Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.

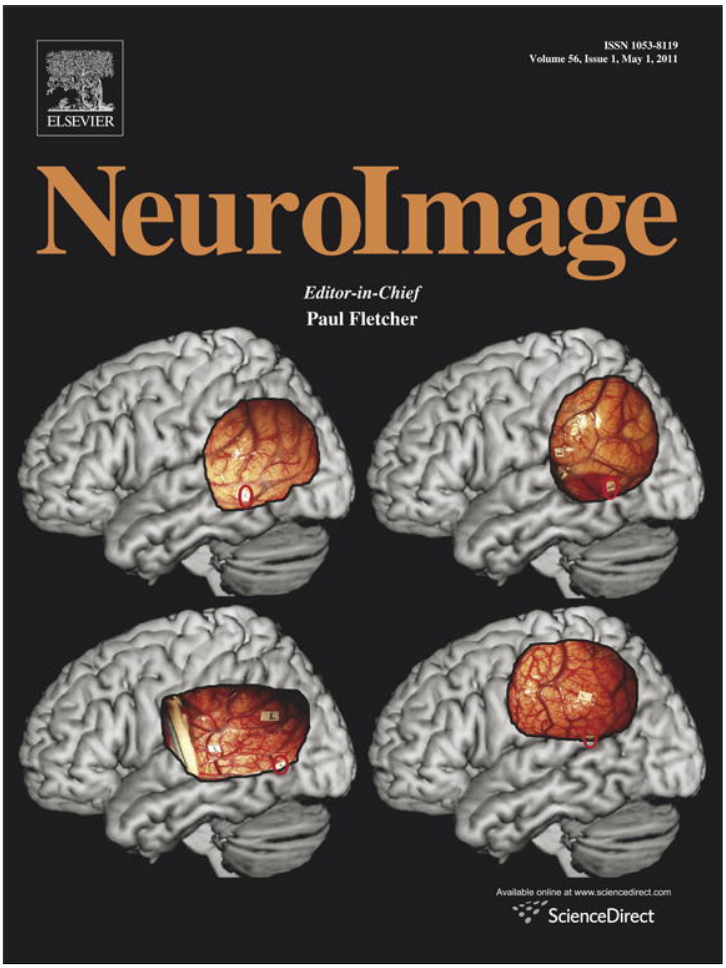

This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

http://www.elsevier.com/copyright

Author's personal copy

NeuroImage 56 (2011) 69–77

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

A wavelet-based method for measuring the oscillatory dynamics of resting-state functional connectivity in MEG

Avniel Singh Ghuman *, Jonathan R. McDaniel, Alex Martin

Laboratory of Brain and Cognition, National Institute of Mental Health, USA

ARTICLE INFO

Article history:
Received 19 September 2010
Revised 29 December 2010
Accepted 13 January 2011
Available online 21 January 2011

Keywords:
Resting-state
Magnetoencephalography
Functional connectivity
Phase-locking
Neural oscillations
Oscillatory dynamics

ABSTRACT

Determining the dynamics of functional connectivity is critical for understanding the brain. Recent functional magnetic resonance imaging (fMRI) studies demonstrate that measuring correlations between brain regions in resting-state activity can be used to reveal intrinsic neural networks. To study the oscillatory dynamics that underlie intrinsic functional connectivity between regions requires high temporal resolution measures of electrophysiological brain activity, such as magnetoencephalography (MEG). However, there is a lack of consensus as to the best method for examining connectivity in resting-state MEG data. Here we adapted a wavelet-based method for measuring phase-locking with respect to the frequency of neural oscillations. This method employs anatomical MRI information combined with MEG data using the minimum norm estimate inverse solution to produce functional connectivity maps from a "seed" region to all other locations on the cortical surface at any and all frequencies of interest. We test this method by simulating phase-locked oscillations at various points on the cortical surface, which illustrates a substantial artifact that results from imperfections in the inverse solution. We demonstrate that normalizing resting-state MEG data using phase-locking values computed on empty room data reduces much of the effects of this artifact. We then use this method with eight subjects to reveal intrinsic interhemispheric connectivity in the auditory network in the alpha frequency band in a silent environment. This spectral resting-state functional connectivity imaging method may allow us to better understand the oscillatory dynamics underlying intrinsic functional connectivity in the human brain.

Published by Elsevier Inc.

Introduction

Measuring how neural regions interact is critical for understanding the dynamics of the normal and disordered brain. "Functional connectivity" is thought to reflect these interactions and is defined as "the correlation between spatially remote neurophysiological events" (Friston et al., 1993). Correlation is used as a measure of functional connectivity based on the principle that if two neuronal populations fire together, they are likely to be part of the same functional circuit. Traditionally, changes in the correlations between neural populations are measured across tasks or cognitive states. Recently, it was discovered that the activity in regions forming task-critical networks (for example the networks associated with visual, auditory, memory, and sensorimotor functions Biswal et al., 1995; Cordes et al., 2000; Vincent et al., 2006); correlate even when those tasks are not being performed (when the subject is at "rest," in light sleep, or even sedated; see Boly et al., 2008; Fox and Raichle, 2007 for reviews). Resting-state functional connectivity has garnered a great deal of interest as a method for examining functional networks in a "natural" state. This interest arises from two aspects of resting-state functional connectivity. First, because these networks arise without being driven by a task, this method has the potential to illustrate fundamental aspects of the brain's intrinsic functional organization (Fox and Raichle, 2007). Second, this lack of task can allow us to examine functional connectivity in clinical populations where behavioral responses may be abnormal (Greicius, 2008).

The majority of resting-state functional connectivity studies have been performed using functional magnetic resonance imaging (fMRI). Because of the relatively poor temporal resolution of fMRI (~.5–1 Hz and below) these studies have been restricted to examining correlations in slow resting-state oscillations. These studies have demonstrated that there are strong correlations within functional networks in the very low frequency (<.1 Hz) aspects of the resting-state activity (Fox and Raichle, 2007). Correlations in this low frequency band are surprising since most electrophysiological aspects of neural activity occur at a much faster time scale. One likely possibility that has emerged is that these very slow fluctuations in large part reflect slow changes in underlying higher frequency neural activity (He et al., 2008; Leopold et al., 2003; Mantini et al., 2007; Nir et al., 2008). However, little is known regarding the exact nature of correlated activity in these higher frequencies. To explore these faster aspects of resting-state functional connectivity requires methods to

* Corresponding author at: Laboratory of Brain and Cognition, NIMH, 10 Center Drive, Room 4C116, MSC 1366, Bethesda, MD 20892-1366, USA.
E-mail address: ghumana@mail.nih.gov (A.S. Ghuman).

1053-8119/$ – see front matter. Published by Elsevier Inc.
doi:10.1016/j.neuroimage.2011.01.046

Author's personal copy

examine interactions using electrophysiological measures of neural activity that have a higher temporal resolution, such as magnetoencephalography (MEG).

Methods for measuring functional connectivity in non-invasive electrophysiological data have generally examined task/condition differences in the correlation of the time–frequency response of the neural activity arising from disparate brain regions (Jerbi et al., 2007; Lin et al., 2004; Tass et al., 2003). These studies were designed to allow researchers to make inferences about whether the task or cognitive state modifies functional connectivity. For example, the spectrum of interregional correlations has been examined before and after learning (Duzel et al., 2005; Ghuman et al., 2008), compared across attentional states (Gross et al., 2004), and for different visual stimuli (Bar et al., 2006; Kaiser et al., 2004). Many of these approaches were adapted from methods that have been used successfully to examine functional connectivity using electrode recordings in animals (Engel et al., 2001; Roelfsema et al., 1997; Varela et al., 2001).

