacrifices they made for me all have helped me get to where I am today. Beyond where am I, they have also helped me become who I am today. My parents were the first in their families to go to college; my father also being the first to receive a law degree. They have worked tirelessly to both push and support my brother and me. They gave us both a foundation to become good and caring people and the tools to become successful as well. As a child, I am pretty sure they thought I would burn down the house with one of my "experiments," but they never tried to restrain my curiosity; in fact, they helped it flourish. My brother has and will always be there for me. It was also his words that motivated me to go to college. My wife has been with me through this entire process. Without her, this would have not been possible. She is always in my corner. My sons, well, they are my greatest inspiration. They bring enumerable joy to my life. I hope that I can do for them, what my parents did for me.

TABLE OF CONTENTS

ACKNOWLEDGMENTS

The work presented here was, in part, funded by grants from the National Institutes of Heath (R33DA022768, R33DA026109, and R21MH096329) and by the Staglin IMHRO Center for Cognitive Neuroscience. It was also made possible with the help and support of many people, including the members of my committee, Mark S. Cohen, PhD, Committee Chair, Daniel Ennis, PhD, Martin Monti, Phd, and Daniel Jiong Jion Wang, PhD; my former adviser, Dario Ringach, PhD; the person who taught me pyschtoolbox, Abtine Tavassoli, PhD; the person behind much of the data this work is based upon, Agatha Lenartowicz PhD; and many others of the UCLA community, Susan Bookheimer, PhD, Austin Head, Edward Lau, Andrew Cho, Shruthi Chakrapani, Marlo Duran, Jun Dizon, and Malina Revett Beatrice. Additionally, beyond being the chair of my committee, I would like to thank Mark Cohen, PhD for his advise and counsel, being my mentor and my advocate, and all of his help in editing this document.

CURRICULUM VITAE

Education:

2011 M.S. University of California, Los Angeles Bioengineering
2005 B.S. University of Vermont, Physics

Additional Training:

2012: Neuroimaging Training Summer Program
2006: Georgetown University, Post Baccalaureate Pre Medical Program

Grants and Fellowships:

2011 Neuroimaging Training Program Fellowship
2004 Hughes Endeavor for Life Sciences Fellowship

Teaching:

2013-2014: Neuroimaging Training Summer Program, UCLA, Lecturer / Lab Instructor
2012-2014: Principles of Neuroimaging, UCLA, Lecturer / Lab Instructor
2004-2005: Astonomy, UVM, Teaching Assistant
2004-2005: Physics Department Tutor, UVM, Teaching Assistant
2003: Physics Lab, UVM, Instructor

Honors:

2005 Sigma Xi Research Honors Society
2005 Sigma Pi Sigma Physics Honor Society
2003 Tau Beta Pi Engineering Honor Society
2002-2005 Deans List, UVM

Patents in Process:

CR Rodriguez, MS Cohen, "Fully automated localization of Electroencephalography (EEG) electrodes." USPTO. Assigned to The Regents of the University of California (Oakland, CA, US).

CR Rodriguez, MS Cohen, A Lenartowicz "A method for measuring cardiac timing from a ballistocardiogram." USPTO. Assigned to The Regents of the University of California (Oakland, CA, US).

Abstracts:

Rodriguez, CR., Lenartowicz, A, Cohen MS "An EKG-free method for extracting BCG timing from the EEG signal." 2013, Organization For Human Brain Mapping

Rodriguez, CR, Cohen, MS, "A method for fully automated localization and identification of EEG electrodes from magnetic resonance images." 2013, Organization For Human Brain Mapping

Suthana NA, Yap, N., Rodriguez, C., Wong, W., and Knowlton BJ. "Hippocampal CA3DG activity during encoding is associated with successful pattern separation." 2013, Society for Neuroscience

Suthana, N., Yap, N., Rodriguez, C., Wong, W., Knowlton, B. Hippocampal CA3DG activity during encoding is associated with successful pattern separation. Society for Neuroscience, San Diego, CA. Poster Presentation. Nov, 2013

Papers:

Lenartowicz A, Lu S, Rodriguez C, Lau EP, Walshaw PD, McCracken JT, Cohen MS, Loo SK (2016) Alpha desynchronization and frontoparietal connectivity during spatial working memory encoding deficits in ADHD: A simultaneous EEGfMRI study. NeuroImage: Clinical 11:210-223.

Rankin JM, Rodriguez C, Wright GA (2006) Bistable profile illumination in pulsars B0919+ 06 and B1859+ 07. Monthly Notices of the Royal Astronomical Society 370:673-680.

Commercial Products:

- MRI Compatible Projection System – Designed in conjunction with SMRT Image
- USB to TTL Device
- MRI Hardware Monitoring System

Other Select Professional Positions:

2013-Current: Director of fMRI and Technical Services, SMRT Image LLC
2013-Current: Technical Development, Center for Cognitive Neuroscience, UCLA
2007-2009: Senior Scientist, Quality Control and Assurance, Osiris Therapeutics INC.

INTRODUCTION

Since its inception roughly 25 years ago functional magnetic resonance imagining (fMRI) (Bandettini et al., 1992, Kwong et al., 1992, Ogawa et al., 1992) has opened up a new realm of possibilities with neuroimaging. It has confirmed (Yacoub et al., 2008), debunked (Nielsen et al., 2013) as well as opened new understandings (Fox et al., 2005) of how the mind works. When taking the whole family of non-invasive neuroimaging modalities into account it has by far the highest spatial resolution with some studies using resolutions on the order of sub millimeter (Yacoub et al., 2008). With all of its accolades and the beautiful pictures it creates, fMRI still has highly limited temporal resolution with the current upper limit of its sampling rate being on the order of hundreds of milliseconds (Feinberg and Setsompop, 2013).

Electroencephalography (EEG) fMRI's much older cousin has been around for over a century (Helmholtz, 1853) and gave some of the first views into the active mind. Even now it is still the key measurement for sleep staging (Kales and Rechtschaffen, 1968, Hori et al., 2001) and for diagnosing many neurological disorders such as epilepsy (Smith, 2005) and schizophrenia (Elbert et al., 1992). EEG, whose instrumentation is technologically primitive when compared with that of fMRI, leaves it in the proverbial dust in that EEG has essentially limitless temporal resolution. EEG spatial resolution on the other hand is more of an inference than a certainty in that in its native domain it a measurement of electrical potential differences on the scalp. Transforming these measures across the scalp into the space of the brain is done through a process referred to as source localization. Source localization's resolution is on the order of one to two centimeters (Akalin Acar and Makeig, 2013).

It then seems natural to combine the two and their individual strengths to create the optimal picture of the active brain. Some of the uses for simultaneous EEG & fMRI (EfMRI) are for the location of epileptic foci and to study: sleep and its disorders (Czisch et al., 2002, Caporro et al., 2012); attention mechanisms (Lenartowicz et al., 2016); and the interconnectedness of the origins of each of its signals (Ritter and Villringer, 2006). EfMRI not only holds the promise of combining the true strengths of each of the components, but it also has the potential to benefit from the fact that its two signals arise from two separate physiological phenomena, EEG from the ion gradients created from the depolarization of neurons, where the fMRI signal arises from vascular changes that occur to accommodate the increased metabolic rates of active neurons. EEG is able to realize transient events; ones that are so narrow in their temporal breadth that they do no evoke a vascular response as long as there is temporal and spatial synchronicity amongst thousands of firing neurons. fMRI can measure events from amongst a disorganized and asynchronous group so long as they are sustained or large enough to create a vascular response.

While the disparate strengths and signal origins of each modality are complimentary of one another, the methods by which they are measured makes them highly incompatible. It is the desire to merge the two and the difficulty in doing so that has driven much of the work of Mark Cohen, Rachel Goldman, Karen Mullinger, Phillip Allen, Rami Niazy, as well as many others. While there are seemingly endless promises for combining EEG and fMRI the fact is that methods by which each of the signals are acquired make the two rather incompatible with one another – each acts as a source of

noise and artifact for the other. It is the MRI machine's high static magnetic field as well it extreme sensitivity that makes it highly incompatible with most other electronic devices.

The issues between the two can be broken up into two parts, non-signal and signal related (Meriläinen, 2002). The non-signal issues being safety issues and device communication issues. These have mostly been solved with careful material selection, electrical isolation, careful wiring and radio frequency shielding (Nöth et al., 2012). The signal related problems, however, have been the most vexing. The effects that the EEG induces on the MRI signal result largely from the electrodes and the amplifier / digitizer (Mullinger et al., 2008, Rosenkranz and Lemieux, 2010). These can act as an antenna both to siphon the RF energy away from and introduce electronic noise into the MR system. The effects MRI induces on EEG come both from its high static magnetic field as well as its rapidly switching gradient fields that introduce signals into the EEG recording (Allen, 1998).

How MRI Works in brief?

Magnetic Resonance (MR) or Magnetic Resonance Imaging (MRI) works by imaging "packets" of magnetic dipoles (M_0) tipped from their parallel alignment with an external magnetic field (B_0) and oscillating at their resonant frequency (i.e. Larmor frequency) about B_0. The net magnetic dipole moment (μ) of an atom arises from the sum of all of the spin angular moments of the charged particles that comprise it. The net μ measured in the case of fMRI is from hydrogen molecules bound or unbound, typically those found in water. An atom / molecule with a net spin experiences the Zeeman effect/splitting when in the presence of an external magnetic field (B_0 in the case of MRI). This is the splitting of energy levels of the atom / molecule. The difference of the energy

between the levels is equal to the energy of a photon with the Larmor frequency, which is in the radio frequency range at typical MR scanner field strengths (1.5T to 3T).

Boltzmann statistics predicts at room temperature that the number of spins at lower energy will slightly outnumber the spins with higher energy. It is due to both Zeeman splitting and Boltzmann statistics that the sum of a collection of molecules (packet) ends up having a net (bulk) magnetization (M), which at equilibrium is parallel to B_0 (longitudinal axis) and has a magnitude M_0.

The bulk magnetization (M) is tipped (excited) from the longitudinal axis by an external RF field (B_1). The frequency of the RF pulse is matched to the Larmor frequency of the bulk magnetization of interest (for fMRI molecules containing hydrogen). As the magnetic dipoles transition from their induced excited state to their equilibrium state they emit RF photons whose frequency is the Larmor frequency and is determined by the magnetic field strength an the gyromagnetic ratio of the species of interest. To encode location, magnetic gradient fields alter the B_0 magnitude across space. In doing so they change the frequency of the bulk magnetization's precession. It is this encoding that is received as the signal in the receivers of the MR machine. The received signal then goes through phase sensitive detection to isolate the frequency and phase of interest from the noise. It is this cleaned up signal that is used to populate what is know as k-space, a spatial frequency representation of the image. The k-space "image" is then run through an inverse Fourier transform to create the MR image. Gradient fields can be used to encode much more than location when used in combination with various pulse sequences. They can be used to encode motion, change contrast be tissues, suppress undesired signals, as well as a multitude of other things.

In the case of fMRI, the functional part is typically a measure of local vascular changes, which are deemed to be the result of changes in local neural activity. Typically, the measure is of blood volume or blood oxygenation, the latter being the most common. Specifically, blood oxygenation level dependent (BOLD) imaging measures increases in MR signal that are the result of increased oxygen content in the blood, which is a corollary for an increase in local neural activity.

How EEG works in brief?

Electroencephalography works by measuring small voltage differences across the scalp that is a reflection of the neural activity beneath. The voltage differences arise from the ion gradients created from depolarization and repolarization of active neurons. If the aggregate clusters of active neurons form what is referred to as a secondary current. It is this summation, this secondary current, that gives rise to the EEG signal on the scalp. How these secondary currents are realized on the scalp is a function of it proximity to scalp, the orientation with respect the scalp, and synchronicity of the neuronal population that generates it. Other sources of scalp voltage differences are from muscle activity, primarily those of the facial muscle and eyes. These non-EEG signals typically can be distinguished from the true EEG signal by their spatial distribution across the scalp, amplitude and time course.

What is EEG source localization in brief?

Source localization in EEG is the process of taking measured electrical potential differences on the scalp and estimating the location of their sources inside the head. Sources sometimes referred to as current generators or current sources are the secondary

currents from synchronous of clusters of neurons. Locating these sources involves essentially two things, the forward problem and the inverse problem. The forward problem is modeling how a given source would appear on the surface of the head. This is typically accomplished using a MR structural scan subjects head and employing things like boundary element or finite element models to estimate the conductive properties of the head. The inverse problem takes the conductivity model from the forward problem and uses it to project inwards to their sources the EEG measurements from the surface of the head. The forward problem has a unique solution. The inverse problem is ill defined, underpowered and with infinite solutions. It is only with a priori constraints that the inverse problem can be solved. The assumptions include things such as: where they can originate from, e.g. the brain; the maximum number of concurrently active sources; and models for how the lead fields of the sources propagate, e.g. as a dipole, quadrapole, current distribution, etc. (Michel et al., 2004).

Safety Concerns with combining EEG and MRI

The magnetic field is problematic for many reasons, but first and foremost for safety. It can be the source of burns and injuries from projectiles. Any ferromagnetic material within the standing static field of the MRI has the potential for becoming a projectile. Conductive materials in general can also serve as sources of danger in they can heat up due from induction or energy absorption due to the rapidly switching field magnetic field gradients or acting as an antenna for radio frequency signals of the MRI machine.

These safety concerns can sometimes be dealt with using careful planning, material and wire length selection, and device positioning. It is the signal related

incompatibilities, the artifacts that each introduced into one another's recording that have proved most daunting to deal with.

EEG signal issues caused by MRI

In the EEG recording most of the induced signals (artifacts) are the result of changes in the magnetic flux in the circuit loops within the EEG system. In an overly simplified trace of this loop, it goes from the measurement point (electrode) to recording instrument (EEG) through a conductive wire, then from the EEG to the reference point (a second electrode) through a wire, and finally from electrode to electrode through the subject's tissue. Signal can be introduced into the EEG recording from any of the possible perturbations of magnetic flux over time. According to Faraday's Law these include: 1) magnetic field strength, 2) the above mentioned conductor loop area and 3) loop orientation said loop with respect to magnetic field. In short any change of the position of the wires / electrodes or change in the magnetic field strength through the loop will result in a signal in the EEG recording.

Perturbations to the first of these lead to the largest error in terms of energy of these artifacts in the EEG recording. This artifact is know as the gradient artifact (GA) whose magnitude is in the range of two to three orders of magnitude greater than the EEG signal itself when acquired using a modern 3T scanner. The MRI system uses what are referred to as gradient fields to perturb the magnetic field in order to create an image. Specifically the gradient fields are used in the selective excitation of nuclei of interest, to create various forms of image contrast, and to be able to localize the emitted signals from the aforementioned nuclei. These changes to the magnetic field cause changes in the

magnetic flux through the EEG circuit loops, which turn induces signal in the EEG recording.

The next artifact in terms of energy in the EEG recoding is the ballistocardiogram (BCG), a signal introduced into the EEG from cardiac related motion. The ballistocardiogram is believed to originate from three sources (Mullinger et al., 2013) all of which find their origins in the cardiac cycle, cardiac induced head rotations and translations (HRT), scalp pulsations (SP) and voltage differences across the head arising from the Hall effect (HE) phenomena. HRT is simply the translations of the ballistic motion of the heart to the head. Scalp pulsation is simply that, the movement of the scalp from the pulsing vessels below. Both HRT and SP cause changes in loop area and orientation. These changes lead to changes in magnetic flux over time, which induces a signal in the EEG recording. The other component, the Hall effect, is the separation of moving charges as they are passed through a magnetic field. In this case the charges are particles in the blood. As the blood circulates through the body and magnetic field the positive and negative charges separate. This charge separation when aggregated creates voltages differences on the scalp that the EEG system measures.

Respiratory Motion (RM) and General Subject Motion (GSM) are similar to HRT but neither is time locked to the cardiac cycle. It is true that rate of RM and the cardiac rate are correlated in that when the heart rate changes so does the other, however one is not the trigger for the other. That is, a heartbeat does not trigger an inhalation or exhalation. Likewise an inhalation or an exhalation does not trigger a heartbeat. The periodicity of RM is approximately one tenth of that of the cardiac cycle, having frequency in the sub-Hertz range. The physical displacement of the RM tends to be larger

than HRT as it is the whole torso's motion being translated to head motion rather than just the heart. However, the motion takes place over a longer interval, resulting in a more dispersed signal. The slower rate of RM leads to a lower amplitude signal when compared to the HRT. This is because the induced signal is directly proportional to the rate of change (Faraday's Law). GSM on the other hand is sporadic, rapid and comparatively large in overall displacement. This make it the largest, in terms of amplitude, of signals induced in the EEG recording. The GSM can be caused directly by a motion of the head and face or indirectly from motions of, neck, shoulder, arm, leg, etcetera that eventually translate into head motion. While the analog to digital converters of MR compatible EEG systems are designed to work linearly over many orders of magnitude (μV to mV), GSM can induce signals that are large enough that they can lead to clipping. Clipping by the analog to digital converters results in an unrecoverable loss of signal.

Vibrations from the MRI system's cryopump (CP) also induce signal in the EEG recording. The cryopump is a device used to maintain the pressure of the liquid helium in the MRI system, which keeps the superconducting MRI magnet cool. The cryopump motions are translated from the pump, which is attached to magnet to the patient table to the subject themselves. These motions, like SP and HRT, perturb the areas and orientations of the conductive loops that lead to changes in the magnetic flux, which in turn leads to signal being induced in the EEG recording.

MRI signal issues caused by EEG

MRI is an emission rather than transmission imaging modality. It other words the signal used to create the image originates rather than is absorbed or reflected by the

subject. Both the endogenous nature, as well as the prediction by Boltzmann statistics that higher energy will only slightly outnumber lower energy spins (the difference being the signal origin), greatly limit the amplitude of the signal, this in turn requires the sensitivity of the recording machine to be very high. BOLD imaging takes this small signal and looks for subtle changes in its amplitude to access neural activity levels. In the MR signal the artifacts are the result of perturbations to the small MR signal or to image acquisition and reconstruction constraints.

In the category of signal perturbation there are three main problems. First is introduction of external signal (noise) into the MR signal, second is the decrease of true MR signal, and third is distortion of the signal. The introduction of external noise can occur two ways, either through faulty radio frequency (RF) shielding and or RF filtering of the EEG system in the magnet room, or the cabling of the EEG system acting as a path / antenna for electronic noise from other devices in the magnet room to enter into the scanner. The reduction of the MR signal comes from RF parasitic losses, which occur from the EEG system acting again as an antenna, though in this case the direction is reversed. Here it is the excitatory B_1 pulse that is diverted from the tissue and spins that it was to excite. As such the net signal is reduced. The distortion of the MR signal comes from mismatch in the magnetic susceptibility (μ) of materials in the EEG system, particularly the electrodes and the cabling, and that of the tissue being imaged. Susceptibility mismatches distort $B_{0,}$ $B_{1,}$ and the gradient fields. This occurs because magnetic field lines bend into materials that have higher magnetic susceptibility than their neighbors (Mullinger et al., 2008).

Image acquisition & reconstruction constraints are brought around by the additional bulk that is created by the EEG electrodes and cap. As overall subject volume increases the possibility of wrap around increases, this is where one part of the object being imaged wraps (is incorrectly located) to the opposite side of image. Wrap around occurs in the phase encoding direction of the image where MR signal producing material exists outside of the selected field of view (FOV) of the image. The wrap occurs because phases of the spins here are either higher than π or lower that $-\pi$ leading to their incorrect assignment in k-space. In other words if the object being imaged is larger in the phase encoding direction than the FOV prescribed by the operator the parts of the object outside the FOV will be wrapped around to the other side of the image. Increasing the object size, in this case the adding of the EEG electrodes to the subject, can lead to image wrap around if the increased object size is not taken into account when assigning the FOV.

What is done to fix the EEG recording?

To recover the highest fidelity EEG signal from EfMRI typically the following operations and order are performed: careful wiring, cyropump shutoff, GA cleaning, BCG cleaning.

Close attention must be paid to careful wiring in order to minimize the induced signals as much as possible. This includes minimizing the loop area, stabilizing the cabling from vibrations and keeping the cabling as parallel to the B_0 as possible. Magnetic flux is directly proportional with loop area. By reducing the loop area all induced artifacts are reduced. Stabilizing the wiring comes into play in that preventing as much of the loop area as possible from moving minimize the changes in flux which minimize the effects of the cyropump, BCG and GSM.

To eliminate the signals that come from the cyropump vibrations, the simplest solution is to just shut it off during the duration of the EfMRI recording. For short durations and with adequate fill levels of liquid helium the MRI hardware can tolerate the cryopump shutoff. If, however, the cryopump is left off for an extended duration there can be a significant boil off of liquid helium, which is needed to cool the superconducting coil that generates B_0. This can lead to the self-quenching (shutdown) of the MRI scanner. Scanner quenching is both very costly in scanner time and money, as such much care must be taken in turning the cyropump off.

GA cleaning is primarily accomplished using template subtraction (Allen, 1998, Allen et al., 2000). Template subtraction relies on the fact that the GA is highly repeatable between repetition times (TR) and that it is an entirely additive noise source. Where the GA is highly repeatable across TRs, the EEG is not, if fact it is has a zero mean. By taking a windowed average of the EEG signal across TR segments the EEG portion of the signal cancels out leaving a high fidelity template of the GA artifact. Since the GA is an additive noise source simply subtracting this template from the EEG recording removes the GA. Because of perturbations to loop areas and orientations from subject motion as well as baseline drifts and amplitude scaling from changes that occur at the electrode interfaces, the GA does vary slightly over time. To compensate for this, rather than including all possible TR segments to create a GA template a moving windowed average is used to create the template. The moving window allows the template to adapt and have a higher fidelity to the local GA structure.

After the GA has been removed BCG cleaning is then performed. There have been several methods developed to accomplish this. They fall into two camps: template

subtraction and component subtraction. Template subtraction works much like that of the GA cleaning in that a window average is performed about the cardiac events to generate a template of the BCG. This template is then subtracted from the EEG recording. Component subtraction involves decomposing the EEG recoding through processes such as principal component analysis (PCA), independent component analysis (ICA), or convex analysis. The components that show a high temporal correlation with the cardiac events are zeroed out. In addition to the temporal correlation, spatial distribution patterns of the components can be used to aid in the identification of BCG components. After the BCG components have been removed (zeroed out) the remaining components are then used to reconstruct the EEG recording without the BCG.

What is done to fix the MR recording?

While most of the EEG cleaning is done either in real time or in post processing all of all of the induced artifacts in the MR recording are handled in a preventative nature.

The EEG system has robust RF shielding and filtering to prevent signals from its hardware from contaminating MR signal. To prevent the EEG system from acting as a source of noise or as an antenna for RF noise generated by other equipment in the MR suite, RF noise tests are regularly run to check for RF shielding and filter failures. Problem sources are identified and remedied before performing EfMRI. To prevent parasitic RF losses the EEG cable is tuned to minimize RF absorption by controlling wire length and cable path (Mullinger et al., 2008). The use of RF baluns to mitigate this is also currently under investigation. Susceptibility artifacts are generally not an issue in that the effected area is quite local to the location of the susceptibility mismatch. In this case the surface of the head. This is far enough away from the brain as to cause no real

issue (Mullinger et al., 2008). That being said magnetic susceptibility is still taken into account in material selection order to minimize the mismatch with tissue. The wrap around artifact is taken care of by either using radio translucent (invisible to the MRI) material to make the EEG net/cap or by increasing the FOV to include the net/cap.

Limitations of the current methods & how they can be improved

Analysis of Moving Windowed Average Template Subtraction

There are many options on how to implement a moving windowed average template subtraction (MWATS) filter for the purpose of removing the gradient artifact from the recorded EEG. The options include: the region of support of the filter; whether or not to include the segment to be cleaned in the creation of the template, how to weight the template, i.e. equally within a TR segment or on a point-by-point basis; what type of weighting to apply to the windowed average, e.g. none, Gaussian, Hamming, etc.; and how many TR segments to include in the windowed average. The issue is that the implications of each of these options for both the fidelity of the cleaned EEG signal to the true EEG signal and to the separability of the two signals, i.e. the quality of GA removal, are poorly understood and have up and to this point not been thoroughly examined. As a result there is little consensus across the EfMRI research community about how to best implement the MWATS filter.

The work that follows in Chapter 1 contains experimental results complemented by an analytic analysis of the implications of the various options for the implementation of the MWATS. Specifically what is looked at it how the various options affect: the fidelity of the cleaned EEG to the actual EEG signal; how well the GA is separated from the EEG signal; and the spectral content of the cleaned EEG. For the experimental

portion, in order to properly quantify fidelity and separability, a ground truth signal was created. To make this signal as realistic as possible an actual EEG recording from a human subject made in the MRI scanner with the cyropump and gradient fields off and used for subsequent analysis. This signal was then transmitted, replayed, through the entire EEG acquisition pipeline so as to make it as true to an actual EEG recording as possible. Custom circuitry was design and built for this purpose, the design of which has been open sourced. A link to the design, the build of materials, and recordings made with the ground truth signal can be found in Chapter 1. Finally recommendations and guidelines for implementing the MWATS filter using the various options are put forward. The guidelines take into account the practical aspects of implementing a MWATS filter depending the subject population and the type of EEG analysis being performed.

An ECG-free Method for Establishing Ballistocardiogram Timing

Current methods for suppressing the ballistocardiogram artifact all rely on knowing the time points at which it occurs. This timing is established indirectly through the correlated and simultaneously acquired electrocardiogram (ECG) signal. The ECG is used because its signal magnitude is on the order of twice that of the EEG and its R-wave peak is distinct and narrow. Using the ECG has it drawbacks though. First, it requires sacrificing recording channels in the EEG system to record the ECG signal that otherwise could be used to record EEG. Second, it adds complexity to the setup of an already complicated experimental setup. ECG acquisition requires extra: subject preparation time for cleaning and abrading of the chest; and experiment setup time from securing leads from vibration. In addition finding the optimal lead placement on subjects with large chests whether from: muscle, breast tissue, or fat can be difficult. Finally, it is prone to

failure (Srivastava et al., 2005). The leading manufactures of MR compatible EEG systems only use a simple difference measurement between two points for the ECG recording. This is likely to minimize the number of EEG channels sacrificed to record the ECG, however it then only requires a single failure at either one of the ECG contacts for it to fail. These failures can be brought around in part from subject movement that can lead to a torqueing of contacts and issues with the lead placement. Attempts have been made previously to extract the BCG timing directly from the EEG recording (Srivastava et al., 2005). In these attempts the way in which the BCG signal was accentuated over the EEG blurred it's timing and made it susceptible to artifacts such as those that would be caused by eye blinks.

The work that follows in Chapter 2 describes a method for establishing the BCG timing directly from the EEG recording: 1) the timing of which is time locked to that of the ECG; 2) provides equivalent peak detection performance to that of the ECG; 3) is more robust in that there is not a total loss signal with the loss of a single electrode contact; and 4) can be applied retrospectively to EEG recordings in which the ECG recording failed. In addition to providing a single time point for each BCG event, the ECG R-wave peak equivalent, the method also provides the window about the time point over which the BCG event occurs.

Automatic Identification of EEG Electrodes in a 3D Volume

The accuracy of the source localization is directly tied to the accuracy of the forward model (Akalin Acar and Makeig, 2013), the total number of electrodes and the accuracy of the knowledge of their locations with respect to the brain (Michel et al., 2004). The former gives the head conductivity model used by the various inverse solution

algorithms to back project the scalp EEG back into the head. The latter has to do with the accuracy and resolution of the sampling of the spatial frequencies of the EEG on the scalp (Spitzer et al., 1989, Gevins et al., 1991, Srinivasan et al., 1996, Srinivasan et al., 1998). In their 2004 paper, Michel et al. suggest a minimum of electrode count of at least 60, but more on the order of 100+ with an inter-electrode spacing in the range of 2-3cm to avoid distortions in the measurements scalp potential distribution.

As mentioned above, to accurately model subject specific head conductivity a high resolution structural MR scan of the subject's head is required. On this structural scan a boundary element or finite element models are used to measure the distributions of the various conductive elements of the head, e.g. white matter, grey matter, cerebrospinal fluid, the skull and the scalp.

To locate the scalp potential measurements with respect to the head model the electrode positions are typically measured using either an electromagnetic digitizer (Polhemus Inc.,) or photogrammetry system (Electrical Geodesics Inc.,) in reference to three or four fiducial points on the scalp. The fiducial points that are typically used are the nasion and the two preauricular points. The inion is sometimes included in addition to these points. These same fiducial markers are also located in the MR scan and used to bring the measured electrode positions into register with the head model from the MR scan.

Existing systems for measuring the electrode positions are time consuming and require a significant amount of manual input and are still prone to many sources of error. Notably, the fiduciary marks are linked to external markers, not to the brain. Typical shifts in subject position (e.g., changing from a sitting to a prone posture) displace the

brain substantially with respect to these surface markers. Electromagnetic digitizers are time consuming on the front end in that the fiducial points as well as each electrode position must be identified and measured. On the Electrical Geodesics Inc.'s 256 channel EEG net this is at least 261 individual measurements: one for the common mode electrode, one for the reference electrode, 256 for each EEG channel, and at least three for the surface fiducial markers. While these positions are being measured, the subject has to remain still to prevent movement of the EEG electrodes with respect to the scalp and the digitizer apparatus (Russell et al., 2005). In contrast to electromagnetic digitizers, photogrammetry systems are time consuming on the back end. To avoid the issues with measuring individual points for their high channel count nets EGI developed a photogrammetry system to measure the electrode on the scalp. Their system takes 11 photos around the surface of the head simultaneously. This allows for a rapid measurement time for all of the electrode positions, however the measurements do not include the individual channel identifications. This requires the operator to manually identify 20 marker electrodes across the 11 pictures. Using these marker electrodes their system attempts to identify or estimate the locations of the remaining electrode channels (Russell et al., 2005). The photogrammetry system though can be easily confounded by the obscuring in the photos of the electrodes by the subject hair and by the improper placement of the subject head in terms of the fields of view of the cameras (Russell et al., 2005).

In the work that follows in Chapter 3 is a method to automatically locate and identify electrode position from a volume image. In the case where the volume image is the MR structural scan used to create the subject specific head conductivity model a

measurement step and a registration step are eliminated from the source imaging process. The measurement step eliminated being the separate measurement of electrode positions. The registration step eliminated being the registering of the measured electrode positions to the head conductivity model through the use of the surface fiducial markers, as the two are already in the same coordinate system. Beyond eliminating these two steps, both of which can be potential sources of error, the algorithm is fully automatic, eliminating the tedious manual input of locating and identifying the electrodes and the possibility of operator bias.

References:

Akalin Acar Z, Makeig S (2013) Effects of forward model errors on EEG source localization. Brain Topogr 26:378-396.

Allen PJ, Josephs O, Turner R (2000) A method for removing imaging artifact from continuous EEG recorded during functional MRI. Neuroimage 12:230-239.

Allen PJ, Polizzi, G., Krakow, K., Fish, D. R., & Lemieux, L. (1998) Identification of EEG Events in the MR Scanner: The Problem of Pulse Artifact and a Method for Its Subtraction. Neuroimage 8:229-239.

Bandettini PA, Wong EC, Hinks RS, Tikofsky RS, Hyde JS (1992) Time course EPI of human brain function during task activation. Magnetic resonance in medicine 25:390-397.

Caporro M, Haneef Z, Yeh HJ, Lenartowicz A, Buttinelli C, Parvizi J, Stern JM (2012) Functional MRI of sleep spindles and K-complexes. Clinical neurophysiology 123:303-309.

Czisch M, Wetter TC, Kaufmann C, Pollmächer T, Holsboer F, Auer DP (2002) Altered processing of acoustic stimuli during sleep: reduced auditory activation and visual deactivation detected by a combined fMRI/EEG study. Neuroimage 16:251-258.

Elbert T, Lutzenberger W, Rockstroh B, Berg P, Cohen R (1992) Physical aspects of the EEG in schizophrenics. Biological psychiatry 32:595-606.

Feinberg DA, Setsompop K (2013) Ultra-fast MRI of the human brain with simultaneous multi-slice imaging. Journal of magnetic resonance 229:90-100.

Fox MD, Snyder AZ, Vincent JL, Corbetta M, Van Essen DC, Raichle ME (2005) The human brain is intrinsically organized into dynamic, anticorrelated functional networks. Proceedings of the National Academy of Sciences of the United States of America 102:9673-9678.

Gevins A, Le J, Brickett P, Reutter B, Desmond J (1991) Seeing through the skull: advanced EEGs use MRIs to accurately measure cortical activity from the scalp. Brain topography 4:125-131.

Helmholtz HV (1853) Ueber einige Gesetze der Vertheilung elektrischer Ströme in körperlichen Leitern, mit Anwendung auf die thierisch-elektrischen Versuche (Schluss.). Annalen der Physik 165:211-233.

Hori T, Sugita Y, Koga E, Shirakawa S, Inoue K, Uchida S, Kuwahara H, Kousaka M, Kobayashi T, Tsuji Y (2001) Proposed supplements and amendments to 'A manual of standardized terminology, techniques and scoring system for sleep stages of human subjects', the Rechtschaffen & Kales (1968) standard. Psychiatry and clinical neurosciences 55:305-310.

Kales A, Rechtschaffen A (1968) A manual of standardized terminology, techniques and scoring system for sleep stages of human subjects: US Department of Health, Education and Welfare, Public Health Service, National Institutes of Health, National Institute of Neurological Diseases and Blindness, Neurological Information Network.

Kwong KK, Belliveau JW, Chesler DA, Goldberg IE, Weisskoff RM, Poncelet BP, Kennedy DN, Hoppel BE, Cohen MS, Turner R (1992) Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. Proceedings of the National Academy of Sciences 89:5675-5679.

Lenartowicz A, Lu S, Rodriguez C, Lau EP, Walshaw PD, McCracken JT, Cohen MS, Loo SK (2016) Alpha desynchronization and frontoparietal connectivity during spatial working memory encoding deficits in ADHD: A simultaneous EEGfMRI study. NeuroImage: Clinical 11:210-223.

Meriläinen V (2002) Magnetic resonance imaging with simultaneous electroencephalography recoding: safety issues [diploma thesis]. Helsinki University of Technology, Department of Electrical and Communications Engineering.

Michel CM, Murray MM, Lantz G, Gonzalez S, Spinelli L, Grave de Peralta R (2004) EEG source imaging. Clin Neurophysiol 115:2195-2222.

Mullinger K, Debener S, Coxon R, Bowtell R (2008) Effects of simultaneous EEG recording on MRI data quality at 1.5, 3 and 7 tesla. International Journal of Psychophysiology 67:178-188.

Mullinger KJ, Havenhand J, Bowtell R (2013) Identifying the sources of the pulse artefact in EEG recordings made inside an MR scanner. Neuroimage 71:75-83.

Nielsen JA, Zielinski BA, Ferguson MA, Lainhart JE, Anderson JS (2013) An evaluation of the left-brain vs. right-brain hypothesis with resting state functional connectivity magnetic resonance imaging. PloS one 8:e71275.

Nöth U, Laufs H, Stoermer R, Deichmann R (2012) Simultaneous electroencephalography‐functional MRI at 3 T: An analysis of safety risks imposed by performing anatomical reference scans with the EEG equipment in place. Journal of magnetic resonance imaging 35:561-571.

Ogawa S, Tank DW, Menon R, Ellermann JM, Kim SG, Merkle H, Ugurbil K (1992) Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. Proceedings of the National Academy of Sciences 89:5951-5955.

Ritter P, Villringer A (2006) Simultaneous EEG-fMRI. Neurosci Biobehav Rev 30:823-838.

Rosenkranz K, Lemieux L (2010) Present and future of simultaneous EEG-fMRI. Magnetic Resonance Materials in Physics, Biology and Medicine 23:309-316.

Russell GS, Jeffrey Eriksen K, Poolman P, Luu P, Tucker DM (2005) Geodesic photogrammetry for localizing sensor positions in dense-array EEG. Clin Neurophysiol 116:1130-1140.

Smith S (2005) EEG in the diagnosis, classification, and management of patients with epilepsy. Journal of Neurology, Neurosurgery & Psychiatry 76:ii2-ii7.

Spitzer AR, Cohen LG, Fabrikant J, Hallett M (1989) A method for determining optimal interelectrode spacing for cerebral topographic mapping. Electroencephalography and clinical Neurophysiology 72:355-361.

Srinivasan R, Nunez PL, Tucker DM, Silberstein RB, Cadusch PJ (1996) Spatial sampling and filtering of EEG with spline laplacians to estimate cortical potentials. Brain topography 8:355-366.

Srinivasan R, Tucker DM, Murias M (1998) Estimating the spatial Nyquist of the human EEG. Behavior Research Methods, Instruments, & Computers 30:8-19.

Srivastava G, Crottaz-Herbette S, Lau KM, Glover GH, Menon V (2005) ICA-based procedures for removing ballistocardiogram artifacts from EEG data acquired in the MRI scanner. Neuroimage 24:50-60.

Yacoub E, Harel N, Uğurbil K (2008) High-field fMRI unveils orientation columns in humans. Proceedings of the National Academy of Sciences 105:10607-10612.

ANALYSIS OF MOVING WINDOWED AVERAGE TEMPLATE SUBTRACTION

Chapter 1

Abstract:

In simultaneous (EEG-fMRI) Electroencephalography (EEG) and functional Magnetic Resonance Imaging (fMRI), the Gradient Artifact (GA) proves a significant challenge. GA induced into the EEG recording by the magnetic resonance imaging (MRI) system's gradient fields is in today's 3 Tesla field MR systems on the order of 2 to 3 magnitudes greater than the EEG signal itself (Yan et al., 2009). The accepted standard for cleaning or removing the GA from the EEG record is through the use of moving windowed average template subtraction (MWATS). MWATS has many implementation options: causal or centered; time-point weighting; whether to weight TR window by TR window or sample point by sample point; whether to include the TR window to be cleaned in the template; and the number of TR windows to include in the template. There appears to be little consensus in how to select among, or to apply, these variations. To bring a better understanding of MWATS, and each of its options, we examined the method both from an analytic standpoint, and through the study of experimentally recorded data. To set a ground truth for the data analysis, we developed a means to create a synthetic EEG signal that could be carried through the entire analysis pipeline. Based upon both the analytic and experimental results, we have developed guidelines on how to best apply the MWATS, based both upon the subject population, and on the type of analysis being performed.

Keywords:

Simultaneous EEG-fMRI, EEG-fMRI, functional Magnetic Resonance Imaging, fMRI, EEG, Gradient Artifact, Moving Windowed Average, Template Subtraction

Introduction:

Simultaneous collection of Electroencephalography and functional Magnetic Resonance Imaging was adopted for two reasons. The first was to merge the imaging strengths of constituents: temporal resolution for EEG, and spatial resolution for fMRI. The second was to combine signals from two different origins: changes in scalp potential resulting from the depolarization of neurons beneath for EEG (Helmholtz, 1853), and vascular changes caused by changes in the metabolic need of neurons due to actively level for fMRI (Bandettini et al., 1992, Kwong et al., 1992, Ogawa et al., 1992) in the interest of better understanding human brain activity and physiology, and to improve precision in such challenges as brain reading. Melding EEG with fMRI, unfortunately, is technically challenging, because the nature by which each of the respective signals is captured makes the two modalities almost mutually incompatible. The incompatibilities exist in two domains: safety and signal contamination. The safety issues have been mitigated by careful materials selection and proper subject preparation. The signal issues arise because each modality serves as a source of signal artifact for the other (Goldman et al., 2000, Mullinger et al., 2008).

In this chapter, we focus on the gradient artifact (GA), which is induced in the EEG recording by the switching gradient fields of the MRI system. Specifically, we focus on the options for applying the community standard for cleaning/removing the GA from the EEG record: template subtraction (TS), or moving windowed average template

subtraction (MWATS) (Allen et al., 2000). The options include: causal or centered templates; the moving weight type; whether to weight the moving window respecting the repetition times (TR) boundaries of the MR acquisition; whether to include the TR window to be cleaned in the template; as well as the number of TR segments to include in the template. How MWATS is implemented affects both the cleaning of the GA and the remaining EEG signal – the latter being the subject of surprisingly little attention. In this chapter we provide both an analytical analysis and experimental results. Based upon these we put forward recommendations for the implementation of MWATS.

The GA arises from the perturbations to the standing magnetic field (B_0) of the MRI machine by the gradient fields (G_x, G_y, & G_z). Such field gradients are necessary both to create the image and to introduce contrast within it. The EEG system can be considered as a collection of conductive loops, *Figure 1*. There are at least as many loops as measurement electrodes (recording channels). Faraday's law of induction states that an electromotive force (*e.m.f.*) is generated in a conductive loop that is equal to the negative of the change of dot product of the loop area and the magnetic field with respect to time, *Equation 1*. In other words, any change in the magnetic flux over time generates a corresponding *e.m.f.* in the loop. The sum of impedances of the elements in the loop determines the current through the loop and the voltage measured by the EEG system; that is, the voltage across the input to the EEG amplifier. Here, it is the time-varying magnetic field (from both magnetic gradients and radio frequency pulses) that inductively induces the voltage difference that is measured by the EEG system, *Equations 2, & 3*. Over the course of a recording, B_0 is considered to be static and thus no time dependence is associated with it. B_1, the radio frequency field, is used to bring the spins to their

excited state , but the induced artifact can be managed relatively simply, as. the typical B_1 frequencies of tens to hundreds of megahertz allow it to be eliminated by passive hardware filters in the EEG isolation system prior to the instrument amplifier. These filters serve the purpose of keeping radio frequency noise from both entering and exiting the EEG system.

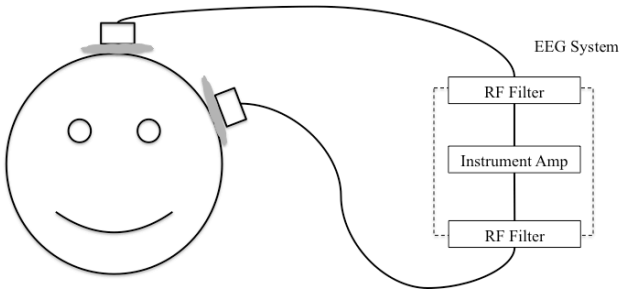

Figure 1: A single conductive loop of the EEG System - Starting at the scalp, the loop passes through a conductive medium (typically a silver chloride gel) to the measurement electrode. Next, it passes through the electrode lead, then through the EEG hardware before returning back through the reference electrode lead, through the reference electrode, then again through a conductive medium to the scalp and through the body tissue back to its starting point in the scalp.

