Spontaneous and task-related changes in resting state connectivity

Kaitlyn Casimo

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

University of Washington

June 1, 2018

Reading committee:
Jeffrey G Ojemann, Chair
Kurt E Weaver
Rajesh PN Rao

Program authorized to offer degree:
Neuroscience

University of Washington

Abstract

Spontaneous and task-related changes in resting state connectivity

Kaitlyn Casimo

Chair of supervisory committee:
Jeffrey G Ojemann
Neurological Surgery, Graduate Program in Neuroscience

Resting state brain connectivity is thought to represent ongoing cognitive processes, including memory consolidation, that occur outside of the context of a specific task. While the existence of synchronous brain activity at rest is well documented, its specific properties have not been thoroughly investigated with electrocorticography, which is notable for its high resolution in fast frequencies thought to represent local neuronal firing rates, the fixed placement of electrodes that do not shift over the course of days, and a variety of other properties that support unique insights into brain function relative to other electrophysiological methods in humans. Our understanding of resting state connectivity depends on thoroughly characterizing its properties, such as over the course of minutes, between multiple sessions, and after a demanding behavioral task. The goal of my work is to use electrocorticography to evaluate various features of resting state connectivity, including multiple connectivity methods to capture different features of the signals over a wide range of frequency bands.

I identified regionally specific patterns of phase synchrony over the course in a single session of the resting state, including lower frequencies dominating the patterns in frontal regions and higher frequencies dominating in parietal regions, and a frontal-to-parietal directional flow of synchrony. I examined how patterns of connectivity vary between sessions of the resting state, and found that strong connectivity in phase locking and amplitude correlation predicts high stability across the sessions, but high stability is not exclusively associated with strong connectivity. I also found that within-region connectivity was more stable than connectivity between different regions, and individual regions expressed different relationships between connectivity and stability. Finally, I found that changes in connectivity patterns of inferior parietal lobule, and to a lesser extent primary motor cortex, following the execution of a novel task especially differentiated individuals who successfully performed the task from those who did not. Surprisingly, I did not find any pre-task connectivity patterns that were predictive of

task performance, in any frequency band or brain region, with any of the connectivity measures I tested.

My findings contribute to a deeper understanding of the specific properties of resting state connectivity, the way it spontaneously varies over time, and the changes that arise specifically from learning a new task. The methods I used to evaluate connectivity and change over time can be applied with ease to a wide range of other conditions or to data from additional brain regions in order to develop a wide and robust profile of resting state connectivity features. Understanding the varying properties of the resting state over space and time will enable us to better understand the functions of the resting state specifically, and the brain independent of task execution more broadly.

TABLE OF CONTENTS

1. INTRODUCTION AND SCOPE 8

1.1 Functional connectivity – measurement, computation, and roles 9

1.2 Research approach 13

1.3 Electrocorticography and resting state connectivity 15

1.4 Spontaneous change in resting state connectivity in electrocorticography 16

1.5 Task-linked changes in resting state connectivity in electrocorticography 17

1.6 Overview of dissertation structure 18

2. MEASURING BRAIN CONNECTIVITY AT A SINGLE TIME POINT 19

2.1 Motivation and role in broader project 19

2.2 Introduction to section 20

2.3 Methods 22

2.4 Results 27

2.5 Discussion 33

3. ESTIMATING SPONTANEOUS CHANGE IN BRAIN CONNECTIVITY OVER TIME 39

3.1 Motivation and role in broader project 39

3.2 Introduction to section 40

3.3 Methods 42

3.4 Results 46

3.5 Discussion 54

4. ESTIMATING CHANGE IN BRAIN CONNECTIVITY ATTRIBUTABLE TO LEARNING 60

4.1 Motivation and role in broader project 60

4.2 Introduction to section 61

4.3 Methods 63

4.4 Results 69

4.5 Discussion 78

5. GENERAL CONCLUSIONS **83**

5.1 Review of findings **83**

5.2 Conclusions **85**

5.3 Limitations **86**

5.4 Future work **88**

6 AND 7. SIDE PROJECTS **90**

6. CHANGES IN CONNECTIVITY ATTRIBUTABLE TO LEARNING OF A BRAIN-COMPUTER INTERFACE DEVICE **91**

6.1 Introduction to section **92**

6.2 Review of literature **94**

6.3 Current research **100**

6.4 Conclusions **104**

7. COMPARING BRAIN CONNECTIVITY IN HUMANS, MONKEYS, AND SHEEP **106**

7.1 Introduction to section **107**

7.2 Methods **110**

7.3 Results **115**

7.4 Discussion **120**

8. ABBREVIATIONS **125**

9. ACKNOWLEDGEMENTS **126**

10. REFERENCES **127**

TABLE OF FIGURES AND TABLES

Figure 2.1 24
Table 2.1 26
Table 2.2 27
Figure 2.2 28
Figure 2.3 29
Figure 2.4 31
Figure 2.5 33
Figure 3.1 43
Figure 3.2 44
Table 3.1 45
Figure 3.3 47
Figure 3.4 49
Figure 3.5 52
Supplementary Figure 2.1 58
Supplementary Figure 3.2 59
Figure 4.1 63
Figure 4.2 64
Figure 4.3 66
Table 4.1 66
Table 4.2 68
Table 4.3 69
Figure 4.4 71
Figure 4.5 73
Figure 4.6 75
Figure 4.7 77
Figure 6.1 99
Figure 6.2 102
Figure 6.3 103
Table 7.1 112
Figure 7.1 112
Figure 7.2 115
Figure 7.3 117
Figure 7.4 120