There are many methodological issues that must be considered in resting-state functional connectivity before adapting task-related approaches to resting-state data. These issues appear because of the lack of a within-subjects comparison condition in resting-state studies. For example, when examining resting-state functional connectivity between brain areas, one must take into account the fact that some of the activity projected to each area originates from common sensors. In task-based studies, crosstalk is somewhat mitigated because it is present to some degree in all conditions and is reduced when a comparison across conditions is performed.

Here we introduce a novel wavelet-based method for measuring resting-state phase-locking between electrophysiological signals that are measured non-invasively, but mapped onto the human brain, and introduce normalization to reduce a major crosstalk artifact. Wavelet-based analyses have the advantage of not requiring the data to be stationary (Percival and Walden, 2000) and therefore are likely to be more appropriate for non-stationary neural data than Fourier-based methods. It should be noted that while we primarily discuss this method for MEG, in principle it could be used for any non-invasive electrophysiological measure of brain activity, such as electroencephalography (EEG). After describing the method, we use simulations to test the spatial sensitivity and specificity of the functional connectivity on the cortical surface. We use these simulations to examine a key artifact that appears in non-invasive resting-state functional connectivity analyses of electrophysiological data and demonstrate a procedure for reducing much of this problem. Finally, we apply the method to examine the spectral functional connectivity in the left (LH) and right hemisphere (RH) auditory network using MEG with eight subjects. While these MEG results are primarily used to validate the method, this is also the first demonstration of connectivity between the LH and RH auditory cortices in a true resting state. Previous studies have used fMRI to examine connectivity in this network (Cordes et al., 2000), but fMRI cannot be considered a true resting state for the auditory cortex because of the noise the MRI machine produces; in contrast, MEG is silent.

Methods

The method for calculating resting-state functional connectivity was adapted from the dynamic statistical parametric mapping method developed by Lin et al. (2004). Specifically, the process involves six steps: 1) artifact removal 2) selection of "seed" region of interest (ROI) 3) calculation of the inverse solution and projection onto the brain 4) wavelet transformation of the signal 5) calculation of the phase-locking values (PLVs; Lachaux et al., 1999) between the seed ROI and every other location in the brain and 6) repeat steps 3, 4, and 5 with empty room data and normalize the original PLVs by the empty room noise PLVs to reduce crosstalk induced by the imperfect inverse solution. Note that when a linear inverse operator is used, the order of steps 3 and 4 does not matter.

Artifact removal

Removing artifacts due to heartbeats, eye blinks, eye movements, and other non-neural sources is critical because some of these signals can dwarf the neural component of the MEG data. The magnitude and ubiquity of these distortions can result in spurious phase-locking and correlation over large portions of the brain.

For artifact removal, we first visually inspected the data for any respiratory artifacts. These artifacts are generally uncommon in MEG data, particularly with the use of third-order gradiometer compensation. However, if subjects have metal on or in their persons that they have not removed, or are unaware of, these artifacts can be substantial (one subject in our study was excluded for this reason). We then used a short-time Fast Fourier filter to band-pass the data 1–50 Hz. This removed low frequency drift, any residual respiratory artifact, 60 Hz line noise, and any DC offset.

Cardiac artifacts were removed using an independent component analysis-based procedure following Liu et al. (2010). Briefly, the MEG sensor data were decomposed into a number of independent components (ICs) using EEGLAB (Delorme and Makeig, 2004). The ICs was identified as being a cardiac artifact if the IC had a peak in its autocorrelation coefficient corresponding to between .6 and 1.5 Hz, the IC had a time course that contained periodic features that were similar to those seen on an electrocardiogram, and had power in the MEG sensors that experience suggested most often contained cardiac features. Across our 8 subjects, between 1 and 3 cardiac ICs were rejected and the remaining ICs were reassembled for further processing.

Typically, ocular-motor artifact rejection is accomplished by removing trials where eye blinks are evident and when there are spikes in the MEG signal, which is impossible for a resting-state analysis because of the lack of trials. Alternatively, artifacts are removed by removing the component of the signal that corresponds to eye blinks or an electrooculogram (EOG) channel using independent components analysis (ICA) or principle components analysis (PCA). However, ICA/PCA methods leave some residual MEG signal that corresponds to the artifact. Because the MEG signal from eye blinks can be up to 10 times the magnitude of the brain signal this residual artifact signal can result in false-positive functional connectivity.

To monitor ocular muscle activity we measured EOG along with the MEG measurements. The EOG time courses were pseudo-Z transformed into standard deviation (SD) units (by subtracting the mean across the epoch and dividing by the SD). These data were visually inspected and the minimum size of each subject's eye blinks in these units was determined (mean = 1.64 [in SD units], min 1.5, max 2.0). Similarly, the MEG signal at each sensor was transformed into SD units and points that exceeded a threshold of 5 SD from the mean across the epoch were found. Centered about each point that exceeded these thresholds, the data 300 ms before and 500 ms after these points were excluded from further analysis (on average this removed approximately 1200 ms per eye blink). These additional windows of data were excluded so that sufficient pre- and post-artifact data were removed to ensure no artifacts remained in the analyzed time courses. These conservative thresholds and windows removed on average 36% (SD = 26%) of the time points across our eight subjects.

Removing these data points creates discontinuities in the waveforms because some contiguous points in the resultant data are not actually contiguous in time. We tested whether these discontinuities cause any undue negative consequences for the analysis by comparing the error caused by removing 35% of simulated 12 minute scans two different ways. Specifically, we either removed 1200 ms windows randomly spread across the 12 min (i.e. simulating eye blinks and creating discontinuities) or randomly selected a continuous 7.8-minute window out of the 12 min (i.e. 65% of the data) for analysis. Comparing the error caused by these two different methods relative to the entire 12-minute window allows us to examine whether the discontinuities in and of themselves were detrimental. We found that neither method for

Author's personal copy

removing 35% of the data caused any bias in the phase-locking (beyond the bias towards higher phase-locking due to the reduced degrees of freedom which was very small at all frequencies), though both methods did cause errors. Comparing the magnitude of the errors for the different methods of removing 35% of the data, we found that below 3 Hz the discontinuities caused significantly larger error compared to the continuous 7.8 min. However, at higher frequencies no substantial difference in the amount of error was seen. It should be noted that if smaller windows than the very conservative 1200 ms for each eye blink is removed, the cutoff frequency at which substantially greater error is present for the discontinuous data is higher than 3 Hz. This analysis demonstrates that, for frequencies above 3 Hz, the discontinuities caused by this method of eye blink removal does not cause substantial negative consequences to the data.

This method of removing the timepoints when the EOG channel exceeds a threshold is preferred to an ICA/PCA approach under most circumstances because of the potential for false-positive phase-locking due to residual eye blink signal with the ICA/PCA approaches. However, there are two cases where the ICA/PCA approach may be preferred. First, if the scientific question relates to oscillatory dynamics below 3 Hz, then the discontinuities will cause substantially excess error. Second, if a subject has excessive blinks resulting in removal of too much of the data and the subjects are too valuable to eliminate from the analysis (e.g. clinical populations).

Seed selection

Many methods exist for selecting the seed ROI. Seeds can be chosen based on functional localization, anatomical landmarks, or some combination of both. One critical factor is to ensure that ROIs are chosen independently of the analysis of interest. For example, as described in the example with real MEG data below, we chose our seed ROIs based on an independent functional localizer to find LH and RH auditory cortex. These locations were used as seed ROIs. The MEG signal is averaged across each location in the seed ROI for the seed time course.