Equation 1: Faraday's Law of Induction

$$\mathrm{EMF}_{\mathrm{electrode}} = -\frac{d}{dt}\int_{surface} B(t,x,y,z)\bullet dA(t),$$

where $B(t,x,y,z) = B_0 + B_1 + G_X(x,t) + G_Y(y,t) + G_Z(z,t).$

Equation 2: Serial Inductance Path of a Conductive Loop of the EEG system. The serial inductance in (a) is broken up into (b) static and (c) dynamic components

(a) $Z_{\mathrm{Electrode\ Loop}}(t,w)=Z_{\mathrm{Dynamic}}(t,w)+Z_{\mathrm{Static}}(w).$

(b) $Z_{\mathrm{Static}}(t)=Z_{\mathrm{Subject}}(w)+Z_{\mathrm{Channel\ Electrode}}(w)+Z_{\mathrm{Channel\ Lead}}(w)+Z_{\mathrm{Instrument\ Amp}}(w)+Z_{\mathrm{Reference\ Lead}}(w)+Z_{\mathrm{Reference\ Electrode}}(w).$

(c) $Z_{\mathrm{Dynamic}}(t,w)=Z_{\mathrm{Channel\ Electrode\ Interface}}(t,w)+Z_{\mathrm{Reference\ Electrode\ Interface}}(t,w).$

Equation 3: Approximating the series inductance of the conductive loop as a series resistance over the narrow band frequencies of EEG.

$$Z_{\mathrm{Electrode\ Loop}}(t,\omega) = Z_{\mathrm{Dynamic}}(t,\omega) + Z_{\mathrm{Static}}(\omega) \rightarrow R_{\mathrm{Electrode\ Loop}}(t) = R_{\mathrm{Dynamic}}(t) + R_{\mathrm{Static}}.$$

Equation 4: Appling Ohm's Law to calculate the measured voltage of GA by the EEG system.

$$sGA_{\mathrm{Electrode}} = R_{\mathrm{Instrument\ Amp}} \times \frac{EMF_{\mathrm{Electrode\ Loop}}}{R_{\mathrm{Electrode\ Loop}}}.$$

To separate/clean the GA signal from the EEG signal, a domain must be found in which the two are distinct and identifiable. In the spectral domain, there is overlap in the content of the GA and that of the EEG signal (Bénar et al., 2003). This overlap is large enough that it all but eliminates the possibility of using a simple spectral filter. In the time domain, however, the two can be teased apart. This is because the GA ideally is identical in each TR segment — the time between MR measurements, i.e. the time between radio frequency excitation pulses — while the EEG is not. In fact, through averaging, the EEG, with its zero-mean distribution, should asymptotically approach zero. TS works by exploiting this difference, the time locked repeatability in the GA and lack thereof in the EEG, to tease the two apart. The GA signal (sGA) can be removed from the true EEG (sEEG) by constructing a TR based windowed average, and then subtracting this averaged waveform from each TR. By averaging the raw EEG signal across TR windows, the GA component should remain unchanged while the EEG component should approach zero, leaving a perfect temporal domain template of the GA signal. This GA template then can be subtracted from the raw EEG recording (sRec) leaving the sEEG, *Figure 2*.

While for the most part the sGA is largely the same between TR segments, there are perturbations in the signal both in endogenous and exogenous forms. In addition, the typical assumption that sRec, *Equation 5*, is composed of only three signals is violated in real world data. The signal at the input to the EEG amplifier is actually a linear combination of at least 5 signals: the sGA; the sEEG; the ballistocardiogram (sBCG) (Mullinger et al., 2013); the cryopump vibrations (sCryo) (Bénar et al., 2003); and multiple other sources including gross subject movement (GM), expressed as ε, *Equation*

6a. To illustrate the relative magnitudes of all of the signal components of sRec, the amplitudes of each component are expressed by font size in *Figure 3*.

We consider endogenous perturbations to the GA signal to include changes in electrode impedance, signal drift, and drift between the clocks of MR and EEG over the recording interval.

Exogenous changes, on the other hand, have to do largely with changes in the conductive loop area and its orientation. Exogenous changes can take the form of oscillatory motions, such as the scalp pulsation portion of the sBCG and sCryo where the start and end positions of the motions are likely in the same place; and displacement motions, such as (GM) and head rotation and translation portions of the sBCG where the start and end positions of the motions are in different locations.

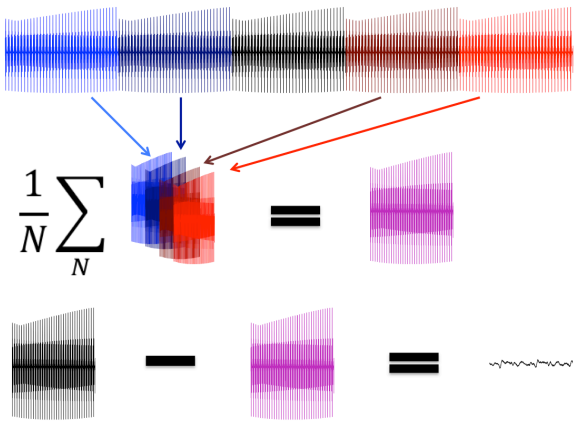

Figure 2: A cartoon representation of TS – Top Row: Here, the TR window being cleaned of the GA is the center TR segment (black section). Middle Row: The average of flanking TR windows, two on either side is used to create the artifact template (magenta). Bottom Row: The artifact template then is subtracted from the TR window being cleaned, leaving behind only the EEG signal.

The oscillatory motions by returning the electrode loop to its approximate starting position induce only a transient signal superposed onto the already existing signals, but do not change the long-term GA signal. The displacement motions are problematic in that

they both induce a large transient signal, enough so that it can cause clipping in the analog to digital converter of the EEG system, and because they can lead to altered GA signals before and after head motions. Clipping is a problem, because when it occurs the superposition of artifact and signal components in the EEG becomes non-linear once the receiver becomes saturated. The EEG signal during time points in which clipping occurs cannot be recovered. The change in sGA pre and post motion is because the GA is specific to both the loop area and its orientation with respect to the magnetic field. If the starting and stopping positions of the loop differ, so to the sGA on before and after the motion. The overall signal *timing* of the GA is essentially independent of motion, but its morphology will change. Variance in the shape of the GA as a function of the electrode contact positions can be seen in *Figure 4*.

Equation 5: Simplistic form of the components of the recorded EEG signal during an MR aquisistion. The recorded signal by the EEG system (sRec) is composed of the gradient artificat signal (sGA), the true EEG signal (sEEG), and all other noise (ε).

$$\mathrm{sRec}[n] = \mathrm{sGA}[n] + \mathrm{sEEG}[n] + \varepsilon[n].$$

Equation 6: Detailed form of the components of the recorded EEG signal durring an MR aquisistion and the process for template subtration. In (a), the recorded signal by the EEG system (sREC) is composed of the gradient artificat signal (sGA), the ballistocardiogram signal (sBCG), the true EEG signal (sEEG), the induced signal from vibrations from the cyropump (hellium pump) of the MRI scanner, and all other noise (ε). In (b) and (c), the recorded signal is broken up into two parts, sGA and the other signals (sO). In (d), the formulation of the sO is in terms of sRec - sGA

(a) $\mathrm{sRec}[n] = \mathrm{sGA}[n] + \mathrm{sBCG}[n] + \mathrm{sEEG}[n] + \mathrm{sCryo}[n] + \varepsilon[n].$

(b) $\mathrm{sO}[n] = \mathrm{sBCG}[n] + \mathrm{sEEG}[n] + \mathrm{sCryo}[n] + \varepsilon[n].$

(c) $\mathrm{sRec}[n] = \mathrm{sGA}[n] + \mathrm{sO}[n].$

(d) $\mathrm{sO}[n] = \mathrm{sRec}[n] - \mathrm{sGA}[n].$

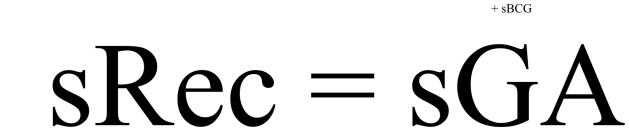

Figure 3: Cartoon figure of the respective component signal magnitudes represented by font size.

Signal amplitude scaling occurs as a result of changes in electrode impedance that cause a slow modulation of the recorded EEG signal. In *Equations 2, 3 & 4* note that as the interface impedance increases, the GA signal magnitude measured by the EEG system is reduced. The interface impedance increase likely can be attributed to drying of the electrode gel or salt solution (Ferree, 2000). Scaling poses a problem in the creation of the GA template because the amplitude of the GA becomes unstable over time.

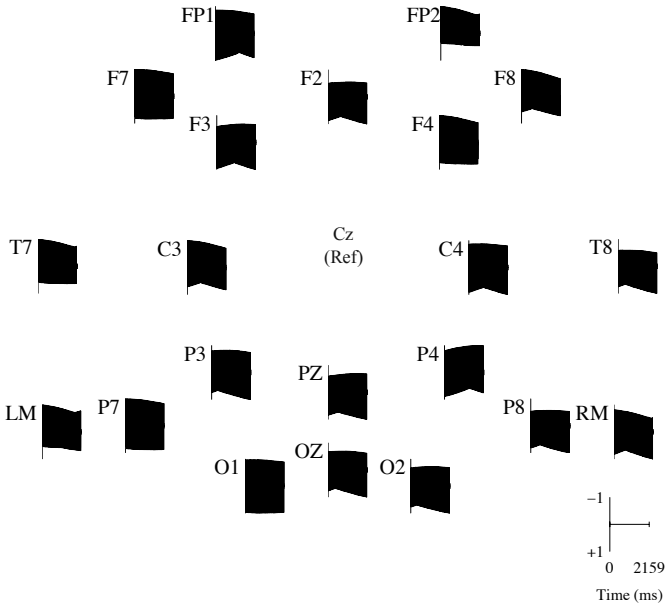

Figure 4: Topographic plot of the envelopes of a normalized average GA artifact across the scalp. Note that for the most part, adjacent electrode's envelopes differ. This is likely the result of differences in the electrode lead paths and conductive paths through the subject volume.

Signal drift occurs as the mean signal about which upon which the EEG is superimposed changes over time (Simons et al., 1982). This is referred to often as the "battery effect." It is presumed due to a polarization of electrodes from an electrochemical effect of metallic substance (the electrode itself) being in contact with an electrolyte (the gel or salt solution). The ions in the electrolyte migrate to the metallic surface of the electrode causing charge to accumulate. The drifting occurs because the

charge rate accumulation is not necessarily (and most likely not) the same for the channel electrode and the reference electrode. This difference creates the "battery" in battery effect, the slow drift of the baseline. A simple centered high-pass filter in post-processing can prevent very slow drifts from causing problems with the GA template creation. There is loss of sub Hz frequencies in the true EEG; Teasing these signals apart from impedance drift is beyond the scope of our current work.

The recording of clock drift consists of two components. The first is that if the duration of 1 TR is not an integer multiple of the EEG sampling dwell time, then the sGA within one TR windowed segment will not match exactly the sGA in the temporally adjacent TR window segments. Second, even TR is an integer multiple of the dwell time, there is also the requirement that the clock of the MR system, and the clock of the EEG analog to digital converter are entirely synchronous. The peak error in the GA from TR to TR due to clock slippage is equal to the product of the maximum value of the derivative of the sRec, *Equation 7a*, with the size of the clock slippage, *Equation 7b*. The slippage error can lead to differences in the GA between TR windows, which will translate into errors in the GA template. This mismatch will cause a degradation in the separation of the sGA and so The clock slippage issue is handled in one of two ways, either by the Cohen (Cohen, 2000) or Allen (Allen, 1998) methods. The Cohen method uses a phase locked loop driven by the MR scanner's master clock to drive the timing of the analog to digital (A/D) converter of the EEG system. By altering the EEG's A/D clock, a phase shift is introduced into each TR. This phase shift however is small, and for most purposes ignored. The Allen method uses both high sampling rates (small sample periods) coupled

with up-sampling (interpolation) and Fourier space phase alignment to line up the TR windows.

Equation 7: Formulas for (a) calculating the maximum rate of signal change of the EEG signal and (b) the potential maximum error from slippage.

(a) Maximum Rage of EEG Signal Change = $\operatorname{argmax}\left(\frac{d}{dn} s \operatorname{Re} c\right)$.

(b) Peak Value of Slippage Error = slippage time $\times$ $\operatorname{argmax}\left(\frac{d}{dn} s \operatorname{Re} c\right)$.

To deal with short timescale modulations to the GA from exogenous changes and electrode impedance changes, a moving windowed average, rather than the average of all TR segments, is used to create the template. A cartoon illustrating this can be found in *Figure 2*. This allows the GA template to adapt to the local sGA. The caveat, however, is that although the EEG has "zero mean", it converges to this value as the square root of the number of averaged segments. There therefore is a tension between the GA template's fidelity to the local sGA structure and the template's independence and separability from the sO and sEEG with it. All of this is linked to the number of TR segments included in the moving average (Ritter et al., 2007). The as more TR segments are used to create the GA template, the template becomes more independent of sO, resulting in better preservation of the EEG signal. When fewer TR segments are used the slow variations in the sGA are not embodied in the template, leading to a better removal of the sGA. In other words, the smaller the number of included TR segments, the better the sGA is removed from sRec, but at the same time, more non GA signal is added back into the sRec. The more TR segments included, the worse the template subtraction is able to adapt to the changing sGA and the less it is removed from sRec, but also the less non GA signal is added back into the sRec.

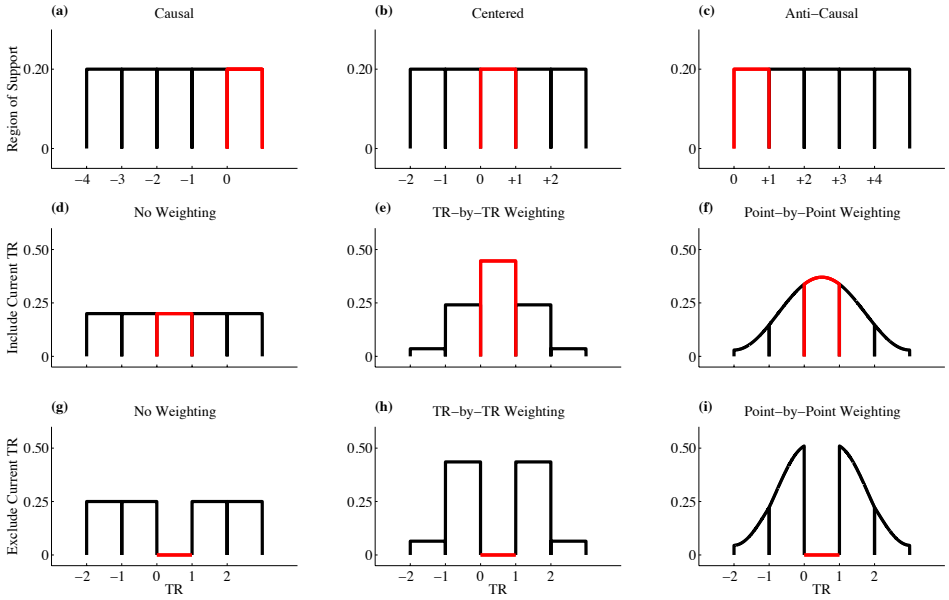

Figure 5: Compilation of the many different options for constructing the GA template. The red regions represent the TR segment to be cleaned and the black regions represent adjacent TR segments. The y-value indicates the relative contribution of each point to the GA template. In all scenarios, these points sum to 1. Top Row: The possible regions of support for the MAWTS filter: (a) causal; (b) centered; and (c) anti-causal. Middle Row: Centered filters with the various weighting schemes, which include the TR segment to be cleaned in the template: (d) no weighting; (e) TR-by-TR weighting; and (f) Point-by-Point weighting. Bottom Row: Centered filters with the various weighting schemes, which exclude the TR segment to be cleaned from the template: (d) no weighting; (e) TR-by-TR weighting; and (f) Point-by-Point weighting. All of the filters depicted are defined as 5-segment, 5-point, or length 5 filters regardless of whether they include the TR segment to be cleaned.

We present here a careful analysis of the many options for implementing a moving windowed average template subtraction filter (MWATS). The options we consider are: whether to implement a causal or centered filter, *Figure 5a & b respectively*; the type of weighting that should be applied to the filter, *Figure 5d-i*; how the weighting should be applied, *Figure 5e,f,h & i*; whether to include the to-be-cleaned TR in the template, *Figure 5d-i*; and the number of TR segments to be included in the template.

- A *causal* filter is a filter uses only data collected prior to, or including, the current time point. By contrast, an anti-causal filter is a filter that uses future inputs and it,

too, may also use the current input. A symmetric centered filter depends equally on past and future inputs and, this too may also include the present.

- The moving average is weighted for two reasons. Moving average filters are essentially a convolution in the time domain, decreasing the weighting of more distant events in time tends to make such filters smoother in the spectral domain (because abrupt changes in the weighting at the filter edges has the effect of adding energy at high temporal frequencies). The second reason is to put a stronger emphasis on the local structure of the GA when creating the template. Weighting typically is accomplished in one of two manners: point-by-point or TR-by-TR. In a point-by-point method, the weighting is applied on a sample-by-sample basis ignoring TR window bounds. In TR-by-TR methods, the weighting is applied equally across all points within a TR segment, but varies between TR segments.
- The inclusion in the template of the to-be-cleaned TR segment, *Figure 5a-f*, is intended to put increased emphasis on the local structure of the GA. The closest gradient artifact segment to the one being cleaned is itself. By adding it into the moving windowed average, it is thought, by some, to allow for a better matching of the template to the TR being cleaned. As of the publication of this chapter, some commercially available software implementations of the MWATS use this approach.
- Changing the filter length—the number of TR segments included in the moving average—is done much like weighting to change how much emphasis is placed on the local structure of the GA. In all of the examples in *Figure 5*, the window

length is defined as a five segment or five point filter regardless of whether they include the TR to be cleaned.

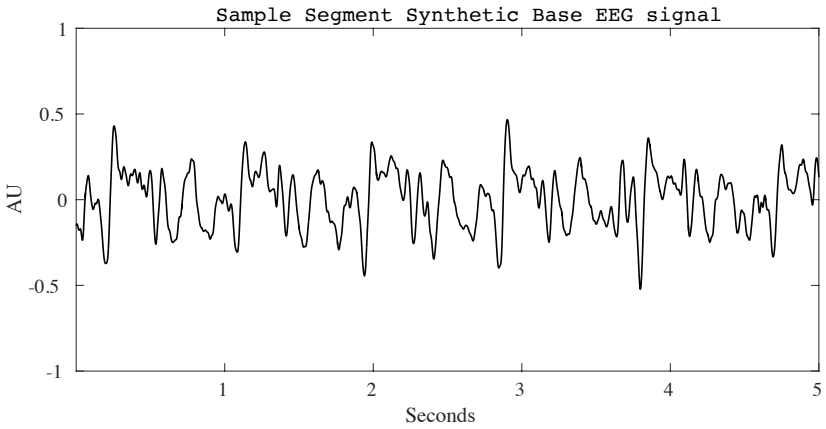

Figure 6: A five second sample of the synthetic EEG base signal, 1 - 50Hz band-pass filtered and normalized. The original EEG recording was collected from a human subject at isocenter in the MRI scanner with a sampling rate of 1kHz. The cryopump was turned off, and there was no active MR scan during this recording. The total length of the recording was 5 minutes 47.6 seconds. The base signal was taken from the Oz channel with a Cz reference. The band-pass filter was applied to avoid introducing the signal drift and line noise from this recording into any new recording. The amplitude units are arbitrary because all of the signal scaling happens in the attenuation circuit.

To test the various MWATS options for their effects on both the fidelity of sO and the separability of sGA from sO, we created a ground truth EEG signal. In their 2007 paper, Ritter et al., attempted to create a ground truth by adding standing sine waves into an already acquired EEG recording that was contaminated with GA. To build upon this, and to better represent sO, we developed a means to generate a synthetic signal, which we collected on a phantom placed in the MR scanner. The synthetic signal was recorded by the EEG system and had to travel through the entire EEG acquisition pipeline. This ensured that sGA and sO passed through the same filters, (RF and anti-alias) and analog to digital conversion. More importantly both signals were affected equally by changes do to baseline signal drift and signal scaling that are the result of impedance and polarization changes at the electrode interfaces. The content of our synthetic signal was made to be more like that of EEG acquired in an MR scanner. The

reference (base) of the synthetic signal was generated from an EEG recording from a subject in-magnet with the gradients and cryopump switched off. This recording was then band-pass filtered from 1-50Hz to avoid reintroducing recorded line noise and signal drift in the synthetic EEG recordings, *Figures 6 & 7*.

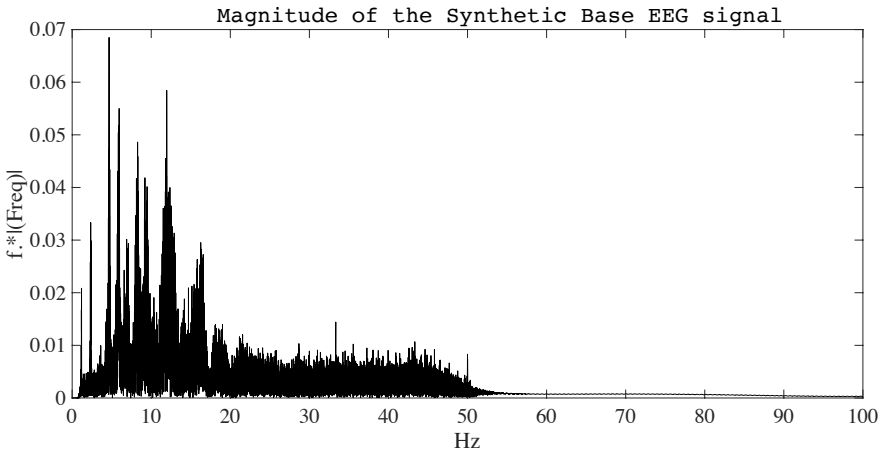

Figure 7: Spectral amplitude plot for the base signal for the synthetic EEG post 1-50Hz band-pass filtering. The original EEG recording made using a human subject at isocenter in the MRI scanner with a sampling rate of 1kHz. The cryopump was turned off and there was no active MR scan during this recording. The total length of the recording was 5 minutes 47.6 seconds. The base signal was taken from the Oz channel with a Cz reference. The band-pass filter was applied to avoid introducing the signal drift and line noise from this recording into any new recording. The spectrum is corrected for 1/f effects by multiplying the magnitude value by the corresponding frequency.

We created the ground truth to test various cleaning options, by making three consecutive recordings each using the base signal above. The two flanking recordings were acquired using the base signal with no gradients and no cryopump. These recordings then were detrended and averaged to create a ground truth signal for the middle recording. This was done to compensate for any signal amplitude scaling from impedance changes at the electrode interfaces. The middle recording was acquired using the base signal with the cryopump off but during an MR acquisition (the gradients were on).

We evaluated the various MWATS options for sO signal fidelity post cleaning and the signal separability of sGA and sO using the synthetic signal recordings. We quantified the fidelity of the sO post MWATS to the ground truth signal as the root mean

squared error (RMSE) between the two signals. The smaller the RMSE measure the closer the cleaned signal matched to ground truth signal. Separability of the signals, the effectiveness of the removal of the GA from the recorded signal, was measured by calculating the mean energy per sample point in a TR windowed average over the entire cleaned signal. Ideally, any component not time locked to the TR should converge to zero in this windowed average. The smaller the value of this residual energy per sample (REsamp) the better the MWATS performed at separating sGA and sO.

In addition to the experimental results we performed an analytic analysis of the MWATS. We did this to better understand: its frequency response characteristics, and the signal pollution from the template subtraction operation. This aided us in the interpretation of the synthetic signal cleaning results.

Materials & Methods:

All of the code mentioned below for implementing the MWATS filter, and for generating the synthetic EEG signal; the phantom datasets with the synthetic signal; the information for constructing the signal attenuator circuit (circuit schematic, board file, a build of materials) are available at www.brainmapping.com/LINT/crodriguez.

EEG Recording System

All EEG recordings were performed with Electrical Geodesics Inc.'s (EGI, Eugene, OR) GES300 MR system amplifier and recorded using Net Station v4.54 recording software. All recording were made with a sampling rate of 1kHz and Cz referenced signal. EGI's MR compatible Hydrocel™ 256 channel net was the electrode system used for all recordings. Nets were positioned on the human participants by

aligning the vertex electrode with the vertex of the head, identified at the midpoint of the anion-to-inion and left/right preauricular landmarks. Electrode impedances were <50 kΩ, as recommended by the manufacturer. Two additional external electrocardiogram (ECG) electrodes were applied to the chest of the human participants (positions: left midclavicular line in 5th intercostal space and 4th left intercostal space at sternal border).

MRI System and Scan Settings

All in-magnet EEG recordings were collected at isocenter in a 3 Tesla Siemens Trio MRI scanner (Erlangen, Germany). The following are the settings for recording where the MRI scanner was active: The receive Siemens 12 Channel Head Coil was used as a receiver, and the transmit signal was generated using the body coil. 36 slices in a 64x64 matrix of 3mm isotropic voxels were acquired at a TR = 2.160sec echo time (TE) = 30msec, and at a flip angle of 90°. Slices were prescribed to allow for full volume coverage.

Data Processing

All signal processing and analysis was performed on a 2012 MacBook Pro Retina, OSX 10.10.2, 2.7 GHz Intel Core i7 with 16Gb of ram (Cupertino, CA), running Matlab™ 2012a with psychtoolbox (Brainard, 1997) version 3.0.12, and EEGLAB (Delorme and Makeig, 2004) version 13.4.4b.

Creating the Synthetic EEG + BCG Base Signal

We generated the Synthetic signal from an EEG recording with the subject at isocenter, head first and supine, in the bore of Siemens 3 Tesla Tim Trio MRI scanner. The cryopump was switched off, and no MR scanning took place during the EEG acquisition. The subject comfort fan, as well as any unnecessary electronic devices, were also turned off or unplugged. The recorded signal from the Oz channel with a Cz reference was extracted. This extracted signal was then zero-phase band-passed filtered 1-50Hz, *Figure 5 & 6*. The low-pass was implemented as a 12^{th} order Butterworth filter and the high pass portion was a 6^{th} order Butterworth filter. Frequencies below 1Hz were removed to eliminate from the synthetic signal. The removal of frequencies above 50Hz likewise was used to eliminate line noise of 60Hz as well as its harmonics from the synthetic signal. The recording was made with a sampling frequency of 1kHz and had a recording length of 5 minutes 47.6 seconds.

Subjects

Participants received verbal and written explanations of study requirements, and prior to any study procedures, were provided written informed permission/assent as approved by the UCLA Institutional Review Board (respective Protocol IDs Follow). An EEG recording from a male subject, age 27 was used to generate the base for the synthetic signal, IRB#11-002946 An EEG recording from a female subject, age 21, was used for the analytic analysis, and to generate the GA template used in the figures, IRB#12-000266.

Synthetic EEG + BCG Signal Generation

We generated the signal using the data processing computer listed above. The base signal recorded with a 1kHz sampling frequency up-sampled to 48kHz. The up-sampled signal was then low-pass filtered at 100Hz with a 12^{th} order Butterworth filter. The filtered signal was played through the headphone jack of the computer. The audio signal was carried over a fiber-optic Toslink repeater system from the console room to scanner room and then into a custom-built attenuator circuit. The output of the attenuator circuit was connected to a phantom. The phantom consisted of a honeydew melon (Ralphs's grocery, Santa Monica CA) with a towel placed on top of the melon, and an EEG net paced on top of the towel. We selected the melon for its approximate match to the size of a human head. The circumference of the melon with a towel on top was measured to be 58.7cm. The was towel soaked in KCl at the concentration used for the setup of EGI's EEG net. The connection of the attenuator circuit to the phantom was made using a 20' coaxial copper cable terminated on the phantom end with a pair of nickel plated brass buttons, size 2, (Prym Consumer USA, Inc., Spartanburg, SC). The mating portions of the buttons were attached to the towel using 24 AWG solid core aluminum wire. One lead was placed near Oz and the other near Cz on the phantom. The nominal signal level measured by the EEG system ±22μV with a maximum peak-to-peak of -168.23 to 143.44μV.

Recording the transmitted Synthetic EEG + BCG Signal

We recorded the synthetic signal with the same settings used in collecting the reference signal. Three recordings were made back-to-back. In all three the cryopump was turned off. In the second recording an MR scan was performed with the acquisition

settings listed above. We aligned the first and third (flanking) recordings, with no gradient artifact, using cross correlation and transistor-transistor logic (TTL) pulses, then calculated their average. We used this average signal as an estimate of ground truth for the recording made during the MR acquisition. We took the average to compensate for any impedance changes in the electrodes between the acquisitions with and without the MR scanner running.

Analyzing the Synthetic EEG + BCG Signal

We analyzed the cleaned signal using 2 metrics: *separability* and *fidelity*. Separability is defined here as how effectively the sGA is removed from the recording, and fidelity is defined as how closely the cleaned sO (of which sEEG is a component) approaches the true sO. On the extremes, a perfect separability would result from zeroing out the sRec, removing sGA completely, but would also result in zero fidelity of the EEG signal. On the other extreme leave sRec unaltered, in which case there would be zero separability but perfect fidelity. To characterize the separability of the two signals (sGA and sO) we measured the total energy in the all TR windowed average after GA cleaning. We then divided the residual energy by the number of samples in a TR in segment, to determine the residual energy per sample (REsamp). This normalization allows this measure to be independent of the TR length. Ideally, with enough averages, anything that is not time locked to the TR approach zero. The components that do not are likely the residuals left over from imperfect cleaning. To characterize fidelity we measured the root mean squared error (RMSE) between the cleaned signal and ground truth signal above.
A causal filter with window length of five TRs would not start cleaning until the 5^{th} TR segment and would clean until the very last TR segment. Likewise a centered filter with a

window length of 17 would not start cleaning until the 8th TR segment and would clean only until the 8th to last TR segment. To account for this, we used only the TR segments that would be cleaned by all the considered filters in forming the RMSE and REsamp metrics.

Results from the Synthetic EEG acquisition:

Overall data quality

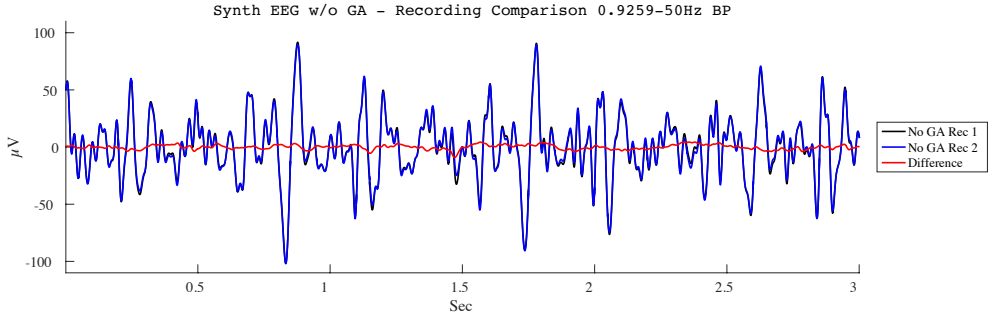

Figure 8: A three second snippet of the two synthetic EEG recording made without aligning GA. Note the slight magnitude differences in the two signals. The first recording, on average, had a slightly larger amplitude than the second. Drying of the electrodes, and consequently the increase loop impedance likely accounts for some of this mismatch.

The overall data quality and reproducibility of the synthetic EEG recordings without GA was good. The RMSE between the two, over the entire recording in the frequency band of 0.9259-50Hz, was 2.3741 μV or 7.9% of the mean of the standard deviation of each of the signals. The reason for using this restricted spectral band, rather than the entire frequency band of the recording, was, as stated before, to avoid issues with asynchronous line noise and baseline drift between the two recordings. Using the peak-to-peak measurements there was an overall decrease in signal amplitude of 1.7963% between the early and late recording. This decrease was likely due to impedance changes at the electrode interface, e.g., electrode drying. The increased impedance likely can account for the decrease in magnitude of the signal, and for some the mismatch between the two recordings, red trace, Figure 8. In the difference signal between the synthetic

EEG recordings without GA, Figure 9d, the spectral energy appears more or less unforminly distributed in the 1/f corrected spectrum. There are distinct peaks though at 12.66, 17.51, 30 and 60Hz. The former three peaks only appear distinct in the difference signal and the 60Hz peak is distinct in both the average, and difference signals, Figures 9c & d. None of these peaks are distinct in either of the synthetic EEG recordings without GA, Figures 9a & b. This likely points to the fact that they are induced signals from a standing noise source(s) who's asynchrony between the two recordings accentuates them in the difference signal. The signal scaling, standing noise asynchrony, as well as differences in the random noise, likely account for the majority of the difference in the two recordings.

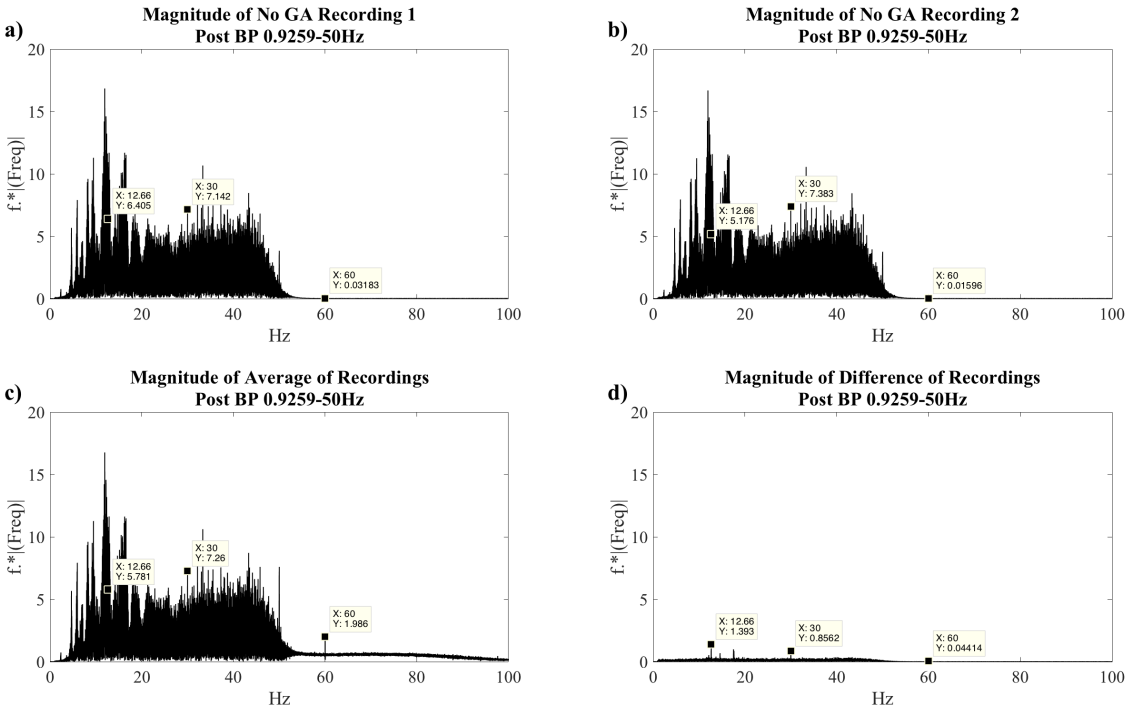

Figure 9: Magnitude spectrum of the difference between the first and second synthetic EEG recordings without GA post band-pass filtering from 0.9259-50Hz. Here like in Figure 5 the spectrum is corrected for 1/f effects by multiplying the magnitude value by the corresponding frequency. Note the distinctive peaks at 12.66, 17.51, 30 and 60Hz. The former three are only distinct in difference signal while the line noise (60Hz) is distinct in both the difference and average signals. They are likely the result of asynchrony in a standing noise source(s) between the two recordings.

Comparison of Causal Versus Centered Cleaners

Table 1: Simple moving average filter comparison of causal and centered filter as well as the effect of including the TR to be cleaned in the formation of the template. For comparison the RMSE between the two EEG recordings with no GA over the same region used in the analysis was 2.3769 µV.

Filter Type	Causal	Causal	Center	Center
Window Length	4	5	5	5
In TR to be cleaned in template	Yes	No	Yes	No
# TRs Windows in Template	4	4	5	4
RMSE (µV)	15.6911	16.3336	13.9386	15.8221
REsamp (µV^2/Sample)	0.1003	0.2296	0.0318	0.0497

In making a comparison of causal versus centered filters, *Table 1*, we found the centered filters to outperform the causal filters in terms of residual energy per sample whether or not they included the TR to be cleaned in the template. In terms of the RMSE, the fidelity measure of the cleaner, the centered filters also outperformed their causal counterparts. In the case of the centered filter that included the TR to be cleaned in the template, the effective number of TR averages was five, likely leading to an artificially lower RMSE as compared to the other filters, each of which all 4 averages. Using the centered filter which did not include the TR to be cleaned in it template, a like-for-like comparison with regards to the number of averages can be performed. Here the RMSE is close to, but not as good as, its causal counterpart which included its TR to be cleaned in the template but outperformed the causal filter which did not. Measured by the residual energy per sample, a measure of separability, the centered filter far outperformed its causal counterparts. In this case there was a 50.45% reduction in REsamp over the causal filter which included the TR to be cleaned in the template, and a 78.35% reduction in the one that did not. In summary the centered filters outperform their causal counterparts in terms of separability. In terms of fidelity the centered filter which did not include the TR to be cleaned in its template outperformed its causal counterpart, however this was not

the case when comparing it to the causal filter which included the TR to be cleaned in its template.

Comparison of Various Window Lengths and Weighting Types

Table 2: Root mean square error (RMSE) versus the number of TR segments included in the template and the weighting type applied. All calculations were made on a centered filter with a TR-by-TR weighting scheme. The TR to be cleaned was not included in the template formation. For comparison the RMSE between the two EEG recordings with no GA over the same region used in the analysis (cleaned) was 2.3684 µV. The weighting types analyzed were: simple moving average (SMA), exponential (EXP), Gaussian (GAU), Hanning (HAN), Hamming (HAM), Blackman (BLK) and a flat top (FTP).

RMSE (µV^2)								
Weight		SMA	EXP	GAU	HAN	HAM	BLK	FTP
Window Length	5	16.084	17.627	22.268	22.268	19.747	22.268	22.107
	7	13.571	16.579	17.091	17.858	16.376	19.031	30.709
	9	12.177	16.281	14.684	15.302	14.271	16.361	27.331
	11	11.217	16.189	13.240	13.711	12.917	14.613	23.582
	13	10.461	16.159	12.255	12.626	11.970	13.404	20.907
	15	9.885	16.150	11.523	11.829	11.264	12.517	18.970
	17	9.445	16.147	10.949	11.209	10.708	11.833	17.510

Table 3: Residual energy per sample point versus the number of TR segments included in the template and the weighting type applied. All calculations were made on a centered filter with a TR-by-TR weighting scheme. The TR to be cleaned was not included in the template formation. The weighting types analyzed were: simple moving average (SMA), exponential (EXP), Gaussian (GAU), Hanning (HAN), Hamming (HAM), Blackman (BLK) and a flat top (FTP).

REsamp (µV^2/Sample)								
Weight		SMA	EXP	GAU	HAN	HAM	BLK	FTP
Window Length	5	0.055	0.050	0.052	0.052	0.050	0.052	0.052
	7	0.075	0.052	0.051	0.050	0.052	0.050	0.066
	9	0.089	0.053	0.058	0.055	0.059	0.052	0.060
	11	0.104	0.053	0.066	0.063	0.068	0.058	0.054
	13	0.123	0.053	0.075	0.071	0.077	0.064	0.051
	15	0.140	0.053	0.084	0.080	0.087	0.071	0.050
	17	0.154	0.053	0.093	0.089	0.096	0.079	0.052

Following the causal versus centered analysis we measured the separability and fidelity for various filter lengths and weighting types. Here all the filters were centered and excluded the TR to be cleaned from the template. The simple moving average for a

given window length provided the lowest RMSE values, *Table 2,* but had the poorest performance in terms REsamp, *Table 3*. Conversely the flat top and exponential weightings provided some of the poorest performance in terms of maintaining the fidelity of sO, but some of the best performance in terms of separating the sGA and sO signals. Plots corresponding to the measured RMSE and REsamp values in Tables 2 & 3 can be found in Figure 10a and b respectively. Stated differently, if the goal was to maintain the fidelity of the sO and the sEEG, then the simple moving average provided the best results. If the goal was separation of the sGA and sO signals, i.e. the removal of sGA, then the exponential and flat top weightings provided the best results. The other weighting types threaded the middle with regards to separability and fidelity. As an example, the Hamming weighting trades slightly decreased performance in terms of fidelity with increased performance in terms of separability in comparison to the flat weighting of the simple moving average.

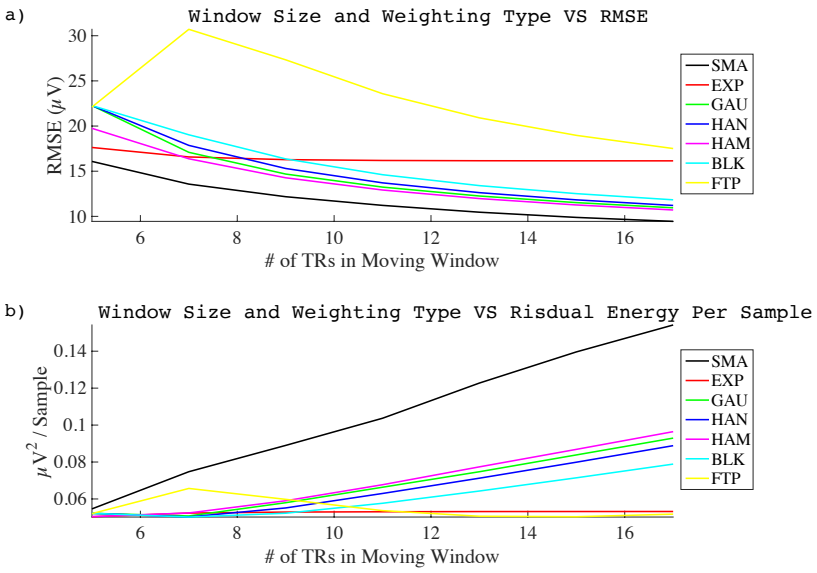

Figure 10: Plots of The RMSE (a) and RE per Sample Point (b) Versus the Number of TR Segments Included in the Template and the Weighting Type Applied. All calculations were made on a centered filter with a TR-by-TR weighting scheme. The TR to be cleaned was not included in the template formation.