Inverse solution

The location of the cortical current sources cannot be precisely determined using the measured magnetic fields from outside the head. Thus, we estimate the location of these sources with the cortically constrained minimum norm estimate (MNE; Hämäläinen et al., 1993) using the MNE™ software suite v2.7. Briefly, a linear inverse operator W is applied to the measured signal to calculate the MNE

$$y(t) = Wx(t)$$

where $x(t)$ is a vector that represents the MEG channel data at time t and $y(t)$ is vector representing the corresponding current projected onto the cortical surface. The estimated activity at each source location is a weighted combination of all the data arising from all sensors (as it is for most distributed source models). Thus, each source shares the same underlying data to some degree; substantially so for sources that are close together in the brain. This leads to "imperfections" due to the inverse solution where the activity in nearby sources appears to be more similar than it actually is in the brain. The artifact caused by these imperfections are discussed and addressed in detail later. The expression of W is calculated using the L_2 norm, which yields

$$W = RA^T\left(ARA^T + \lambda^2 C\right)^{-1}.$$

where A is the free source orientation solution of the forward problem calculated using the boundary element method. C and R are the noise and source covariance matrices respectively. λ is a weighting factor that is used to avoid the magnification of errors in the data and $\lambda^2 \approx (1/SNR)$. We used a value of 3 for this, as is often done in MEG analysis (Hämäläinen, MNE software user's guide version 2.7, 2009). R was depth-weighted to overcome the superficial bias of the MNE with a depth factor of .8 (Lin et al., NeuroImage, 2006). Furthermore, because cortical neurons are known to be preferentially oriented perpendicular to the cortical surface, we used a loose orientation constraint. Specifically, the component of R normal to the surface was multiplied by 1 and the components transverse to the surface was multiplied by .4. Typically in non-resting-state scans, MEG measurements that are taken while the subject is in the scanner (but not performing the task) are used to calculate the noise covariance matrix (Hämäläinen et al., 1993). However, this would reduce any spontaneous covariance in the data, which is precisely what we wish to examine in studies of resting-state functional connectivity (Lin et al., 2004). Thus, the noise covariance matrix was calculated from 12 min of continuous empty room MEG measurements collected immediately prior to putting the subject in the scanner.

Here we use the MNE inverse solution to project our MEG data onto the brain, however in principle any distributed source model for calculating the inverse solution could be used in this procedure. For example, beamformers, which are another type of inverse solution method, could be used (Sekihara et al., 2001). However, in practice, beamformers assume minimal covariance among sources for constraining the inverse solution. To implement this assumption they use the neural data itself to build the covariance matrix and the process removes spatial covariances in the data. This will remove precisely the coupling between sources that we are trying to find when examining functional connectivity.

The wavelet transform

The MEG data at each source location were spectrally decomposed using a continuous wavelet transform by temporally convolving the signal with a complex Morlet wavelet centered at each frequency of interest f and at each time in the scanning run t, after artifact removal (Lachaux et al., 1999). The Morlet wavelet is used because it has a Gaussian window shape in both time and frequency while maintaining a sinusoidal underlying structure. This wavelet structure yields easily interpretable results in the time and frequency domain because they yield qualitatively similar data as when a time-frequency analysis is done with a Fourier transform (though the wavelet is better suited for non-stationary data). The Morlet wavelet is described by the equation:

$$G(t,f) = \frac{1}{\sqrt{2\pi f}} e^{\left(\frac{-t^2}{2\sigma^2}\right)} e^{i2\pi f t}$$

where σ is the SD of the Gaussian in the time domain. To ensure stability of the wavelet transform σ must be at least $\frac{5}{2\pi f}$ and here we set σ to be $\frac{7}{2\pi f}$. Because the wavelet convolution introduces Gaussian temporal blurring with an SD of σ, the effective number of independent samples (degrees of freedom for statistical tests) of the transformed time course is $\frac{N-1}{\sqrt{2\pi(f_s\sigma)^2}}$ (Friston et al., 1995) where f_s is the sampling frequency of the data and N is the number of time points in the sample.

One question we examined was whether a different inverse solution (particularly a different noise covariance matrix) was required for different frequencies. We found that using a broadband signal (band passed to 1–50 Hz) or using the same frequency as we were examining in the PLV calculation made negligible difference (the correlation coefficient across source locations between these two different techniques was greater than .95 across many frequencies and subjects). Thus, for

Author's personal copy

simplicity, we used the same inverse solution, based on the broadband signal, for all frequency bands.

Phase-locking

Resting-state phase-locking measures the variability over time of the phase difference between the seed and the other cortical locations (Lachaux et al., 1999; Lin et al., 2004). Specifically, the phase-locking value (PLV) is defined as:

$$PLV_l = \frac{1}{N}\left|\sum_{n=1}^{N} e^{i(\theta_{seed}(n)-\theta_l(n))}\right|$$

where N is the number of time points in the sample and $\theta_{seed}(n)$ and $\theta_l(n)$ are the phase of the wavelet convolved data in the seed ROI and the cortical locations "*l*" respectively. The PLV varies between 0 (a random phase relationship) and 1 (a consistent phase difference at all time points).

Statistics

The statistical significance of the PLV in individuals can be calculated using the Rayleigh statistic (Fisher, 1993), where the degrees of freedom are the effective number of independent samples (as calculated above). When examining the significance across subjects for PLVs between a seed location and every other location in the brain, we first use a normalization appropriate for Rayleigh distributed data. Specifically, we took the square root of the PLVs, which we confirmed normalized Rayleigh distributed data using D' Agostino's K^2 test. A *t*-test was used to test for significance and the critical values were determined after Bonferroni correcting for multiple comparisons (specifically, 273, the number of functional sensors and the maximal number of independent data sources according to information theory). When examining the significance between ROIs a *t*-test was performed at each frequency of interest. The critical values were determined after using a cluster-level correction for multiple comparisons (described in detail in Maris and Oostenveld, 2007). Briefly, all frequency points that were $p<.05$ were found, clusters in adjacent frequencies were found, and the sum of the t-values in each cluster was determined (cluster "mass"). A distribution was created by permuting the data and the statistical significance of the cluster mass was determined. This method inherently controls for multiple comparisons across frequencies because it has a global null hypothesis.

Subjects

Ten subjects (6 males, mean age = 26.8, SD = 4.4) participated in the MEG experiment. One subject was excluded due to unusually large cardiac and respiratory artifacts and another subject was excluded due to head movement in excess of .5 cm, thus eight subjects are included in the analysis (5 males, mean age = 24.7, SD = 6.3). All subjects were naïve to the goals of the experiment. The Institutional Review Board of the National Institutes of Health approved all procedures and written informed consent was obtained for all subjects. Subjects were compensated for their participation.

Recording

Neuromagnetic responses were recorded at 600 Hz using a 275 channel whole head MEG system in a shielded room (VSM MedTech, Ltd., Canada). The MEG is equipped with 275 radial gradiometers and synthetic 3rd order gradient noise cancellation was used. Head position coils were placed at the nasion and left and right preauricular points to coregister the anatomical MRI and the MEG sensors. Head position was determined at the beginning and end of each run to ensure that head movements did not exceed .5 cm for any subject. Eyeblinks and eye movements were recorded using a bi-polar EOG electrode placed about each subject's left eye.

Structural MRI

Structural MRI images were obtained separately using a 3-Tesla whole-body scanner (GE Signa, USA). A high-resolution T1-weighted 3D volume was obtained for each subject. The MEG data were coregistered with anatomical MR images using fiduciary headpoints. Freesurfer™ was used to create a cortical surface model for each subject using an automatic reconstruction algorithm. The cortical white matter was segmented providing a topologically correct representation of the surface with approximately 150,000 vertices per hemisphere. The cortical surface was then decimated to approximately 4000 source dipoles per hemisphere, approximately 1 dipole every 10 mm along the cortical surface.

Behavioral tasks

Each subject first participated in a 12-minute rest-state scan where their task was to fixate on a centrally presented cross. Subjects were not given any details about later tasks during the rest-state scan. Having subjects fixate is important in MEG (as opposed to subjects having their eyes closed or eyes open unfixated), because this reduces eye movement artifacts. Following the rest-state scan, the subjects were presented with a series of button-press trials. Every 2 s, a white fixation cross in the middle of the screen would change to green or purple for 200 ms. The subjects were instructed to press a button with their left index finger when the cross turned green and a different button with their right index finger when the cross turned purple. The subjects were presented with 70 trials in each condition. Finally, using MEG compatible ear buds, the subjects were presented with a series of auditory clicks trials binaurally. Each trial consisted of 500 ms of a 1910 Hz pure tone amplitude modulated at 40 Hz followed by 1000 ms of silence. The subjects were presented with 100 trials and asked only to fixate on a centrally presented cross throughout the experiment.

MEG analysis for functional localizer

The noise-normalized, dynamic statistical parametric mapping procedure described in Dale et al. (2000) was used to estimate task-related neural activity for each individual. This technique yields f-distributed data that correspond to the significance of the activity at each time point relative to the interstimulus period.

To localize LH primary motor cortex (M1), the locus of activity about the central sulcus occurring between approximately 0 to 150 ms after subjects pressed the button with their right hand was calculated. The portion of this activity anterior to the central sulcus was used as the location of the LH M1 ROI for each subject.

To localize LH and RH auditory cortex, the activity between 50 and 550 ms after the onset of the auditory clicks was averaged. The portion of the activity corresponding approximately to Heschl's gyrus was used as the location of the LH and RH primary auditory cortex ROI for each subject.

Results and discussion

Spatial sensitivity and specificity simulations

One major concern when employing a linear inverse solution with electrophysiological data is that most inverse solutions act as a spatial filter where the data projected onto each point in the brain is a combination of the data derived from each sensor location (Hämäläinen et al., 1993). Therefore, some artifactual phase-locking will be introduced because each source location shares some data with each

Author's personal copy

other source location. An example of artifactual phase-locking introduced by the imperfect inverse solution can be seen in Fig. 1a. To generate this image, 8 min of empty room data were projected onto a subject's cortical surface using the MNE inverse operator. Phase-locking was then calculated for this empty room noise between a seed ROI chosen in the LH sylvian fissure (near Heschl's gyrus; the green dot in Fig. 1a) and every other source location on the cortical surface. As can be seen, the PLV analysis of these empty room noise data identified a cloud around the seed of significant phase-locking ($p<.05$, corrected for multiple comparisons), even though there was no neural signal between the seed and these other locations. To illustrate this point, we graphed the PLV values for each dipole on the surface in terms of their distance from the seed in the sylvian fissure (Fig. 2). As can be seen, even out to 7 cm from the seed, false-positive PLVs are seen in many dipoles ($p<.05$ falls within 1 standard deviation of the mean PLV). One point to notice is that, though there is a general falling off of the PLV by distance, the spatial distribution is extremely heterogeneous (the distribution is extremely non-uniform in Fig. 1a with many local minima within the false-positive cloud and the standard deviations are large in Fig. 2). This is because the amount of data shared between sources is dependent on both the distance between the sources and their relative geometries.

Can these false-positive PLVs be accounted for? Here we show that much of the crosstalk arising from the imperfections in the inverse solution can be removed if the PLVs calculated on empty room data are subtracted from the PLVs calculated on neural data. To demonstrate this we ran a series of simulations where the same data were projected onto a seed source in the LH sylvian fissure and onto a series of second target sources pseudo-randomly chosen in the left hemisphere. The form of this simulated signal was

$$q = a \cdot \sin(f \cdot 2\pi \cdot t)$$

where t is time. The amplitude of the sine wave was chosen such that the PLV between the seed and a RH auditory cortex ROI was approximately .14 (similar to that seen in our MEG data below), this corresponded to $a = 2.6 \cdot 10^{-8}$ A m. The frequency of the sine wave was chosen as 10 Hz, also corresponding to the peak frequency seen in the MEG data below. Using the forward solution derived from the subject's structural information and the boundary element model, these data were then projected back onto the sensors. Noise was then added which had the same covariance structure and amplitude as measured empty room data. These data were then projected back onto the cortex using the MNE inverse solution and the phase-locking between the seed and target sources from which the data arose was calculated. An example of this simulation is seen in Fig. 1b where the sine wave signal is present in both the sylvian fissure and a frontal pole location. We then compared the phase-locking values seen for

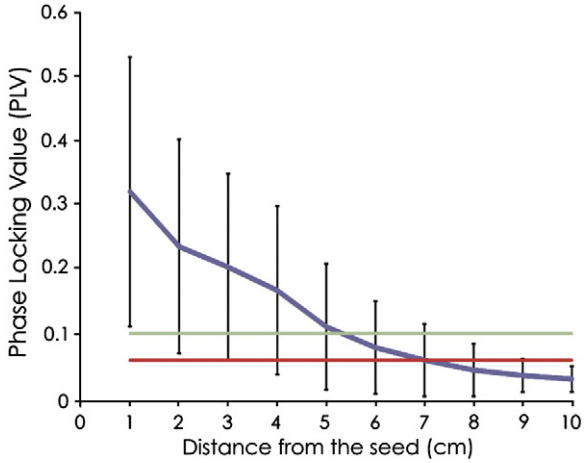

Fig. 2. A systematic analysis of the artifact that results from the inverse solution. Phase-locking values in empty room data as a factor of the Euclidean distance (rather than the distance on the cortical surface) from a seed in the LH sylvian fissure. The means and standard deviations of the empty room phase-locking data shown in this figure are plotted for all dipoles in an annulus 0–1 cm from the seed, 1–2 cm from the seed, etc. The red line is the PLV that corresponds to a p-value of .05 and the green line corresponds to a p-value of .05 corrected for 273 comparisons (the number of functioning sensors). As can be seen, there is a general trend for lower artifactual phase-locking with increasing distance from the seed. Also note however that the standard deviations across the dipoles at each distance are quite large demonstrating that this artifact is not completely spatially uniform about the seed.

these simulated data with the values seen between the same sources for empty room data without any simulated data. Fig. 1c shows an example of the difference between the simulated and empty room PLVs.

The difference between the simulated and empty room PLVs was closely related to the PLV value seen in empty room data. Specifically, where the phase-locking of empty room noise is the highest, the difference between the PLVs for the simulated signal and the empty room data is near zero. However, when the empty room PLVs dropped below approximately .15 (which occurs when the sources are separated by approximately 4–6 cm; Fig. 2), the phase-locking was significantly larger for the simulated data than for the empty room data. Thus, after normalizing by the empty room data, there is a cloud of sources close to the seed location where the PLV is suppressed (i.e. false negatives were seen in this cloud), after which the measure becomes sensitive to PLV at approximately 4–6 cm. In these simulations the false-positive rate fell to chance everywhere in the brain except for adjacent to the target location. Note that, in task-based studies of functional connectivity, comparing across conditions