Discussion:

We will be using three time series in discussing the examined options for MWATS, they are: TR-by-TR (TR #), point-by-point ($TP_{m,n}$), and sample number (S_k). These series are illustrated in Figure 11. The moving windowed average (MWA) can be thought of operating in each of these time series as follows: In the TR-by-TR frame MWA operates by creating an average TR window, *Figure 2*. In the point-by-point frame MWA operates on a single time point within the TR across windows, *Figure 11c-f*. In the sample number frame MWA operates across the section of the recording containing the GA, *Equations 10 & 11*.

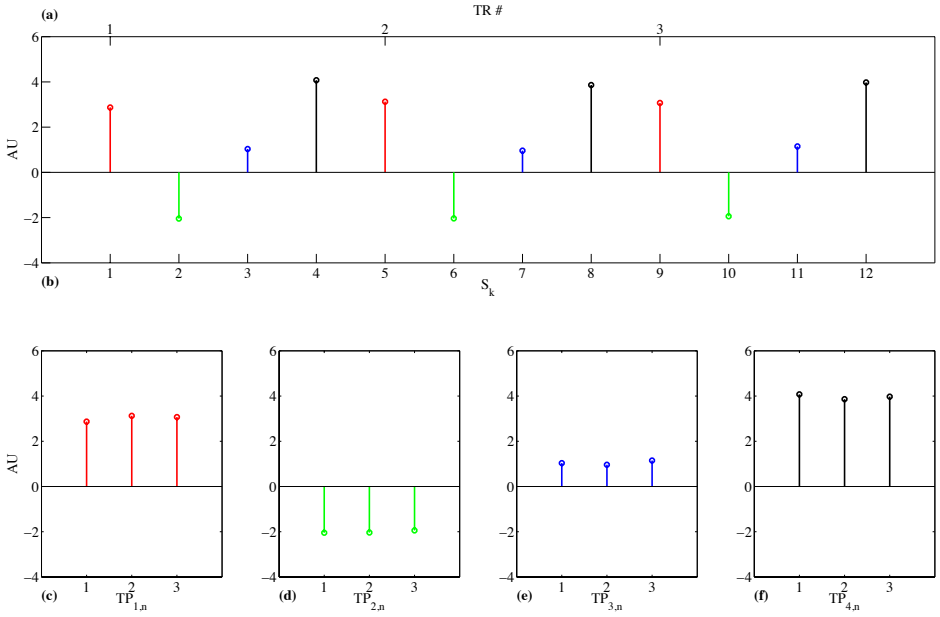

Figure 11: An illustration of three different time series and how they are related: (a) TR-by-TR (TR #), (c-f) Point-by-Point ($TP_{m,n}$), and (b) sample number (S_k). TR-by-TR looks at the entire TR windows as a block consisting of fixed amount of samples (M = TR X Sampling Rate). Each of these samples constitutes a time point within a TR segment ($TP_{m,n}$). Here "m" denotes the time point within a TR window and ranges from 1 to M where there are M sample points per TR. The "n" denotes the TR from which the time point came and ranges from 1 to N, N being the number of TR windows fully captured in the recording. Point-by-Point looks a specific time point "m" across TRs. The last time series is the total sample series S_k. K denoting the sample position which runs from the first to last sample in the recording.

Should a Causal or Centered Filter be used?

In considering a causal versus centered filter it is worth mentioning there are two scenarios to be considered, real time and post processing cleaning. For real time cleaning the filter must be a causal as there is no knowledge of future events, thusly what we consider here is in terms of post processing cleaning.

In *Figure 12* causal and centered filters are considered in the point-by-point series using a simple 3-point filter. X(t) being the $TP_{m,\_}$ going into the filter and H(t) being the impulse response of a simple moving average filter. Here X(t) should be thought of as modulations to the GA as a whole. An example of a modulation to GA would be increasing impedances of the electrode interfaces over time from drying out. This will increase the overall loop impedance, *Equations 2 & 3*, and therefore decrease the overall magnitude of the GA over time, *Equation 4*. A simple moving average causal filter is *center top* and a simple moving average centered filter is *center in Figure 12*. The result, Y(t), is on the right hand side. The centered and causal filter types both preserve amplitude, but the output after the causal filter is shifted to the right in time, i.e. to the future, of both the centered filtered data and X(t) itself. This shifting to the future, lagging the output, creates mismatch between the causal filtered data and X(t) which plays out in TS as an error in the GA template, bottom left plot of *Figure 12*. While neither the causal or centered filter give a perfect fit to X(t) as both act as a form of low pass filter, the causal filter performs worse as it creates a lag. In the sample number series the derivations of both causal and centered MWATS can be found in *Equations 8 & 9* respectively. The takeaway message from the two derivations is that there is no phase angle, *Figure 13d*, associated with the centered filter thus no lag.

In considering the results from *Table 1*, a sizable difference can be seen in the separability (REsamp) between the causal and centered filters which did not include the TR to be cleaned in the template, 0.2296 and 0.0497 μV/sample. The more effective reduction in GA in the centered filter likely accounts for its slight improvement in RMSE over the causal filter as well. In summary the centered the outperforms the causal filter.

Equation 8: Mathematical formulation of a causal comb filter. (a) In the time domain a comb filter adds weighted lagged samples to the current sample. Here the lag "*k*" equals the length of a TR segment in samples and the number of lags added equals the width of the window in TRs, *α* is the weighting function applied, e.g. for a simple moving average α = -1/(WL-1) for all lags. It is also worth noting that this is almost an exact representation of MWATS. The difference is in "*k*". This is because the lag changes from one TR to the next due to the clock synchronization of the EEG system to MRI system. The variability of k will be on the order of 1e-3 for sampling rate of 1kHz. In (b) is the z-transform of (a). In (c) is the formulation of the impulse response function "H[z]" of (b). By substituting the "z" in (c) for "$e^{j\omega}$" and applying Euler's formula the impulse response function is formulated for the frequency domain (d). In (e) the formulas for the calculation of magnitude and phase angle for the impulse response function are shown.

(a) $y[n] = x[n] + \alpha x[n-k] + \alpha x[n-2k] + \ldots .$

(b) $Y[z] = X[z] + \alpha z^{-k} x[n] + \alpha z^{-2k} x[n] + \ldots .$

(c) $y[n] = x[n] \otimes h[n] \rightarrow H[z] = \frac{Y[z]}{X[z]} = 1 + \alpha z^{-k} + \alpha z^{-2k} + \ldots .$

(d) $H[e^{j\omega}] = 1 + \alpha \cos(j\omega k) - j\alpha \sin(j\omega k) + \alpha \cos(j\omega 2k) - j\alpha \sin(j\omega 2k) + \ldots .$

(e) $\left|H[e^{j\omega}]\right| = \sqrt{H[e^{j\omega}]^2}$ and $\measuredangle H[e^{j\omega}] = \arctan\left(\Im\left(H[e^{j\omega}]\right) / \Re\left(H[e^{j\omega}]\right)\right).$

Equation 9: Mathematical formulation of a centered comb filter. (a) In the time domain a comb filter adds weighted lagged samples to the current sample. Here the lag "*k*" equals the integer multiples of length of a TR in samples and the number of lags added equals the width of the window in TRs, α is the weighting function applied, e.g. for a simple moving average α = -1/(WL-1) for all lags. As in Equation 8 this is almost exact representation of MWATS due to the issue of system clock syncing. In (b) is the z-transform of (a). In (c) is the formulation of the impulse response function "H[z]" of (b). By substituting the "z" in (c) for "$e^{j\omega}$" and applying Euler's formula the impulse response function is formulated for the frequency domain (d). In (e) the formulas for the calculation of magnitude and phase angle for the impulse response function are shown.

(a) $y[n] = x[n] + \alpha x[n-k] + \alpha x[n+k] + \alpha x[n-2k] + \alpha x[n+2k] + \ldots .$

(b) $Y[z] = X[z] + \alpha z^{-k} x[n] + \alpha z^{k} x[n] + \alpha z^{-2k} x[n] + \alpha z^{2k} x[n] \ldots .$

(c) $y[n] = x[n] \otimes h[n] \rightarrow H[z] = \frac{Y[z]}{X[z]} = 1 + \alpha z^{-k} + \alpha z^{k} + \alpha z^{-2k} + \alpha z^{2k} \ldots .$

(d) $H[e^{j\omega}] = 1 + \alpha \cos(j\omega k) - \alpha \cos(j\omega 2k) + \ldots .$

(e) $\left|H[e^{j\omega}]\right| = \sqrt{H[e^{j\omega}]^2}$ and $\measuredangle H[e^{j\omega}] = \arctan\left(\Im\left(H[e^{j\omega}]\right) / \Re\left(H[e^{j\omega}]\right)\right) = 0.$

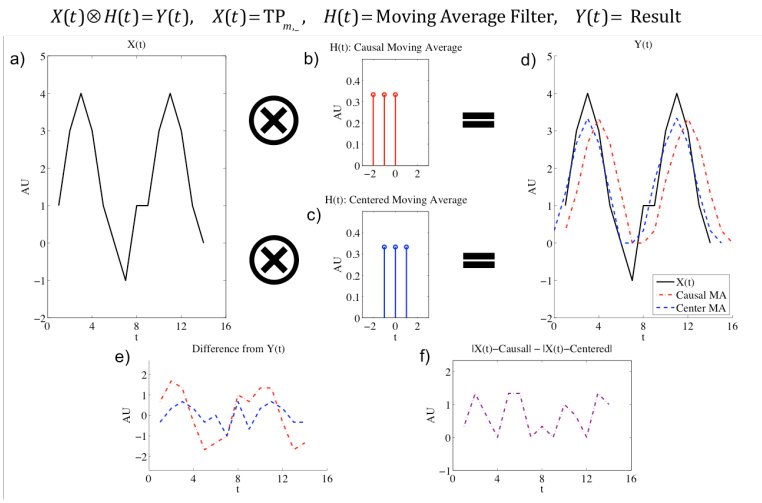

Figure 12: In the point-by-point frame the first time point in a window ($TP_{1,\_}$) is averaged only with its respective counterparts in other windows, likewise for the other m points. So in essence there are actually M averages taking place in MWA. One of these M averages will be used in considering a causal or centered for the MWA. Here X(t) is a collection of a single time point across several TR windows (a), H(t) is either a 3-point causal (b) or 3-point centered (c) simple moving average filter and Y(t) is the filtered result (d). (e) is the difference between X(t) and the filtered results Y(t). (f) is a comparison of the fits of the two filters, positive values are where the centered outperforms the causal and vice versa for the negative. Note that at no point does the causal outperform the centered filter for matching the true value of X(t). Also note that abrupt changes, the peaks & valleys, in X(t) are not well captured by a moving average, as it is by its very nature a low pass filter.

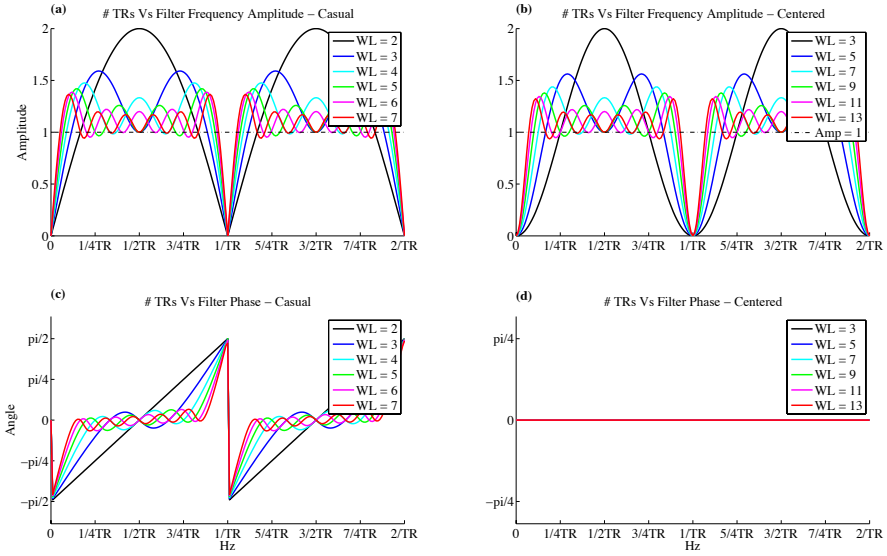

Figure 13: Frequency and phase responses of causal and centered simple moving average filters of various window lenghts. The frequency response of the MWATS causal filter can be seen in a) at window lenghts of 2-7 and likewise for the centered in b) at window lenghts of 3, 7, 5, 7, 9, and 13. The respective phase response plots can be found in c) and d) respectively. As the window lengths increase the notch at the 1/TR harmonics becomes narrower a) & b) and the phase response of the causal filter becomes flatter, c). The phase response of the centered filter remains flat regardless of window length.

Should the TR to be Cleaned be Included or Excluded the from the Template?

In deciding whether or not to include the TR segment to be cleaned into the moving window average it is worth considering again what the average template does. The point of creating the average template is to separate the sGA from all other signals not time locked to the TR, i.e. sO. In doing so, a template is made of sGA, which then is subtracted from the sRec, leaving OS, *Equation 6d*. Conceptually, adding the TR to be cleaned to the template adds no information to the template subtraction in terms of how to separate the sGA and OS within that TR. Consider the limit of a moving window size of one. Using this as the template will result in a zero signal after cleaning, a perfect removal of the sGA but a total loss of the sO. Expanding on this the result, including the current TR in the template will only result in the scaling of the OS within current TR post cleaning by 1 – the weight applied to the current TR. In the case of a simple moving average this would equal 1 – (1 / Window Length), a formal proof follows in *Equation 10*.

Reconsidering the results of filters in Table 1 which included the TR to be cleaned in the template outperforming, in terms of RMSE, excluding it from the template, the benefits seen in RMSE and REsamp are a misleading, in that the gains are the result of scaling of the cleaned signal. In the case of the causal filter the signal was scaled down by 25%, and by 20% in the case of the centered filter. The different scaling amount between the two is due to the different number of effective averages in these conditions. As an example, the equivalent filter to a causal filter of length 4 that included the TR to be cleaned, 4 averages, would be a length 4 filter that excluded the TR to be cleaned from the template, 3 averages, whose resultant signal magnitude was scaled down 25%. The REsamp with this filter would be 0.1782 μV/sample. Taking the same all-cleaned TR

windowed average and scaling the magnitude down by 25%, then calculating the average residual energy per sample, gives a new REsamp value 0.1003 μV/sample, the same as measured in Table 1 for the causal filter which included the TR to be cleaned in the template. In short the TR to be cleaned should not be included in the template as it adds nothing with regards to separability and results in the down scaling of the cleaned signal.

Equation 10: The mathematical process behind MWATS. Note the indices are ambiguous to whether this is a causal or centered filter as the same logic applies to both. A more detailed formula of Equation 6d is displayed in (a). Here TREV(*i*) is the starting sample of the TR window to be cleaned, *i* is the index of the TR window to be cleaned (current TR), TRint is the number of samples per TR window, *j* is the channel to be cleaned and WL is the length of the cleaning window in TRs and α is the weighting function. (b) is the mathematical construction of the gradient artifact template ($MR_{Template}$). (c) is the gradient artifact template without the current TR ($MR'_{Template}$). (d) is the reformulation of (b) using $MR'_{Template}$. (e) is the reformulation (a) using $MR'_{Template}$. Note that including the current TR in the template only scales the resulting cleaned signal by 1-α(*i*), (i.e. one minus the weight applied to the current TR). To prevent scaling of the cleaned EEG signal α(*i*) = 0.

(a) $\mathrm{EEG}_{Clean}[j, TREV(i):TREV(i)+TR\,\mathrm{int}-1] = \mathrm{EEG}_{Raw}[j, TREV(i):TREV(i)+TR\,\mathrm{int}-1] - \mathrm{MR}_{Template}$,
where l =k+WL-1.

(b) $\mathrm{MR}_{Template} = \sum_{k}^{l} \alpha(\theta) \times \mathrm{EEG}_{Raw}[j, TREV(\theta):TREV(\theta)+TR\,\mathrm{int}-1]$,

where $\alpha(\theta)$ = weighting function, e.g. SMA = $\frac{1}{WL}$, for all θ.

(c) $\mathrm{MR}'_{Template} = \sum_{k}^{l} \alpha(\theta) \times \mathrm{EEG}_{Raw}[j, TREV(\theta):TREV(\theta)+TR\,\mathrm{int}-1],\ \theta \neq i.$

(d) $\mathrm{MR}_{Template} = \mathrm{MR}'_{Template} + \alpha(i) \times \mathrm{EEG}_{Raw}[j, TREV(i):TREV(i)+TR\,\mathrm{int}-1].$

(e) $\mathrm{EEG}_{Clean}[j, TREV(i):TREV(i)+TR\,\mathrm{int}-1] = (1-\alpha(i))\mathrm{EEG}_{Raw}[j, TREV(i):TREV(i)+TR\,\mathrm{int}-1] - \mathrm{MR}'_{Template}.$

How Should the Weighting be Applied?

In understanding whether to apply a weighting on a TR-by-TR or a point-by-point basis it is worth the revisiting the region of support for the filter. A causal filter relies only on past and present inputs, *Figure 5a*, and has an output shifted to the right in samples. A centered filter relies equally on past and future inputs and can include the present, *Figure 5b*, and does not have a shift in its output. An anti-causal relies only on future and present inputs, *Figure 5c*, and shifts the output to the left in samples. Semi-causal and semi-anti-causal are similar to their namesakes in that their output is shifted

but they can rely on past, present and future events. What makes them "causal" or "anti-causal" is that the heavier weighting is applied to the past for the former or to the future for the latter in terms of inputs. If the weighting is done on a TR-by-TR basis, Figure *14b*, all time points within each TR window experience the weighting equally. A point-by-point weighting, *Figure 14c*, on the other hand causes heavier weighting on future events for points that appear early in the TR window, *Figure 14g*, and vice versa for events that occur late in the TR window, *Figure 14i*. This has the effect of making the first half of the TR widow experience a semi-anti-causal filter, where the second half experience a semi-causal filter. Introducing a causal or anti-causal nature to the filter introduces lag in the MWA which degrades the templates ability to match the TR window to be cleaned as seen the REsamp values in *Table 1*. Consequently any weighting applied to the filter should be done respecting the boundaries of the TR windows, i.e. on a TR-by-TR basis.

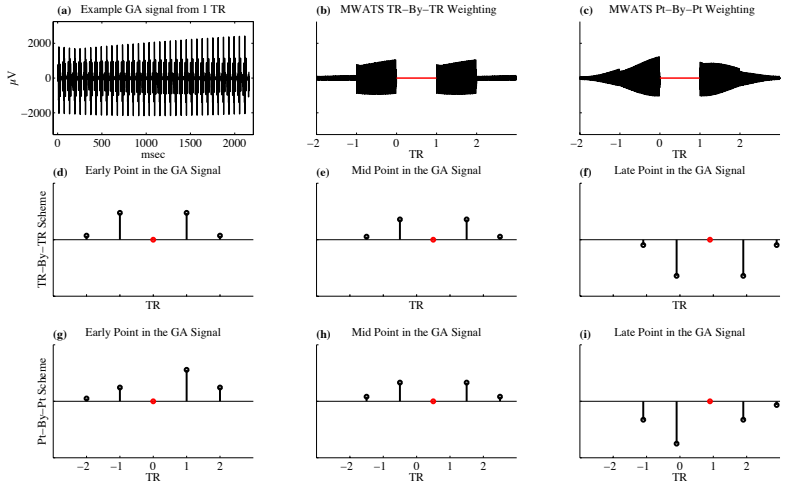

Figure 14: Weighting schemes versus filter regions of support. Here an example of a single TR segment of gradient artifact is displayed in (a). In (b) and (c) are examples of window length five centered moving filter which exclude the TR to be cleaned from the template which have a TR-by-TR and Point-by-Point weighting scheme respecively. In the middle row are points taken from early (d) the middle (e) and late (f) in the TR window across segments for the TR-by-TR weight scheme. Likewise in the bottom row (g-f) for the Point-by-Point weighting scheme. In the middle note how all of the filters are centered, i.e. symmetric weighting. In the bottom row note that early points in the GA will have a heavier weighting toward future events (anti-causal) where the opposite is found for late points in the GA (semi-causal).

What type of Weighting Should be Applied and How Many TRs Should be Included?

Weighting and the number TRs segments to include in the template are considered here in tandem, because both act to change the size of the filter. In doing so, both affect the suppression of the non-GA signals (sO) in the template, the separability of sGA and sO, and the frequency characteristics of the filter, *Figures 15 & 16*.

Weighting the template affects both its time and frequency domain representations. In the time domain, weighting the template puts more emphasis on the local structure of the GA but also on the EEG, BCG, etc. as well. In the frequency domain, weighting acts to smooth the ripples in the frequency spectrum, *Figure 15a*. The GA occupies narrow bands at integer multiples of 1/TR. The ideal filter would be one with unit amplitude at all frequencies that are not integer multiples of 1/TR, and zero otherwise, *Figure 15a & b* black dashed line. That is, it would only remove the GA signal bands while leaving the remaining spectrum untouched. Filter amplitudes below the dashed line are the sO template frequencies that are subtracted from the current TR and vise versa for the amplitudes above 1. The subtracted frequencies can be thought of as adding negative signal, thus a MWATS with deviations in either the positive or negative pollute the current TR with erroneous signals. In characterizing how the different weights affect the filter there are two values to look at. The first being the narrowness of the notch, measured by its full width at half of it maximum value (FWHM). To allow the FWHM to be comparable across different TR values, the FWHM value is normalized by 1/TR(sec). Furthermore, to give context to this measure, it is converted into percent of the spectrum retained post filtering, Equation 8a which is 1 minus the FWHM normalized by width of its spectral repeat. The second measure is the percent non-ideality of the filter, measured

by the area between the filter and the ideal, *Figure 15c*, normalized by the area of the ideal, *Equation 8b*. Both of the measures in *Equation 11* we made on a centered MWATS filter with a window length of 13 using different weights, the results of which can be found in *Table 4*. The weighting types we analyzed were: simple moving average, exponential, Gaussian, Hanning, Hamming, Blackman and flat top.

Equation 11: Formulas for Percent Spectrum Retained Post Filtering and Percent Non-Ideality of a filter.

$$\text{(a) Percent Spectrum Retained} = 100 \times \left(1 - \frac{\text{FWHM}}{\text{TR(sec)}}\right).$$

$$\text{(b) Percent Non-Ideality} = 100 \times \sum_{0}^{1/TR} \frac{|1 - H(z)|}{\Delta z}.$$

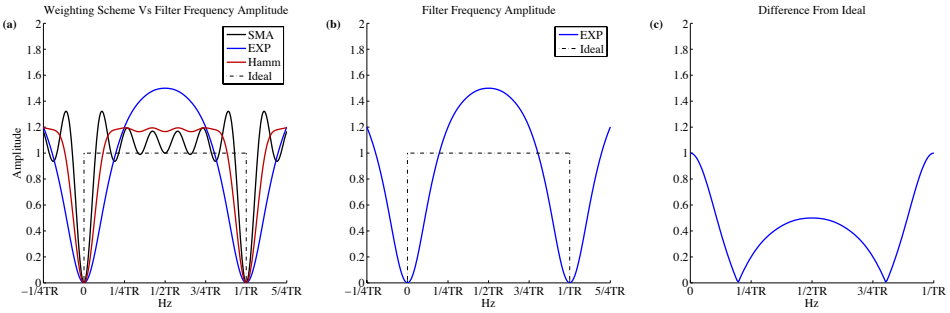

Figure 15: Difference from ideal by weighting type. (a) The frequency response of as a function of the weighting for a centered simple moving average (SMA), exponential (EXP), Hamming (Hamm) and ideal filter. All filters of plotted are of window length 13. (b) The frequency response for the exponentially weighted filter, blue line, is plotted with the ideal filter, black dashed line. (c) The difference between the two filters in (b) is plotted.

Table 4: Difference from Ideal Filter Across Weighting Schemes. All calculations were performed on a centered MWATS filter with a length of 13. TR to be cleaned was not included in the template formation. The weighting was performed on a TR-by-TR basis. The resolution in the frequency domain was ΔF = 0.0030660Hz. The time resolution (1/ ΔF) was matched to the recording length (number of TR segments cleaned) and TR (sec) from the synthetic recording.

Weighting Scheme	% Spectrum Retained 100*(1-FWHM (Hz) / TR sec)	% Difference From Ideal
Simple Moving Average	89.404	18.7473
Exponential	69.5364	43.9709
Gaussian	84.106	28.0636
Hanning	82.7815	29.7643
Hamming	84.106	26.9013
Blackman	80.1325	33.2898
Flattop	61.5894	61.6811

The simple moving average, i.e. a flat weighting, performed best among the group tested: it had the narrowest notch and the smallest difference from the ideal filter. In other words, for a given widow length, the flat weighting both retained the largest amount of spectrum and maintained the highest fidelity for the sO.

In thinking about the weighting in the time domain, by putting increased emphasis on the local structure of the GA in the window across which the filter is acting in essence is being narrowed. Narrowing the filter has two drawbacks. The first is that narrowing the window in effect reduces the number of averages. Reducing the number of averages reduces the suppression of sO in the template. The second drawback is that as the window becomes narrower, the more non-ideal the filter becomes, *Figures 15 & 16*. The TR segment being cleaned is contaminated by the combination of whatever non-GA signal, sO, is not suppressed in the template and the non-ideality of the filter. The larger each of these, the more contamination the MWATS introduces into the TR segment being cleaned. Beyond the contamination issue though, there a problem of discarding portions of the frequency spectrum, the FWHM of the notches normalized by 1/TR. This is inversely proportional to the window size. The simple moving average gives the maximum possible window size for a given length and therefore has narrowest notch or retains the largest amount of spectrum.

Expanding on weighting, window length works similarly. The longer the window the more separable the sGA and OS are, and the more spectral content is retained in the cleaned signal. Using the same tools applied above, Equation 8, we measured the percent spectrum retained, and the non-ideality of a simple moving average template subtraction filter, *Table 5*. Expanding on this analysis, the suppression of the non-GA signal in the

template was measured as a function of filter length. Assuming a normal distribution for non-GA signals, in a simple moving average template subtraction these signals should fall of as $1/\sqrt{\#\text{ of Averages}} = 1/\sqrt{\text{Window Lenght -1}}$. The pollution from the non-GA signal in the template falls off even faster though as the non-ideality of the filter falls off with window length as well, *Table 5 & Figures 16a, 17, 18e*. In Figure 18e the pollution from template in the Non-GA bands introduced by MWATS is measured as a percent of the energy in the TR segment being cleaned. To obtain pollution level of 5%, the window length would need to be at least 13 TR segments long and for a 1% level the window length would need to be at least 45 TR segments, *Figure 18e*.

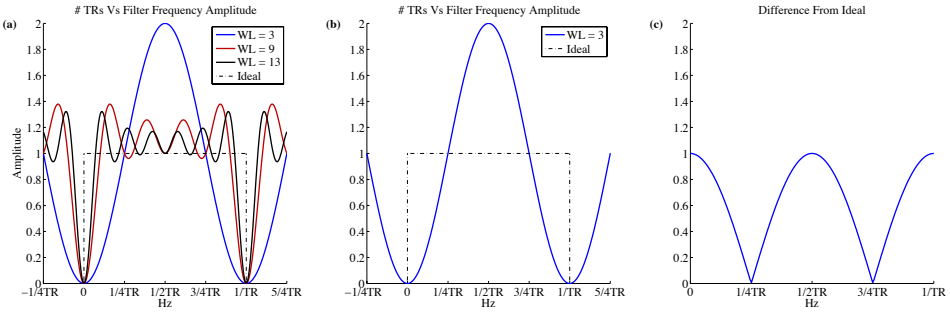

Figure 16: Difference from the ideal filter as a function of window size. (a) The frequency response of as a function of the window length for a centered simple moving average with length of with 3, 9 and 13 TR segments as well as the ideal filter. (b) The frequency response for the filter with a window length 3, blue line, is plotted with the ideal filter, black dashed line. (c) The difference between the two filters in (b) is plotted.

Table 5: Difference from Ideal Filter Across Window Lengths. All calculations were performed on a centered MWATS filter with a flat weighting. TR to be cleaned was not included in the template formation. The resolution in the frequency domain was ΔF = 0.0031494Hz. The time resolution (1/ ΔF) across all window length was matched to the number of TR segments cleaned by the length 17 filter and TR (sec) from the synthetic recording. The difference in the ΔF used for the calculations in this table and Table 4 accounts for the slight difference in values seen for the SMA filter of length 13.

SMA Window Length	% Spectrum Retained 100*(1-FWHM (Hz) / TR sec)	% Difference From Ideal
5	71.4286	42.0362
7	79.5918	31.5082
9	83.6735	25.5060
11	86.3946	21.5645
13	89.1156	18.7832
15	90.4762	16.6540
17	91.8367	15.0413

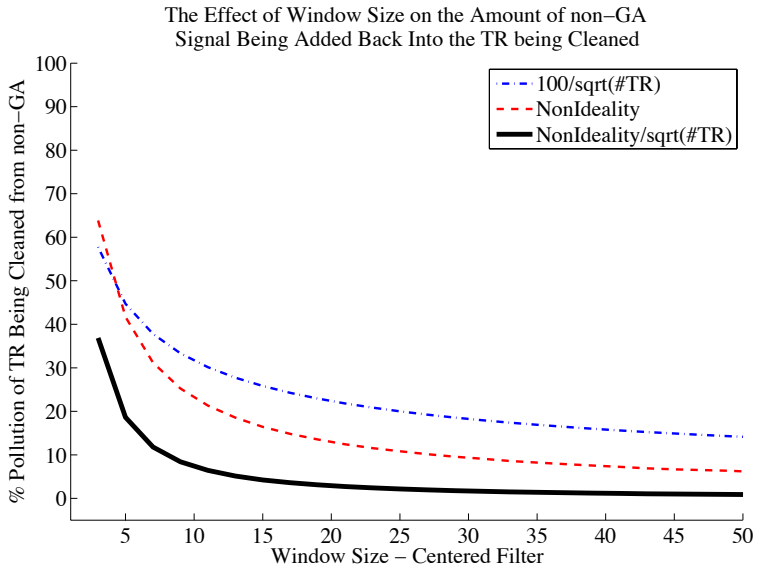

Figure 17: Percent signal pollution as a function window length. The filter percent non-ideality as a function of window length is plotted in red-dashed line. The percent of the non GA signal in the template as a function of window length is plotted in the blue-dashed line. Their combination as a function of window length is plotted in the black line.

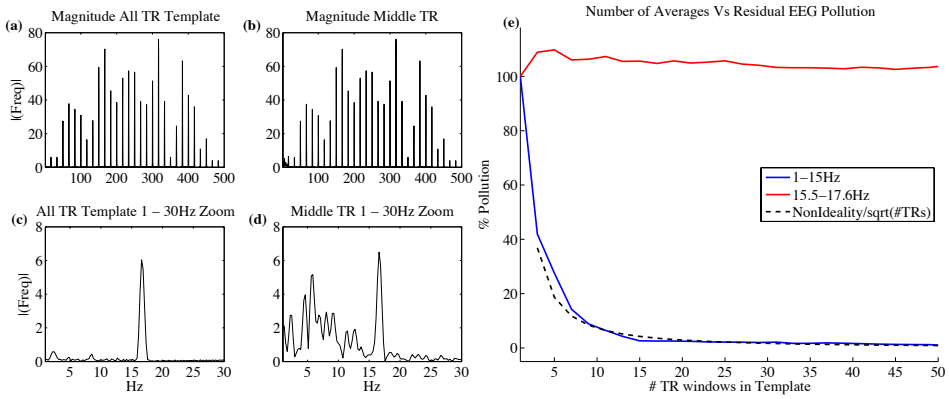

Figure 18: The measurement of percent signal pollution by energy from a subject recording. (a) Example of a magnitude spectrum of the all TR segment windowed average from a subject recording. (b) The magnitude spectrum of the center most TR segment in the same recording. (c) and (d) 0-30Hz zooms of the magnitude spectrums of (a) and (b) respectively. (e) Plot of the percent of energy in the template compared to the energy center TR segment as a function of window length, the 1-15Hz band, blue line, 15.5-17.6Hz band, red line. A trend line of the non-ideality over the number of averages is displayed in the dashed black line. There were 36 slices in this MR acquisition with a TR of 2.160 sec resulting in a slice timing frequency of 16.67Hz. Note the energy in the band that contains the slice timing artifact remains constant regardless of the number of TRs included in the template, red line. The energy in the band not occupied by the GA signal (1-15Hz) falls of as the Non-Ideality over the square of the number of averages, blue line. Putting numbers to the trend line, at a window length of 3 the residual pollution from the non GA signal in the template into the true sO from cleaning is 36.9% by energy, at a window length of 5 is 18.6%, at a window length of 9 is 8.5%, at a window length 13 is 5.1% and at a window length of 45 is 1%.

The findings in *Tables 4 & 5* follow the trends found using the synthetic data in *Table 2*. In *Table 4* if listing weighting types in terms of either the amount of spectrum discarded or the percent non-ideality of the filter the order is: simple moving average (i.e. none or flat), Hamming, Gaussian, Hanning, Blackman, exponential, and flat top. This is the same order that is found *Table 2*. Similarly in *Table 5*, the amount of spectrum discarded and the percent non-ideality both decrease with the filter length, following the findings in *Table 2*.

In the process of filter selection some practical factors come into play beyond separability and fidelity. The first of these being the subject population. Populations prone to movement, e.g. children, where subject movement in the scanner is likely to be increased will have displacement motions of the electrode loop where the GA is likely to be different pre and post motion. The more TR segments that are included in the template the longer the GA needs to be stable. One drawback of the synthetic recording was that it was too good of a subject. Specifically the phantom did not move over the recording, this meant that the GA was stable in terms of displacement changes over the recording duration. In their 2008 paper, Herrmann and Debener, described lack of head motion over a recording as "rather unrealistic". Secondly, the longer the window the more of the EEG recording will need to be discarded. That is, a centered cleaner will need a padding of half the total number of TR segments averages (Window Length / 2 -1) on either end of cleaning. The window length can be adapted, reduced, to clean these padded regions, but frequency content will be decreased and the signal pollution increased in these regions as a result. Finally a factor to consider is how much spectrum is acceptable to remove.

Longer filters will result in narrower notches, *Figure 16a,* and thus more retained spectrum content .

Beyond subject population and the amount of data that can be discarded, the type analysis being performed, e.g. spectral, single trial, evoked response potential (ERP), should inform the decisions on how many TR segments to include and what type of weighting to apply to the template. In spectral analysis retaining as much of the spectrum as possible will likely drive these decisions. The weighting that retains the most spectrum is a flat weighting. The number of TR segments to include in the window will be driven by the subject population, how much data can be discarded and the retaining of the spectrum. The more TR segments included the more of the spectrum will be retained, however the more of the data will be corrupted by subject motion and the larger the padding regions will be on either end of the recording. In single trial analysis maintaining the fidelity of the sEEG will likely drive these decisions. Here again, a flat weighting would work best. The number of TR segments to include in the template will be driven here by the subject population and how much data can be discarded. In ERP analysis fidelity can be traded off somewhat for gains in separability. ERP analysis has the benefit of averaging trials. Slight errors in fidelity can be averaged away. In terms of GA though, the poorer the separability is, the more GA signal remaining in the cleaned signal there will be. As the amount GA in the cleaned signal increases the more averages will be required to suppress it. Here the tradeoffs of a Hamming or Gaussian weighting may reduce the number of averages in the analysis needed to suppress the residual GA where really not costing much in terms of fidelity of in the EEG. In ERP analysis the data sets tend to be long to account for the number of averages need. As such the likely driving

factor for the number of TR segments to include in the template will be primarily the subject population.

Conclusion:

Apart from perfect fidelity, doing nothing, or perfect separability, subtracting the recording from itself, anything that is done to the sRec will affect both the fidelity of the true EEG signal by way of sO and how well two signals, sO and sGA, are separated. In deciding on the options in implementing MAWTS to remove sGA, choices must be made based in part on what is needed for the analysis. Most of the choices are clear cut however: in post-hoc analysis use a centered filter, do not include the current TR in the template, and apply any weighting by respecting TR segment boundaries (equally to all points with in a TR segment). Decisions regarding the appropriate window length and weighting should both be informed by the makeup of the subject population as well as the type of analysis on the EEG data being performed.

In the literature the settings used for the MWATS are either not noted or vary wildly. As examples: in their 2008 paper, Herrman and Debener, suggest a window length of 30; Liu, et al., (2016) used a window length of 41; Hoppstädter et al., (2015), Sockeel (2016), and Xiao, et al., (2016) did not list the settings they used; and Zich, et al., (2015) grew their window with each successive TR. This is problematic in that spectral content and fidelity of the EEG cleaned of GA are affected by the cleaner choices. As a general default, a window length of 13 with a flat weighting may be good starting point, but consideration should be given to: the subject population and the frequency of the their movement, and to the type of analysis being performed on the cleaned EEG data. In addition the decisions on window length should be made in part based upon the amount

of the EEG data that is going to be discarded. A total of the window length – 1 segments of TR data of the EEG will not be the cleaned by the filter. Over the 5 minutes 47.6 second recording of the synthetic data, we discarded 17 – 1 TRs worth of data for the analysis in order to keep the regions analyzed consistent across different cleaner window lengths. This meant that a total of 34.56 seconds (16*2.160sec) of data were discarded. Putting it other terms, in the case of this recording 9.94% of the recorded EEG data was discarded. Had we used our recommended default this would have been 7.46%. In addition to discarding data in the temporal domain, data also is discarded in the frequency domain as notches are introduced in the spectrum by the MWATS. Using our recommended default filter, 10.6% of the spectral content is eliminated by the notches introduced by the MWATS.

Beyond the loss of both spectral and temporal content, MWATS introduces contamination into the EEG signal. Using our recommended default filter there is approximately 5%, by energy, contamination of the retained spectrum from non-GA signals in the template being added into the clean EEG signal.

Due to the spectral losses, the temporal losses and the signal contamination the MWATS should not be thought of as a final solution to GA cleaning. New methods are needed to reduce or eliminate these losses and pollution. Until then however, the work we presented here allows for more informed decisions regarding how moving windowed average template subtraction should be applied.

References:

Allen PJ, Josephs O, Turner R (2000) A method for removing imaging artifact from continuous EEG recorded during functional MRI. Neuroimage 12:230-239.

Allen PJ, Polizzi, G., Krakow, K., Fish, D. R., & Lemieux, L. (1998) Identification of EEG Events in the MR Scanner: The Problem of Pulse Artifact and a Method for Its Subtraction. Neuroimage 8:229-239.

Bandettini PA, Wong EC, Hinks RS, Tikofsky RS, Hyde JS (1992) Time course EPI of human brain function during task activation. Magnetic resonance in medicine 25:390-397.

Bénar C-G, Aghakhani Y, Wang Y, Izenberg A, Al-Asmi A, Dubeau F, Gotman J (2003) Quality of EEG in simultaneous EEG-fMRI for epilepsy. Clinical Neurophysiology 114:569-580.

Brainard DH (1997) The Psychophysics Toolbox, Spatial Vision. Spatial vision 10:433-436.

Cohen MS (2000) Method and apparatus for reducing contamination of an electrical signal. (USPTO, ed): Regents of the University of California.

Delorme A, Makeig S (2004) EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. J Neurosci Methods 134:9-21.

Ferree TC, Luu, P., Russell, G. S., & Tucker, D. M. (2000) Scalp electrode impedance, infection risk, and EEG data quality. Clinical Neurophysiology 112:536-544.

Goldman RI, Stern JM, Engel J, Cohen MS (2000) Acquiring simultaneous EEG and functional MRI. Clinical Neurophysiology 111:1974-1980.

Helmholtz HV (1853) Ueber einige Gesetze der Vertheilung elektrischer Ströme in körperlichen Leitern, mit Anwendung auf die thierisch-elektrischen Versuche (Schluss.). Annalen der Physik 165:211-233.

Herrmann CS, Debener S (2008) Simultaneous recording of EEG and BOLD responses: a historical perspective. Int J Psychophysiol 67:161-168.

Hoppstädter M, Baeuchl C, Diener C, Flor H, Meyer P (2015) Simultaneous EEG–fMRI reveals brain networks underlying recognition memory ERP old/new effects. NeuroImage 116:112-122.

Kwong KK, Belliveau JW, Chesler DA, Goldberg IE, Weisskoff RM, Poncelet BP, Kennedy DN, Hoppel BE, Cohen MS, Turner R (1992) Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. Proceedings of the National Academy of Sciences 89:5675-5679.

Liu Y, Bengson J, Huang H, Mangun GR, Ding M (2016) Top-down modulation of neural activity in anticipatory visual attention: control mechanisms revealed by simultaneous EEG-fMRI. Cerebral Cortex 26:517-529.

Liu Z, Kecman F, He B (2006) Effects of fMRI-EEG mismatches in cortical current density estimation integrating fMRI and EEG: a simulation study. Clin Neurophysiol 117:1610-1622.

Mullinger K, Debener S, Coxon R, Bowtell R (2008) Effects of simultaneous EEG recording on MRI data quality at 1.5, 3 and 7 tesla. International Journal of Psychophysiology 67:178-188.

Mullinger KJ, Havenhand J, Bowtell R (2013) Identifying the sources of the pulse artefact in EEG recordings made inside an MR scanner. Neuroimage 71:75-83.

Ogawa S, Tank DW, Menon R, Ellermann JM, Kim SG, Merkle H, Ugurbil K (1992) Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. Proceedings of the National Academy of Sciences 89:5951-5955.