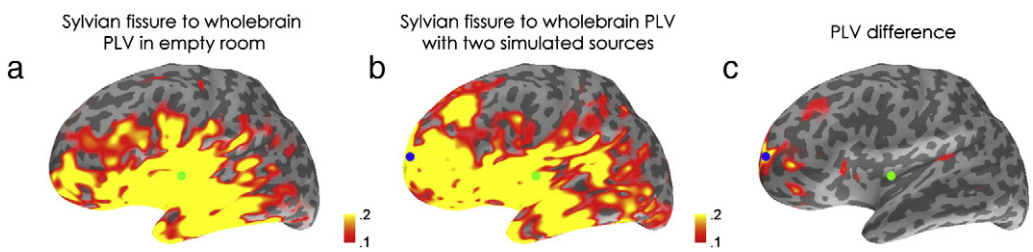

Fig. 1. Examples of the artifact caused by the imperfect inverse solution and simulation of the correction for this artifact. a) Phase-locking for empty room data projected onto the brain. A seed was placed in the LH sylvian fissure (green dot) and PLVs were calculated between this seed and every other location in the brain. A large cloud of significant PLVs is seen even though there is no phase-locked signal emanating from the brain. This false-positive phase-locking occurs because, in most distributed source inverse solutions, the data in each dipole on the brain is a linear combination of the data from all of the sensors (Hämäläinen et al., 1993). Therefore, a large cloud of dipoles shares a significant portion of their data resulting in false-positive phase-locking. A threshold of .1 was used because this corresponds to a p-value of approximately .05 corrected for multiple comparisons. b) Simulated resting-state phase-locking. A sine wave was placed in both the LH sylvian fissure (green dot) and the LH frontal pole (blue dot) and PLVs were calculated between the LH sylvian fissure seed and every other location in the brain. Due to the false-positive phase-locking, it is difficult to differentiate the true phase-locking from the artifact. c) Subtracting panel a from panel b demonstrates that when the false-positive phase-locking is accounted for, the true phase-locking can be detected.

Author's personal copy

similarly reduces much of the false-positive PLVs induced by the imperfect inverse solution.

When examining functional connectivity using fMRI, generally the correlation values are compared to the correlation expected for random data. For electrophysiological data, however, significant crosstalk between sources is introduced by the inverse solution, as shown above. Thus, for these data, if the PLVs expected from random noise were used as the null hypothesis the false-positive rate would be high. However, by using the PLVs derived from empty room noise as the null hypothesis, we trade this large cloud of false-positives (7–8 cm) for a smaller false negative cloud (4–6 cm).

One potential concern is that the empty room data may not have the same properties as the resting-state data. This is critical because if the phase distribution of the empty room data were different than the resting-state data the empty room data may not accurately reduce the crosstalk inherent in the inverse solution. The critical question is whether the normalizing data and the resting-state data have the same phase distribution, because the phase-locking method removes the amplitude information. To address this concern we compared the PLV for empty room data to phase shuffling the resting-state data. We phase shuffled recorded resting-state data for 4 of our subjects and calculated PLVs for 3 different frequencies (5, 10 and 20 Hz) and 2 different seed regions (a total of 24 comparisons). This phase shuffling was done such that the phase distribution of the data was not changed after the shuffle. The average correlation between the phase shuffled data and the empty room data was .90 (standard deviation of .05). As a comparison, the average correlation between different phase shuffles within the same subjects was .92 (standard deviation of .04). This demonstrates that the empty room noise has very similar properties as the resting-state data. This is not surprising because over a 12 minute scan, both the resting-state data and the empty room noise have a nearly uniform phase distribution. Indeed, any random data with a uniform phase distribution could be used for this normalization. Empty room data is a convenient choice because it needs to be recorded for the noise covariance matrix and contains many of the sources of noise present in MEG data. The fact that any random data with a uniform phase distribution could be used for normalization provides a potential way of using this method for examining resting-state functional connectivity in EEG data where there is no direct equivalent of empty room data. More problematic for EEG is that empty room data is not available for the noise covariance matrix in the inverse solution, therefore the forward solution alone may have to be used in EEG for producing the source estimate.

Resting-state functional connectivity in the auditory network

To demonstrate the efficacy of this resting-state PLV method on real data, we first examined phase-locking between LH and RH primary auditory cortex and for a control comparison between LH motor cortex and RH auditory cortex (Fig. 3a). We found larger PLVs between LH and RH auditory cortex than between LH motor cortex and RH auditory cortex in the alpha (~7–15 Hz) frequency range ($p = .006$ corrected for multiple frequency comparisons). In contrast to the PLV results, greater broadband power was seen in the LH motor cortex than in RH and LH auditory cortex in all eight subjects. This excludes the possibility that the greater PLV between RH and LH auditory cortex was due to greater signal-to-noise ratio in the auditory cortex than the motor cortex (see Ghuman et al., submitted for publication for a full discussion of the effect of signal-to-noise ratio on correlation-based measures of functional connectivity). The PLVs between LH and RH auditory cortex were also greater for the resting-state data than for the empty room data in the same frequency range ($p = .034$). Interestingly, alpha frequency band activity in the primary auditory cortex has previously been shown to be modulated following auditory stimulation (Basar et al., 2001; Lehtela et al., 1997; Schurmann et al., 1997; Tiihonen et al., 1991; van Dijk et al., 2010). Additionally, patients with tinnitus have been shown to have abnormal alpha phase-locking at rest (Schlee et al., 2009). Thus,

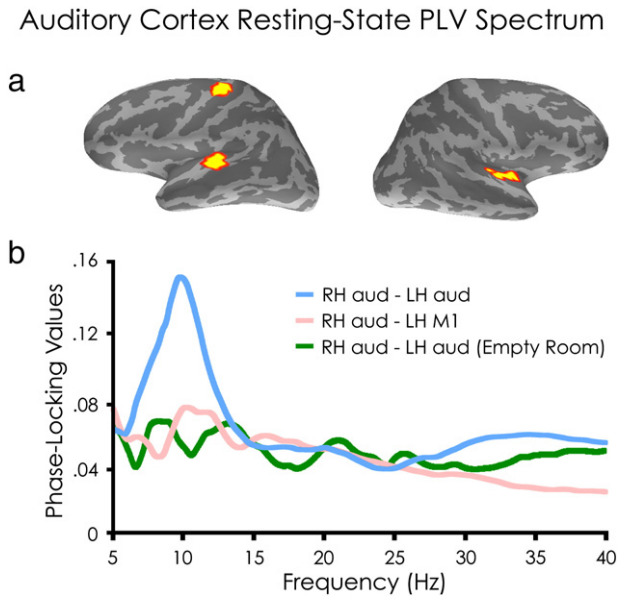

Fig. 3. a) Map of the extent of the LH auditory cortex, RH auditory cortex, and LH motor cortex ROIs across the 8 subjects. b) Average phase-locking values with respect to frequency. The average PLV between RH auditory cortex–LH auditory cortex during the resting-state, RH auditory cortex–LH motor cortex during the resting-state, and RH auditory cortex–LH auditory cortex for empty room data are plotted.

previous studies have demonstrated that activity in the alpha frequency bands is important to processing auditory information and here we show that the same frequency band is also important to the intrinsic functional connectivity in the auditory network.

We then examined the resting-state phase-locking between the LH/RH primary auditory cortex and every other region in the brain. We found that at 10 Hz (the peak of intrinsic phase-locking in the auditory network), significant PLVs were seen between primary auditory cortex and the contralateral auditory network. Specifically, when the seed was in the LH primary auditory cortex, significant PLVs were seen in contralateral primary auditory cortex (supporting the result in Fig. 3b) as well as in the adjacent superior and middle temporal gyri, the ventral part of the RH somatomotor strip, and RH inferior frontal cortex (Fig. 4 left). When the seed was placed in RH primary auditory cortex the PLVs were weaker and sparser for the RH auditory seed than the LH auditory seed. However, significant PLVs were still seen in contralateral primary auditory cortex, as well as the adjacent superior and middle temporal gyri, and the ventral part of the LH somatomotor strip (Fig. 4 right). We also examined the whole brain functional connectivity to the LH/RH auditory cortex at 24 Hz, a frequency that is not expected to show significant phase-locking (Fig. 3b). Indeed, no significant resting-state phase-locking was between RH or LH auditory cortex and any other region of the brain (Fig. 5).