Ritter P, Becker R, Graefe C, Villringer A (2007) Evaluating gradient artifact correction of EEG data acquired simultaneously with fMRI. Magnetic resonance imaging 25:923-932.

Simons RF, Miller GA, Weerts TC, Lang PJ (1982) Correcting baseline drift artifact in slow potential recording. Psychophysiology 19:691-700.

Sockeel S, Schwartz D, Pélégrini-Issac M, Benali H (2016) Large-Scale Functional Networks Identified from Resting-State EEG Using Spatial ICA. PloS one 11.

Yan WX, Mullinger KJ, Brookes MJ, Bowtell R (2009) Understanding gradient artefacts in simultaneous EEG/fMRI. Neuroimage 46:459-471.

Zich C, Debener S, Kranczioch C, Bleichner MG, Gutberlet I, De Vos M (2015) Real-time EEG feedback during simultaneous EEG–fMRI identifies the cortical signature of motor imagery. NeuroImage 114:438-447.

AN ECG-FREE METHOD FOR ESTABLISHING BALLISTOCARDIOGRAM TIMING

Chapter 2

Abstract:

Second to the gradient artifact, the three components of the ballistocardiogram (BCG): scalp pulsation; head rotations and translations; and the Hall effect, are the largest contaminants of an electroencephalography (EEG) recording acquired within the magnetic resonance imaging (MRI) environment. All current methods of suppressing the ballistocardiogram artifact rely upon temporal correlation with the artifact. There are two ways to establish this timing, either through a simultaneously acquired electrocardiographic (ECG) recording or through the EEG itself. The most common method at present is to infer the timing of the BCG based upon the timing of the R-wave complexes in the ECG recording. Acquisition of the simultaneous ECG signal has several drawbacks: first and foremost, it adds an additional point of technical failure; secondly, it increases set up time and experiment complexity; and third, it can serve as additional path for the introduction of artifacts into both the EEG and MRI recordings. Alternative methods extract the timing of the BCG directly from the EEG, avoiding these drawbacks. However these methods typically depend on the mean of the absolute signal across channels (MAS) (Srivastava et al., 2005). By averaging across the head and across artifact components MAS temporally blurs, i.e. broadens the peaks of, the BCG signal which leads to variance between its timing and the ECG timing, this in turn leads to reduced artifact suppression efficacy in the various cleaning algorithms. Presented here is a method to extract the BCG timing directly from the EEG record; it both avoids the pitfalls of previous EEG based methods, and the complications of using an ECG based

method. Our new method: gives timing that is locked to the ECG; has peak detection performance equivalent to the ECG; is more robust against total failure from signal loss; and it can be applied retrospectively on datasets where the ECG was lost. It accomplishes this through anatomically based channel selection, a reference scheme designed to suppress signal from head and facial motion, a constrained peak detection algorithm, and automatic error detection and correction algorithms. Beyond providing a single time point for each BCG event (the ECG R-wave peak equivalent) our method provides also a window about the time point over which the BCG event occurs

Keywords:

Simultaneous EEG-fMRI, Electroencephalogram, EEG, functional Magnetic Resonance Imaging, fMRI, Ballistocardiogram, BCG, Ballistocardiogram Timing, Electrocardiogram, ECG, EKG, Electrocardiogram Timing, Shannon Energy Envelope, Peak Detection

Introduction:

In simultaneous (EEG-fMRI) Electroencephalography (EEG) and functional Magnetic Resonance Imaging (fMRI) the Ballistocardiogram artifact (BCG) proves a significant challenge. While in today's 3 Tesla field MRI scanners, the BCG is an order of magnitude smaller than the gradient artifact (GA) the BCG still may be one to two orders of magnitude greater than the EEG signal itself. This is of particular import for EEG signals in the sub 50Hz frequency range, as this is where the majority of the spectral power of this BCG artifact lies (Weissler, 1974). BCG is the result of three different processes (Mullinger et al., 2013) all of which originate with the heartbeat: scalp

pulsation (SP), head rotations and translation (HRT), and changes to the electrical potentials across the scalp from the Hall effect (HE). The currently popular methods to remove the BCG from the EEG recording all rely on correlating artifact events to cardiac timing from a simultaneously recorded electrocardiogram (ECG). These methods include template subtraction (Allen, 1998), principal component fitting (Niazy et al., 2005), independent component analysis (Nakamura, 2006, Koskinen and Vartiainen, 2008), and orthogonal matching pursuit (Xia et al., 2014). Difficulties with the ECG recording/acquisition have lead to attempts to use the "map hilliness" (Lehmann, 1971) of the EEG recording, the mean of the absolute signal (MAS) as a surrogate for the ECG recording (Srivastava et al., 2005). This technique however has never been adopted as the standard method, perhaps due to its inability to yield cleaning results as good as those whose timing was derived from the ECG. Acquiring the ECG has several drawbacks as it creates an additional point of failure, increases set up time and experiment complexity, and serves as a potential a route for the introduction of further artifact into both the EEG and MRI recordings (Mullinger et al., 2008a). Developed here is an EEG-based approach that circumvents the drawbacks of acquiring the simultaneous ECG signal, and the shortcomings of the MAS method. Our new method (Rodriguez et al., 2015) does so by careful *a priori* channel selection, by re-referencing, and through a model-based peak detection algorithm that includes automatic error detection and correction algorithms.

The translations for each of the three sources of BCG signal from heartbeat to artifact are different, but the methods by which their signals are created fall into two general areas: induction and charge separation. Scalp pulsation (SP) and Head rotation/translation (HRT) result from electromotive forces induced in the conductive

loops, the Hall effect (HE) results from separation moving charges in the blood by the magnetic field of the MRI.

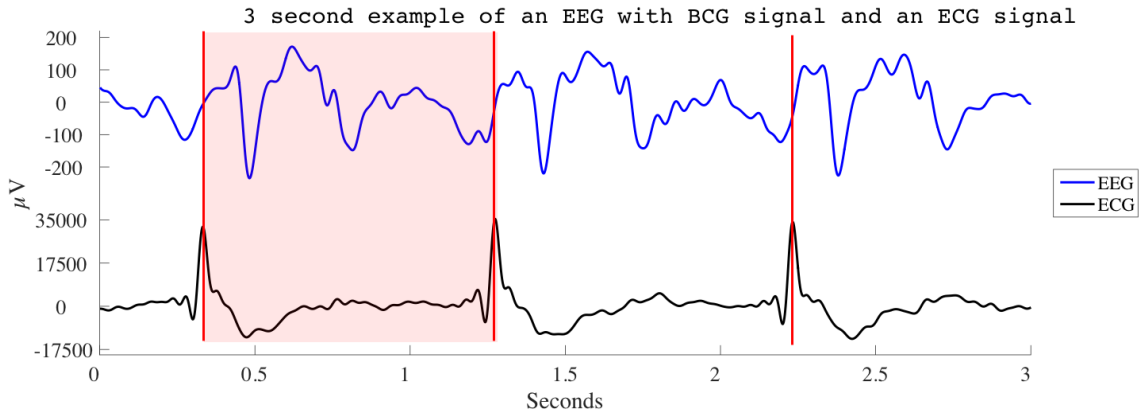

Figure 1: Example of an EEG recording with BCG contamination, and the corresponding ECG recording. The EEG (blue trace) and the ECG (black trace) signals are plotted using different scales. The amplitude of the ECG signal is 2 orders of magnitude greater than that of the EEG signal. The red vertical lines correspond to the R-wave peaks of the ECG signal. The shaded region is the inter-beat interval (IBI) between the first and second heartbeats. Note the pattern that repeats in the EEG, the BCG, with every successive heart beat. Here both the EEG and ECG signals are band-pass filtered from 1-25Hz. The EEG signal is T3-Cz.

Faraday's Law of Induction

The EEG device can be viewed as a collection of conductive loops, *Figure 2*. Faraday's law of induction states that an electromotive force (*emf*) will be induced into these loops equal to the opposite of the time rate of change of the magnetic flux though the loops, *Equation 1*. By extension, a voltage difference signal proportional to the electromotive force will be introduced into the recording due to changes in the magnetic flux, *Equation 2*. The GA arises from changes in the magnetic field with respect to time, the scalp pulsation and head rotation and translation components of the BCG arise principally from time varying changes to the loop geometry.

Scalp pulsation

The scalp pulsation (SP) component of the BCG arises from physical displacement of the electrodes due to outward expansion of the scalp from the expansion

of underlying blood vessels. The vast majority of the scalp expansion will be due to arterial vasculature as its pressure is much higher than that of the venous system. With each heartbeat, specifically during the depolarization of the ventricles of the heart, coincident with the QRS complex in the ECG, blood is forced through the arterial vasculature throughout the body. In the supine position, the standard for brain fMRI, there is minimal movement of the head with respect to the torso. Consequently there is almost no change from heartbeat to heartbeat in the length the blood must travel from the heart to reach the head. The time delay from heartbeat to scalp pulsation, pulse wave time, can be calculated using the pulse wave velocity (PWV) and the distance traveled from the heart, *Equation 3*, called the Moens–Korteweg equation. Examining the Moens–Korteweg equation, it is clear that the delay can be seen to be purely a function of the properties of the blood vessels and the distanced traveled through them. When taken in the context of EEG-fMRI the rate of change in the properties of the blood vessels will be miniscule. As such, the time between the heartbeat and scalp pulsation become a function of distance traveled from the heart. In terms of overall characteristics, SP is the most time locked to the ECG QRS complex timing, the most repeatable, the narrowest in temporal breadth, and in the middle of overall amplitude (Mullinger et al., 2013).

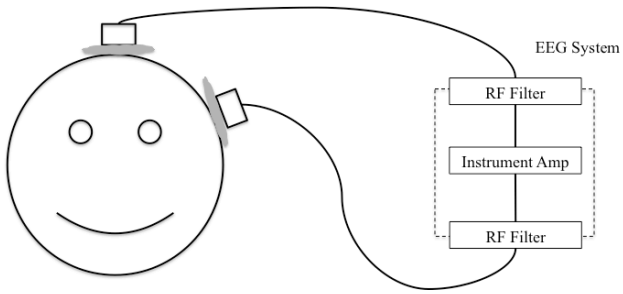

Figure 2: A single conductive loop of the EEG System - Starting at the scalp, the loop passes through a conductive medium (typically a silver chloride gel) to the measurement electrode. Next, it passes through the electrode lead, then through the EEG hardware and returns back through the reference electrode lead, through the reference electrode, then again through a conductive medium to the scalp and through the body tissue back to its starting point in the scalp.

Equation 1: Faraday's Law of Induction, the electromotive force induces in a conductive loop is equal to the negative of the rate of change of the magnetic flux within the loop. With the gradient artifact it is the perturbations to the magnetic field B(t) that cause the in the magnetic flux. The ballistocardiogram artifact arises principally from perturbations in the loop area and orientation, dA(t), resulting in a net magnetic flux.

$$\mathrm{EMF}_{\mathrm{Electrode}} = -\frac{d}{dt}\int_{surface} B(t) \bullet dA(t).$$

Equation 2: The SP and HRT portions of BCG are equal to the the *emf* they induce, multiplied by the ratio of the impedance of the instrument amp, $Z_{\text{Instrument Amp}}$, over the total impeadance of the electrode loop, $Z_{\text{electrode Loop}}(t)$. The electrode loop's time dependence is due to the fact that the electrode interface portions of the loop to not have a static impedance, e.g., the electrodes "dry out" .

$$\mathrm{BCG}_{\mathrm{Electrode}} = \frac{\mathrm{Z}_{\text{Instrument Amp}}}{\mathrm{Z(t)}_{\text{Electrode Loop}}}\mathrm{EMF}_{\mathrm{Electrode}}.$$

Equation 3: The Moens–Korteweg equation for pulse wave velocity. Here E_{inc} is the incremental elastic modulus of the blood vessel wall, h is the constraint ratio of the vessel wall thickness, r is the vessel radius and ρ is the blood density.

$$\text{(a) Pulse Wave Velosity} = \sqrt{\frac{E_{inc} \bullet h}{2r\rho}}.$$

$$\text{(b) Pulse Wave Time} = \text{Pulse Wave Velosity} \times \text{Distance From Heart}.$$

Head Rotations and Translations

Head rotations and translations (HRT) are the ballistic part of the BCG. This component of the BCG results from the transfer of the kinetic energy of the heart motion to the head. Movement of the heart causes the torso to move, the movement of the torso moves the neck, and so on. As the head moves the conductive loops move, generated a signal in the EEG leads. Each stage of the translation can be can be thought of as damped inertial movement. Inertia dampers are devices that transfer short-duration high-energy motion into longer duration lower energy motion. The devices can be found, for example, at the top of skyscrapers as large suspended masses used to dampen the effects of strong bursts of wind. In the present case the high intensity short duration motion is from the

heart, the large masses are the torso, neck and head, and the conductive loops are the equivalent of the moving building. The burst of movement from each heartbeat gets dispersed through each step of translation through the body. The transfer functions of each of these translations are not locked however. As an example, the energy transfers from the heart to the torso is dependent on the current state of respiration (Prisk, 2001) (the chest expansion), the mass of the chest in relation to the heart, and the damping of the motion through air in the lungs. The result of both the damping and the variation in the transfer functions temporally blur this component of the BCG and offset with respect to that of the heartbeat itself. In terms of overall characteristics, HRT is in the middle of the BCG components in terms of its time lock to the ECG QRS complex timing, repeatability and temporal breadth, but it has the largest overall amplitude (Yan et al., 2010).

Hall Effect

The Hall effect simply stated is the separation of moving charges by their sign in a magnetic field. This charge separation creates an electrical potential. When the Lorentz force – the combination of electric and magnetic forces on a charge, *Equation 4* – equals zero the system has come to equilibrium. The voltage at equilibrium is the Hall voltage (V_{Hall}) and is at its maximum when the velocity of the charge is perpendicular to the magnetic field, and at a minimum where they are parallel, *Equation 4c*. V_{Hall} also is proportional directly to the distance by which charges can be separated.

Putting the Hall effect in the context of physiology and the BCG; the charged particles are charged molecules in the blood. Their velocity is determined both by the speed at which the heart is moving the blood, and by the orientation of the blood vessels

with respect to the magnetic field. The separation distance is the distance across the blood vessel in the direction of the cross product of the velocity and the magnetic field. V_{Hall} would be the voltage difference measured across this distance. What is measured by the EEG system it the aggregate of these voltages conducted through the tissues of the head, the summed Hall effect (HES). The Hall effect component of the BCG as measured by the EEG system is subject to the resistivity's of the tissues of the head, the rate at which the blood is flowing and details of its perfusion throughout the head, the latter two being capable of varying over a heartbeat-by-heartbeat time frame. In its overall characteristics, HES is the worst of the BCG components in terms of its poor time lock to the ECG QRS complex timing, and its repeatability and temporal breadth; it has the smallest overall amplitude (Mullinger et al., 2013).

Equation 4: Equilibrium condition for the Lorentz Force, F = 0. Here *q* is the charge, *E* is the electric field, *v* is the velocity of the charge, *B* is the magnetic field vector, and SD is the vector distance of the separation of the charges.

(a) $q\left[\vec{E}+\left(\vec{v}\times\vec{B}\right)\right]=0.$

(b) $\vec{E}=-\left(\vec{v}\times\vec{B}\right).$

(c) $V_{Hall}=\vec{E}\bullet\overrightarrow{SD}=-\vec{v}\times\vec{B}\bullet\overrightarrow{SD}.$

ECG Method

The ECG signal reflects the electrical activity of the heart. It arises from the depolarization of the heart muscle during the heartbeat. To measure the ECG, changes in voltage differences across the skin are measured; a minimum of two points is required. The leading manufacturers of MRI compatible EEG amplifiers do this in one of two manners, either with a channel and a reference placed on the chest across the heart or by placing the ECG channel electrode on the back and using the same reference as the EEG

channels. By acquiring both signals on the same analog to digital clock, timing slips that might otherwise occur between the clocks of different recording devices are avoided. This keeps the ECG signal and EEG signal in lock step with one another. The down side of this is that a channel that could otherwise be used to acquire EEG is instead repurposed for ECG.

From the acquired signal the heartbeat timing is established by one of two components: the onset of the QRS complex, or that of the peak of the R-wave, *Figures 1 and* 12, of the QRS complex.

Acquiring the ECG in EEG-fMRI adds complications, and the potential for errors in the form of an additional point of failure as well as electrode/lead placement issues due to subject size, gender and or weight. In addition to signal acquisition issues, ECG acquisition adds safety hazards in the form of subject stimulation and/or burns.

The EEG system consists of many multiple channel pairs, the loss of any one of which is not generally a critical failure. The typical method by which the ECG signal is captured, as implemented by the two major manufactures of MR compatible EEG systems, Electrical Geodesics, Inc. (Eugene, Oregon) and Brain Products, GmBH. (Gilching, Germany), is through measurement of the voltage difference between two electrodes. Failure of either or both of the electrodes in the pair, whether from a lead popping off (loss of contact), the impedance getting too high, or a failure between an electrode clip and electrode pad, all lead to a critical failure in capturing the ECG signal. The failure rate of ECG in the field is not known, but as example of the problem, in their 2005 paper, Srivastava, et al., had an 80% failure rate of the ECG signal, and in the data sets used in this paper there was a 12.5% failure rate. Such problems have not been

uncommon in our own lab. The failure of ECG in an MR environment is not just limited to EEG-fMRI. Cardiac timing is also used in critical applications such as cardiac MRI, where the timing is used to gate the MR image acquisition. In their 2004 paper, Larson, et al, developed a "self-gated" method for cardiac MRI to avoid the all of the complications, including signal failure, associated with ECG in the MR environment.

The issue of ECG signal failure can arise from improper electrode placement. Placement can be confounded subject body type, size and gender. Subjects with large chests either from muscle, breast tissue or excess fat, complicate the recording as the signal is reduced by the increased distance of the electrodes from the heart. Finding the optimum placement of the electrodes in these situations is both difficult and time consuming.

The safety issues with acquiring ECG arise from the cabling of the leads (Kugel et al., 2003). Large loops from improper wiring can lead to subject stimulation, burns, and/or data corruption (Mullinger et al., 2008a). The larger the loop, the larger the induced electromotive force, Equation 1, this, coupled with the low impedance ECG contacts can cause subject stimulation or electrical shocks (Lemieux et al., 1999). The induced voltage from the switching gradients also can saturate the analog to digital converters of the recording amplifier, leading to clipping and data corruption. In addition, failure or improper cabling insulation can lead to subject burns either directly or indirectly. The direct route is the heating of the subject tissue by the conductor or passage of current between the leads through the skin. The indirect route comes about from capacitive coupling of the subject tissue with the conductor. The coupling serves as a route for excess deposition of the radio frequency energy into tissue.

Existing Non-ECG Method (MAS / GFP)

In recordings where the ECG signal was absent or unusable/corrupted, the MAS of the EEG has been used as a substitute for the cardiac timing signal (Srivastava et al., 2005). Using the MAS, the mean of the absolute signal, is a simple way to combine the BCG artifact across channels. In doing so the signal that is common among most electrodes, the BCG, is retained where the signals unique to individual electrodes, the EEG and noise, are suppressed. The implicit assumption however, that there is zero correlation across channels with the EEG and noise, is false (Makeig et al., 1996). In this case though, the BCG is correlated across the head where the correlation of EEG and that of the noise are more likely to be local. Beyond this however, using the mean of the absolute signal has the effects both of decreasing the signal to noise ration (SNR), *Figure 3 black traces*, and blurring its timing. It is due to this blurring that MAS is unable to lead to cleaning results as good as those whose timing is derived from the ECG.

The temporal blurring is due in part to four reasons: rectifying the signal, integration of the artifact timing across the head, integration of the timing of across artifact components, and variability introduced by the respiratory cycle state on the BCG (Prisk, 2001).

As seen in *Equations 1 & 2,* the BCG signal is the first derivative with respect to time of position of the conductive loops. The SP and the HRT components of the BCG cause an oscillatory motion in these loops. This oscillation leads to a biphasic signal, *Figure 3a & Figure 5, top three traces*, whose sign depends on the relationship between the reference and channel electrodes, and their relationship to magnetic field. Rectifying the signals resolves the sign difference, *Figure 3b and Figure 5, bottom trace*, so that

signals with the same morphology, but opposite sign, do not cancel one another in the average. Taking an example of a center reference, Cz, the signal seen by channels over the right and left temporal arteries would be virtually the same, but have opposite polarity. Rectification turns one positive BCG deflection peak into two. In addition it rectifies the "noise" signal eliminating the zero mean nature of both the EEG and other zero mean "noise" signals. The latter merely just decreases the BCG peak amplitude relative to the background, *Equation 5* and *Figure 3b*.

Equation 5: Consequences to the signal to noise ratio (SNR) when rectifying and taking the windowed average of a signal. Here, the signal is composed of the BCG and normally distributed zero-mean noise (a). After calculating the windowed average of the absolute signal (b) the peak value is still equal to the peak value of the windowed value of the absolute signal (c). However the background signal with a zero-mean distribution pre-rectification is no longer is a zero-mean distribution post-rectification, consequently the windowed average of the rectified EEG is no longer zero (d). Therefore while the overall peak signal amplitude remains the same the noise floor increases, resulting in an overall SNR decrease.

(a) $\text{Signal}_{\text{Raw}} = \text{BCG} + \text{Noise}, \text{Noise} = \varepsilon, \text{N}(\sigma, 0).$

(b) $\text{Signal}_{ABS} = \left|\text{Signal}_{\text{Raw}}\right|.$

(c) Peaks: $\overline{\text{Signal}_{ABS}} = \overline{\left|\text{Signal}_{\text{Raw}}\right|} = \overline{\left|\text{BCG}\right|}.$

(d) Background: $\overline{\varepsilon_{ABS}} = \overline{\left|\varepsilon\right|} \neq \bar{\varepsilon} = 0.$

The temporal blurring that comes from averaging the timing of the rectified signal across the head is more problematic. When considering, for example, the SP component, one path blood travels across the scalp from the heart is through the external carotid artery from which it branches out across the scalp *Figure 4*. The time delay between different points across the head of the pulse wave of the SP is proportional to the difference in their distance to the heart, *Equation 3*. This temporal dispersion creates feathering out and blending of the rectified BCG peaks when averaging across the head, *Figure 3c & d*.

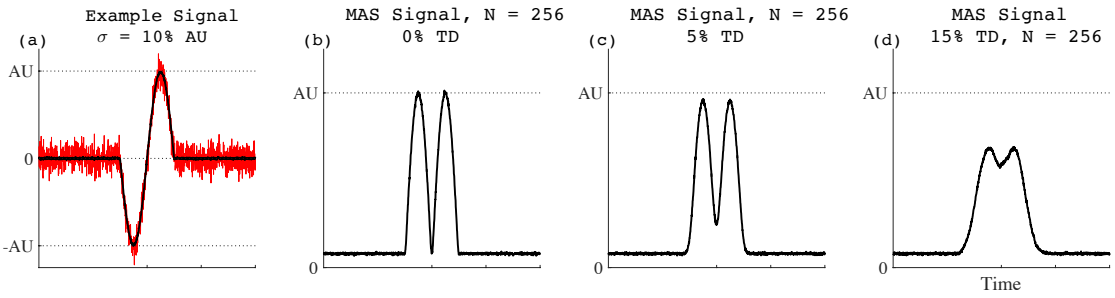

Figure 3: An example of the consequences to both SNR and peak definition from taking the windowed average of signal with background noise (ε = N(σ,0)) after rectification. Here the signal would be the BCG, and the noise represents the actual EEG signal in each channel. An example of the signal from a single channel is indicated by the red trace in (a), where σ =10% of peak amplitude of the peak of the signal. This is an approximation of the ratio of the amplitudes of EEG and BCG when recorded in a 3T MR scanner. The signals in each channel are rectified, then averaged across all of the channels. In (b) the resultant signal post rectification and averaging is plotted. Here it is assumed that the signal occurs concurrently in all of the recording channels. i.e., the standard deviation of the temporal dispersion (TD) = 0% of the period of the signal. Note with even when TD=0 the ratio of signal peak amplitude relative to the noise floor (SNR) has decreased by as compared to signal in the average of the un-rectified signals (black trace in (a)). In (c) a dispersion in the signal timing across the all of the channels is modeled. The standard deviation in the timing across the channels is 5%, TD = 5%, of the period of the signal. For a biphasic signal with a period of 200ms, an approximation from Figure 8, this would a standard deviation in the timing of the BCG events of 10ms. Note both the decrease SNR and the smoothing of the peaks relative to (b). In (d) the TD is increased to 15%. The SNR decrease is problematic in terms of peak detection as a result of the blending of the peaks in the rectified biphasic waveform. In the above examples the noise (EEG) was modeled as uncorrelated across the channels for simplicity. This is of course incorrect. The affect of correlated across the channels would be to decrease the SNR further.

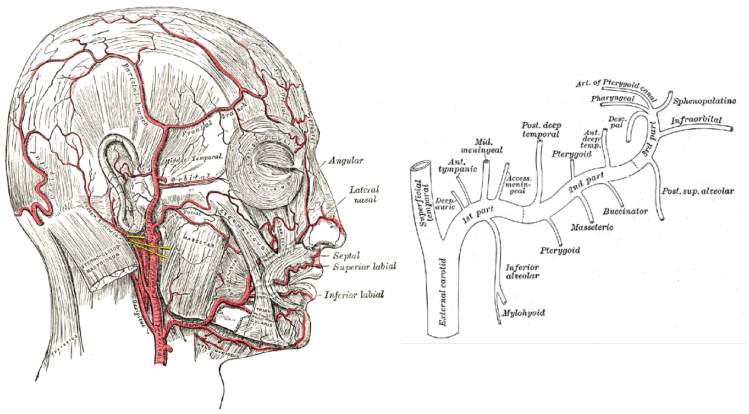

Figure 4: Henry Vandyke Carter - Henry Gray (1918) Anatomy of the Human Body, Bartleby.com: Left: Gray's Anatomy, Plate 508. Licensed under Public Domain via Wikimedia Commons Right: Gray's Anatomy, Plate 511. Licensed under Public Domain via Wikimedia Commons.

By taking the average signal across the head, the three components of the BCG are also blurred. The spatial distribution of the BCG components are different, as are their

timing and amplitude characteristics (Mullinger et al., 2013). The mean of the absolute signal the most weight to the BCG component with the largest amplitude, the HRT. As noted, the HRT has variability in its translation from the QRS to the signal recorded by the EEG system. This leads to variability in the "heartbeat" timing and to a lower quality of artifact cleaning.

As the HRT will be the most heavily weighted component when in using the MAS, it is important to consider how the respiratory cycle affects the HRT. On a simplified level the amplitude of the HRT signal will be at a maximum at the end of inspiration, and at a minimum at the end of expiration (Prisk, 2001). The variability in HRT amplitude will thus vary the contributions of the other components, SP and HE, to the overall timing signal.

New EEG Based Method (LR mean)

Presented here is a new EEG-based method (Rodriguez et al., 2015) to provide cardiac timing that circumvents the complications of acquiring the ECG, yet avoids also the timing variance issues of published methods. The previous attempts to derive cardiac timing from EEG based signals, e.g., MAS, blurred the cardiac timing by mixing the timing of the events across the head and by mixing the of the timing of the individual BCG components, i.e. SP, HRT, HES. To avoid these problems the method presented here, referred here to as LR mean, uses a different base EEG signal that is constructed from careful channel selection and re-referencing. In addition, a new method to extract the ECG R-wave peak equivalent time points is introduced. This algorithm is a robust constrained peak detection that is augmented by pattern matching and multiple layers of automatic error correction.

LR mean Channel Selection

The channels whose signal will comprise the LR mean, *Figure 6*, are chosen based upon the following factors:

- inclusion in the 10-10 EEG montage
- proximity to large scalp arterial vessels
- lateralization
- distance from each other
- distance from facial musculature.

Restricting the subset to the 10-10 EEG system is premised on increasing the compatibility across manufactures and net/cap densities. The distance to scalp arteries is important in ensuring that the SP component will be maximal. The lateralization is to minimize the effects of HRT, and of other artifacts that are symmetrical across the head. Narrowing the spatial spread of the subset is done to maintain narrow peaks, i.e., avoid temporal smoothing. Avoiding facial musculature minimizes the effects of eye blinks and jaw movement on the signal.

Our choice to focus on the SP component was based on three factors: temporal stability, early onset, and predictability of location. The temporal stability of the offset of the SP peak to the ECG R-wave peak, as discussed, is a result of the hydro-mechanical coupling of the two. This coupling also causes the temporal offset from the R-wave peak for the SP signal to be less than that of the HRT. This latency difference has been demonstrated both by Debener et al. (2008) and Mullinger et al. (2013). The predictability comes from *a priori* knowledge of the 10-10 EEG system (Luu and Ferree,

2000, Oostenveld, 2001, Jurcak et al., 2007) channel locations with respect to scalp artery locations (Marano et al., 1985).

Lateralization comes into play for multiple reasons: to minimize HRT, to minimize signals from rigid body motion, and to minimize the effects of symmetrical facial movement. HRT arises in part from avoiding the channels most affected by rotation and translational motion (Jurcak et al., 2007, Mullinger et al., 2013). Lateralization, combined with the reference scheme discussed below, reduces signals that are similar between the sides of the head such as eye blinks, jaw movement, and lateral head movement, *Figure 5*.

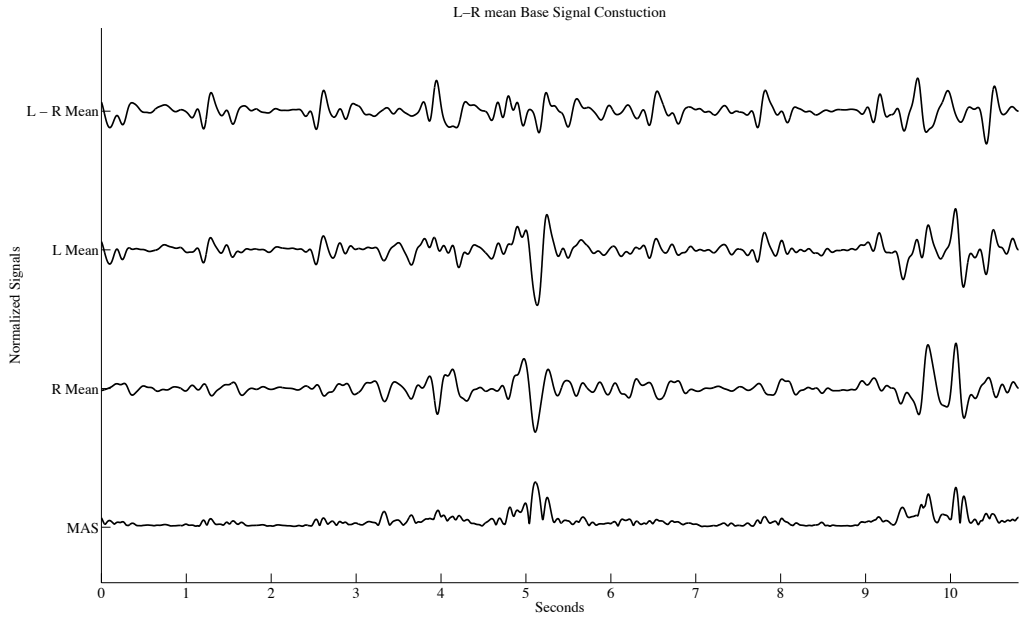

Figure 5: LR mean signal deconstruction and facial artifact suppression. The top trace is an example of the LR mean signal. The left mean and right mean signals and middle top and middle bottom respectively. The bottom trace is equivalent MAS signal. Note the large deflection at about 5 seconds. This is likely to eye blink. In the left and right mean traces the artifact is sizable (more that twice that of their biphasic SP peaks) and fairly symmetric. This same artifact is suppressed greatly in the left – right mean signal (LRmean). The same artifact by contrast is enhanced relative to the background in the MAS. Another likely muscle or motion artifact can be seen at about 10 seconds. Here again the artifact is suppress in the LRmean and enhanced in the MAS. All traces are plotted on the same amplitude scale.

Keeping the spacing between the leads small minimizes the difference in arterial distance between the channels. This is important to minimize the temporal variance of the SP peak onsets. The arterial distance from the heart to channel dictates the SP onset time, *Equation 3*. When the channels differ in distance from the heart, the peak onset times will differ as well, leading in turn to a broader peak, thereby reducing the temporal accuracy of the detection. It is important to keep in mind that the spatial spread (scalp) and arterial spread are related but not in a proportional manner, as it is important to take into account the branching of the arteries, *Figure 4*.

By choosing channels that avoid facial musculature we minimize spurious artifacts from jaw and eye movement. While lateralization, and the reference scheme discussed below, aid in the minimization of these signals the displacement of the electrodes due to facial movement is much greater that that from SP. Unless there is near complete symmetry between the left and right sides, the signals from channels over the facial musculature would dominate the LR mean.

Appling all of the above criteria leads us to select the subset of channels found in *Table 1*. The sets can be reduced further to accommodate subjects with excessive jaw movement and smaller head sizes by removing those channels with the furthest spread and closest proximity to jaw musculature.

LR mean Reference Scheme

Table 1: Left and Right channel subsets. These subsets can be reduced further if cases of excessive jaw movements and small head sizes to only the channels in BOLD.

Left Channels	Right Channels
F9, **FT9**, **TP9**, F7, **T7**, **P7**, C3	F10, **FT10**, **TP10**, F8, **T8**, **P8**, C4

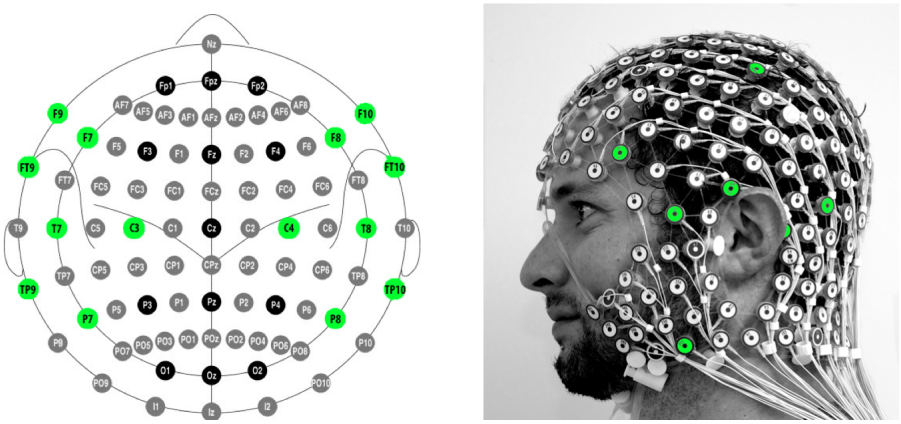

Figure 6 - Left: 10-10 system layout with left and right mean channels hilighted in green. Right: subject wearing EGI's 256 channel MR compatible EEG net with the left mean channel subset highlighted in green.

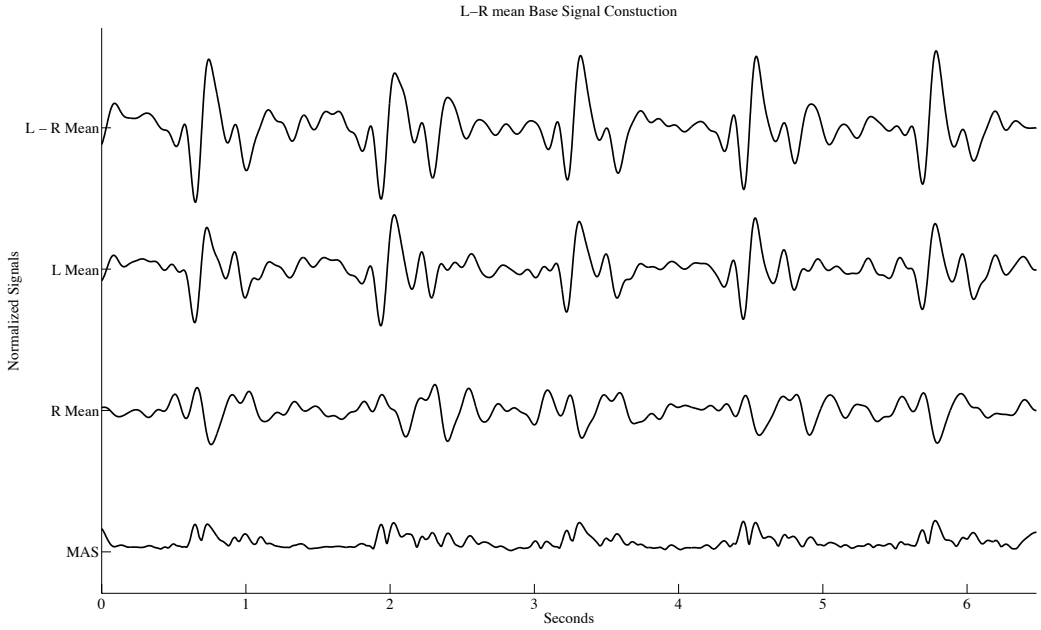

Figure 7: LR mean signal construction and biphasic peak enhancement. The top trace is an example of the LR mean signal. The left mean and right mean signals and middle top and middle bottom respectively. The bottom trace is the equivalent MAS signal. All traces are plotted on the same amplitude scale. Note the enhancement in amplitude of the biphasic peaks relative to baseline of the LR mean signal when compared to it's component signals, the left mean and the right mean signals. The opposite is the case for the MAS.

As mentioned above the selected subset of channels are re-referenced to left right dipole arrangement. Specifically, the mean of the right subset of channels is subtracted from the mean of the left subset of channels. This is done for two complementary

reasons. The first is that this arrangement emphasizes similar signals whose polarities are reversed across the head, e.g. SP, *Figure 7*. The second is that this suppresses similar signals whose polarities are the same on left and right, e.g. HRT, eye blinks and some types of jaw movement, *Figure 5*. Specifically in terms of movement, simultaneous motion in the left and right channels, in the same direction, is suppressed whereas movement in opposite directions is emphasized.

Peak Detection Algorithm

Once we have constructed the base signal, we apply a multilayered peak detection algorithm to identify the BCG timing. We developed a new algorithm, *Figure 8*, for peak detection for two reasons. The first being that the peak amplitude relative to the background of the LR mean signal, *Figure 10, bottom trace*, was much less than that of the ECG based signal, *Figure 1, bottom trace.* In other words the "SNR" for peak detection in the ECG is higher than that of LR mean. The second reason was that the peak detection algorithm must operate on a signal more complex than the EEG. The recorded ECG signal is, for the most part, composed of the true ECG signal, the GA, and other noise. Peak detection is therefor fairly straightforward, in that there is only one dominant peak per cardiac cycle, the R-wave. The LR mean signal, however, is composed of the BCG signals (SP, HRT, HES), the GA signal, the EEG signal, and other noise. The EEG component and the HRT component of the LR mean signal, while minimized, complicate peak detection by adding additional biphasic peaks to the signal.

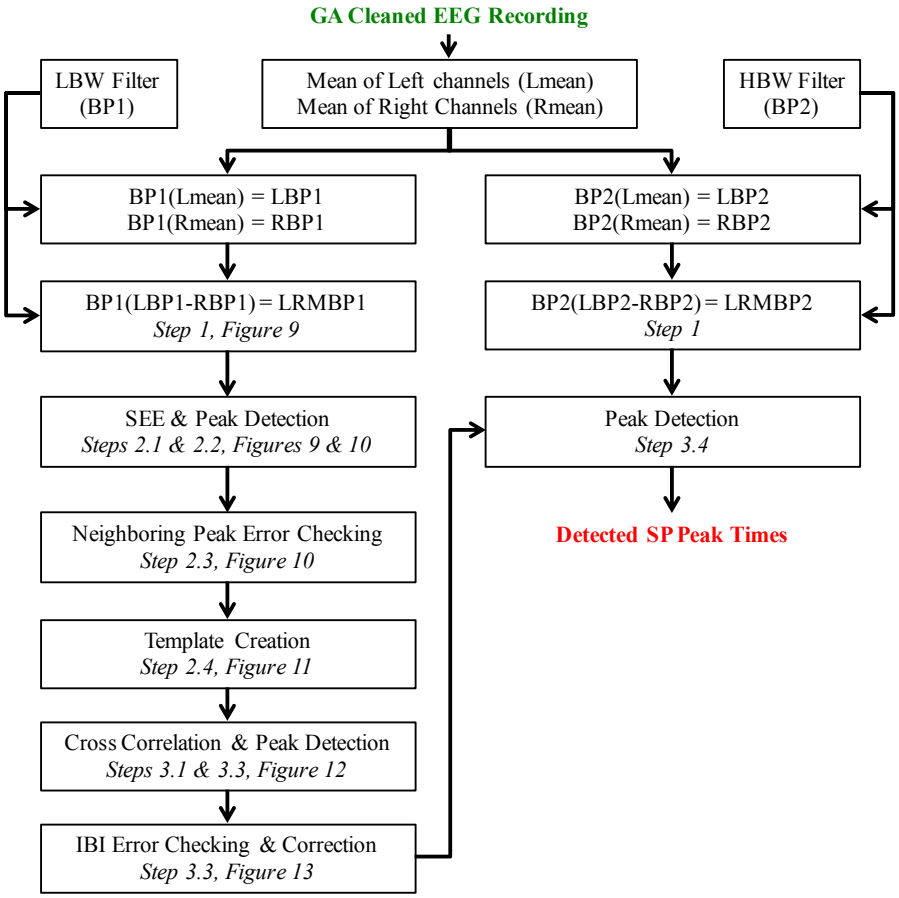

Figure 8: Flow chart of the progression of the algorithm annotated with corresponding steps and figure numbers.

The layers of the peak detection algorithm consist of signal conditioning, Shannon Energy Envelope (SEE) (Choi and Jiang, 2008, Manikandan and Soman, 2012, Zhu and Dong, 2013) peak detection and template creation, and lastly template matching and error checking. The first step of the process entails zero-phase band pass filtering of the LR mean to those bands where the majority of the spectral power of the ECG/BCG exists. In the second step we apply a SEE operator to the filtered LR mean signal to emphasize the biphasic peak of the SP over the background noise and spurious signals, generating the SEE signal. Constrained peak detection is then run on the SEE signal. We use the times detected for the peaks in the SEE signal as reference points in the filtered LR mean signal

for a first pass detection of the BCG peaks. These BCG peaks are then analyzed for erroneous detections based upon their relation to one another as well as their relation to the peaks from the SEE signal. The LR mean peaks that we determine to be correct detections then are used to create a windowed average of the LR mean, a BCG template. In the third step of the process this BCG template is cross-correlated with the filtered LR mean signal to create a pattern-match signal. A second round of constrained peak detection is applied to the pattern-match signal, the peaks of which are run through error checking and correction. The corrected peaks from the pattern match signal become our reference points for peak detection in a wider bandwidth band-pass filtered version of the LR mean. The peaks of this last detection are used as the BCG timing.

Step 1: LR mean Signal Conditioning / Filtering

The first step of the BCG timing detection is to create two different band-pass filtered versions of the mean of the left channels and mean of the right channel subsets as well as their respective difference. That is the left mean is filtered as well as the right mean, then the difference of the filtered left and right means is filtered once more. The two filters are the same in order and type, but their bandwidth differs. To avoid time shifts/lags both filters are implemented in a centered (i.e. zero-phase) fashion. The first filter, the low bandwidth filter (LBW) is a 0.75Hz, 6^{th} order, high-pass Butterworth with a 10Hz, 12^{th} order low-pass Butterworth filter. The second filter, the high bandwidth filter (HBW) is the same as the LBW except the low-pass 3dB point is set to 50Hz instead of 10Hz.

Step 2.1: Shannon Energy Envelope

In the next step of the timing detection, a SEE operator, *Equation 6c*, is applied to the normalized difference *Equation 6b*, of the LBW filtered LR mean signal. Applying the SEE puts emphasis on the middle amplitude values of a signal, i.e. attenuate both low and high amplitude values of a signal (Choi and Jiang, 2008). The difference operator changes the signal from an amplitude signal to a rate of change signal. By doing so the large bipolar deflection, *Figure 7 top trace*, of the LR mean signal is highlighted. The corruption of the difference signal by noise is reduced by the LBW filter. The SEE operator, when operating on a normalized distribution, creates a positive envelope *Figure 9, bottom trace*. By doing so it simplifies greatly the first round of peak detection.

Equation 6: Shannon Energy Envelope Equation of the Normalized Derivative. Here x[n] is time domain representation of the signal, in this case the LR mean. (a) is the difference operation, (b) is the normalization of the difference and (c) SEE[n] is the transformed version of the normalized difference of x[n]. In this formal form of the equation the "*" indicates the complex conjugate, while in reality, (b) would be entirely real-valued. Also, note that the "×" in (c) denotes taking the inner product.

(a) $x'[n] = \frac{d}{dn} x[n].$

(b) $x'_N[n] = x'[n] / \max\left(\left|x'[n]\right|\right).$

(c) $SEE[n] = -\left(\left(x'_N[n] \times x'^{*}_N[n]\right) \times Log\left(x'_N[n] \times x'^{*}_N[n]\right)\right).$

Step 2.2: LR mean Peak Detection Round 1 – SEE

A constrained peak detection algorithm is applied to the SEE signal. It is based on the assumption of 120 beats per minute (BMP) upper limit for a supine resting heart rate. The peaks detected from the SEE signal then are used as reference points for the SP peak search in the LR mean LBW signal. The first stage of this search involves finding the local maxima and minima in the LR mean LBW signal nearest to the SEE peak reference

point. Once these local maxima, minima, and SEE peak sets have been identified they are run through a series of error checks.

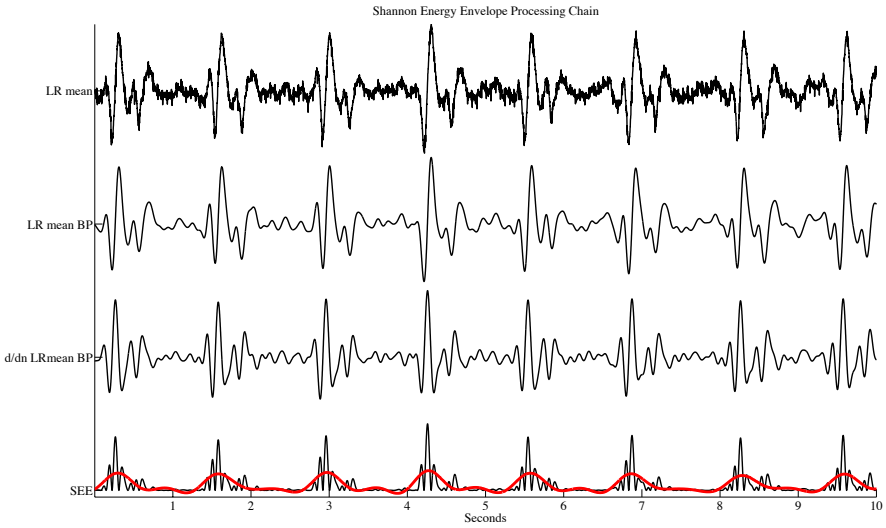

Figure 9: Shannon Energy Envelop Processing Pipeline. The unfiltered LR mean signal can be seen in top trace. The LBW band-pass filtered LR mean signal and its difference signal can be seen in the middle top and middle bottom traces respectively. The SEE signal can be seen in the black bottom trace. The red trace at the bottom is the 2Hz envelope (low pass filtered signal) of the SEE signal. All signals are plotted with normalized amplitude.

Step 2.3: Error Checking Round 1

Ideally, the Shannon Energy Envelope of the of the differentiated LR mean LBW signal yields peaks at the center of the large bipolar deflection of the BCG, *Figure 7*, as such a local minimum peak should be on one side of the SEE peek time point and a local maximum on the other. The temporal order of the minima, SEE peak time points and the maxima will be the same across all BCG events. Whether the maxima or minima occur first depends on the orientation of B_0, which varies by manufacturer. The order the peaks appear in gives the first method for error checking. We do this by checking whether the order of peaks in the set is correct, e.g. minimum before SEE peak time point before maximum or vice verse (depending on the B_0 orientation.) In addition to the order error checking, a second layer of error checking is done by comparing the amplitude and

temporal distances within a peak set (local minima, SEE peak and local maxima) against the averages across sets. Here peak sets outside more than 1.5 standard deviations away from the mean are ignored. Finally, the inter-beat-interval (IBI) between peak sets is considered. Here IBI's of sets above 1.5 seconds (40 BPM) and below 0.5 seconds (120 BPM) are ignored.

Step 2.4: IBI calculation and Mean BCG Pattern Template

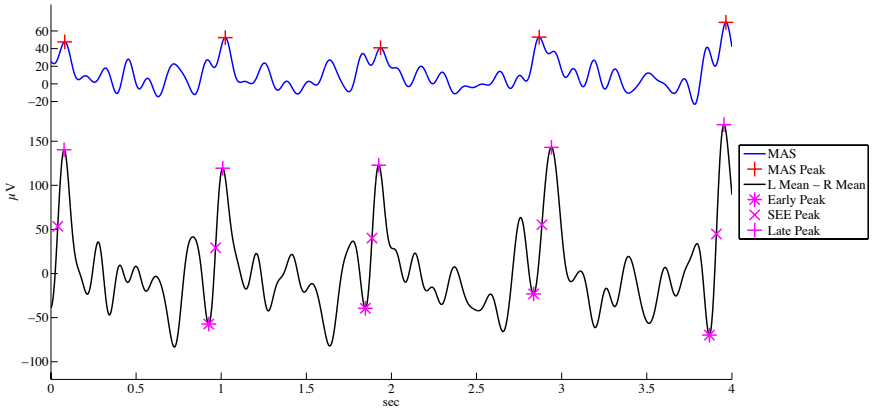

Figure 10: Example peak time points from the MAS signal and LR mean signal. In the MAS signal, blue trace, the correct peak is ambiguous. The late peak is selected peak in 80% of BCG events in the MAS signal. In the LBW LR mean signal, black trace, the SEE peak time point is denoted by a magenta "x". The early peak (the local minimum) is marked with a magenta asterisk, and the late peak (the local maximum), is marked with a magenta cross. Note the nearly 4 times larger peak signal amplitude in the LBW LR mean signal as compared to the MAS signal.

The early peaks from the sets that make it through are then used to create a windowed average of the LR mean LBW BCG, *Figure 11*. The window size of which is determined by taking whichever is larger, the mean IBI minus 2 standard deviations or 0.5 seconds.

Step 3.1: LR mean Cross Correlated with BCG Pattern Template

The windowed average BCG template is then cross-correlated with the LBW LR mean signal to find any missed peaks. The cross-correlation signal has large peaks

corresponding to maximum alignment between the BCG template and LBW LR mean signal, *Figure 12, bottom trace*. These peaks are then run though a second constrained peak detection.

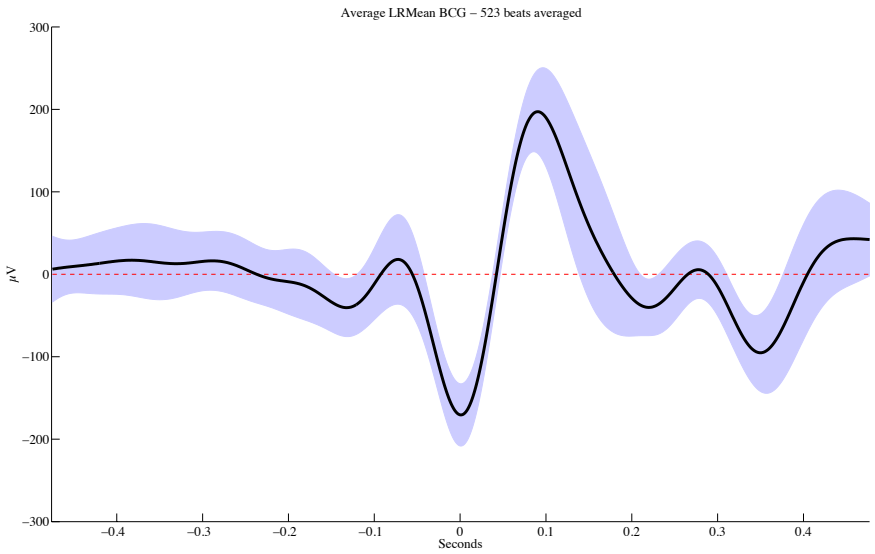

Figure 11: Extended BCG Template. Here the solid black line is the windowed average LBW LR mean BCG waveform centered on the early peak at 0 seconds. The tinted region represents the standard deviation in the BCG waveforms. The dashed red line is zero amplitude for reference. The width of the template in this figure is the mean IBI rather than the mean IBI – 2 standard deviation. This is done for ilustraitive purposes.

Step 3.2: LR mean Peak Detection Round 2 – Cross Correlation

To detect the peaks of the cross correlation a constrained peak detection similar to that in step 2.3 is run. The algorithm though is no longer blind to the data at this step as it was in step 2.3. As such the criteria for minimum separation of the peaks in this step set is based upon the data. Here the minimum separation criteria is the temporal width of the average BCG event. This width is established based up the standard deviation in the windowed average. Using the trace in *Figure 11* as an example the width would be approximately 0.6 seconds or about -125 to 375ms about the early SP peak. The -125ms and 375ms values come from that these are the farthest points away from the SP peak

where 0, *Figure 11 dashed red trace*, is not included within the bound of a standard deviation away from the average waveform, *Figure 11 shaded region*. In the case from *Figure 11*, the 0.6 second minimum beat separation would lead to maximum heart rate of 100 beats per minute.

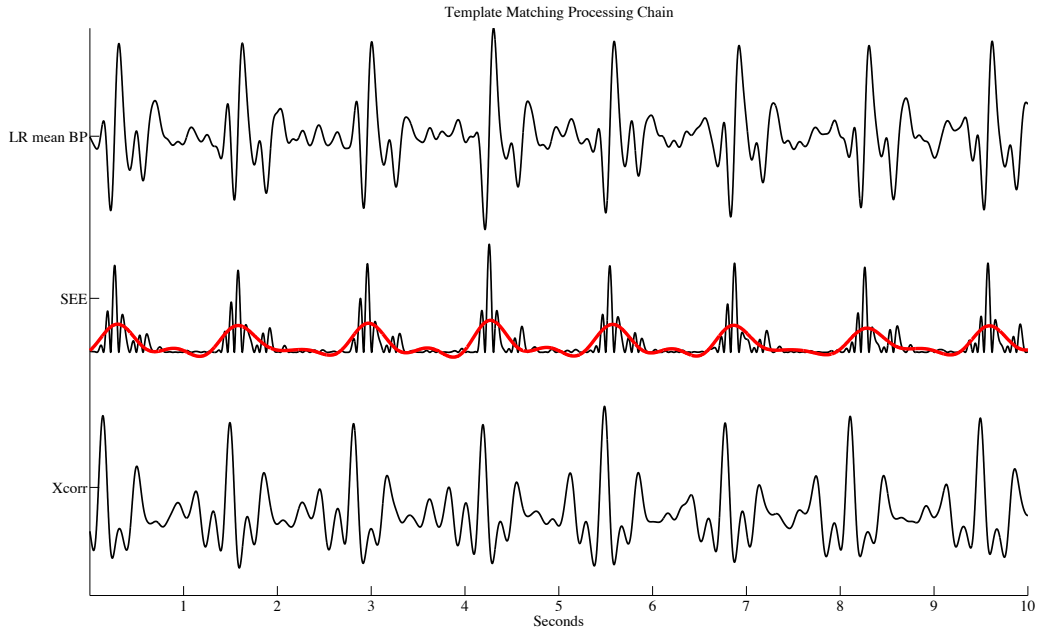

Figure 12: Template Matching Processing Pipeline. The unfiltered LR mean signal can be seen in top trace. The SEE signal can be seen in the black middle trace with a 2Hz envelope red trace. A template from a windowed average of the LBW LR mean signal. The windows for the average are centered around the local minima (or maxima, depending on B0 orientation) preceeding the time points corresponding to the SEE signal. The resultiung signal from the template and LBW LRmean signal cross coreelation can be seen in the bottom trace.

Step 3.3: Error Checking and Correction - Inter-beat Interval Evaluation

With the peaks detected in the cross correlation signal error checking and correction is then performed using the instantaneous inter-beat interval (IBI). Fences of ± 0.1 seconds of a 4-point (beat) centered moving average of the IBI are used to flag outlier IBI's, *Figure 13 red traces*. Errors in detection are found when two successive IBIs are outside the fences. Specifically one of the pair above the positive fence and the other below the negative fence, *Figure 13 indices 330 and 331*. In cases with this successive

bipolar deviation the largest peak within the fenced region of the cross correlation signal is used to replace the erroneous detection.

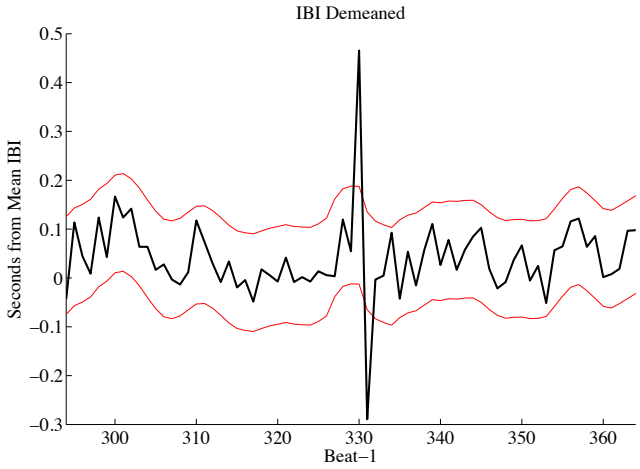

Figure 13: An example of a successive bipolar deviation in the Inter-beat Interval (IBI). Here the black trace is the deviation of individual IBI's from the mean across all IBI's in seconds. The red curves/fences above and below represent ± 0.1 seconds of a 4 point centered moving average of the individual IBI, black, trace. The x-axis is in terms of the index minus 1 of detected peak/beat, as this is a difference measure. Here note the positive followed immediately negative outliers, indexes 330 and 331 respectively. This bipolar arrangement indicates the detected beat at index 331 was likely too far to the right (late) in time.

Step 3.4: Applying timing to Wider Bandwidth

After the IBI error correction the peaks in the HBW LR mean signal closest to those found in the LBW LR mean signal are located. The peak times from the HBW signal are then returned as the SP component BCG peaks times.

Testing the New EEG Based Algorithm

This new EEG based method to extract cardiac timing was designed to avoid the complications and failures associated with acquiring ECG in an MR environment. In addition the algorithm was designed to be robust against artifacts such as eye blinks and subject head movement that would otherwise corrupt an EEG signal. To test the robustness of this method an EEG-fMRI and differential ECG dataset with considerable

known subject movement was used. Here the subject pool was made up of adolescent males with and without (controls) attention deficit hyperactivity disorder.

Methods:

Subjects

A total of 16 boys (10 with ADHD, 12-16 years old) were recruited from the Los Angeles community through flyers, community organizations (CHADD; www.chadd.org), and the UCLA ADHD clinic. Exclusion criteria included: IQ < 80, history of learning disabilities, co-morbid Axis I diagnoses other than oppositional defiant disorder, and current use of psychotropic medications other than psychostimulants. Parents/participants received verbal and written explanations of study requirements, and prior to any study procedures, provided written informed permission/assent as approved by the UCLA Institutional Review Board (IRB#12-000266). No participants, including those diagnosed with ADHD, were on medication during clinical assessment or testing, withholding use 24-48hrs prior to session, consistent with medication half-life.

EEG Recording System

All EEG recordings were performed with Electrical Geodesics Inc.'s (EGI) NetAmps 300 EEG amplifier and recorded using Net Station v4.54. All recordings were made with a Sampling Rate of 1kHz and Cz referenced signal. EGI's Hydrocel 256 MRI compatible net was the electrode system used for all recordings.

MRI System and Scan Settings

All in magnet EEG recordings were performed at isocenter in a Siemens Tim Trio 3 Tesla MRI scanner. The receive coil was Siemens 12 Channel Head Coil and the transmit coil was the Body Coil. 36 slices in a 64x64 matrix of 3mm isotropic voxels were acquired at a TR = 2.160sec and TE = 40msec. Slices were prescribed for full volume coverage, and tilted to minimize contribution of fronto-orbital sinus.

Data Processing

All signal processing and analysis was performed on a 2012 MacBook Pro Retina, OSX 10.10.2, 2.7 GHz Intel Core i7 with 16Gb of ram (Cupertino, CA), running Matlab™ 2012a with psychtoolbox version 3.0.12 (Brainard, 1997), and EEGLAB version 13.4.4b (Delorme and Makeig, 2004).

GA cleaning

MR gradient artifact cleaning was performed using the moving windowed average template subtraction algorithm described in the previous chapter. The settings for the moving windowed average template subtraction algorithm were a centered filter with a window width of 13 TR intervals, and no amplitude weighting was applied across TR intervals.

ECG Peak Timing Extraction

The method described by Zhu and Dong (2013) was used to extract the R-wave peak timing from the ECG recording. The R-wave peak of the QRS complex was used instead of the complex onset time due to signal to noise limitations in the ECG recordings.

BCG Peak Timing Extraction

The method described above for LR mean peak detection was used to extract the BCG SP component time points.

Determining Corrupted Recording Sections

The following process was used to determined areas of possible signal corruption, spikes, by excessive subject movement or by the pre-scan pulses of the MRI in either the EEG of ECG traces. Fences of the mean ± 5 standard deviations values of GA cleaned in the LRM signal and mean ± 7 standard deviations in the ECG signal were used to identify corrupted recording spikes. A padding ± 200msec about each of the spikes was used to set a buffer zone. The EEG and ECG corrupted recording sections were combined to demark exclusions areas for lag comparison.

ECG & BCG Missed and False Detection

Equation 7: Calculations for both sensitivity (a) specificity (b) and accuracy (c). Here DP_{True} is the number of correctly detected peaks, MP is the number of missed peaks, and DP_{False} is the number of incorrectly detected peaks

$$\text{(a)}\ \ S_E = 100 \times \frac{DP_{True}}{DP_{True} + MP}.$$

$$\text{(b)}\ \ S_P = 100 \times \frac{DP_{True}}{DP_{True} + DP_{False}}.$$

$$\text{(c)}\ \ A = 100 \times \frac{DP_{True}}{DP_{True} + DP_{False} + MP}.$$

Manual inspection was performed on both the ECG and LR mean traces across the whole recording. All peaks were counted either true, false, or missed outside the above mentioned exclusion areas. In the exclusion areas if the true peak could be identified by manual inspection in either of the traces, true, false or missed peak tallies

were made. Calculations of sensitivity (S_E), specificity (S_P) and accuracy (a) using Equations 7a-c respectively.

BCG-ECG Peak Lag Comparison

The temporal lag between the early SP peak of the LR mean and the R-wave peak of the ECG was calculated by simple subtraction. The mean and standard deviation of which where then calculated.

Results & Discussion:

Overall data quality

In both data sets both peak detection algorithms worked well. The peaks detected being the early SP peak for EEG data and the R-wave peak for the ECG data. For EEG the overall sensitivity was 98.63% and specificity was 96.95% across all data sets. For ECG it was 98.19% and 99.53% respectively across 14 of the 16 data sets. Two sets, numbers 8 and 9, were omitted from the calculation due to failures in the ECG recording. Most likely the failure was due to the loss of an ECG lead contact. The results of both the algorithms and manual inspection can be found summarized in Table 2. The EEG also had trouble in 2 sets, 11 and 15, in that the majority of data in the right channels in both sets was unusable. This could have been due to the channels being disturbed by the padding placed about the subject's head or from the placement of headphones over these electrodes. In these two sets the left channel mean referenced to Cz was used rather LRM. If these two sets are excluded from the overall sensitivity and specificity calculations the EEG peak detection has 98.58% and 97.84% respectively.

Table 2: Tabulated Results the detection algorithms and from manual inspection. Here bold font indicates a case subject, gray highlight indicates subjects where the ECG acquisition failed, and gray font indicates BCG peak detection where only one side (left in both cases) of the EEG was used.

SET	BCG						ECG						Useable Data %			Lags (msec)	
	DP_T	DP_F	MP	Se	Sp	A	DP_T	DP_F	MP	Se	Sp	A	EEG	ECG	Both	μ	σ
1	441	19	6	98.66	95.87	94.64	460	5	1	99.78	98.92	98.71	95.45	98.98	95.45	126.39	4.57
2	**814**	**3**	**10**	**98.79**	**99.63**	**98.43**	**817**	**0**	**10**	**98.79**	**100.00**	**98.79**	**97.99**	**99.11**	**98.62**	**184.34**	**5.04**
3	**664**	**41**	**9**	**98.66**	**94.18**	**93.00**	**700**	**4**	**10**	**98.59**	**99.43**	**98.04**	**95.72**	**92.84**	**90.01**	**129.04**	**4.88**
4	**519**	**11**	**6**	**98.86**	**97.92**	**96.83**	**519**	**4**	**17**	**96.83**	**99.24**	**96.11**	**94.84**	**97.51**	**94.52**	**136.67**	**3.34**
5	799	3	8	99.01	99.63	98.64	779	1	30	96.29	99.87	96.17	98.65	98.91	98.65	152.18	3.66
6	704	8	19	97.37	98.88	96.31	677	5	49	93.25	99.27	92.61	97.35	99.18	97.35	123.11	5.40
7	618	18	5	99.20	97.17	96.41	624	3	14	97.81	99.52	97.35	90.59	97.68	89.41	132.38	6.74
8	614	7	12	98.08	98.87	97.00	N/A	N/A	N/A	N/A	N/A	N/A	97.28	N/A	N/A	N/A	N/A
9	**812**	**8**	**19**	**97.71**	**99.02**	**96.78**	**N/A**	**N/A**	**N/A**	**N/A**	**N/A**	**N/A**	**96.24**	**N/A**	**N/A**	**N/A**	**N/A**
10	**679**	**36**	**12**	**98.26**	**94.97**	**93.40**	**708**	**6**	**13**	**98.20**	**99.16**	**97.39**	**93.43**	**98.64**	**92.92**	**121.83**	**6.33**
11	895	48	8	99.11	94.91	94.11	937	7	4	99.57	99.26	98.84	95.53	97.73	96.24	63.78	7.09
12	**789**	**3**	**8**	**99.00**	**99.62**	**98.63**	**784**	**2**	**14**	**98.25**	**99.75**	**98.00**	**97.21**	**99.08**	**97.21**	**141.56**	**5.71**
13	**838**	**32**	**9**	**98.94**	**96.32**	**95.34**	**879**	**0**	**0**	**100.00**	**100.00**	**100.00**	**96.63**	**98.68**	**96.46**	**116.03**	**4.82**
14	**730**	**16**	**8**	**98.92**	**97.86**	**96.82**	**730**	**5**	**19**	**97.46**	**99.32**	**96.82**	**95.44**	**97.76**	**95.31**	**130.83**	**5.51**
15	832	130	10	98.81	86.49	85.60	954	3	0	100.00	99.69	99.69	94.86	98.93	94.76	130.87	8.17
16	507	1	7	98.64	99.80	98.45	514	0	1	99.81	100.00	99.81	97.93	98.18	97.92	136.76	2.53
μ	702.07	26.36	8.93	98.63	96.95	95.65	720.14	3.21	13.00	98.19	99.53	97.74	95.83	98.09	95.35	130.41	5.27
σ	139.12	33.55	3.43	0.51	3.37	3.23	153.49	2.33	13.49	1.85	0.35	1.94	2.10	1.62	2.89	25.39	1.52

ECG results

The two most prominent features of the ECG data are its magnitude and its acquisition failure rate. The ECG signal magnitude is on the order of at least ten times greater than the BCG artifact and a hundred times greater than the EEG itself. It is to the point where its signal magnitude becomes comparable with that of the GA. With the equipment and field strength in this study this is the expected range, it is also the reason for the different fence values between the EEG and ECG signals in marking so-called bad regions. The failure rate of the ECG, 12.5%, in this study may be somewhat inflated due to the nature of the subject population, adolescent males, 62.5% of which have a clinical diagnosis of ADHD. Even with this subject population 98.09% of the captured ECG signal was useable. In addition to the demographics of the population in this study, subject body size and shape, may have played a role in the failure of the ECG recoding in subject 8 as well. Larger subject size, as mentioned earlier, reduces the signal quality of the ECG as well, and can lead to increased discomfort in the scanner. Increased

discomfort leads to increased subject motion that can lead to "popping" one or both of the ECG leads causing a failure in the ECG recording altogether. The lack of redundancy, i.e. a simple difference measure, in the ECG system of the current EEG MR compatible amplifiers means that only one electrode needs to fail for the ECG recording to fail. Coupling the motion prone population with the less than robust ECG recording likely increased the ECG failure rate. While the exact cause of the failures cannot be ascertained definitively, the lack of redundancy / robustness of the ECG is known to be a problem in EEG-fMRI (Srivastava et al., 2005).

An additional known issue related to measuring ECG inside in MRI environment is the apparent inversion of the T-wave, *Figures 1, 14 & 16*. When measured outside of large standing magnetic field inversion of the T-wave signal in ECG can be indicative of coronary ischemia, Wellen's Syndrome, a central nervous system disorder, left ventricular hypertrophy, as well as many other maladies. This is not the case here. The "T-waves" seen in the above mentioned figures are actually the superposition of 3 signals. The first signal is the end of the S-wave, the completion of the depolarization of the Purkinje fibers, at which time the aortic value opens and blood rushes from the heart. The second signal is the voltage difference across the subject chest resulting from the Hall effect as charged particles in the blood become separated as they flow over the aortic arch. The last signal is the true T-wave, the repolarization of the ventricles. This superposition of signals from here on will be referred to as the T-wave trough (TWT).

Another notable feature regarding the ECG signal is the success rate for the Shannon Energy Envelope based peak detection on the in magnet recording. The ECG peak detection used here was based upon the work done Zhu and Dong (2013). They

reported in their paper an average peak detection specificity of 99.92% and sensitivity of 99.92% for recordings made out of magnet. The in-magnet ECG recordings from this data set, all of which went through gradient artifact cleaning, had an average sensitivity and of 98.19% and specificity of 99.53%. These numbers also were comparable to the performance reported by Fisher, et al, in their 1999 paper, in which they evaluated the use of a 12-lead vector cardiogram and found sensitivity of 100% and specificity 99.8% respectively. This is not withstanding the fact that a 12-lead vector cardiogram ECG recording has a comparatively higher signal to artifact ratio than a 2-lead differential ECG recording (Mullinger et al., 2008b). In addition, Fisher, et al, had to simulate the GA artifact in their data as their ECG system was not MR compatible. In their simulation the GA was unchanging between TRs, which is an unrealistic assumption (Allen et al., 2000), and their subject population consisted of only a single healthy male volunteer.

An interesting side note from the data in this study is the difference in the regularity of the IBI between the case and control subjects in this study. It appeared that the case subjects had less variability in their IBI overall as compared to the control subjects. On a few of the case subjects there appeared to be an increase in the high frequency content of the IBI. These increases have been correlated to attention level in subjects with ADHD (Börger et al., 1999). While a clear narrow band respiratory sinus arrhythmia (RSA) (~0.1Hz variability in the IBI) was observed in the most of the control subjects, in most of the case subjects, this variability was not detectable, not as tightly banded or shifted to a higher frequency.

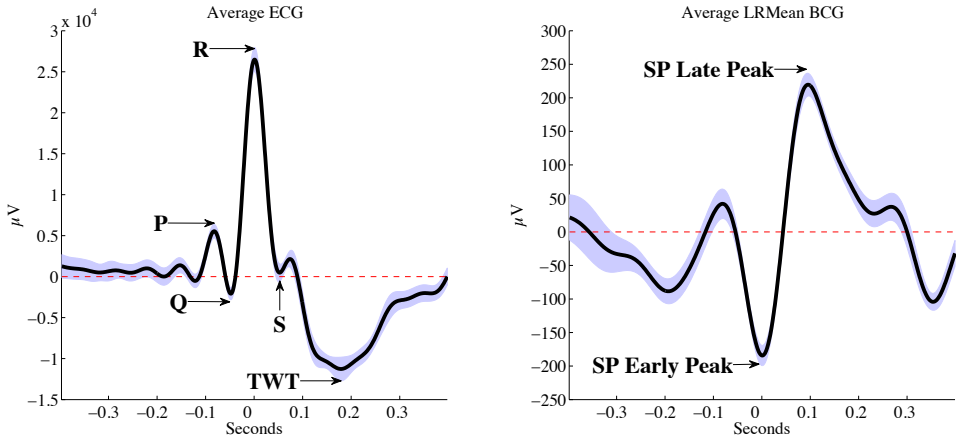

Figure 14: Example of average heartbeat and BCG event from Set 5. Here is the left plot is a R-wave peak centered average of ECG and the right plot is a negative (early) BCG peak centered average. For reference the P, Q, R, & S waveform peaks are labeled on the ECG trace. In addition on the ECG trace the T-wave trough (TWT) is also label. The TWT is a blending of the end of the S-wave, the T-wave and the signal from the Hall effect as the blood ejected from the heart passes over the aortic arch. In the BCG trace the early and late peaks of the SP biphasic waveform labeled. In both plots the shaded region represents one standard deviation away form the mean signal (black traces).

In a final note on the ECG peak detection algorithm used in this study, the only peak used was the R-wave. There might be potential for improving performance by using more of the waveform peaks, e.g. P, Q, & S. This would be performed by finding each of the peaks, then using their relation to one another both inter and intra-beat as an additional means of error checking.

BCG results

The LRM data had two notable features, contamination from subject motion and robustness to complete signal loss. Here the LRM signal of the EEG showed greater motion contamination compared to the ECG. Even so, the cardiac timing could be established in in all 16 datasets.

Applying the amplitude thresholds (5σ and 7σ respectively) to LRM and ECG signals the algorithm labeled 95.83% of the LRM data motion or MR prescan artifact free, compared to 98.09% of the ECG data. However, the periods considered corrupted

by the algorithm from motion or prescan pulses in the ECG data overlapped more than 99.5% of the time with the periods considered corrupted by the algorithm in the LRM data. The different fences, 5σ for EEG and 7σ for ECG, may have contributed to the differences in the amount of data being considered corrupted by the algorithm. However, the magnitude difference of approximately 10 fold between the two signals, BCG and ECG necessitated the different thresholds. Another factor that likely played a role in these differences, was that head moves more than the torso while in the scanner. As the head is a fairly rigid body it is safe to assume though that motion contamination seen in the LRM likely is common across all electrodes. The fact that the LR mean minimizes the signal from these motions probably means that the EEG data as a whole was effected as well. As such the unusable data in terms of the LR mean is more or less a moot point in that BCG cleaning unusable EEG data does nothing in that the EEG data are already unusable.

The robustness of the LRM method was demonstrated in subjects 11 and 15 where the signal in the channels on the right side subset in both cases was unusable. Upon manual inspection of the signal it appeared that there was a loss of contact between the scalp and electrodes. One possible cause of this is that these electrodes where perturbed as a set either during the padding of the subject's head against the MR head coil or during the placement of the headphones on the subject. Even with the loss of half of the electrodes, which turned the LRM to just a LM with CZ reference, the sensitivities for each set were 99.11% and 98.81%, and specificities were 94.91% and 86.49% respectively. The reduced specificity was likely due to the loss of the benefits of

lateralization. While the specificity is significantly lower than the group average it is better than no signal at all.

ECG specificity, sensitivity and accuracy compared to the BCG

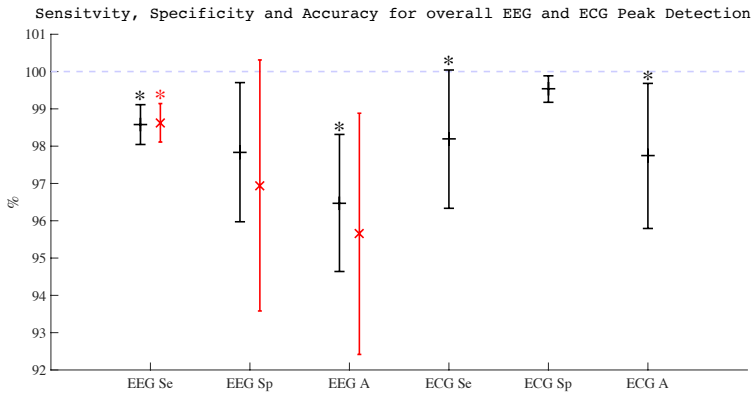

Figure 15: Measures of the EEG and ECG based peak detection. The error bars are displayed as μ ± σ for sensitivity (Se), specificity (Sp) and Accuracy (A). The black bars for the EEG values were calculated using all 16 data sets. The red bars for the EEG values were calculated using only the data where both sides (left and right) of the EEG were used (Black Font, Table 2). The ECG error bars were only calculated using the ECG data where a signal was present, i.e. not including subjects 8 & 9. An asterisk above an error bars indicates no significant difference in the measures was found between the EEG and ECG based methods. Significant difference was calculated at a P = 0.05 using two sided t-test.

Comparing the measures of sensitivity, specificity, and accuracy between ECG and EEG based peak detection a two-sided t-test with a significance level of 0.05 was used, *Figure 15*. No significant difference was found however in the sensitivity measures. The null hypothesis could not be rejected however with regards to the mean of the specificity measures. The lower specificity value for the measure for the EEG based method is likely due to the waveform shape between the ECG and LRM BCG. Rather than having a single large distinct peak, R-wave peak, the LRM BCG waveform has several. These several peaks while rare can lead to incorrect peak identification by the algorithm. A possible solution to this is to add an additional layer of error checking based upon intra-beat intervals within each BCG waveform. The accuracy if including the two

subjects where only the left side EEG channels were used, no gain of the bilateral reference, the null hypothesis again could not be rejected. However, when only comparing the LRM with the ECG, even with the lower specificity measure there is no significant difference in the accuracy measure between the methods.

ECG to BCG lag

Performance-wise the EEG based peak detection was comparable to the ECG based method. Beyond specificity and sensitivity, the lag between the EEG and ECG was also stable within subject with an average offset of 130.41 ± 5.27 milliseconds (SD) between the R-wave peak and the early peak in BCG of the LRM. This variance is most likely attributable to five sources: the sampling frequency, clock synchronization between the EEG system and MR scanner, the temporal smoothing of the BCG peak from the spatial spread of the EEG channel subset, the temporal smoothing from the HBW band-pass filter, and finally using the R-wave peak as a correlate for the opening of the aortic valve.

The sampling frequency affects the lag calculation in so much as the certainty of the location of a detected peak is at its best only good ±0.5 sample periods of where it was detected. This limit on peak location accuracy when factoring the distance between two peaks gives a certainty of ±1 sample periods in the lag measure. In the case of the data here which was acquired here the gives a lower limit on the accuracy of ±1 millisecond.

Synchronization between the EEG's A/D clock and MR scanner clock and is performed to assure optimal matching between the GA across volume periods (TR's). In doing so a slight, sub sample period jiggering occurs. If the modification to the timing

occurs between the ECG and BCG peak this will lag measure. No record however is kept by the EEG system of when and how much of a time shifts was made. The ratio of the average temporal offset between the peaks (130.41 milliseconds), to the TR of these recordings (2.160 seconds), is about 6%. This means on average a clock synchronization would have occurred between the two peaks, R-wave and early SP, 6% of the time.

The HBW band-pass filter, which is necessary to suppress the imperfect GA cleaning of high frequencies, unfortunately results in temporal broadening of both the peaks of the ECG R-wave and the BCG SP peaks. However, as mentioned earlier, the majority of both the ECG and BCG spectral power lies below the 50Hz cutoff of the HBW band-pass filter.

As discussed in the channel selection section above, the locations of the left and right channels was chosen to minimize temporal dispersion of the peaks and facial movement artifacts while maximizing robustness and suppression of the EEG signal. The resultant tradeoff does leave some temporal blurring of the peak, though this no where as bad as averaging across the head.

Finally, calculating the lag is done here using the R-wave peak, early ventricular depolarization, as a marker for the ejection of blood from the heart rather than at the end of S-wave, the end of the depolarization of the Purkinje fibers and contraction of the rim of the ventricular muscle. It is after the Purkinje fibers have depolarized that the aortic valve opens and blood is ejected from the heart. The reason for using the R-wave peak over the S-wave end is two fold. Firstly, the R-wave peak is much, and therefore easier to detect. and robust against error. Secondly, and more importantly, where the S-wave ends and the T-Wave starts is difficult to determine, see *Figures 14 & 16*. In the right plot of

Figure 16, the RST waveform zoom, a clear variance can be seen in timing of the S-wave peak and S-T wave transition zone with respect to the R-wave peak. As the scalp pulsation relies on the ejection of blood from the heart this variance is reflected in the lag between the R-wave and BCG.

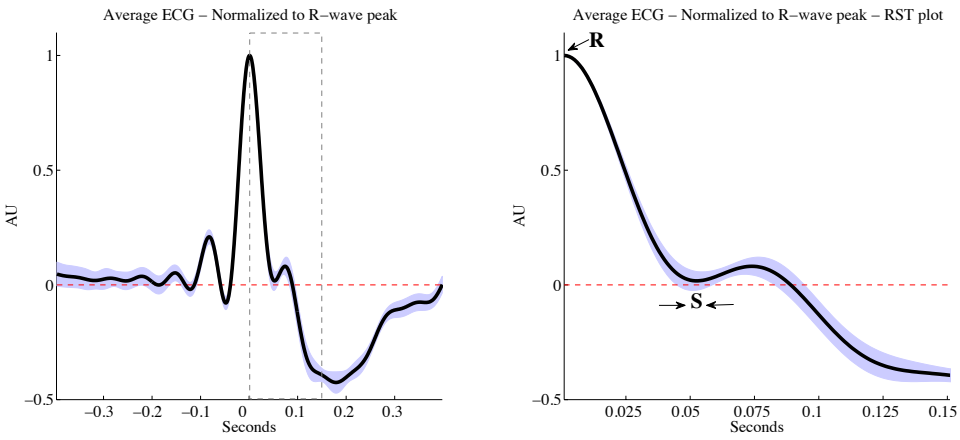

Figure 16: Averaged ECG from Set 5 normalized and aligned to R-wave peak value. Left: average PQRST waveform trace, black, with ± 1 standard deviation shaded in violet, dashed box represents zoom region of plot on right. Right: zoom of left plot, specifically the R-S-T transition. Note the temporal variance of the S-wave peak with respect to the R-wave peak. Also note the temporal variance of the S-T transition zone (right of S-wave peak) with respect to the R-wave peak.

Moving from the within subject standard deviation of the lag measurement to the across subject average lag, a value of 130.41 milliseconds was measured. This number is much different from the accepted standard lag (Vanderperren et al., 2010) of 210 milliseconds between the QRS and BCG. Allen, et al, in 1998 derived this value from the largest peak value in the BCG which corresponds to the HRT component (Yan et al., 2010). In their 2013 paper, Mullinger, et al, found that by isolating peak components that the SP appears earlier, and at approximately 140 milliseconds after the R-wave peak, which is much more in line with the value found here.

A quick comment on the outlier value for lag of 63.78 milliseconds in set 11. This is the result of the peak detection algorithm of using positive peak just prior to the early SP. It was the lack of the LR reference in this set that lead to the peak detection algorithm choosing the incorrect peak.

ECG and BCG windowed averages

Lining up the windowed BCG and ECG averages based upon the average lag between the two reveals several interesting features, *Figure 17*. Firstly, the difference in temporal breadth between the two signals is substantial. Secondly, the SP wave of the BCG and the T-wave trough (TWT) look very similar in temporal breadth. Lastly, there is a biphasic wave in the BCG that appears before the SP biphasic wave whose initial peak occurs around the time of the Q-wave of the ECG.

The overall temporal dilation of the BCG, compared to the ECG signal, is due to the means by which the two signals come about. The PQRST waveforms of the ECG arise from temporally discrete electrical potential changes, polarization and depolarization, of tissues in the heart. The BCG signal, as mentioned earlier, arises from three sources, two of which are from the displacement of the EEG electrodes in the magnetic field, SP and HRT, and the other from electrical potentials being generated on scalp from the separation of charge in the blood due to its motion in a magnetic field, HES. As discussed earlier the mechanics of the translation the heartbeat to each of the respective component signals has differing amounts of both temporal offset stability and temporal dilation from the original heartbeat signal.