One thing to note is that these data also support the simulations shown in Fig. 1 and 2. Specifically, if the raw PLVs are examined without comparing to the empty room data, it is nearly impossible to draw any conclusions due to the large cloud of false-positive PLVs (Figs. 4 and 5). It is only when the statistical values of the resting-state PLVs vs. the empty room PLVs are shown that the true resting-state functional connectivity is seen.

One difference between the real data and the simulation is that some residual significant phase-locking was seen near the seed for the real data. This effect is clearly seen in Fig. 5 where the only significant patches of phase-locking are around the seed locations. This result is surprising given that the simulations suggested that, if anything, there should be false negative PLVs near the seed. Thus, this phase-locking is due to some factor that was not taken into account in the simulations.

Author's personal copy

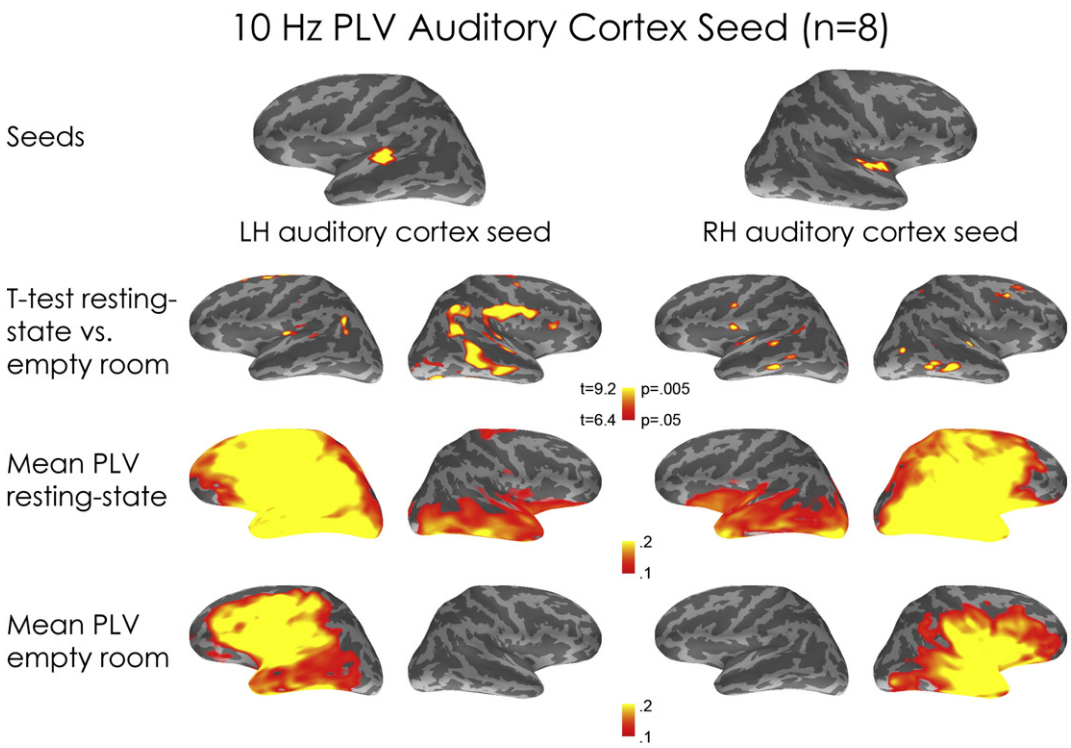

Fig. 4. Resting-state phase-locking between LH and RH auditory cortex and the rest of the brain at 10 Hz. A *t*-test across 8 individuals of resting-state vs. empty room PLV reveals that the LH auditory cortex shows significant (corrected for multiple comparisons) functional connectivity to RH primary auditory cortex as well as the adjacent superior and middle temporal gyri, RH inferior frontal cortex, and the ventral part of the somatomotor strip. Functional connectivity starting from a seed in the RH auditory cortex is weaker and sparser (though still significant), compared to the LH auditory cortex seed. Specifically, RH primary auditory cortex demonstrated significant functional connectivity to LH auditory cortex, the adjacent superior and middle temporal gyri, and a region that straddles the ventral part of the somatomotor strip. Much as in the simulated and empty room data (Figs. 1 and 2), the mean resting-state PLV across 8 subjects is difficult to interpret because of the large artifact due to the imperfect inverse solution. Much of this cloud remains for the mean empty room PLVs across the 8 subjects (empty room data was collected prior to each subject's session). The auditory network connectivity is only revealed for the statistical test of the resting-state vs. empty room phase-locking.

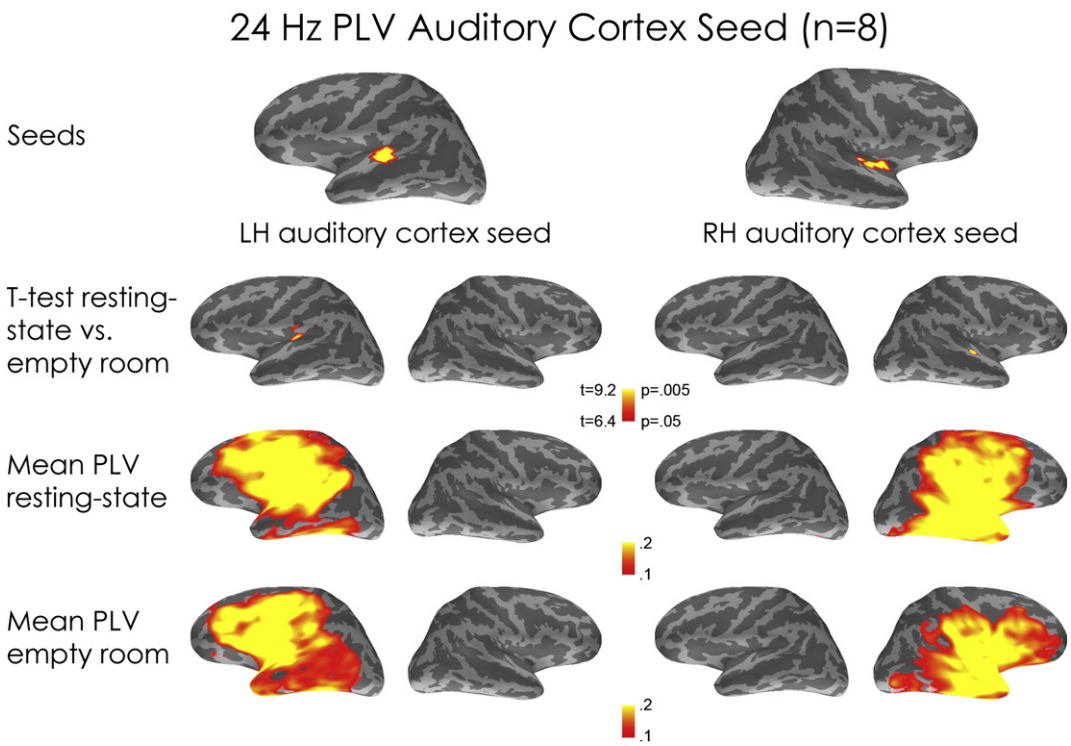

Fig. 5. Resting-state phase-locking between LH and RH auditory cortex and the rest of the brain at 24 Hz. No regions showed significant resting-state phase-locking vs. empty room phase-locking except for small regions close to the seeds. This is because the mean PLV at 24 Hz between the RH and LH auditory cortex was similar for the resting-state and for empty room data (Fig. 3b). Thus, almost the entire resting-state PLV cloud at 24 Hz is due to the artifact caused by the crosstalk in the inverse solution.

Author's personal copy

One possible source of discrepancy is inaccuracies in the forward model that cause the simulation to deviate slightly from the real data. The fact that the patch around the seed location appears in both Figs. 4 and 5, but not in the simulation, makes an inaccuracy in the forward solution a likely possibility, rather than an artifact in the signal that would likely be frequency specific. Regardless, adhering to a rule of thumb of excluding PLVs within 4–6 cm of the seed will eliminate this concern.