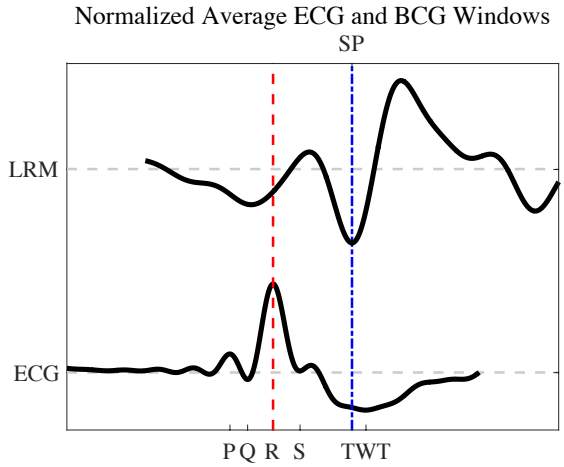

Figure 17: Subject 5 peak value normalized windowed averages of the ECG and BCG waveforms. The waveforms are aligned by their temporal lag. Gray dashed lines represent the zero voltages of the respective windowed averaged signals. The red dashed line corresponds to the timing of the R-wave peak and the blue dashed line corresponds to the early scalp pulsation peak (SP) of the BCG. Here the average lag is distance between the red and blue dashed lines.

One way the timing of the transfer of signal from heart and blood motion to BCG can be seen is in the temporal breadth of TWT and SP biphasic peaks. The TWT signal is generated by the flow of blood from the heart. As blood flows from the heart it creates the depression seen in following the S-wave in the ECG. While the ECG signal is depressed there is an active flow of blood out of the heart, thus the temporal breadth of the signal is the time over which blood is moving from the heart. In the LRM signal this is seen as the biphasic signal whose early peak is demarcated at "SP" in *Figure 17*. As blood is pushed out from the heart it fills the arteries, eventually reaching the scalp, where it displaces the electrodes outwards creating the downward signal in the LRM. As the blood flow starts to wane the electrodes move inwards creating the upward deflection in the LRM. As both the TWT and SP signals rely on active blood flow from the heart it is not surprising that the two seem to match one another closely in temporal duration. The offset of these two signals is the travel time of the blood pressure wave from the heart to the head. Unfortunately, there is not a clear marker for the TWT start and the best

approximation of the location, the downward zero crossing of the ECG post S-wave, will be readily affected by the smallest noise in the ECG signal.

The early biphasic waveform visible in the average LRM trace in *Figure 17* appears roughly to reach its first peak around the Q-wave of the ECG trace thus before the ventricular contraction. The question then is whether this earlier waveform is part of the same heartbeat as the SP waveform on which the window is centered, or an issue with the how the window is framed. By expanding the window in *Figure 11* to include the surrounding BCG events, *Figure 18a*, the specific heartbeat in the BCG to which these peaks belong. Specifically, this is done by minimizing the summed variance across a window width that is equal to mean IBI, *Figure 18b*. That is to say, the summed variance will be minimized in a window whose peaks are most tightly tied in both amplitude and timing with the window reference point, the early SP peak, and thus the same BCG event/heartbeat as the reference point. This method can also be used to determine ideal offset/lag for using the SP time points in BCG cleaning. The leading and lagging edges, green and red lines respectively, mark the region most tightly tied to the SP peak at the center of *Figure 18a*. The early waveform in question lies within this region. This peak then must be the result of the atrial contraction of the heart, the P-wave. The method by which the P-wave is translated from the heartbeat into this early waveform is likely primarily through HRT as both HE and SP are driven by blood flow in the head of which atrial contraction has a minimal impact. The temporal dilation, ECG to BCG, in the case of subject 5 approximately 30%, coupled with the peak being on the leading edge of the BCG event window likely means that this early waveform is influenced by the previous heartbeat. The nature by which this early waveform arises as well as it's the likelihood of

it being influenced of the previous beats makes it an ill candidate for a reference point for BCG cleaning. Even so, it should be included within the BCG window for cleaning.

Case Versus Control Peak Detection Performance

Based on a two-sided t-test, p = 0.05, no significant difference between was found in the sensitivity, specificity, and accuracy measures between the case and control subjects.

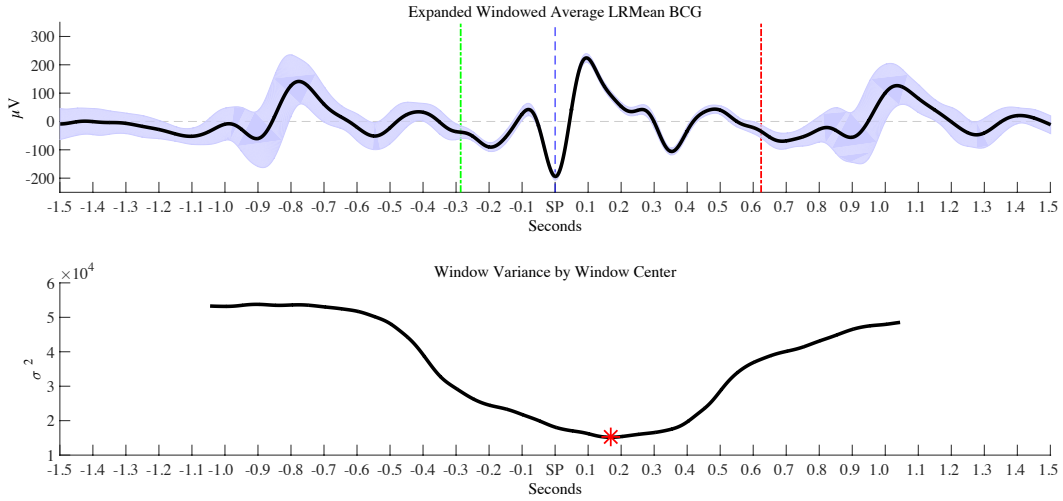

Figure 18: Measuring the window extents of the BCG artifact in the LR mean signal. In the top plot, (a), a 3-second windowed centered on the early detected SP peaks. The data in each window is normalized to the early SP peak amplitude. The average signal is represented by the black trace with region of ± 1 standard deviation shaded in violet. In the bottom plot the summed variance in a sliding window moving across the expanded windowed average. The width of the sliding window being equal to the average inter-beat interval. The location of the minimum value in the bottom trace indicates the location of the center of window that minimizes the variance in the windowed average, top trace. The leading and lagging edges of this window are denoted by the green (early) and red (late) dashed lines respectively in the top plot. Bleed over from previous events as well as the simplification of using a fixed window width leads to some flaring in the variance at the inside edges of the window seen here in the flair up of variance <-0.25 and >0.55 within the bounded region, top plot.

Conclusion:

Using the ECG-free method and algorithm, LRM, BCG events were detected successfully and their event timing determined. No significant difference in sensitivity or accuracy between the LRM and ECG method was found. A slight decrease, 1.69%, in the

specificity was found in EEG based method as compared to the ECG based method. The overall latency between the R-wave peak and the SP peak was very stable within individual subjects. The majority of the variance was likely due to the variability in timing of the R-Wave peak and of the ejection of the blood from the heart, the opening of the aortic value at the end of the S-wave. Beyond just providing a point to window about the algorithm also provides the extents both forward and backward of the window. By providing the window extents the method presented here can be used by all of the various existing BCG cleaning algorithms. In addition to providing a means to avoid acquiring the ECG, the method presented here can be used retrospectively on EEG acquisitions where the ECG data was corrupted or not acquired.

Future Improvements

The accuracy of peak detection in the LRM and ECG signals might be improved by using multiple points for validation. That is, in the case of the ECG signal, to use the measures of inter-beat and intra-beat intervals as well as the respective amplitudes between the P, Q, R, S waveforms as additional means of validation / error checking. Likewise, for the LRM signal the inter-beat and intra-beat intervals as well as the respective amplitudes of the HRT peaks and SP peaks can possibly provide additional means of validation / error checking.

Future Study

Beyond providing a means to eliminate ECG, and to recover EEG recordings without useable ECG, the detection of BCG events and estimation of their temporal extent, opens up many possibilities for future study.

First and foremost is how various BCG cleaning algorithms outputs are changed based upon the new data points. While the ECG R-wave peak and early SP peak of the BCG show a fairly locked lag the variance introduced by using the ventral contraction, as in current ECG based methods, rather than blood ejection, will most likely worsen cleaner performance. Inferring timing from a correlated source, the ECG, rather than extracting the timing directly via the LR mean, will never provide as accurate as a result, as it relies on an estimate. However, it is still worthwhile to studying the difference between the two.

Secondly the time points provided by this algorithm serve as a means to determine the BCG event temporal extent within each channel. That is, the timing of the BCG is different across the head. As such, the temporal extents of each event across the head will be different with respect to the SP peak of the LRM. This is true whether using ECG or BCG based timing. The method of looking at the variance across a wide windowed average discussed above and demonstrated in *Figure 18* can be used to determine the temporal extents within channel. It is this timing that should be used for the BCG cleaning. Moreover, with the temporal extents of the BCG events within each channel it would be worth exploring methods for measuring the pulse wave velocity from the EEG recording. The electrode locations with respect to one another can be determined directly from structural MRI scans, see the following chapter. With the addition of an angiography scan to establish atrial distance between electrodes accurate measurements of pulse wave velocity should be possible.

Third it seems reasonable and likely that there might be overlap in BCG event windows. With the temporal dilation from the translation of the cardiac motion into the

BCG coupled with a short IBI / fast heart rate one BGC event most likely will run into the next. This lack of discreteness to the BCG event windows can be thought of as a form of system memory. Specifically, it is worth exploring what effect this system memory has on BCG cleaning using the existing cleaning algorithms.

Finally it would be interesting to study the IBI and intra-beat interval (IaBI) of the ECG, the BCG, and how the two relate. The ECG intra-beat timing, that is the timing between the P, Q, R, S, & T waves can be markers for cardiac problems. For example, extended QT intervals, the depolarization to repolarization of the verticals, can be a marker for tachyarrhythmias. Another example is PR intervals less than 120 milliseconds or greater than 200 milliseconds, are marker for a disconnect of the conduction path from the AV node or an atrioventricular blockage respectively. Beyond the static distances in the IaBI it would be interesting to see is how the variance in these distances relates to cardiac pathophysiology. Variation on in IBI has two know sources in healthy subjects, respiratory sinus arrhythmia, the variation of the heart rate due to location in the respiratory cycle, and low frequency oscillations due to Mayer Waves, the cyclic changes in blood pressure brought about by variations in the baroreceptors and chemoreceptors of the circulatory system. Changes in the variability of the IBI have also been found post myocardial infarction (Kleiger et al., 1987), as well as with but not limited to, diabetic neuropathy, sepsis, and quadriplegia. The route by which the IBI is changed in these cases is a from of perturbation to the autonomic nervous system either in a decrease of parasympathetic activity, an increase in sympathetic activity or both. Here again, these are global changes, but local changes can also be seen in the IBI from mental workload (Nickel and Nachreiner, 2003), and anxiety (Jonsson, 2007). Expanding on this as

mentioned above the case subjects appeared to have a noticeable different IBI's compared to the controls. Specifically, the spectral content of the case subjects IBI traces seemed shifted higher and less well defined / more dispersed in frequency. It would be interesting to one, see if the inter-beat interval can serve as a diagnostic tool for ADHD, and two, to see how the level of inter-beat interval variability changes with attention.

Final Thoughts

The EEG based method presented here serves a robust alternative to the ECG method for determining BCG event timing. Not only are the confounds introduced by the ECG recording avoided but a likely more accurate estimate of the event timing is given. Finally, beyond just being a means to avoid ECG this method also provides a method to recover recordings in which the ECG data was not recorded or was corrupted.

References:

Allen PJ, Josephs O, Turner R (2000) A method for removing imaging artifact from continuous EEG recorded during functional MRI. Neuroimage 12:230-239.

Allen PJ, Polizzi, G., Krakow, K., Fish, D. R., & Lemieux, L. (1998) Identification of EEG Events in the MR Scanner: The Problem of Pulse Artifact and a Method for Its Subtraction. Neuroimage 8:229-239.

Börger N, van Der Meere J, Ronner A, Alberts E, Geuze R, Bogte H (1999) Heart rate variability and sustained attention in ADHD children. Journal of Abnormal Child Psychology 27:25-33.

Brainard DH (1997) The Psychophysics Toolbox, Spatial Vision. Spatial vision 10:433-436.

Choi S, Jiang Z (2008) Comparison of envelope extraction algorithms for cardiac sound signal segmentation. Expert Systems with Applications 34:1056-1069.

Debener S, Mullinger KJ, Niazy RK, Bowtell RW (2008) Properties of the ballistocardiogram artefact as revealed by EEG recordings at 1.5, 3 and 7 T static magnetic field strength. Int J Psychophysiol 67:189-199.

Delorme A, Makeig S (2004) EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. J Neurosci Methods 134:9-21.

Fischer SE, Wickline SA, Lorenz CH (1999) Novel real‐time R‐wave detection algorithm based on the vectorcardiogram for accurate gated magnetic resonance acquisitions. Magnetic Resonance in Medicine 42:361-370.

Jonsson P (2007) Respiratory sinus arrhythmia as a function of state anxiety in healthy individuals. Int J Psychophysiol 63:48-54.

Jurcak V, Tsuzuki D, Dan I (2007) 10/20, 10/10, and 10/5 systems revisited: their validity as relative head-surface-based positioning systems. Neuroimage 34:1600-1611.

Kleiger RE, Miller JP, Bigger JT, Moss AJ (1987) Decreased heart rate variability and its association with increased mortality after acute myocardial infarction. The American journal of cardiology 59:256-262.

Koskinen M, Vartiainen N (2008) Removal of ballistocardiogram artifact from eeg data acquired in the mri scanner: selection of ica components. In: Engineering in Medicine and Biology Society, 2008 EMBS 2008 30th Annual International Conference of the IEEE, pp 5220-5223: IEEE.

Kugel H, Bremer C, Püschel M, Fischbach R, Lenzen H, Tombach B, Van Aken H, Heindel W (2003) Hazardous situation in the MR bore: induction in ECG leads causes fire. European radiology 13:690-694.

Larson AC, White RD, Laub G, McVeigh ER, Li D, Simonetti OP (2004) Self‐gated cardiac cine MRI. Magnetic Resonance in Medicine 51:93-102.

Lehmann D (1971) Multichannel topography of human alpha EEG fields. Electroencephalography and clinical neurophysiology 31:439-449.

Lemieux L, Allen P, Krakow K, Symms M, Fish D (1999) Methodological issues in EEG-correlated functional MRI experiments. IJBEM 1:87-95.

Luu P, Ferree T (2000) Determination of the Geodesic Sensor Nets' average electrode positions and their 10–10 international equivalents. Technical note 1-11.

Makeig S, Bell AJ, Jung T-P, Sejnowski TJ (1996) Independent component analysis of electroencephalographic data. Advances in neural information processing systems 145-151.

Manikandan MS, Soman K (2012) A novel method for detecting R-peaks in electrocardiogram (ECG) signal. Biomedical Signal Processing and Control 7:118-128.

Marano SR, Fischer DW, Gaines C, Sonntag VK (1985) Anatomical study of the superficial temporal artery. Neurosurgery 16:786-790.

Mullinger K, Debener S, Coxon R, Bowtell R (2008a) Effects of simultaneous EEG recording on MRI data quality at 1.5, 3 and 7 tesla. International Journal of Psychophysiology 67:178-188.

Mullinger KJ, Havenhand J, Bowtell R (2013) Identifying the sources of the pulse artefact in EEG recordings made inside an MR scanner. Neuroimage 71:75-83.

Mullinger KJ, Morgan PS, Bowtell RW (2008b) Improved artifact correction for combined electroencephalography/functional MRI by means of synchronization and use of vectorcardiogram recordings. Journal of Magnetic Resonance Imaging 27:607-616.

Nakamura W, Anami, K., Mori, T., Saitoh, O., Cichocki, A., & Amari, S. I. (2006) Removal of Ballistocardiogram Artifacts From Simultaneously Recorded EEG and fMRI Data Using Independent Component Analysis. Biomedical Engineering, IEEE Transactions on 53:1294-1308.

Niazy RK, Beckmann CF, Iannetti GD, Brady JM, Smith SM (2005) Removal of FMRI environment artifacts from EEG data using optimal basis sets. Neuroimage 28:720-737.

Nickel P, Nachreiner F (2003) Sensitivity and diagnosticity of the 0.1-Hz component of heart rate variability as an indicator of mental workload. Human Factors: The Journal of the Human Factors and Ergonomics Society 45:575-590.

Oostenveld R, & Praamstra, P. (2001) The five percent electrode system for high-resolution EEG and ERP measurements. Clinical neurophysiology 112:713-719.

Prisk GK, Verhaeghe, S., Padeken, D., Hamacher, H., & Paiva, M. (2001) Three-dimensional ballistocardiography and respiratory motion in sustained microgravity. Aviation, space, and environmental medicine 72:1067-1074.

Rodriguez C, Cohen M, Lenartowicz A (2015) Systems and methods for measuring cardiac timing from a ballistocardiogram. Regents of the University of California.

Srivastava G, Crottaz-Herbette S, Lau KM, Glover GH, Menon V (2005) ICA-based procedures for removing ballistocardiogram artifacts from EEG data acquired in the MRI scanner. Neuroimage 24:50-60.

Vanderperren K, De Vos M, Ramautar JR, Novitskiy N, Mennes M, Assecondi S, Vanrumste B, Stiers P, Van den Bergh BR, Wagemans J (2010) Removal of BCG artifacts from EEG recordings inside the MR scanner: a comparison of methodological and validation-related aspects. Neuroimage 50:920-934.

Weissler A (1974) Noninvasive cardiology; clinical cardiology monographs. Grune & Stratton Inc, NY.

Xia H, Ruan D, Cohen MS (2014) Removing ballistocardiogram (BCG) artifact from full-scalp EEG acquired inside the MR scanner with Orthogonal Matching Pursuit (OMP). Front Neurosci 8:218.

Yan WX, Mullinger KJ, Geirsdottir GB, Bowtell R (2010) Physical modeling of pulse artefact sources in simultaneous EEG/fMRI. Human brain mapping 31:604-620.

Zhu H, Dong J (2013) An R-peak detection method based on peaks of Shannon energy envelope. Biomedical Signal Processing and Control 8:466-474.

AUTOMATIC IDENTIFICATION OF EEG ELECTRODES IN A 3D VOLUME

Chapter 3

Abstract:

In electroencephalographic (EEG) source localization accurate knowledge of electrode localizations with respect to the underlying anatomy is necessary to project the measured scalp potentials properly onto their sources in underlying brain (Dalal et al., 2014). Historically, this process used system in which simple surface fiducial markers (inion, nasion, left and right preauricular points) are identified, and a proportional measurement grid was created based on these points. The location of the brain with respect to the surface anatomy was inferred based on models of average brain locations. The current state-of-the-art generally involves two measurements. The first is of the underlying brain anatomy using a structural magnetic resonance imaging (MRI) scan. The second is of the electrode positions themselves using a digitizer or photogrammetry system. As the measurements are made separately they must then be brought into alignment. This is generally done by locating fiducial markers on the surface of the head in each measurement. These markers then are aligned between the measurements to bring the two into register. This requires manual input, is time consuming and a potential source of error. Beyond this though, the methods to measure the electrode positions are also time consuming and tedious. In addition to the tedious nature of both tasks, the manual input allows for operator bias and variability. Presented here is a new method that measures the electrode positions using the same structural MRI scan used to measure the anatomy, turning two measurements into one. In addition this is done without the use of exogenous or fiducial markers. Exogenous markers are not needed because the electrode

positions are measured directly. Fiducial markers are not needed as the electrodes and anatomy are measured in the same coordinate system. As such no *post-hoc* registration of any kind is needed – eliminating a source of error. Beyond these advantages the method presented here is fully automatic both for locating the electrodes in the MRI image but also identifying them by channel number – removing the potential for operator bias.

Keywords:

EEG-fMRI, EEG, Source Localization, Inverse Problem, Photogrammetry

Introduction:

Source localization is the process of projecting electroencephalography's (EEG) measured signals from their locations on the surface of the head onto their presumptive sources in the underlying brain. It consists of a forward problem and an approximated inverse solution. The forward problem entails the modeling of how signal sources in the brain project outward on the surface of the head. To do so, typically, a subject specific head conductivity model is made by measuring the distributions of the different conductive elements in the head from a structural magnetic resonance imaging (MRI) scan. Methods such as finite (Gençer and Acar, 2004) or boundary (Akalin-Acar and Gençer, 2004) element analysis are applied to the MRI image to locate and identify the distributions of the different conductive elements of the head (e.g. scalp, skull, cerebrospinal fluid (CSF), and gray and white brain matter). The inverse solution takes the forward model, the measured EEG signals, and then applies an a priori model for the source's electric field to project the signals from the surface of the head back onto their origins in the brain. The a priori models for the electric fields of the sources include an

equivalent current dipole (Henderson et al., 1975) and a distributed source model (Pascual-Marqui et al., 1994, Grave de Peralta and Gonzalez Andino, 2002). These inverse solutions are not unique, as the recorded scalp potentials – at any spatial resolution – could be formed by an infinite number of different source distributions. Therefore, all current inverse solvers use additional spatial priors requiring, for example, that the source dipoles are distributed smoothly and normal to the cortical surface. Beyond the a priori and forward models, the inverse solution is reliant on the accurate knowledge of location of the measured EEG signals with respect to the underlying anatomy (Akalin Acar and Makeig, 2013). The current primary means of measuring this relationship is through either electromagnetic digitization (He and Estepp, 2013) or photogrammetry (Russell et al., 2005). The issues with these methods are four fold. First, both are time consuming and require a considerable amount of manual input (Russell et al., 2005), electromagnetic digitization on the front end, the measuring of the positions, (Khosla et al., 1999) and photogrammetry on the back end, the identification of the electrode channels (Russell et al., 2005). Second, both methods register their measured electrode positions to the underlying anatomy through the use of surface fiducial markers (SFMs). These markers on the surface of the head serve as proxies for the underlying anatomy and allow the electrode positions measured outside of the MRI environment to be brought in to registration with the head anatomy measured by the MRI (Tadel et al., 2011). This registration serves as a potential source of error (Akalin Acar and Makeig, 2013). Third, due to the fact they require manual input they are prone to intra and inter-operator variability (Brinkmann et al., 1998, Koessler et al., 2007). This manual input is for either locating or identifying: the electrodes, the fiducial markers, or both. Fourth, the

position of the brain with respect to the scalp is variable even within the individual, because changes in head position result in considerable shifts in brain position. Presented here is a new method (Rodriguez, 2014) that resolves these issues by automating the processes of electrode position measurement and identification while combining the separate measurements of the electrodes positions and the underlying anatomy into a single measurement. This new method does so by measuring the electrode positions with same structural MRI scan used to measure the anatomy. In contrast to other methods that have done this in the past (Lagerlund et al., 1993, Yoo et al., 1997, Brinkmann et al., 1998, Sijbers et al., 2000) no exogenous markers are used. Fiducial markers are also not needed as the electrodes and anatomy are measured in the same coordinate system. The fact that no markers whether exogenous or fiducial are used means that no registration of any kind is needed – eliminating a source of error. Beyond eliminating a measurement step and a registration step, this new method is fully automatic both for locating the electrodes in the MRI image but also identifying them by channel number, removing the potential for operator bias.

Current methods for the measurement of electrode positions

The current methods for measuring electrode positions fall into three camps, manual measurement, measurement through photogrammetry, and the measurement of exogenous markers.

Manual measurement includes the use of calipers (De Munck et al., 1991) or digitizers whether electromagnetic (He and Estepp, 2013) or ultrasonic (Steddin and Bötzel, 1995). All of the processes entail identifying and measuring the locations of each individual electrode, one by one, in reference to the SFMs. When considering that the

recommended minimum electrode channel count to correctly sample the scalp electric field ranges from 60 (Michel et al., 2004) to 128 (Srinivasan et al., 1998) manual measurement is both time-consuming and tedious especially if using calipers. If extrapolating from the 15 minutes it took Khosla et al. (1999) to measure 64 electrodes, it would take about 60 minutes to measure the 260 points (256 channels, 1 reference electrode, and 3 fiducial points) needed for Electrical Geodesics Inc.'s (EGI) (Eugene, Oregon) high density 256 channel MR compatible EEG net.

Do to the "practical limitations of individual point digitizing" on dense-array (128+ channel) EEG systems Russell et al. (2005) developed a photogrammetry based system. The photogrammetry method measures the electrode positions simultaneously, rather than one by one, using 11 cameras places in a fixed geometry around the head. This greatly reduces the time to measure the positions. To identify the channel associated with each electrode the operator must manually identify 20 marker electrodes across the pictures. Using these marker electrodes, the Russell et al. (2005) algorithm attempts to identify or estimate the locations of the remaining electrode channels.

Measuring exogenous markers placed on the electrodes (Yoo et al., 1997, Sijbers et al., 2000) or the surface of the head (Lagerlund et al., 1993, Brinkmann et al., 1998) with MRI came about in order to eliminate a experimental step and a coordinate system transformation. The elimination of an experimental step comes about by merging the step of measuring the electrode positions with the step of measuring the distributions of the conductive elements of the head by using the same MRI scan for both. Because the markers and head are captured in the same MRI scan they are in a common coordinate system, eliminating the need to calculate a transformation from the SFMs measurements

to the MRI coordinate system.

Drawbacks of the current methods for the measurement of electrode positions

As noted, manual measurement is not practical on dense-array EEG systems due to the time it takes to perform them. This is especially true in clinical and young populations where subject fatigue and compliance are problems (Russell et al., 2005). The benefit of the digitizer systems though is that locating the electrodes, and identifying each measure position by electrode channel number, are done in tandem.

The drawback with photogrammetry is that while the measurement step of the electrode positions is quick, on the order of seconds (Russell et al., 2005), the identification by channel step is not. It requires locating the 20 marker electrodes across the 11 photos. This is not simply locating 20 points, as each marker must be captured in multiple images to create identify its locations in three dimensions. In addition to the identification step drawback, the photogrammetry system can be confounded easily when electrodes in the photos are obscured by the subject hair and by the improper placement of the subject head with respect to the fields of view of the cameras (Russell et al., 2005, Koessler et al., 2007).

The drawback with the use exogenous makers is that they serve as proxies for the electrode positions. As such the electrodes must be registered to the exogenous markers, which in turn are registered to the underlying anatomy. While a coordinate system transformation is eliminated one registration is just changed for another. Here the registration to fiducial markers is simple replaced by registration to the exogenous makers. The markers only serve to denote the location of the electrodes but do not identify them. This, like the photogrammetry method, requires a considerable amount of

work by the operator on the back end to identify the individual electrode channels.

New method for automatic location & identification of EEG electrodes in an MRI scan

Presented here is a new method for locating and identify EEG electrode positions on the surface of the head that: combines the step of measuring electrode positions with the step of measuring the underlying anatomy by using the same MRI scan to measure both and in doing so eliminating coordinate transformations as well; in addition the electrode positions are measured directly, eliminating registration to either fiducial or exogenous markers as the positions are already in register with the underlying anatomy. Finally our method automatically identifies the electrodes by channel eliminating this tedious task as well as inter and intra operator variance of measuring the positions and or the fiducial markers manually (Brinkmann et al., 1998, Koessler et al., 2007). This new method has two steps, *detection* and *identification*. In the detection step the electrodes are located directly in a high resolution structural MRI based upon their physical characteristics including size, shape and spacing. Once the electrodes are detected they are then identified by channel number through registering them to a predefined 3D template that is EEG manufacturer and model (cap or net) specific.

Step 0: The Electrode Template

Independent of both detection and identification steps is the process of constructing a template (3D model) of the electrode distribution being used. This needs to be determined only once per model distribution (cap or net). The same template can in general be used across various sizes of the model. In addition, once a template is made, it can be shared between users and sites. Beyond the model specific distribution the

template also contains information about the physical characteristics of the electrodes themselves such as their size and shape.

Step 0.1: Template Creation

There are two methods to easily create a 3D representation of the channel layout of an electrode array. The first is to modify an existing 3D digital representation such as the ones found in the publically available Brainstorm database (Tadel et al., 2011). The second is to create a new 3D digital representation using one of the electromagnetic digitizer systems. In both cases all place holder electrodes as well as the common and reference electrodes positions must be included in the layout.

Once the 3D model has been created its center of mass is placed at the origin *Equation 1a & b,* and it size is normalized, *Equation 1c*. These two steps simplify the registration process and allow the same template to be used across sizes for a given electrode model array.

Equation 1: Method for centering the template at the origin and normalizing its size. Measuring the center of mass (TP_{COM}) of the template distribution (TP), (a). Moving the location of the center of mass of the template distribution (TP') to the origin, (b). Normalizing the size of TP' by stretching the distribution by the inverse of the magnitude of the vector of the furthest template point from the origin, (c).

$$\text{(a)}\quad \mathrm{TP}_{COM}(x,y,z) = \frac{1}{N}\sum_{i}^{N} \mathrm{TP}_i(x,y,z).$$

$$\text{(b)}\quad \mathrm{TP}'(x,y,z) = \mathrm{TP}(x,y,z) - \mathrm{TP}_{COM}(x,y,z).$$

$$\text{(c)}\quad \widehat{\mathrm{TP}'}(x,y,z) = \mathrm{TP}'(x,y,z) \Big/ \max\left(\left|\mathrm{TP}'\right|\right).$$

Step 1: Electrode Detection

In the electrode detection stage, the algorithm must determine what is, and is not, an electrode in the image. To do so it applies a priori knowledge of where the electrodes are placed on the subject and the physical characteristics of the electrodes and their distribution. The information about the physical characteristics and distribution are taken

from the model specific template, *Step 0*. Knowledge of placement location tells the computer where to search. Specifically, that the electrodes are placed on surface of the head protruding outwards and that they are not in certain locations on the surface such as the ears, neck, tip of the nose and chin. Knowledge of the physical characteristics, i.e. size and shape of the electrodes, tells the algorithm what to look for. For example, in the case of an Electrical Geodesics Inc.'s (Eugene, Oregon) MR compatible HyrdoCel™ Net, the algorithm will search for electrodes pedestals that are cylinders 13mm tall with a diameter of 13mm. Lastly, the knowledge of distribution, specifically the relative distance between electrodes, is applied as one of the means for automatically correcting false detections. Here the algorithm removes detections in which their relationship to their neighbors does not match ones found in the template.

Step 1.1: Preprocessing, Applying the knowledge of location

In the first stage of electrode detection, rather than having the algorithm find where to look, the algorithm decides where not to look. It does so by systematically eliminating regions by applying logical constraints based upon the properties of the image. Each voxel is binned into one of two categories, either a potential electrode voxel (PEV) or a non-electrode voxel (NEV). The algorithm starts with the naive assumption that every voxel in the image is a PEV. Through the following steps the algorithm categorizes each voxel: background signal removal, shell extraction, orientation detection and shell trimming. In the background signal removal step, the algorithm decides which voxels are purely noise and moves them into the NEV category, removing them from further consideration. Next, the exterior surface of the subject volume, the voxels comprising the subject and electrodes, is isolated and retained. The rest of the voxels are

moved to the NEV category. After this, the algorithm detects how the subject volume is oriented within the larger image volume. Measuring the orientation has two benefits. The first is that it aids in the identification and removal of voxels that will confound the shape detection portion of the algorithm. The second is that it aids the channel registration in the identification portion of the algorithm. In the last preprocessing step ,and with the orientation of the subject volume determined, the algorithm identifies and removes the voxels in the shell comprising the nose, ear, chin, and neck.

Step 1.1.1: Background signal removal

Removing the background signal is different than removing the background noise. The three volumes discussed here are the image, subject and background volumes. The image volume is all of the volume elements (voxels) that make up the MR image. The subject volume is made up of all of the voxels with intrinsic MR signal in the image volume. The background volume is made up of all of the voxels with no coherent signal in the image volume. These background voxels have signal in the form of the omnipresent, additive, noise in the MR image. This measurement noise itself is Johnson noise, otherwise known as background noise (Nishimura, 2010). Our algorithm assumes that it is dispersed evenly throughout the image volume. In the MRI raw data this noise is modeled as zero-mean, Gaussian-distribution. In converting the raw signal to an image a Fourier transform in two or more dimensions is applied, and the only the magnitude signal is preserved (phase information generally is not informative to the structural images). The noise has a random *temporal* distribution during signal measurement, which becomes a random *spatial* distribution after the Fourier transform. The magnitude conversion is a rectification step that changes the distribution of the noise to Rayleigh

(Nishimura, 2010). Whether Gaussian or Rayleigh-distributed, voxels containing this equally dispersed additive noise generally will be the elements of the image volume with the lowest intensity signal.

This background region is the first to be excluded from the search, classified as NEVs. To do so the algorithm must decide what is and is not purely noise in the image. Our algorithm applies three thresholds sequentially. In order of application they are based on: intensity, PEV connectivity, and bounding box size.

Step 1.1.1.1: Background signal removal – Appling an Intensity Threshold

The majority of the background signal is from the lowest intensity voxels in the image volume as they contain. These are the voxels whose measured signal arises solely from the Rayleigh-distributed background noise as they have no intrinsic MRI signal. To remove these voxels the algorithm applies a lower-limit threshold on the intensity values of the image. The value for this threshold is set by looking at the entire distribution of voxel intensity values and finding the first local minimum in a histogram of intensity values, *arrow, Figure 1*. All voxels having intensities values below this point will be ignored by the algorithm from this point on, i.e. moved from the PEV to the NEV category, *red line, Figure 1*. This threshold value is set in a conservative manor as the tails of Rayleigh distribution may extend beyond this point, resulting in incomplete removal of the background signal, *Figure 2, middle row*. However, in order to avoid removing the subject volume as well, rather than raising the threshold to remove the entire background signal additional steps are taken. To remove the remaining voxels in the background volume the algorithm next applies connectivity and bounding box filters.

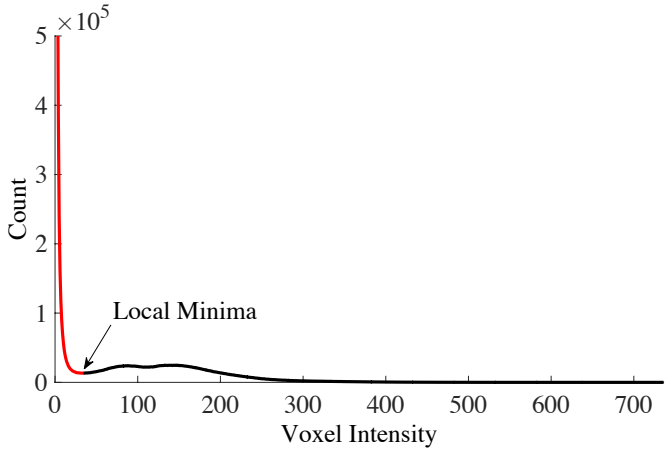

Figure 1: Image volume intensity distribution. The black trace is the intensity values of voxels remaining in the PEV category. The red trace is the intensity values of voxels removed from the PEV category by the threshold.

Step 1.1.1.2: Background signal removal – Appling a Connectivity Threshold

Using the only the voxels remaining in the PEV category the algorithm compares the intensity of the target voxel to its 26 adjacent neighbors. PEVs with fewer than 3 PEV neighbors are removed from the category.

Step 1.1.1.3: Background signal removal – Appling a Bounding box Threshold

Once the low intensity and low connectivity voxels are eliminated the algorithm looks at the sizes of all remaining contiguous PEVs regions. To do it circumscribes each region with a bounding box. The largest of the boxes will contain both the subject head and the electrodes. The regions in all other boxes are removed.

Step 1.1.2 Shell Extraction

With the background volume removed, the exterior surface of the subject volume, *background images, Figure 3,* is then extracted. This shell is where the electrodes will be located and is determined by ray tracing: every possible line parallel to the image volume dimensions is drawn. If a line passes through PEVs, the first and last spatial indices of the crossings are found. These points are used to construct the shell, *red lines, Figure 3*.

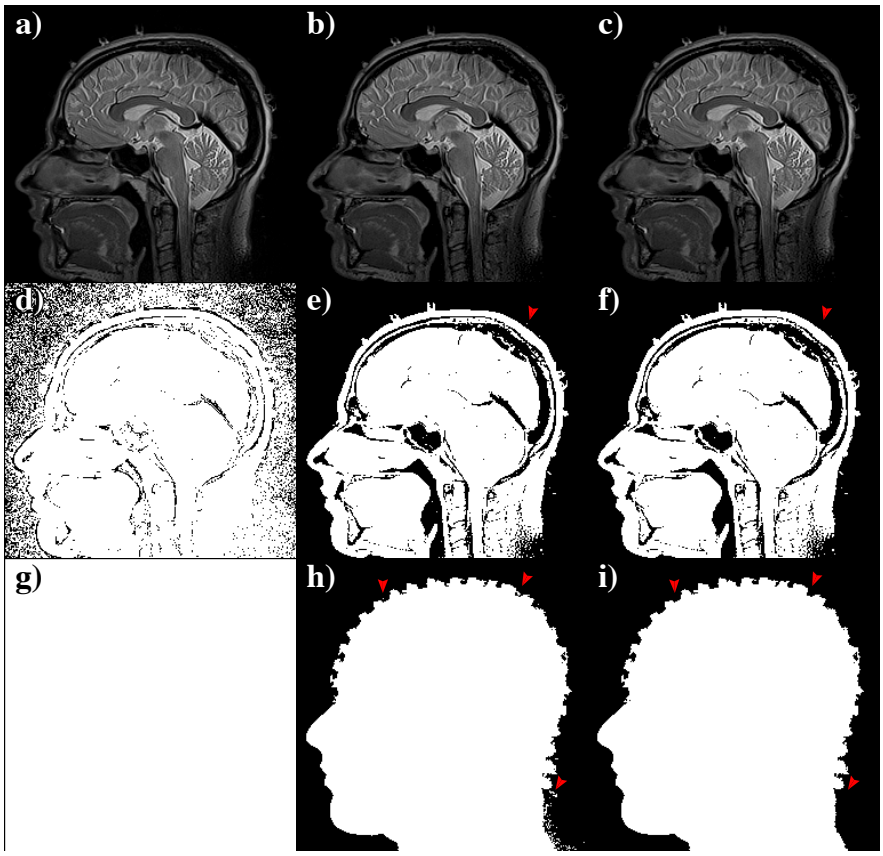

Figure 2: Progression of the removal of the voxels in the background volume from the PEV category. Row 1, (a)-(c), contains the center sagittal slices of the image volume. Row 2, (d)-(f), contains center sagittal slices of the image with the intensities binarized. Row 3, (g)-(h), contains the summed image along the sagittal axis of binary intensities of the image volume. Column 1, (a),(d),(g), contains the unprocessed data. Column 2, (b),(e),(h), contains the data after applying the intensity threshold. Column 3, (c),(f),(i), contains the data after applying the intensity, connectivity and bounding thresholds. The binary images, row 2, are shown to accentuate the signal from the background volume. The summed images, row 3, are shown to further accentuate this signal. Note the vast majority of the background volume is removed by the intensity threshold (d) & (g) versus (e) & (h). Red arrows in (e) and (h) mark residual background signal after the application of the intensity threshold. Notice in (f) and (i) the regions no remain after the connectivity and bounding box thresholds.

Step 1.1.3: Orientation Detection

With the background volume removed and the surface of the subject volume extracted, the orientation of the subject volume in the overall image volume is then determined. Rather than relying on information found in the MRI data (DICOM or NIFDI) file header, the algorithm establishes the orientation directly from the images. This is done for two reasons. The first is that the labeling may be incorrect in the file

header. The second is that a finer level of detail than can be found in the file header is needed to remove from the PEV category the voxels that make up the subject's ears, nose, chin and neck. The orientation is done to two stages, course and fine. In the course stage the cardinal directions (superior – inferior, anterior – posterior and left – right) for the image volume are established. In the fine stage the algorithm establishes how the subject volume is oriented within the greater image volume. Our algorithm is based on relatively simple, but highly reliable, priors regarding head shape.

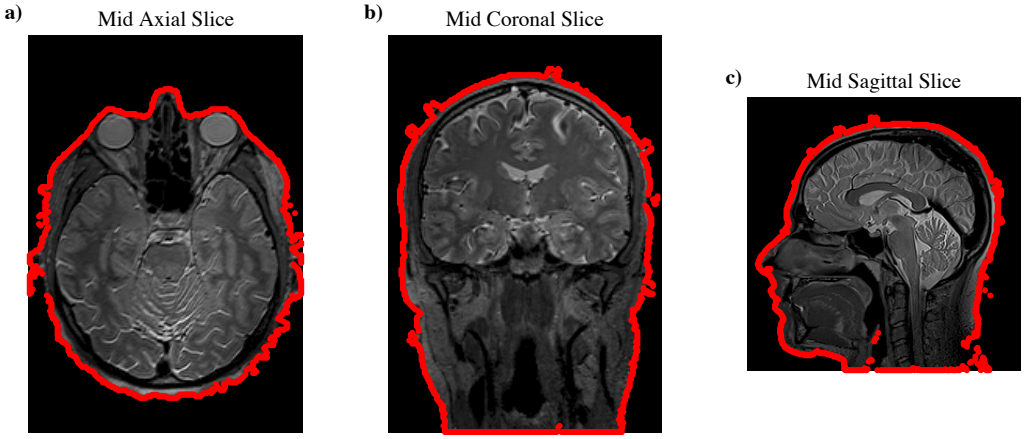

Figure 3: Mid-volume views of subject volume with it shell highlighted. The shell of the subject volume, Red Trace, is shown in overlaid on Mid-axial (a), mid-coronal (b), and mid-saggital slices from the image volume.

Step 1.1.3.1: Orientation Detection – Finding the Principle Image Volume Axes

To establish the principal (cardinal) axes of the image the algorithm first locates the neck of the subject within the image volume. Knowing where the neck is in the image volume establishes the orientation of the all of the principal axes. To find the neck the algorithm examines each of the edges of the image volume, Figure 3. The edge which contains the neck generally will have the largest number of PEVs. This will be the bottom of the image volume. Dividing this edge into four quadrants, *red lines, Figure 4d*, will reveal the back of the image. The side of the bottom edge of the image volume, *bottom,*

Figure 4d, with the quadrants containing the two largest numbers of PEVs will be the back of the image. Then by the process of the elimination the left – right axis of the image volume is established.

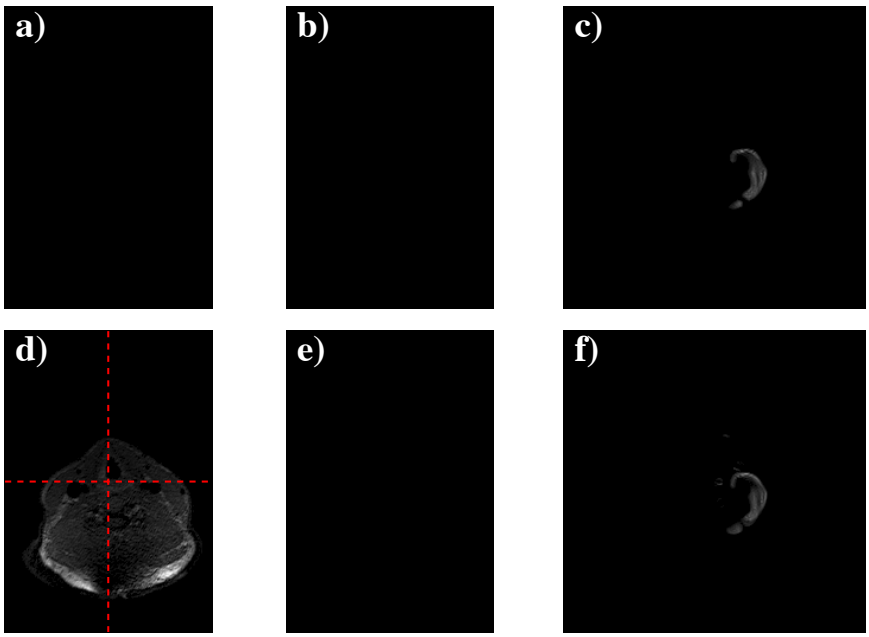

Figure 4: Edges of the image volume. (a) superior most axial slice. (d) inferior most axial slice. (b) anterior most coronal slice. (e) posterior most coronal slice. (c) left most saggital slice. (f) right most sagittal slice. The subject's neck can be seen in (d). Red dashed linse in (d) are the quardant divisions of the bottom edge of the image volume used to establish it's anterior – posterior axis.

Step 1.1.3.2 Finding the Principle Axes of the Subject Volume

Equation 2: Raw moments of a gray scale digital image. In the calculation of the raw moment (a), the sum of "i" and "j" in the moment term, M_{ij} ,equal the order of the moment. The zeroth order moment of the image is its gray level or total pixel mass (b) The 2 – first order moments of the image divided by the total pixel mass give the coordinates of the centroid of the image (c)

(a) $M_{i,j} = \sum_X \sum_Y x^i y^j I(x,y).$

(b) Total Pixel Mass = M_{00}.

(c) Image Centroid, $(\bar{x},\bar{y}) = (M_{10}, M_{01}) / M_{00}$.

With the image volume axes identified, how the subject volume is orientated with respect to these axes is then calculated. To do so, first depth projections through the image volume along its principal axes are made. Here the depths are that of the shell of the

subject volume, *Figure 5*. The image moments are then calculated, *Equations 2-4*, on these projections to determine the centroid and principle axes of the shell, *solid lines, Figure 6*.

Equation 3: Central and normalized central moments of a gray scale digital image. Calculation of the central or second order moments of an image is performed using the equation in (a). The normalized second order central moments (b)-(d) are calculated by normalizing the equation in (a) by the zeroth order central moment. These normalized second order central moments can also be calculated using the raw image moments from Equation 2 using the relationships in (b)-(d)

(a) $$\mu_{i,j} = \sum_X \sum_Y \left(x^i - \bar{x}\right)\left(y^j - \bar{y}\right) I(x,y).$$

(b) $$\mu'_{20} = \mu_{20}/\mu_{00} = M_{20}/M_{00} - \bar{x}^2.$$

(c) $$\mu'_{02} = \mu_{02}/\mu_{00} = M_{02}/M_{00} - \bar{y}^2.$$

(d) $$\mu'_{11} = \mu_{11}/\mu_{00} = M_{11}/M_{00} - \bar{x}\,\bar{y}.$$

Equation 4: Orientation of the Subject Axes with Respect to the Image Axes. Here using the second order central moments from Equation 3 the rotations of the major and minor axes of the depth projections are calculated with respect to the principal axis of the image volume.

(a) $$\theta_{Major} = \frac{180}{2\pi}\arctan\left(\frac{2\mu'_{11}}{2\mu'_{20} - 2\mu'_{02}}\right).$$

(b) $$\theta_{Minor} = \theta_{Major} - 90°.$$

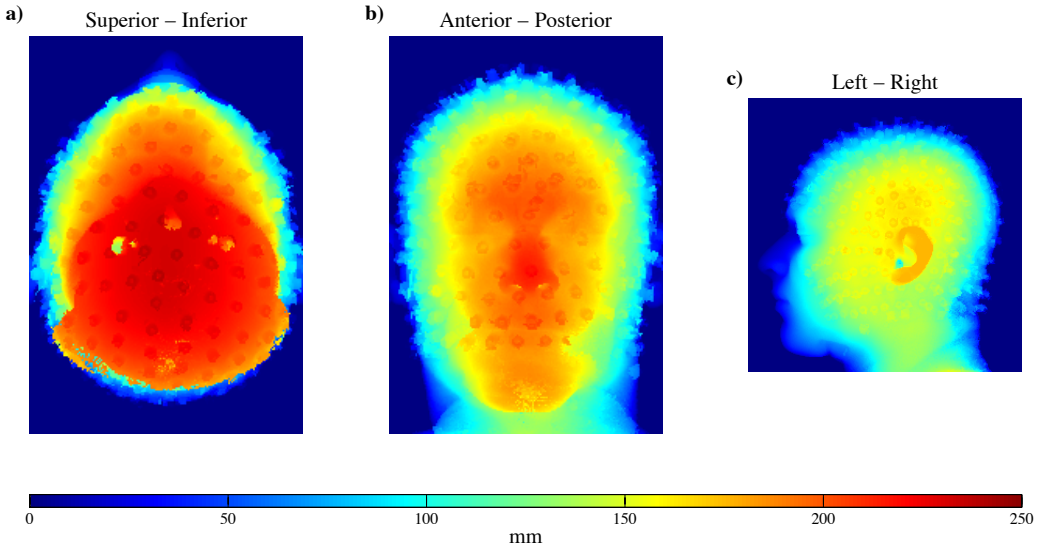

Figure 5: Depth projections of the PEV regions along the principal axes of the image volume. Measurements are made from the same indices used to construct the subject volume shell in step 1.1.3. The difference between each first and last index pair is multiplied by the voxel size. The voxel size is extracted from the NIFDI or DICOM header. If the voxel size is absent the algorithm assumes a 1mm^3 voxel size.

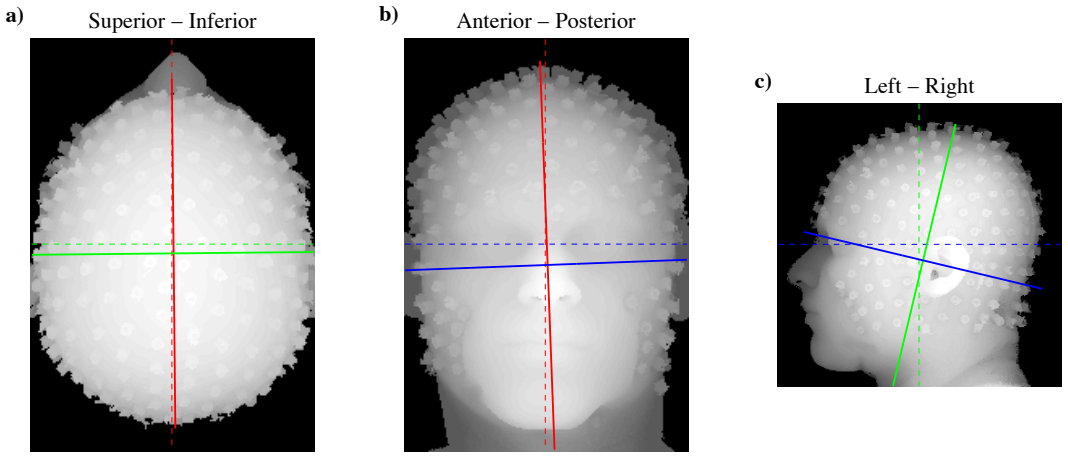

Figure 6: Axial, coronal and sagital depth projection with both image volume and subject volume axes. The subject volume axes (solid) and image cardinal axes (dashed) are as follows: Red = Left – Right Plane; Green = Anterior – Posterior Plane; Blue = Superior – Inferior Plane. Background images are half depth projections. That is in (a) a superior to midline projection, in (b) is an anterior to midline projection, and likewise in (c) a left to midline projection

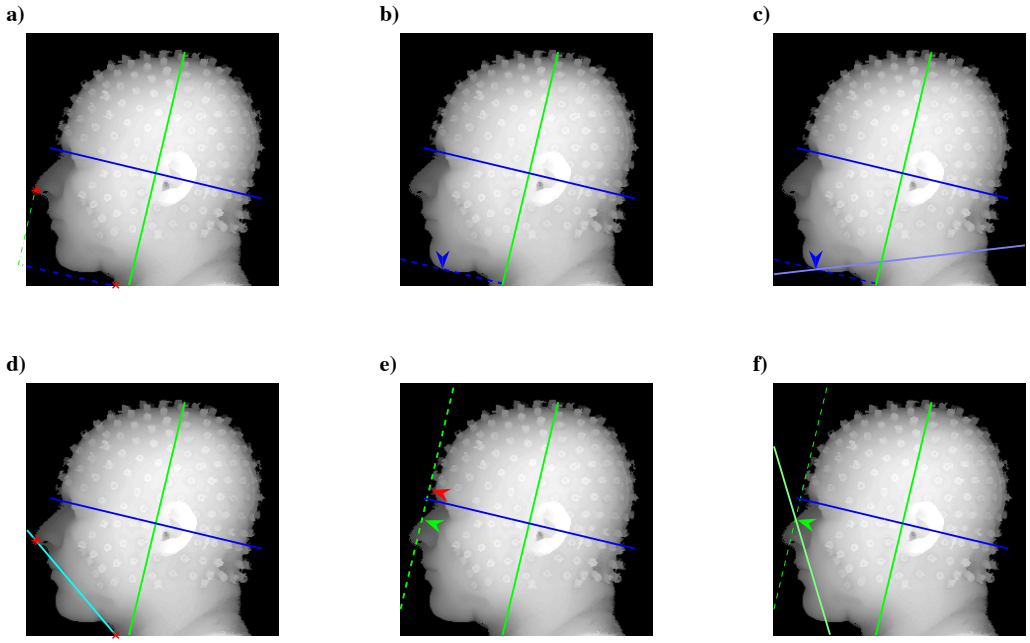

Figure 7: Sagital half-depth projections with fiducial point and trimming plane locations. The subject volume axes (solid) are as follows: Dark Green = Anterior – Posterior Plane; Dark Blue = Superior – Inferior Plane. The point defined as the bottom of the chin is demarcated with a blue arrow in (b) and (c). The bridge of the nose point is demarcated with a green arrow in (e) and (f). The nasion point is demarcated with a red arrow in image (c). The plane parallel to the Superior – Inferior Plane and containing the bottom of the chin point can be seen as a dashed blue line in (b) & (c). The plane parallel to the Anterior – Posterior Plane and containing the nasion and bridge of the nose points can be seen as a dashed green line in (e) & (f). The neck trimming plane can be seen in the solid violet line in (c). The chin trimming plane can be seen as the solid cyan line in (d). The nose trimming plane can be seen in the solid light green line in (f).

Step 1.1.4: Shell Trimming

Once the orientation of the subject volume inside the image volume has been measured, and the dimensions of subject volume have been calculated, the voxels that comprise the neck, ears, nose and chin are identified and removed. These voxels can confound the detection of the electrode positions that will follow in the next step.

The first step in the elimination of these areas is to locate the tip of the nose, *red asterisk, Figure 7a*, and the front of the neck, *red "x", Figure 7a*. The most anterior point on the inferior most image volume plane will be the front of the neck, referred to from here on as the "Adam's apple point". The most anterior point overall in subject orientation will be the tip of the nose.

Using the Adam's apple point and the tip of the nose, the following points are located: the bottom of the chin, blue arrows *Figure 7b & c*, the nasion electrode point, red arrow *Figure 7e*, and the bridge of the nose, green arrow *Figure 7e & f*. The bottom of the chin will be found on the most inferior plane parallel to the superior-inferior subject plane that contains a PEV anterior to the Adams apple point, blue dashed line *Figure 7b & c*. The most posterior and medial point (in subject orientation) on this search plane will be considered the bottom of the chin. The nasion electrode point will be located in a similar manner but slightly modified manner. The search plane here will be parallel to the anterior-posterior subject plane. The criteria will be slightly different, here the nasion electrode point will be found on the most anterior search plane that contains a PEV posterior to the tip of the nose, and is non-contiguous with the tip of the nose, green dashed line *Figure 7e & f*. The nasion electrode point will be the most inferior and medial

point in the non contiguous region with the tip of the nose, the bridge of the nose will be the most superior and medial point in the region contiguous with the tip of the nose.

After the above points have been located next it must be determined if the neck is distinct. That is, whether it was captured in the image volume. To do so the distance between the Adams apple point and the tip of the nose along the subject anterior-posterior axis, blue dashed line *Figure 7a*, is measured. If the distance is greater than 4 cm then the neck is deemed distinct in the subject image. This threshold serves to prevent the chin from being mistaken as the neck.

If the neck and or shoulders are present in the image they are removed by a trimming plane. This cutoff plane is a rotation of the search plane used to locate the bottom of the chin, solid violet line *Figure 7c*. The axis of rotation runs though the bottom of the chin point and is parallel to the left-right subject axis. The rotation is such that the portion of the plane anterior to the bottom of the chin moves downward. The amount of rotation is based upon the layout of the electrode net and will be make and model specific. When using EGI's Hyrdocel 256 net for example the angle will be 20°. All points inferior to this cutoff plane are removed from the PEV category, violet regions *Figure 9b & c*.

The chin voxels are removed after removal of the neck voxels. This is done again with a trimming plane. This plane passes though both the Adams apple point and the tip of the nose. It is a rotation about the left-right subject axis of a plane that parallel to the subject anterior – posterior plane, solid cyan line *Figure 7d*. All points inferior to this cutoff plane are removed from the PEV category, cyan regions *Figure 10b & c*.

Next the nose voxels are removed through the use of a trimming plane. The cutoff plane, much like the neck cutoff plane, is a rotation of the search plane, in this case the one used to locate the bridge of the nose, solid green line *Figure 7f*. The axis of rotation runs though the bridge of the nose point and is parallel to the left-right subject axis. The rotation is such that the portion of the plane superior to the bridge of the nose moves forward. The amount of rotation is based upon the layout of the electrode net and will be make and model specific. When using EGI's Hyrdocel™ 256 net for example the angle will be 30°. All points inferior to this cutoff plane are removed from the PEV category, green regions *Figure 10a, b & c*.

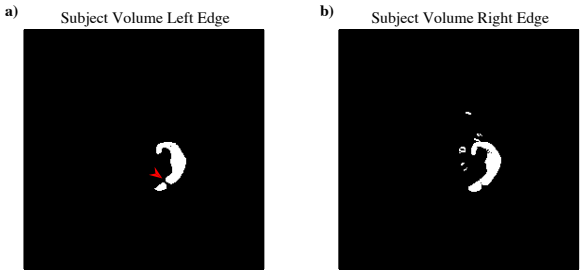

Figure 8: Binarized outermost left and right sagittal slices containing PEV. Bifurcation of the left earlobe region is marked with a red arrow. Note also how the ear is nested amongst electrodes in the slice from the right side (b).

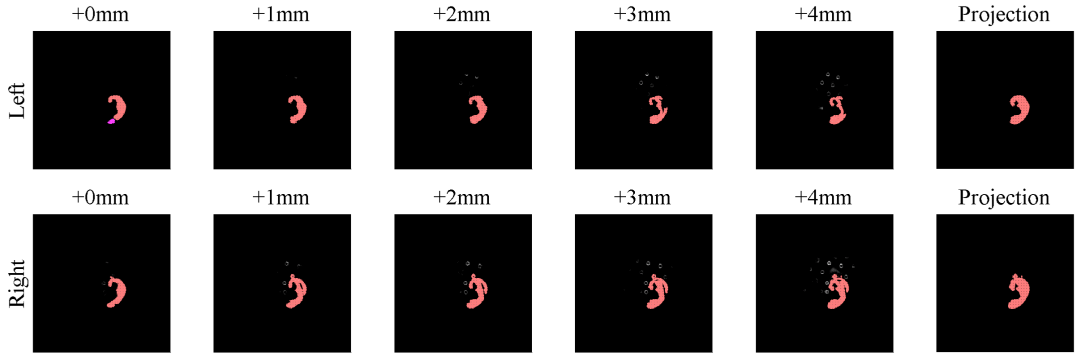

Figure 9: Ear voxels in a slice-by-slice search. The slices from left side of the image volume are shown in the top row and the slices from the right side are shown in the bottom row. The left most column is the outermost sagittal slice containing a PEV. Moving from left to right, the middel columns are adjacent slices moving inward. The right most column is the sagittal projection of all of identified voxels from the slices in the columns to the left. Voxels identified as the ear from the inward growth are colored in red. Voxels identified as the ear from the outward growth are colored in magenta. Note the white regions in both rows. These are electrode voxels. Also note, in the bottom column the algorithm is unable to separate an electrode in contact with the ear from it.

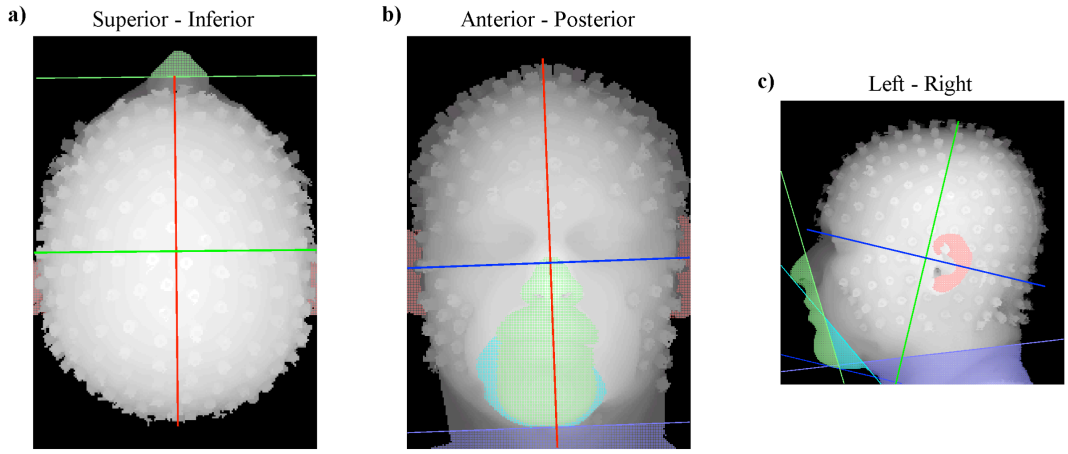

Figure 10: Axial, coronal and sagittal half-depth projections with trimmed region overlays: Background images are half depth projections. That is in (a) a superior to midline projection, in (b) is an anterior to midline projection, and likewise in (c) a left to midline projection. The nose trimming plane can be seen as the light green line in (a) & (c). The neck trimming plane can be seen as the violet line in (b) & (c). The chin trimming plane can be seen as the cyan line in (c). Green patches represent the voxels removed via the nose trimming plane, blue the neck, and cyan the chin. The red patch represents ear voxels removed by the projected expanded seed region.

Finally the ear voxels are removed. In the case of the ears though, a trimming plane approach will not work, due to the fact the voxels comprising the ears are nested amongst the electrodes. Instead, to isolate the ears a seed expansion combined with profile projection approach is taken. The seed expansion portion consists of identifying a seed region and growing to encompass the voxels that have connection to it. The grown region is used to create a profile of the ear, which then is projected along the left-right subject axis inwards. The seed region is located automatically by using the simple criteria that the ears should be most lateral points in the left-right subject orientation. In the case where there are multiple areas in this region, *Figure 8*, the region with the largest contiguous area will be used as the seed. The seed region is grown inwards for a total thickness of 5mm, red regions, *Figure 9, 0-4mm*. To accommodate the bifurcations of the ear within plane, red arrow *Figure 8a*, the algorithm does a second pass moving outward, *Figure 9a* magenta region. The expanded seed region voxels are then flatted along the

left-right axis, *Left and Right Projections, Figure 9*, which are projected inwards to the zero point on the left-right subject axis. The expanded seed regions as well as all voxels within the projections are removed.

Step 1.2: Electrode Identification, Applying knowledge of size, shape and distribution

With the shell of the subject volume extracted and the problematic surface features removed the algorithm then locates the electrodes. To do so it accentuates the contrast in the domain(s) of maximum difference between the electrode and non-electrode voxels. In the case of the EGI's HydroCel™ net the contrasts to accentuate are the shape, size and orientation differences between the subject head and the electrode pedestals. After the contrast is accentuated a threshold is applied to separate the electrode voxels from the remaining PEV. With the electrode voxels separated the algorithm then performs a first round of error checking, looking for false positives by examining the inter-electrode distance.

Step 1.2.1: Shape, Size, & Orientation Accentuation and Extraction

In the example presented here, extracting the electrode positions of EGI's MR compatible HydroCel™ 256 channel EEG net from a high resolution MRI, image intensity proves a poor choice for several reasons. The plastic portions of the electrodes are MR visible, however their image intensities are very similar to those of the scalp tissue, *Figure 2a-c*. More problematic though, is that surface image intestines are not uniform across the scalp, *Figure 2a-c top of the head compared to the back of the neck.* The intensity variability over the surface of the head is due both to susceptibility variations, and to radio frequency (RF) parasitic losses caused from the high density of

cabling in the EEG net (Mullinger et al., 2008). The cabling is packed tightly to minimize loop area. In the case of the HydroCel™ net this cabling is aggregated in the back of the head such that all electrode leads end up in a single bundle at the base of the skull / back of the neck. While minimizing loop area reduces the amplitude of both gradient and ballistocardiogram artifacts in the EEG recording, it introduces susceptibility and RF parasitic loss artifacts in the MRI image (Mullinger et al., 2008). RF Parasitic losses come about from the local reduction in RF field used to excite the protons that are to be imaged. The reduction is due to the absorption of RF energy by other objects in the MR scanner. The parasitic conduction serves also to shunt the emitted RF signal away from the receiver amplifiers. In this case the objects are conductive electrode leads. The amount of absorption increases with the number of cables in the bundle. Susceptibility losses come about from distortion of the magnetic field by objects within it, more specifically at boundaries where there is a mismatch in magnetic susceptibility between two materials causing the magnetic field within that voxel to be inhomogeneous, resulting in reduced T2* relaxation time. In this case the voxels containing both cabling the subject tissue are affected.

Rather than use image intensity, which both varies across the surface of the head and has little or no contrast between the subject and the electrodes, the algorithm uses the shape, the size as well as the relative orientation of the electrode with respect to the head, which are uniform across the surface of the head and have high contrast. The electrodes are oriented such that they protrude outward from the surface of the head. As for the size and shape, thanks in part to the shell trimming performed in *step 1.1.4*, these are distinct from other features on the shell.

To convert the size, shape and orientation differences into image intensities the trimmed shell is binarized and then convolved with a binary image of a 3D rotation of a single electrode, forming a spatial filter. This rotated electrode image is essentially a sphere whose diameter matches the longest diagonal dimension of an electrode. The output of the convolution in restricted to the subject volume by trimming it such that its external surface matches that of the trimmed shell. The result of all of preprocessing steps, and this convolution, is that the parts of the convolved image with the highest voxel intensity, *blue border Figure 11*, will be the electrodes.

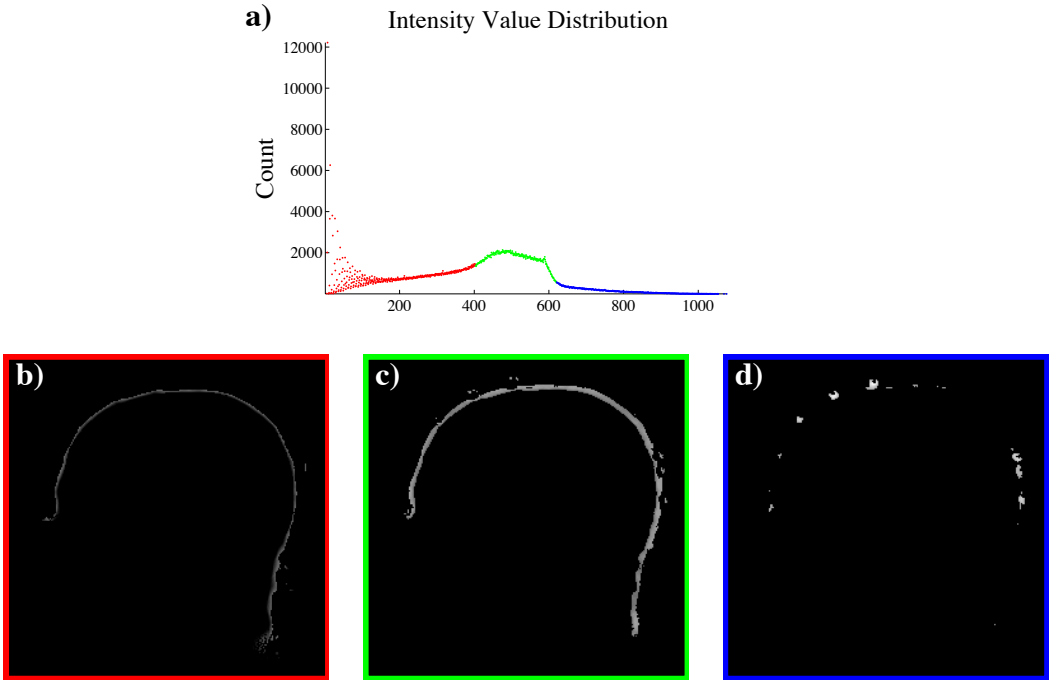

Figure 11: Threshold Value Determination Post Convolution. (a) Intensity value distribution in the contrast accentuated image volume. Inflection points of the smoothed distribution are located at the color boundaries (red – green & green – blue) of the distribution in (a). In (b)-(c) are mid-sagittal images of the voxels from different segments of the distribution in (a). The boarder around each of these images corresponds the region of the intensity distribution in (a) where the voxel are from.

The lower-limit threshold here is automatically set in a similar manner to that in *step 1.1.1.1* . Here the algorithm uses an inflection point rather than a local minima in the intensity distribution to determine the threshold. The inflection point used is the one

corresponding to the highest intensity value *which is the boundary between green and blue regions in Figure 11a*. Beyond using an inflection point rather than a minima, the method by which the algorithm goes about locating the inflection point is slightly different than the method in *step 1.1.1.1* . Rather than using the raw distribution the algorithm applies a 20 point centered moving average filter to smooth the ripples in the distribution. The smoothing is applied because the sample size at this point is relatively small, as about 95% of the voxels of been eliminated from consideration.

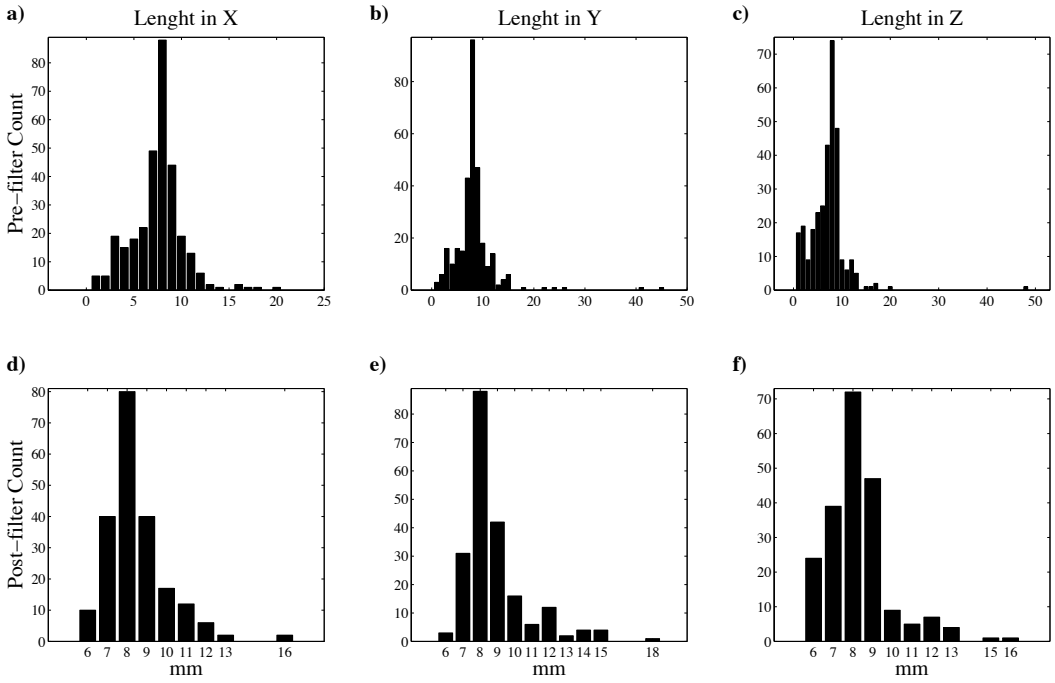

Figure 12: Contiguous region bounding box edge lengths pre and post bounding box filtering. Top row are the distributions of the bounding box edges post connectivity filtering. Bottom row are the distributions of the bounding box edges post bounding box filtering.

Step 1.2.2: Size and Connectivity Verification

After the intensities threshold has been performed, both low connectivity and bounding box filters are applied. The connectivity filter is the same as described in step 1.1.1.2. The bounding box filter here removes any contiguous volumes that too small or too large to be an electrode, *Figure 12*. Too large is any contiguous volume that has a

bounding box edge longer than longest internal diagonal of an electrode. In the case of the HydroCel™ electrode, a 13mm diameter cylinder that is 13mm tall, the longest edge size is $19\text{mm} \approx \sqrt{13\text{mm}^2 + 13\text{mm}^2}$. Too small is set such that no edge is smaller than 25% of the longest internal diameter. The reason the lower limit is so small is because in the areas with signal drop out such as the back of the neck, not all of the electrode voxels may have survived up and through the shell formation process.

Step 1.2.3: Detected Electrode Grid Distance and Angle Measurement

Once the detected regions of electrode voxels have been separated and filtered by connectivity and size the algorithm then checks the spacing between the detected electrodes. It does this to search for false and multiple detections. A false detection being the locating of an electrode where there was none. A multiple detection being a single electrode represented by more than one contiguous group of voxels. Spacing between the detected regions is measured three ways: the distance to the nearest neighbors, the number of nearest neighbors, and the angular separation between the nearest neighbors (internal angle).

The distance is calculated using the center of mass for each contiguous group of voxels, *Equation 8a*. Electrodes within five standard deviations of the average minimum distance between detected regions are considered "nearest neighbors". Contiguous regions, which have center of masses located at less than the length of the longest internal diagonal of an electrode are considered a multiple detections, *Figure 13b*. Contiguous regions that meet this criteria are joined by replacing the individual center of masses of each independent region with that of the group distribution.

Once the nearest neighbors of a detected region are established and any multiple detections are cleared up the angular separation between the nearest neighbors are calculated using the law of cosines, *Equation 5b*.

Equation 5: Calculation for the Angle Between Nearest Neighbors: a) Euclidean distance between 2 points in 3 dimensional space b) calculation of the angle between three points "RP", "a" and "b". The angle calculated is between the three points is adjacent to point "RP".

(a) $\mathrm{Pt}_i\mathrm{Pt}_j = \sqrt{\left(x_i - x_j\right)^2 + \left(y_i - y_j\right)^2 + \left(z_i - z_j\right)^2}$.

(b) $\mathrm{Angle}_{\mathrm{RP}} = \arccos\left(\frac{\mathrm{Pt}_{\mathrm{RP}}\mathrm{Pt}_{\mathrm{a}}{}^2 + \mathrm{Pt}_{\mathrm{RP}}\mathrm{Pt}_{\mathrm{b}}{}^2 - \mathrm{Pt}_{\mathrm{a}}\mathrm{Pt}_{\mathrm{b}}{}^2}{2 \times \mathrm{Pt}_{\mathrm{RP}}\mathrm{Pt}_{\mathrm{a}} \times \mathrm{Pt}_{\mathrm{RP}}\mathrm{Pt}_{\mathrm{b}}}\right)$.

Next the algorithm determines if a detected contiguous region is a *correct detection*, a *potentially false detection*, or a *false detection*. There are three criteria used to determine the classification: the number of nearest neighbors, the distance to the nearest neighbors and the relative angle separation between the nearest neighbors.

Unfortunately clean, hard criteria do not exist in each category due to head shape variability, variability in the placement of the distribution, etcetera. Using EGI's HydroCel™ net to illustrate this, almost all of the interior electrodes are arranged in hexagonal arrangements, *Figure 13a*, which are tiled across the distribution, *Figure 14*. In 2 dimensions with equal spacing between the electrodes all of the internal angles, the acute angles between the dashed lines, are 60 degrees. While nominally 60 degrees, when projecting the hexagonal arrangement into 3 dimensions, along the curve of the scalp, the internal angles decrease slightly. Do to the variability of subject head sizes and shapes this change is subject and electrode placement dependent. Coupled with this is the spacing between electrodes within the arrangement is likely not exactly even. This has to do both with the subject head shape and size as well as how the electrode distribution was applied. The uneven spacing of the electrodes distorts the hexagonal shape slightly

increasing some internal angles while decreasing others, main lobe *Figure 15b*. This blurring likewise occurs with the number of nearest neighbors, *Figure 15a*, and neighbor distance, *Figure 15c*. Due to this blurring a hard threshold based on any individual property is infeasible. Instead the algorithm takes all factors into account simultaneously to classify a detected region into one of the three categories.

In addition to the blurring there is also an issue with where in the distribution the detected region is. The method by which the algorithm applies the criteria is different for the electrodes at the edge of the distribution than and those inside of it. Using the electrode distribution in *Figure 14* as an example, the electrodes at the edge compared to the inside of the distribution have half as many neighbors and a wider distribution in internal angles. The spacing of the edge electrodes with their neighbors is also more variable than those the inside of the distribution. This is because the edges of the distribution are from where the entire distribution is stretched and anchored to the head. This is done by various forms of chin and cheek straps which pull the distribution tight on head by pulling all of the edges forward and down. This frontal anchor point leads to a larger range of inter-electrode spacing at the edges.

Therefore the first step in the classification process is to separate the detected regions into the two categories: electrodes in the inside of the distribution and electrodes at the edges of the distribution. The algorithm does this by estimating of the edges of the electrode distribution. Detected regions within 2cm of the subject volume ears, and 5 cm of the neck, chin and nose trimming planes are considered as the edges of the electrode distribution, the rest of the detected regions are treated as inside the distribution.

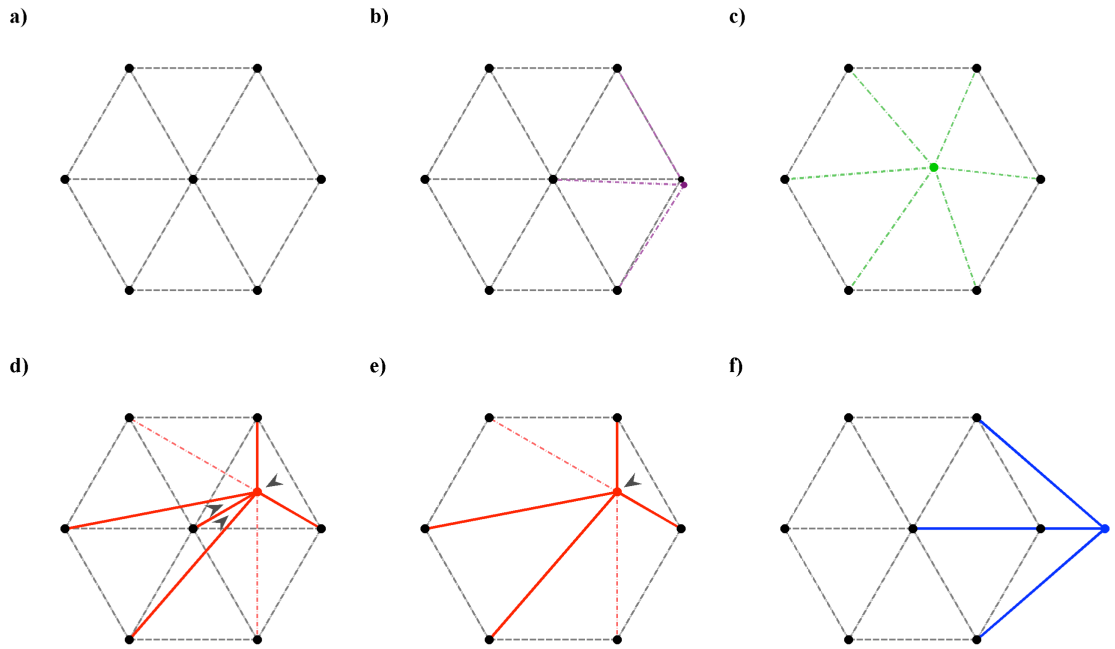

Figure 13: Example Scenarios of Detected Electrode Spacing. Depicted here is a hexagonal electrode arrangement like that found in the EGI HydroCel™ net. Dots represent the centroids of contiguous regions. Lines represent the distance to the nearest neighbors. Only neighbors within 5 standard deviation of the mean spacing of the electrode distribution are considered "nearest". Dashed lines represent being within 2 standard deviations of the mean. Solid lines represent being between 2 and 5 standard deviations of the mean. Gray arrows depict internal angles that are outside 2 standard deviations of the mean of the internal angles. The following scenarios are depicted: (a) normal spacing, (b) multiple detection (c) small shift in one of the electrodes in the distribution (d) a false detection (e) either improperly placed electrode (large shift) or false detection with missed electrode detection, (f) possible edge electrode or false detection.

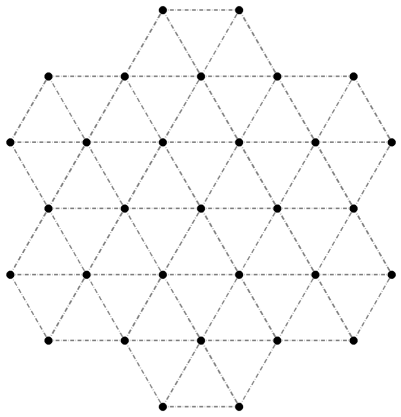

Figure 14: Example of the tiling of hexagonal electrode arrangement. 31 electrode centers of mass (black dots) are interconnected by 72 dashed lines representing the "nearest neighbor" connections. In this cartoon there are 13 internal electrodes and 18 edge electrodes. The edge electrodes have fewer neighbor electrodes (3 or 5) as compared to the internal electrodes (6). The distribution of acute angles is different as well. Here the edge electrodes with only three electrodes have two internal angles of 60° and the edge electrodes with 5 nearest neighbors have four internal angles of 60° and one of 120°.

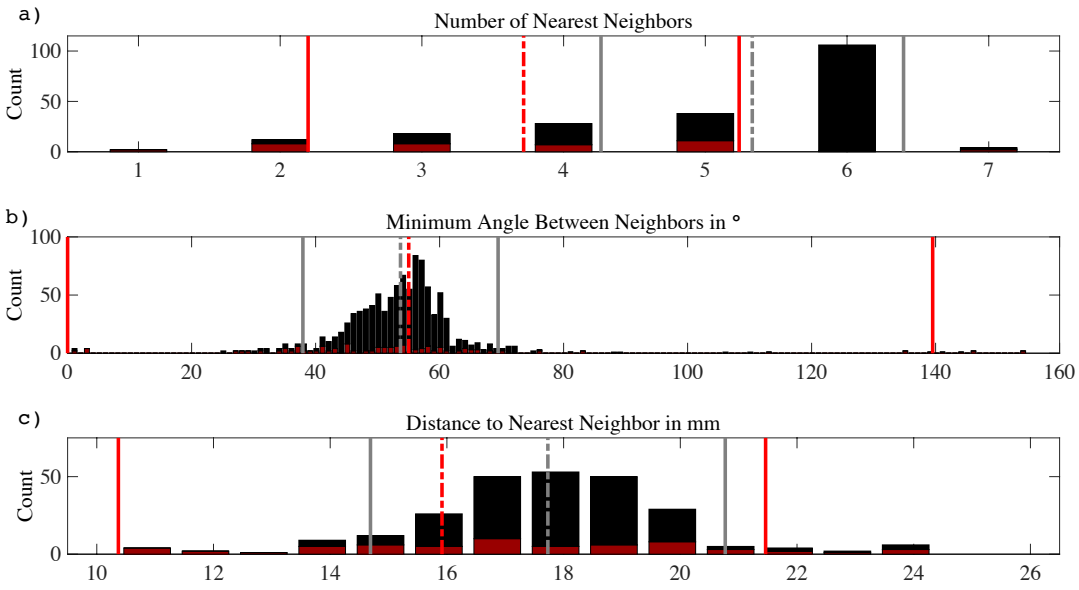

Figure 15 – Distributions of the relationships of the detected electrodes to their nearest neighbors. Distributions of the (a) number of the nearest neighbors, (b) the internal angles, and of (b) the distances to nearest neighbors are displayed with black bars representing the internal regions and red representing edge regions of the electrode distribution. Example data from the detected regions of an EGI HydroCel™ 256 MR compatible EEG net. Dashed lines represent the mean value of value within each of the three distributions, red for the edge electrode and gray from the internal electrodes. Solid lines represent the fences for each distribution, red for the edge electrodes and grey for the internal electrodes.

Fences are then applied to the internal and edge distributions separately for the number of nearest neighbors, internal angles and the distance to nearest neighbors. For internal detected regions the fences are within one standard deviation of the mean for number of nearest neighbors, and two standard deviations of the mean both for internal angles and distances to nearest neighbors. For edge-detected regions the fences are within one standard deviation of the mean for number of nearest neighbors, three standard deviations of the mean for the internal angles and two standard deviations for distances to nearest neighbors. Detected regions containing one or more members outside the fences in all three categories are considered false detections, *Figure 13d*. Regions containing one or more members outside the fences in two of the three categories are considered potential false detections, *Figure 13e*. All other detected regions are considered correct.