More generally, it is critical to note that normalization by the empty room noise does not eliminate all possible sources of crosstalk. There may be other artifacts that cause crosstalk, such as inaccuracies in the forward model, unaccounted for physiological noise such as muscle artifacts or movement artifacts not removed using the noise reduction procedures described above, and other artifacts that arise from the neural data having a different noise covariance structure than the empty room data. These sources of crosstalk were not present in the simulated data (Fig. 2). Thus, although crosstalk is greatly reduced by using the empty room normalization, many sources of crosstalk still must be considered. Because most sources of crosstalk are located spatially close to the seed and target locations, it will be largely removed by excluding all PLVs between sites 4–6 cm or less apart. Nevertheless, it is certainly possible for artifactual crosstalk to occur even beyond 6 cm.

What about amplitude?

So far we have concentrated on examining phase relationships as the measure of functional connectivity. A potentially complementary method for measuring functional connectivity is amplitude correlation. Indeed, coherence, which is a measure of the linear correlation between all properties of waveforms, is a combination of both amplitude correlation and phase-locking. We calculated amplitude correlation by taking the absolute value of the wavelet-transformed data and performing a Pearson's correlation over time between pairs of signals. Significantly greater amplitude correlation was seen between LH and RH auditory cortex than between LH motor cortex and RH auditory cortex in the alpha frequency range (Fig. 6a; $p = .04$ corrected for multiple frequency comparisons), similar to the PLV results (Fig. 3). However, when examining the amplitude correlation at 10 Hz between the LH/RH auditory cortex seeds and every other dipole on the cortex, the auditory cortex dipole contralateral to the seed failed to survive a correction for multiple comparisons based on the number of MEG sensors (Fig. 6b). In fact, most of the dipoles that survived multiple comparisons were close to the seed and, due to their proximity to the seed, are potentially artifactual. Thus, the amplitude correlation results differed from the PLV results in two ways. First, the amplitude correlation in cross-hemisphere auditory network was relatively weak compared to phase-locking. This may be because of the well-described property of weakly coupled oscillatory systems: phase-locking tends to be more sensitive to their coupling than amplitude correlation (Pikovsky and Rosenblum, 2007). It is somewhat surprising that no amplitude correlation was seen as they have been described in the resting-state for data collected using intracranial EEG (He et al., 2008; Nir et al., 2008). However, intracranial EEG has a much higher signal-to-noise ratio than non-invasive measures of neural activity and thus may overcome the relatively poor sensitivity of amplitude correlation. Second, even after normalizing by the empty room data, much greater artifactual crosstalk was seen close to the seed. It is not clear why this occurred, however one possibility is that artifacts other than the crosstalk inherent in the inverse model are greater for amplitude correlation than for phase-locking. Thus, because amplitude correlation is less sensitive to "true" coupling and because greater potentially artifactual amplitude correlation close to the seed is seen, phase-locking may be better suited for measuring functional connectivity in MEG.

Conclusion

We have described a method for detecting and describing the oscillatory dynamics of functional connectivity in MEG resting-state data projected onto the brain. We examined the spatial sensitivity and specificity of this phase-locking technique using simulated data and showed how to use empty room data to account for much of the crosstalk that arises due to the imperfect inverse solution. We then applied this technique to show that the auditory network displays resting-state functional connectivity in the alpha frequency bands, even in a silent environment. This finding further demonstrates the importance of accounting for the crosstalk induced by the inverse solution. These simulations and data suggest that our method may be useful for exploring functional connectivity in electrophysiological data.

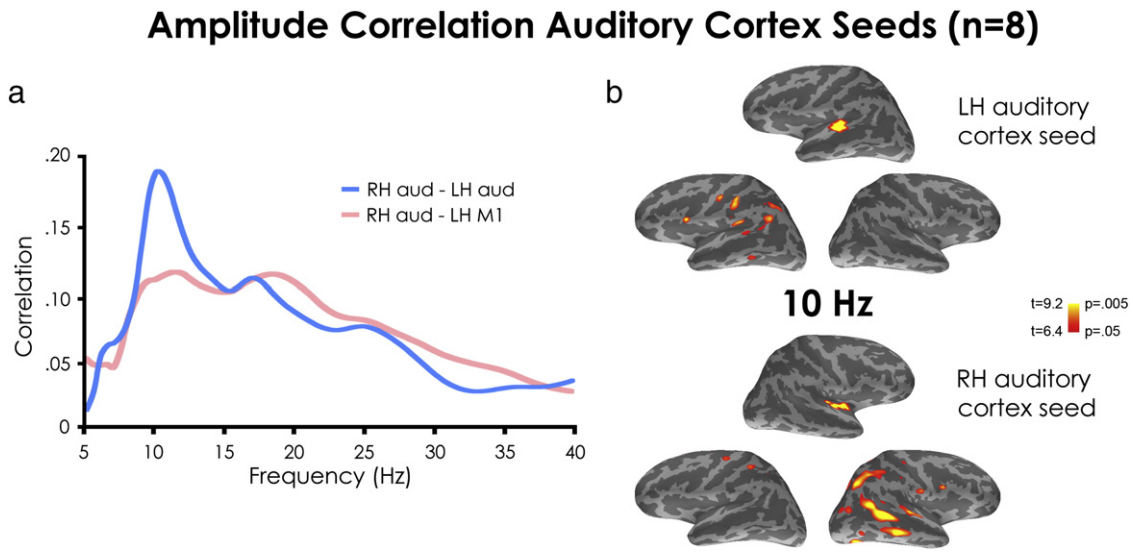

Fig. 6. Resting-state amplitude correlation in the auditory network. a) The average Pearson's correlation between RH and LH auditory cortex and between RH auditory cortex and LH motor cortex as a factor of frequency. Significantly greater amplitude correlation was seen in the alpha frequency range between RH auditory cortex and LH auditory cortex than between RH auditory cortex and LH motor cortex (similar to the phase-locking between these regions; Fig. 3b). b) Resting-state amplitude correlation between LH and RH auditory cortex and the rest of the brain at 10 Hz. No regions of the auditory network contralateral to the seed region showed significant resting-state vs. empty room amplitude correlation (after correcting for multiple comparisons). The only regions of significant amplitude correlation were ipselateral and relatively close to the seeds. Due to the relative proximity of these regions to the seeds, it is difficult to rule out the possibility that this is artifactual amplitude correlation resulting from inaccuracies in the forward model or other crosstalk artifacts.

Author's personal copy

Acknowledgments

We thank Rebecca van den Honert for assistance with data collection and analysis, Steve Gotts, Richard Coppola, and for insightful comments, Tom Holroyd and the staff NIMH MEG core for assistance with data collection, Gang Chen for assistance with statistics, Zhongming Liu for assistance with the cardiac artifact removal procedure, and Ziad Saad for assistance with MRI processing. This work was supported by NIMH-DIRP.

References

Bar, M., Kassam, K.S., Ghuman, A.S., Boshyan, J., Schmid, A.M., Dale, A.M., Hämäläinen, M.S., Marinkovic, K., Schacter, D.L., Rosen, B.R., Halgren, E., 2006. Top-down facilitation of visual recognition. Proc. Natl Acad. Sci. USA 103, 449–454.

Basar, E., Basar-Eroglu, C., Karakas, S., Schurmann, M., 2001. Gamma, alpha, delta, and theta oscillations govern cognitive processes. Int. J. Psychophysiol. 39, 241–248.

Biswal, B., Yetkin, F.Z., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magn. Reson. Med. 34, 537–541.