False detections are removed from further consideration, and potential false detections flagged for further steps of the algorithm.

Step 2: Electrode Channel Identification

After the electrodes regions are extracted from the 3D image and scrubbed from false detections the next step is to register the template from *step 0* to the detected regions. This will identify each detected region by channel number. This process involves six steps: identifying alignment electrodes in the detected distribution; scaling of the template to match the detected electrodes distribution; aligning the template to the detected distribution using an iterative closest point algorithm (Chen et al., 1989, Besl and McKay, 1992); using the channel relations in the template distribution to identify the electrodes in the detected distribution; scaling the template using singular value decomposition (SVD); and average displacement vector warping.

Step 2.1: Identifying Alignment Electrodes

Alignment electrodes are used in two ways. The first is to provide measures for the affine scaling of the template. The second is that they serve as seeds in the following step that identifies the detected electrodes based upon their relation to one another. The electrodes that are used for this purpose are those that are reliably identifiable based upon: (1) their location on the subject volume and, (2) their relation to their nearest neighbors. These will be determined a priori and will be model and manufacturer specific. In the case of the EGI HydroCel™ 256 channel net there are nine electrodes that stand out, *Figure 16*. Of the 269 possible detection points: only six are in the center in a pentagonal arrangement, *Figure 16a*; there is only one electrode on the mid-sagittal

subject plane that has only 2 nearest neighbors and a close to 90 degree internal angle, *Figures 10b & 16 a & b*; and finally there is only one electrode on each of the left and right sides that is closest to both the chin and neck trimming planes, *Figures 10c & 16b*.

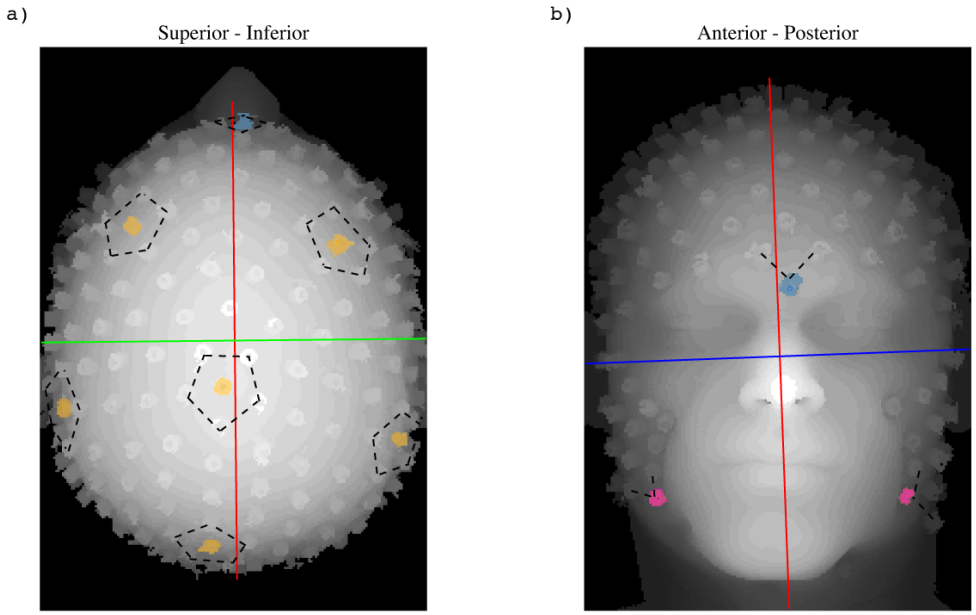

Figure 16: Axial and coronal half-depth projections with the alignment electrodes highlighted. (a) Axial half-depth projections with the electrodes at the center of pentagonal arangements highlighted in orange and the nasion electrode highlighted in blue. Nearest neigbor perimeters of the highlighted eledctrodes are outlined with black dashed lines in (a). Also displayed is the subject volume left – right plane in red and the subject volume anterior – posterior plane in green. (b) Coronal half-depth projections with the chin electrodes highlighted in magenta and the nasion electrode highlighted in blue. Nearest neigbor connections of the highlighted eledctrodes are traced with black dashed lines in (b). Also displayed in (b) is the subject volume left – right plane in red and the subject volume superior – inferior plane in blue. Nearest neigbor perimeters of the highlighted electrodes are outlined with black dashed lines.

The six electrodes at the center of the pentagonal arrangements, *highlighted in orange in Figure 16a*, are identified by the sizes of internal angles, the number of nearest neighbors, and their location on the scalp. In a pentagonal arrangement the electrode at the center should have the majority of its internal angles approximately equal 72°, *Figure 17a*. Likewise each should have five nearest neighbors, *Figure 17b*. In addition, these electrodes will be in the superior half of head and roughly equidistant from one another. These electrodes are identified by channel number based upon their location on head.

The nasion electrode, *highlighted in blue in Figure 16*, is identified in a similar manner to the pentagonal electrodes above. Its internal angle though, should be near to 90°, *Figure 18a*. Likewise it should have only two nearest neighbors, *Figure 18b*. In addition this electrode will be at the most anterior electron on the superior half of head and located near the mid-sagittal line.

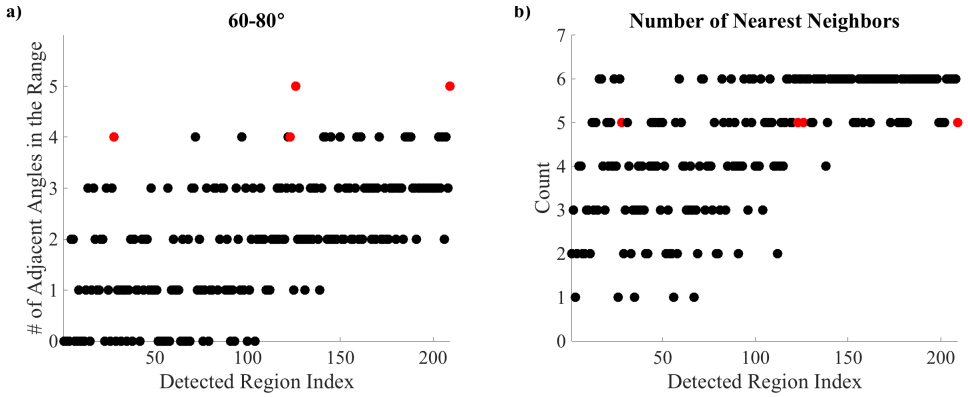

Figure 17: Scatter plots of number of adjacent angles in the 60° - 80° range and number of nearest neighbors by detected region index. (a) Scatter plots of number of adjacent (internal) angles in the 60° - 80° by detected region index. Here the 60° - 80° range encompasses the average 72° internal angle of a pentagonal arrangement. (b) Scatter plots of number of nearest neighbors by detected region index. Red dots in both plots indicate regions that meet the criteria for a pentagonal arrangement.

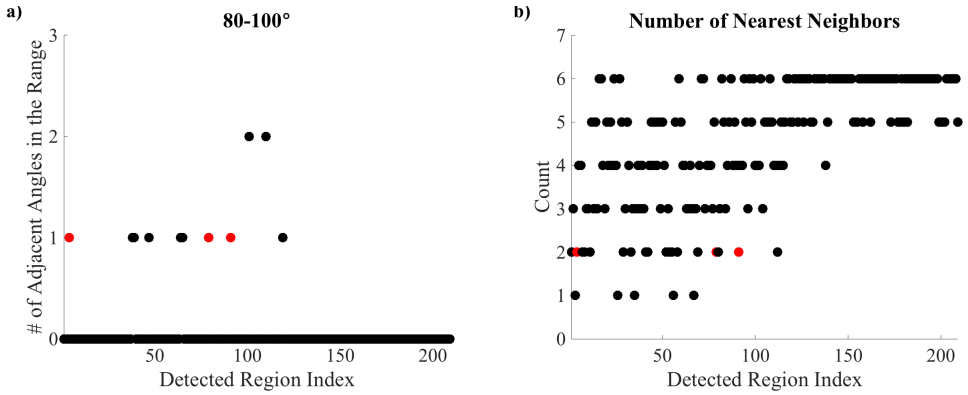

Figure 18: Scatter plots of number of adjacent angles in the 80° - 100° range and number of nearest neighbors by detected region index. (a) Scatter plots of number of adjacent (internal) angles in the 80° - 100° by detected region index. (b) Scatter plots of number of nearest neighbors by detected region index. Red dots in both plots indicate regions that meet the criteria for a pentagonal arrangement.

The two chin electrodes, *magenta highlighted electrodes, Figure 16b*, are located in the same manor. Their internal angle should approximately equal 90°, *Figure 18a*, and

likewise have only two nearest neighbors, *Figure 18b*. In addition these electrode will be the closest of all the electrodes to the intersection of the chin and neck trimming planes, *Figure 10c*. The chin electrodes are identified by channel number based upon which side of the head, left or right, they are on.

Step 2.2: Iterative closest point algorithm alignment

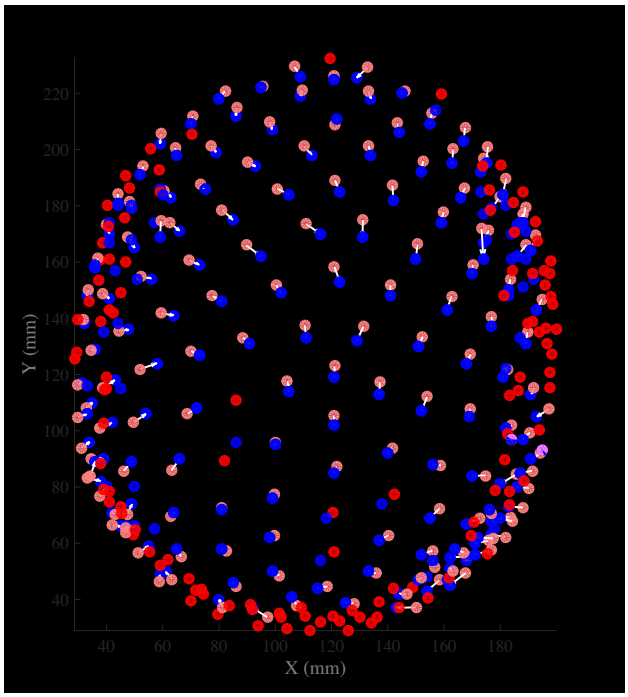

Figure 19: Detected and Template electrode point distributions after iterative closest point alignment. The distributions are displayed in an axial view from superior perspective. The detected distribution COMs can be seen in the blue dots. The template distribution point cloud can be seen in the red dots. The likely point parings between the two distributions are drawn with white vectors. The pairings are based upon the Euclidian distances between the two distributions.

The first step in the first alignment of the electrode template to the detected electrode distribution is to stretch the normalized template. This is done by matching the spatial spread of the alignment electrodes in the template to their spatial spread in the detected distribution using a simple affine transform. One this is done, the iterative closest point algorithm (ICPA) (Chen et al., 1989, Besl and McKay, 1992) is run to perform an initial registration of the template to the detected distribution. The reason that

it is an initial registration is because the ICPA is only a 6-degree of freedom transform, taking care of only translations and rotations.

Step 2.3: Align the template to Detected Electrodes on a Channel-by-Channel Basis

After ICPA the electrodes of the template are brought into register with the detected electrodes on a channel-by-channel basis. This is accomplished through an iterative series of non-linear transformations identifying detected electrodes based upon their Cartesian and Euclidian relationships to already identified electrodes. Using the alignment electrodes determined as above as seeds for this process, the algorithm moves though the detected distribution identifying the detected electrodes based upon on their relationship to one another and by their correspondence to the template.

Step 2.4: Singular Value Decomposition Stretching

Equation 6: Using Singular Value Decomposition (SVD) to stretch the unidentified electrodes in the template to match the identified electrodes in the detected distribution. (a) SVD of the identified electrodes in the detected electrode distribution. (b) SVD of the identified electrodes in the template electrode distribution. (c) multiplying the unidentified template electrodes by the scaling term from (b) and then the inverse of the scaling term from (a) scales the unidentified electrodes in the template to match the identified electrodes detected distribution.

(a) $\mathrm{DE}_{\mathrm{ID}} = U_D \Sigma_D V_D^*.$

(b) $\mathrm{TE}_{\mathrm{ID}} = U_T \Sigma_T V_T^*.$

(c) $\mathrm{TE}_{\mathrm{UID\ Streched}} = \mathrm{TE}_{\mathrm{UID}} \Sigma_T \Sigma_D'.$

With the majority of electrodes identified in *step 2.3*, the algorithm next stretches the remain unidentified electrodes in the template using SVD. This is done by performing SVD separately on both the identified electrodes in the template, *Equation 6b,* and detected distribution, *Equation 6a*. The unidentified electrodes in the template are multiplied by the diagonal matrix of singular values, Σ, from the SVD of the identified template electrodes, then multiplied by the inverse scaling term from the SVD of the

identified detected electrodes, *Equation 6c*. This operation scales the unidentified electrodes to match the scaling of the identified electrodes.

Step 2.5: Identifying by closest point neighbor

Once unidentified electrodes are stretched to match the identified electrodes in the detected distribution, the algorithm then proceeds to identify these remaining electrodes. It does so by finding the closet point pairs between the two distributions, *Equation 5a.* In doing so, the remaining detected regions are identified.

Step 2.5: Filling in the Undetected Channels

After all of the detected regions have been identified by channel number the remaining undetected electrodes are filled in using the template. This is done by applying a local warping to the template. The position of each undetected electrode in the template is warped based upon the average displacement of its nearest neighbors. This displacement is the vector between the template and detected electrode pairs, *white arrows, Figure 19*.

Testing the New Electrode Location and Identification Algorithm

To test the this new method nine MR structural scan were performed across three subjects wearing EGI's HydroCel™ 256 channel MR compatible net. Three of the scans were performed back-to-back to test the algorithms repeatability. The detection rate was measured across all scans to test the algorithms ability to identify the electrodes from the MR scan. Evaluation of the detection of the electrodes was performed by accentuating the intensity values of the detected electrodes and then viewing a surface rendering of the

MR images in Osirix Lite™. Evaluation of the identification of the channels was verified using a custom made graphical user interface (GUI), *Figure 20*.

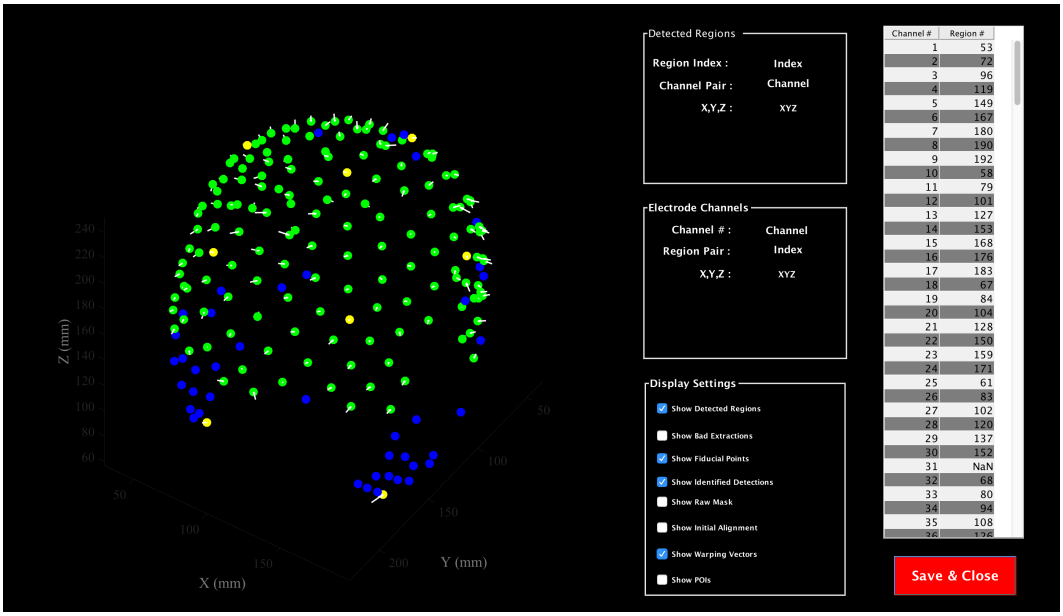

Figure 20: Screenshot of custom built GUI for the inspection of the identification of the detected regions by channel number. This screenshot is post the channel-by-channel identification step. The green dots represent identified detected electrodes. The blue dots represent unidentified detected electrodes. The yellow dots represent alignment electrodes. The white arrows represent the displacement vectors between the template and the detected regions.

Methods:

Subjects

Data from three male, Caucasian participants, mean age of 28, range 25 – 32, were used. Participants received verbal and written explanations of study requirements prior to any study procedures, and were provided written informed consent as approved by the UCLA Institutional Review Board, IRB#11-002946.

EEG Recording System

All MRI scans were performed with the subject wearing Electrogeodesic Inc.'s (EGI) HydroCel™ 256 MRI compatible net. The nets were connected to EGI's NetAmps 300 EEG amplifier. Nets were positioned on the human participants by aligning the

vertex electrode with the vertex of the head, identified at the midpoint of the anion-to-inion and left/right preauricular landmarks. Electrode impedances were <50 kΩ, as recommended by the manufacturer. Two additional external electrocardiogram (ECG) electrodes were applied to the chest of the human participants (positions: left midclavicular line in 5th intercostal space and 4th left intercostal space at sternal border).

MRI System and Scan Settings

All in-magnet EEG recordings were performed at isocenter in a 3T Siemens Trio MRI scanner (Erlangen, Germany). The receive coil was Siemens 12 Channel Head Coil and the transmit coil was the Body Coil. T2-SPACE structural images [TR, 3.2 s; TE 213 ms; FOV, 256 mm; matrix, 256 x 256; sagittal plane; slice thickness, 1 mm; 224 slices were prescribed to allow for full volume coverage.

Data Processing

All signal processing and analysis was performed on a 2012 MacBook Pro Retina, OSX 10.10.2, 2.7 GHz Intel Core i7 with 16Gb of ram (Cupertino, CA), running: Matlab™ 2012a with psychtoolbox (Brainard, 1997) version 3.0.12, and EEGLAB (Delorme and Makeig, 2004) version 13.4.4b; FSL version 5.0 (Jenkinson et al., 2012); and Osirix Lite version 7.0.1 (Rosset et al., 2004).

Brain Extraction

Brain extraction was performed using the FSL BET tool with the neck cleanup option (Smith, 2002)

Detection Confirmation

The detected positions of the electrodes were confirmed visually by setting their intensity value to 1000 and the inspecting the regions using surface rendering in Osirix Lite.

Identification Confirmation

The identification of the detected positions by channel number was done using the custom built GUI shown in *Figure 20*. The reason for the custom GUI was: 1) to allow ready access to EGI's channel map for their 256 channel HydroCel™ electrode net, 2) to aid in the visualization of the distribution detected regions and the template distribution 3) to enable the quick switching between the features displayed, e.g. detected electrode positions, alignment electrode positions, template electrode positions, warping vectors, and registered template positions.

Sensitivity, Specificity, and Accuracy Calculations

Measures of sensitivity, specificity, and accuracy for both the electrode detection and identification stages were made using *Equation 7.*

Equation 7: Calculations for both sensitivity (a) specificity (b) and accuracy (c). Here True is the number of correctly detected or identified electrodes, Missed is the number of missed electrode detections or identifications, and False is the number of incorrectly detected electrodes or identified electrodes.

$$\text{(a)}\ \ S_E = 100 \times \frac{\text{True}}{\text{True} + \text{Missed}}.$$

$$\text{(b)}\ \ S_P = 100 \times \frac{\text{True}}{\text{True} + \text{False}}.$$

$$\text{(c)}\ \ A = 100 \times \frac{\text{True}}{\text{True} + \text{False} + \text{Missed}}.$$

Results & Discussion:

Overall data quality

In all nine MRI datasets the electrodes were extracted and identified. The electrode positions were confirmed by manual inspection to assure correct electrode location and identification using Osirix Lite program and the GUI in shown *Figure 20*. RF parasitic losses were seen in all of the nine data sets. The losses were particularly large posterior to the mandible and inferior to the inion, *left image, Figures 21 & Figure 22a*. These areas coincide with where the lead cabling gathers in the HydroCel™ net, *Figure 22b*.

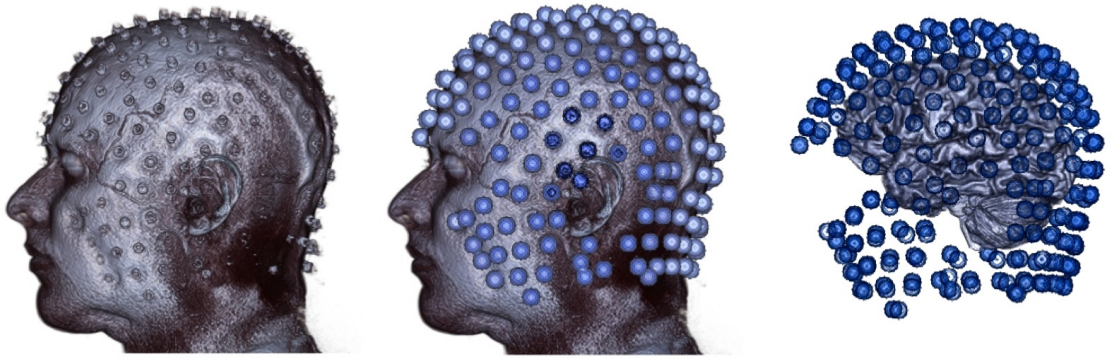

Figure 21: Algorithm progression from MR acquisition through identification. A surface intensity rendering of the raw MR data made using Osirix Lite is displayed in the left image. Surface intensity rendering overlaid with the detected and registered electrode net is displayed in the center image. Brain extracted using the FSL BET toolbox overlaid with the detected and registered electrode net is displayed in the right image. Note the signal dropout from RF parasitic losses posterior to the mandible and inferior to the inion.

Detection step performance

The mean initial electrode extraction, i.e. the sensitivity for all of the data sets, was 69.56%, or 187.1 of the 269 possible detection points. The range was 62.45% – 77.32%, or 168 – 208 points. The 269 points being the 256 channels, the reference electrode, the common electrode and 11 dummy electrodes used to create the distribution shape. In three repeated measures (MR Scans) the mean difference in detected electrode

locations was 0.6117mm with a standard deviation of 0.6294mm. Across all of the scans the specificity was 100%. No detected region other than electrode made it though the error checking in *steps 1.2.2 and 1.2.3*. The accuracy was therefor the same as the sensitivity, 69.56%.

The low sensitivity resulted from three factors: RF parasitic losses, the compression of the electrodes against the surface of the head, and from the conservative thresholds used throughout the algorithm. Areas of parasitic loss where a mixed bag in terms of detection. Depending on the level of signal loss some recovery of the signal was possible by extracting the shell rather relying on signal intensity. As an example, the areas posterior to the mandible and inferior to the inion, where electrodes cannot be seen in the surface intensity image, *left image, Figure 21*, they can be seen in the depth projection image, *Figures 6c, 7, & 10c*. Areas where there was almost a complete signal loss however, *red arrow Figure 21a,* were unrecoverable. Another issue cofounding the detection in these regions was the compression of the electrodes against the surface of the head as a consequence of the subject position, ear protection and head padding. The subjects were in the supine position for the scans, consequently their heads were resting on the occipital channels of the EEG net resulting in compression of these electrodes against their heads. The headphones used to provide hearing protection compressed the electrodes surrounding the ears. Lastly, the padding placed at the sides of the head roughly inline with the mid-coronal plane (just posterior to the mandible), that was used to prevent subject motion in the scanner, compressed the electrodes in these region against the head. The compression of these electrodes changed their profiles such that they became slightly indented, i.e. protruded less from the surface of the head. Lastly the

conservative thresholds used throughout the algorithm were very effective at preventing false detections but they also prevented correct detections in these regions. The coupling of the RF parasitic losses, the change in profile and the conservative thresholds, crippled the detection in these regions. By contrast, detection in superior half of the head and check regions by the algorithm was almost 100%, *green and blue dots, Figure 20*.

The high specificity of the algorithm was the result of three factors: the restriction of the search area by the preprocessing done in *step 1.1*, the conservative thresholds used throughout the algorithm, and the layers of error checking performed by the algorithm in *steps 1.2.2 and 1.2.3.* The preprocessing done in *step 1.1*, especially the surface trimming, helped greatly in the elimination of false detections by removing approximately 95% of the voxels from the electrode search. The surface feature trimming was not perfect though, especially for the ears. As an example in *Figure 9, bottom row* an electrode can be seen being incorrectly identified as part of the ear. Sensitivity also was sacrificed for specificity in the setting of thresholds by the algorithm. In *Figure 11c* an example of this can be seen be seen in voxels that are clearly part of an electrode being removed. In moving the threshold to the left and including these voxels in the detection the sensitivity would increase but there would be an increase in the voxels identified as being part of an electrode that are in fact not, a decrease in specificity. The voxels that are incorrectly identified though have a pretty large barrier to pass from both the bounding box and connectivity filters in step *1.2.2*. Even so, this pales in comparison to the hurdle that must be passed in the error checking performed in step *1.2.3*. The constraint of having to be located where a correct detection should be, essentially eliminates false detections. There is a caveat though with *step 1.2.3*.: the number false detections must be

far less than the number of correct detections. This is because the false detections throw off the statistics for the neighboring distribution calculations. A few likely can be absorbed, as the number grows though, the more the fences set in step 1.2.3 will be affected, thus the needed trade off of sensitivity for specificity.

In the test of repeated measurements, between the three scans acquired back-to-back there was average measurement difference 0.6117mm with a standard deviation of 0.6294mm. This mean lies between the limit of detection (LOD) and the limit of quantization for this scan (LOQ). The LOD being 0.5mm or half a voxel width. The LOQ being the resolution of the scan, 1mm. Registration errors between the scans from subject motion likely accounts for the differences.

Identification step performance

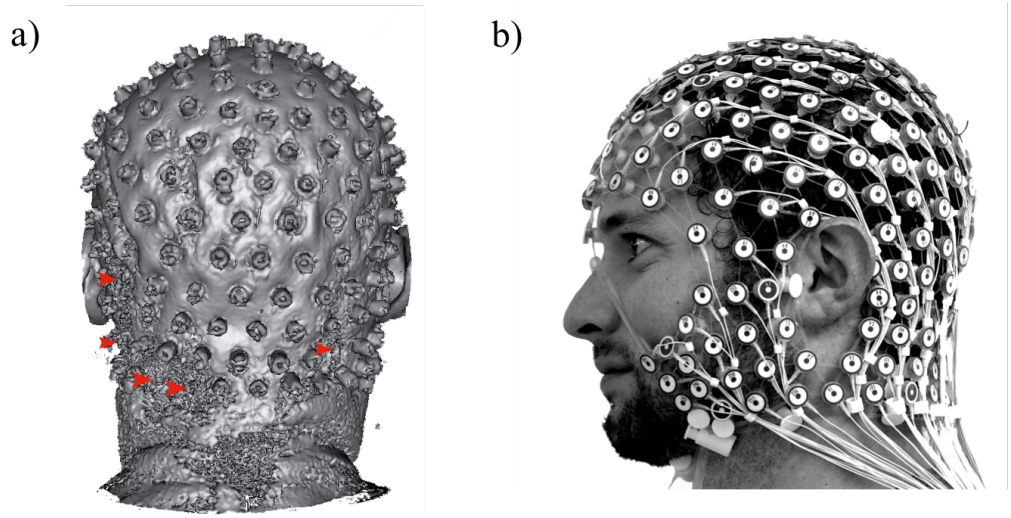

Figure 22: Side-by-side images of a surface rendering from an MR scan and photograph of a subject wearing HydroCel™ 256 channel MR compatible net. (a) Anterior View of a surface intensity rendering of raw MR Data made using Osirix Lite. (b) Photograph of a sagittal view of a subject wearing EGI's MR compatible EEG net. Red markers in (a) indicate electrode positions that were unable to be identified by manual inspection due both to RF parasitic losses and magnetic susceptibility mismatches. Note the areas that coincide with the parasitic losses in (a), match where the electrode lead cabling bundles in (b)

In the identification stage the sensitivity, specificity and accuracy amongst the detected electrodes in all data sets was 100%. While the 100% across the board seems at first glance too good to be true, the reality is that the identification stage of the algorithm is highly constrained and sacrifices made in detection stage resulting in its lower sensitivity were traded for higher performance in this stage. Due to the high specificity in the detection stage the registration (identification) stage was simply aligning two matched point clouds.

Computer Performance

The average runtime of the algorithm on the computer used for the data analysis was under a minute. Much of this speed was achieved by restricting the search regions in the preprocessing done in *step 1.1*, using a Fourier based multiplication rather than convolution in the detection *step 1.2.1*, and using the highly efficient iterative closest point algorithm for the initial registration of the template to the detected electrodes in *step 2.2*.

Comparison to other methods

Comparison of these results to the manual measurement methods is of limited value because the electrodes move between their placement on the scalp outside of the scanner and their measurement inside the scanner (Yoo et al., 1997). That being said the sensitivity and specificity for the identification of the electrodes using a manual method will likely be very high, if not 100%. Measuring the location of the electrodes with respect to the underlying anatomy in these methods is a different story. Location errors in these methods are actually the compounding of the errors of two separate measurements;

errors in the location of the surface fiducial markers in the MR scan and errors in the location of the electrodes with respect to the surface fiducial markers. According to the manufacturers of the Fastrack™ electromagnetic digitization system (Polhemus Inc., Colchester, Vermont) the accuracy of their system is about 8mm (Koessler et al., 2007), which is similar to the accuracy found Khosla et al. (1999). To improve this accuracy Russell et al. (2005) in their study measured the electrode position several times and were improve it to 1.02mm. Though, if measuring the location of 64 electrodes once took Khosla et al. (1999) 15 minutes using the Fastrack™ system, then a single measurement of each position on EGI's 256 channel HydroCel™ net would require about 60 minutes. This time makes multiple measurements for a high density electrode arrays for all practical purposes a non-starter. That being said an 8mm error in the electrode positions from a single measurement is likely to be fairly random in its direction and as such will have little consequence for the localization of sources (Wang, 2001). Systematic errors in electrode positions are much more problematic. They mainly come about by the incorrect location of the surface fiducial markers in the MR scan, or the electrode position measurement, or both. Systematic errors of just 5° lead to large errors (+1cm) in the location of sources (Akalin Acar and Makeig, 2013). For reference, 5° of tilt on a 56cm circumference head, average adult male (Farkas et al., 1992), would be equivalent to a systematic shift on the surface of the head of slightly less than 8mm, or just over half the diameter of one of EGI's electrode pedestals and within the accuracy of Fastrack™ system.

In comparing the 69.56% detection using our method, to EGI's photogrammetry system, similar results were found by Russell et al. (2005) of 74.42%. The detection

criteria for Russell et al. (2005) being that an electrode was visible in more than one camera (Russell et al., 2005). In terms of identification performance measures for EGI's photogrammetry system, manual review of the images is required; false positives, and missed detections, are resolved by the operator (Russell et al., 2005). The estimated time to perform this manual review is 30-40 minutes for the 128 channel EGI HydroCel™ net (Russell et al., 2005), which likely translates to 60-80 minutes for the 256 channel EGI HydroCel™ net. In terms of electrode position errors the mean error in Russell et al. (2005) study was 1.27mm. Like the electrode measurement error in manual method above this is likely to be random in its displacement and as such has little consequence for location of EEG sources. Though, again, similarly to the manual method above, because the electrode positions are measured outside of the scanner they must be brought into register with the anatomical MR scans through the use of surface fiducial markers. There is therefore a possibility for the introduction of systematic errors in the electrode positions with regards to the underlying anatomy from errors in the location of the surface fiducial markers both in the MR scan and photogrammetry data.

A comparison of the use of exogenous markers, whether placed on the electrodes (Yoo et al., 1997, Sijbers et al., 2000), or on the surface of the head (Lagerlund et al., 1993, Brinkmann et al., 1998) with this method is not an "apples to apples" comparison. This is because none of these methods identified the electrodes. The markers used simply to aid in the detection of the electrodes. That being said, in terms of measurement accuracy of the markers placed on the surface of the head, Brinkmann et al. (1998) found the registration error between the centroid locations of the markers to be 2.21mm. either Yoo et al. (1997) nor Sijbers et al. (2000) listed accuracy values.

Comparing the method described here to the electromagnetic digitization and photogrammetry methods, similar performance was found in the specificity of the identification of the electrodes. The detection sensitivity of this algorithm was similar to photogrammetry. The major advantages of this method over the current primary methods for measuring the electrode positions are two fold. The first is that this method eliminates the registration of the electrode positions measured outside of the scanner to the measurement of the anatomy made inside the scanner, eliminating a source of potential systematic error. The second is the huge, at least 60 fold, time advantage of this method has over the current methods. This gain in speed is from the fully automatic nature of the algorithm, i.e. no user input needed, and from the elimination of measurement step and registration steps.

Conclusion:

The method presented and tested here for locating and identifying EEG electrodes from an anatomical MRI scan proved effective and highly time efficient. The time improvements over the current alternative methods were the result of three factors: the elimination of a separate measurement step for the electrode positions, the elimination of a registration step of the electrode positions to the subject anatomy, and through the automatic nature by which the algorithm proceeds. In addition to the time savings gained by eliminating the registration step, doing so also eliminated a potential source of error. Likewise the ability of the algorithm to run automatically, without the need for user input, not only provided a time gain but eliminated possibility for the introduction of variability in measurements within and between operators.

Limitations of this study and method

The study was limited by two things, the subject pool and the lack of ground truth for the electrode positions. The limitation of this method was inability of the detection step to locate electrodes anterior to the mandible and inferior to the inion.

The subject population used here was single gender, single race, and consisted of only three individuals. The effects of various head shapes and hair styles were not thoroughly tested on both the location and identification steps. It is likely that further refinement will be needed to accommodate various head shapes in locating the trimming planes in the preprocessing step. Hair is invisible in the MR translucent and should not introduce additional error in the detection step. In fact, rather than harming the detection, some hairstyles may prove effective in padding certain electrode regions from compression against the head. This padding may actually aid in the detection. Additional hair volume also may be useful for spacing the lead wiring further from the surface of the head and the electrodes. This may reduce the effects of RF parasitic losses in both the head and the electrodes.

The lack of ground truth for the electrode positions was a result of the fact that any electrode position measured outside could not be trusted as a measure for its position inside the scanner. Electrode positions outside the scanner are measured with the subject in a seated position. In the scanner they are measured with the subject in the supine position. This shift in position means that pressure is applied to occipital electrodes likely resulting in some displacement. Padding also is applied around the periphery of the head when the subject is placed in the scanner to minimize their head movement. This padding applies pressure to the electrodes beneath, likely resulting in their movement. Likewise

hearing protection, in the form of head phones are placed over the subject's ears. Like above this likely results in the movement electrodes beneath. The combination of these was likely behind the movement the of the makers noticed be Yoo et al., 1997 in their study.

The reduced detection in the inferior posterior portion of the head was due in part to compression of the electrodes in this region, and from the RF parasitic losses resulting from the accumulated lead cabling in this region. This lack of detection meant the location of the electrodes in these regions had to be estimated by the identification step of the algorithm. There were several areas where this signal loss was so great the correctness of this estimation could not be verified.

Future Improvements

A potential improvement to this algorithm would be to run it iteratively. That is to say the locations of the detected and identified electrodes in the first round could be used as seeds for a second round of detection and identification. Much like the identification step uses the neighbor relationships to identify the detected electrodes by channel number, these same relationships can likely be used to aid in their detection as well. This may be one way to deal with issue of detecting electrodes that have compressed against the surface of the head.

Alternate implementations

While the method presented here used the volume images from MR scans, it should be translatable to other volume imaging methods, e.g. 3D surface scans performed by a laser scanner, Microsoft Kinect™, etc. In addition, while the domains of contrast

used here to highlight the electrodes were, shape, size and orientation, other contrast domains should be applicable as well, e.g. magnetic susceptibility. In the extreme case where the electrodes are not MR lucent, and no difference in things such as magnetic susceptibility can be found, the indentation profile of the electrodes against the head, *left image Figure 21, and Figure 22a,* might serve as a way for locating the electrodes.

Final thoughts

The brain's position in the skull, and how CSF is distributed within it, is known to change with subject position. The shift in brain position is on the order of millimeters and the change of the thickness of the CSF layer can be up to 30% between seated and supine positions (Rice et al., 2013). Errors in the head model are viewed to be the most important source of error in source localization (Akalin Acar and Makeig, 2013). As such it seems that any source localization study should be performed with the subject in the same position that their anatomical scan was measured in. This is regardless of the method used to measure the electrode positions.

References:

Akalin Acar Z, Makeig S (2013) Effects of forward model errors on EEG source localization. Brain Topogr 26:378-396.

Akalin-Acar Z, Gençer NG (2004) An advanced boundary element method (BEM) implementation for the forward problem of electromagnetic source imaging. Physics in medicine and biology 49:5011.

Besl PJ, McKay ND (1992) Method for registration of 3-D shapes. In: Robotics-DL tentative, pp 586-606: International Society for Optics and Photonics.

Brainard DH (1997) The Psychophysics Toolbox, Spatial Vision. Spatial vision 10:433-436.

Brinkmann BH, O'Brien TJ, Dresner MA, Lagerlund TD, Sharbrough FW, Robb RA (1998) Scalp-recorded EEG localization in MRI volume data. Brain topography 10:245-253.

Chen J, Vandewalle J, Sansen W, Vantrappen G, Janssens J (1989) Adaptive method for cancellation of respiratory artefact in electrogastric measurements. Medical and Biological Engineering and Computing 27:57-63.

Dalal SS, Rampp S, Willomitzer F, Ettl S (2014) Consequences of EEG electrode position error on ultimate beamformer source reconstruction performance. Front Neurosci 8:42.

De Munck J, Vijn P, Spekreijse H (1991) A practical method for determining electrode positions on the head. Electroencephalography and clinical neurophysiology 78:85-87.

Delorme A, Makeig S (2004) EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. J Neurosci Methods 134:9-21.

Farkas LG, Posnick JC, Hreczko TM (1992) Anthropometric growth study of the head. The Cleft Palate-Craniofacial Journal 29:303-308.

Gençer NG, Acar CE (2004) Sensitivity of EEG and MEG measurements to tissue conductivity. Physics in medicine and biology 49:701.

Grave de Peralta R, Gonzalez Andino S (2002) Comparison of algorithms for the localization of focal sources: evaluation with simulated data and analysis of experimental data. Int J Bioelectromagn 4.

He P, Estepp JR (2013) A practical method for quickly determining electrode positions in high-density EEG studies. Neurosci Lett 541:73-76.

Henderson C, Butler S, Glass A (1975) The localization of equivalent dipoles of EEG sources by the application of electrical field theory. Electroencephalography and clinical neurophysiology 39:117-130.

Jenkinson M, Beckmann CF, Behrens TE, Woolrich MW, Smith SM (2012) Fsl. Neuroimage 62:782-790.

Khosla D, Don M, Kwong B (1999) Spatial mislocalization of EEG electrodes–effects on accuracy of dipole estimation. Clinical neurophysiology 110:261-271.

Koessler L, Maillard L, Benhadid A, Vignal J-P, Braun M, Vespignani H (2007) Spatial localization of EEG electrodes. Neurophysiologie Clinique/Clinical Neurophysiology 37:97-102.

Lagerlund TD, Sharbrough FW, Jack CR, Erickson BJ, Strelow DC, Cicora KM, Busacker NE (1993) Determination of 10–20 system electrode locations using magnetic resonance image scanning with markers. Electroencephalography and clinical neurophysiology 86:7-14.

Michel CM, Murray MM, Lantz G, Gonzalez S, Spinelli L, Grave de Peralta R (2004) EEG source imaging. Clin Neurophysiol 115:2195-2222.

Mullinger K, Debener S, Coxon R, Bowtell R (2008) Effects of simultaneous EEG recording on MRI data quality at 1.5, 3 and 7 tesla. International Journal of Psychophysiology 67:178-188.

Nishimura DG (2010) Principles of magnetic resonance imaging: Standford Univ.

Pascual-Marqui RD, Michel CM, Lehmann D (1994) Low resolution electromagnetic tomography: a new method for localizing electrical activity in the brain. International Journal of psychophysiology 18:49-65.

Rice JK, Rorden C, Little JS, Parra LC (2013) Subject position affects EEG magnitudes. Neuroimage 64:476-484.

Rodriguez C, Cohen, M. (2014) Fully automated localization of electroencephalography (EEG) electrodes. Regents of the University of California.

Rosset A, Spadola L, Ratib O (2004) OsiriX: an open-source software for navigating in multidimensional DICOM images. Journal of digital imaging 17:205-216.

Russell GS, Jeffrey Eriksen K, Poolman P, Luu P, Tucker DM (2005) Geodesic photogrammetry for localizing sensor positions in dense-array EEG. Clin Neurophysiol 116:1130-1140.

Sijbers J, Vanrumste B, Van Hoey G, Boon P, Verhoye M, Van der Linden A, Van Dyck D (2000) Automatic localization of EEG electrode markers within 3D MR data. Magnetic resonance imaging 18:485-488.

Smith SM (2002) Fast robust automated brain extraction. Human brain mapping 17:143-155.

Srinivasan R, Tucker DM, Murias M (1998) Estimating the spatial Nyquist of the human EEG. Behavior Research Methods, Instruments, & Computers 30:8-19.

Steddin S, Bötzel K (1995) A new device for scalp electrode localization with unrestrained head. J Neurol 242:65.

Tadel F, Baillet S, Mosher JC, Pantazis D, Leahy RM (2011) Brainstorm: a user-friendly application for MEG/EEG analysis. Comput Intell Neurosci 2011:879716.

Wang Y, & Gotman, J. (2001) The infuence of electrode location errors on EEG dipole source localization with a realistic head model. Clinical Neurophysiology 112:1777-1780.

Yoo S-S, Guttmann CR, Ives JR, Panych LP, Kikinis R, Schomer DL, Jolesz FA (1997) 3D localization of surface 10–20 EEG electrodes on high resolution anatomical MR images. Electroencephalography and clinical neurophysiology 102:335-339.