Boly, M., Phillips, C., Tshibanda, L., Vanhaudenhuyse, A., Schabus, M., Dang-Vu, T.T., Moonen, G., Hustinx, R., Maquet, P., Laureys, S., 2008. Intrinsic brain activity in altered states of consciousness: how conscious is the default mode of brain function? Ann. NY Acad. Sci. 1129, 119–129.

Cordes, D., Haughton, V.M., Arfanakis, K., Wendt, G.J., Turski, P.A., Moritz, C.H., Quigley, M.A., Meyerand, M.E., 2000. Mapping functionally related regions of brain with functional connectivity MR imaging. AJNR Am. J. Neuroradiol. 21, 1636–1644.

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

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

Duzel, E., Neufang, M., Heinze, H.J., 2005. The oscillatory dynamics of recognition memory and its relationship to event-related responses. Cereb. Cortex 15, 1992–2002.

Engel, A.K., Fries, P., Singer, W., 2001. Dynamic predictions: oscillations and synchrony in top-down processing. Nat. Rev. Neurosci. 2, 704–716.

Fisher, N.I., 1993. Statistical analysis of circular data. Press Syndicate of the University of Cambridge, New York, NY.

Fox, M.D., Raichle, M.E., 2007. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. Nat. Rev. Neurosci. 8, 700–711.

Friston, K.J., Frith, C.D., Liddle, P.F., Frackowiak, R.S., 1993. Functional connectivity: the principal-component analysis of large (PET) data sets. J. Cereb. Blood Flow Metab. 13, 5–14.

Friston, K.J., Holmes, A.P., Poline, J.B., Grasby, P.J., Williams, S.C., Frackowiak, R.S., Turner, R., 1995. Analysis of fMRI time-series revisited. Neuroimage 2, 45–53.

Ghuman, A.S., Bar, M., Dobbins, I.G., Schnyer, D.M., 2008. The effects of priming on frontal–temporal communication. Proc. Natl Acad. Sci. USA 105, 8405–8409.

Ghuman, A.S., Gotts, S.J., Martin, A., submitted for publication. When differences in correlation do not mandate differences in functional connectivity.

Greicius, M., 2008. Resting-state functional connectivity in neuropsychiatric disorders. Curr. Opin. Neurol. 21, 424–430.

Gross, J., Schmitz, F., Schnitzler, I., Kessler, K., Shapiro, K., Hommel, B., Schnitzler, A., 2004. Modulation of long-range neural synchrony reflects temporal limitations of visual attention in humans. Proc. Natl Acad. Sci. USA 101, 13050–13055.

Hämäläinen, M., Hari, R., Ilmoniemi, R.J., Knuutila, J., Lounasmaa, O.V., 1993. Magnetoencephalography—theory, instrumentation, and applications to noninvasive studies of the human brain. Rev. Mod. Phys. 65, 413–497.

He, B.J., Snyder, A.Z., Zempel, J.M., Smyth, M.D., Raichle, M.E., 2008. Electrophysiological correlates of the brain's intrinsic large-scale functional architecture. Proc. Natl Acad. Sci. USA 105, 16039–16044.

Jerbi, K., Lachaux, J.P., N'Diaye, K., Pantazis, D., Leahy, R.M., Garnero, L., Baillet, S., 2007. Coherent neural representation of hand speed in humans revealed by MEG imaging. Proc. Natl Acad. Sci. USA 104, 7676–7681.

Kaiser, J., Buhler, M., Lutzenberger, W., 2004. Magnetoencephalographic gamma-band responses to illusory triangles in humans. Neuroimage 23, 551–560.

Lachaux, J.P., Rodriguez, E., Martinerie, J., Varela, F.J., 1999. Measuring phase synchrony in brain signals. Hum. Brain Mapp. 8, 194–208.

Lehtela, L., Salmelin, R., Hari, R., 1997. Evidence for reactive magnetic 10-Hz rhythm in the human auditory cortex. Neurosci. Lett. 222, 111–114.

Leopold, D.A., Murayama, Y., Logothetis, N.K., 2003. Very slow activity fluctuations in monkey visual cortex: implications for functional brain imaging. Cereb. Cortex 13, 422–433.

Lin, F.H., Witzel, T., Hamalainen, M.S., Dale, A.M., Belliveau, J.W., Stufflebeam, S.M., 2004. Spectral spatiotemporal imaging of cortical oscillations and interactions in the human brain. Neuroimage 23, 582–595.

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

Liu, Z., Fukunaga, M., de Zwart, J.A., Duyn, J.H., 2010. Large-scale spontaneous fluctuations and correlations in brain electrical activity observed with magnetoencephalography. Neuroimage 51, 102–111.

Mantini, D., Perrucci, M.G., Del Gratta, C., Romani, G.L., Corbetta, M., 2007. Electrophysiological signatures of resting state networks in the human brain. Proc. Natl Acad. Sci. USA 104, 13170–13175.

Maris, E., Oostenveld, R., 2007. Nonparametric statistical testing of EEG- and MEG-data. J. Neurosci. Meth. 164, 177–190.

Nir, Y., Mukamel, R., Dinstein, I., Privman, E., Harel, M., Fisch, L., Gelbard-Sagiv, H., Kipervasser, S., Andelman, F., Neufeld, M.Y., Kramer, U., Arieli, A., Fried, I., Malach, R., 2008. Interhemispheric correlations of slow spontaneous neuronal fluctuations revealed in human sensory cortex. Nat. Neurosci. 11, 1100–1108.

Percival, D.B., Walden, A.T., 2000. Wavelet methods for time series analysis. Cambridge University Press, New York, NY.

Pikovsky, A., Rosenblum, M., 2007. Synchronization. Scholarpedia 2, 1459.

Roelfsema, P.R., Engel, A.K., Konig, P., Singer, W., 1997. Visuomotor integration is associated with zero time-lag synchronization among cortical areas. Nature 385, 157–161.

Schlee, W., Hartmann, T., Langguth, B., Weisz, N., 2009. Abnormal resting-state cortical coupling in chronic tinnitus. BMC Neurosci. 10, 11.

Schurmann, M., Basar-Eroglu, C., Basar, E., 1997. A possible role of evoked alpha in primary sensory processing: common properties of cat intracranial recordings and human EEG and MEG. Int. J. Psychophysiol. 26, 149–170.

Sekihara, K., Nagarajan, S.S., Poeppel, D., Marantz, A., Miyashita, Y., 2001. Reconstructing spatio-temporal activities of neural sources using an MEG vector beamformer technique. IEEE Trans. Biomed. Eng. 48, 760–771.

Tass, P.A., Fieseler, T., Dammers, J., Dolan, K., Morosan, P., Majtanik, M., Boers, F., Muren, A., Zilles, K., Fink, G.R., 2003. Synchronization tomography: a method for three-dimensional localization of phase synchronized neuronal populations in the human brain using magnetoencephalography. Phys. Rev. Lett. 90, 088101.

Tiihonen, J., Hari, R., Kajola, M., Karhu, J., Ahlfors, S., Tissari, S., 1991. Magnetoencephalographic 10-Hz rhythm from the human auditory cortex. Neurosci. Lett. 129, 303–305.

van Dijk, H., Nieuwenhuis, I.L.C., Jensen, O., 2010. Left temporal alpha band activity increases during working memory retention of pitches. Eur. J. Neurosci. 31, 1701–1707.

Varela, F., Lachaux, J.P., Rodriguez, E., Martinerie, J., 2001. The brainweb: phase synchronization and large-scale integration. Nat. Rev. Neurosci. 2, 229–239.

Vincent, J.L., Snyder, A.Z., Fox, M.D., Shannon, B.J., Andrews, J.R., Raichle, M.E., Buckner, R.L., 2006. Coherent spontaneous activity identifies a hippocampal–parietal memory network. J. Neurophysiol. 96, 3517–3531.