1. Introduction and scope

All brain regions exhibit some degree of activity in the absence of a specific task – that is to say, in a behavioral resting state. Not only is unprovoked activity present at rest, it is coordinated between regions, and encompasses the whole resting brain. Co-activation and or temporal synchronization between brain regions reflects the fact that the brain at “rest” is not really resting – it is engaged in a variety of cognitive functions such as memory consolidation (Buckner et al., 2008). The structured, spontaneous activity connecting brain regions at rest is usually studied context of large-scale, long-distance networks, called resting state networks (RSNs) (Buckner et al., 2008; Marcus E Raichle, 2011; Marcus E Raichle & Snyder, 2007). However, resting state connectivity outside of the traditional RSNs is still significant and ordered, and its meaningful origin is validated by the fact that such patterns follow known cytoarchitectural boundaries between brain regions (Gordon et al., 2016).

The ubiquitous presence of resting state activity is highly metabolically demanding. This raises the question of its cognitive function, as the brain does not expend its limited energy resources on activity without cause. Resting state brain activity has been linked to fundamental cognitive functions such as general attention to the internal thoughts (rather than to external stimuli), general readiness to respond to environmental stimuli (rather than attention to a particular stimulus), memory consolidation, idle or self-referential thoughts, and more (Buckner et al., 2008; Fransson, 2005; Hahn et al., 2007; K. Smith, 2012). Studying patterns of resting state brain activity, changes to those patterns over time, and their linkage to specific tasks can expand our understanding of spontaneous activity and its role in cognitive function. In fact, it is connectivity patterns, rather than activity patterns or behavior, that is the most relevant object of study in the resting state (Buckner et al., 2008).

The well-defined RSNs and their postulated cognitive roles are the focus of much research. However, such correlated activity in the absence of a task actually encompasses the entire brain to varying degrees, and in fact is sufficiently structured that cytoarchitectural boundaries between brain regions can be localized on the basis of resting state connectivity patterns (Gordon et al., 2016). This finding suggests that neuronal structure and function across the brain directly contributes to the properties of resting state activity and connectivity (Gordon et al., 2016). The same cytoarchitectural patterns also contribute to the definition of Brodmann areas, which I use to outline regions of interest. Studying patterns of resting state connectivity across the whole cortex, therefore, can illuminate the neural basis of cognitive functions generally associated with the resting state (Buzsáki & Draguhn, 2004). Mapping the brain in the

absence of behavior, further, can illuminate a wide array of physiological properties of the brain that may be obscured by specific task-linked activity.

One of these postulated features of the resting state is in memory consolidation, which occurs over the course of minutes to hours after the remembered event (Buckner et al., 2008; Miall & Robertson, 2006). The resting state may represent a period free of behavioral demands during which learned information distributed widely across the brain can be consolidated and stored (Albert, Robertson, Mehta, et al., 2009; Miall & Robertson, 2006). Indeed, the role of the resting state in memory consolidation is supported by changes to resting state connectivity following learning a task (Danielle Smith Bassett et al., 2011; Darainy et al., 2014; Kelly & Garavan, 2005; F. T. Sun et al., 2007; Taubert et al., 2011; Vincent, 2009), as well as after the consolidation or recall of semantic memories (reviewed in Buckner et al., 2008). This is further supported by observations that artificial disruption to resting state connectivity is associated with impaired recall (Censor et al., 2014; Chan & LaPaglia, 2013).

1.1 Functional connectivity – measurement, computation, and roles

The resting brain is generally studied on the basis of functional relationships between brain regions, usually termed functional connectivity, rather than the activity of individual brain regions. Functional brain connectivity reflects the co-activation of pairs or groups of brain regions, whose activity increases and decreases together, or in opposition to each other. It is purely coincident, and does not postulate on the causes or sources of such synchronization. It is distinct from anatomical connectivity, which reflects the presence of direct white matter pathways between brain regions but not the relative activity of those regions. The presence of a functional relationship between a pair or group of regions does not imply the existence of a direct anatomical connection between those regions. Functional connectivity is also distinct from effective connectivity, which reflects causal effects of one brain region on others. (While I include some effective connectivity in this work, it is not the focus. Resting state connectivity studies certainly have explored effective connectivity, including in the context of cortical stimulation revealing directional influences, but less often than more general functional connectivity.)

Functional connectivity arises from, and is calculated on the basis of, oscillatory brain activity that can be compared between pairs or groups of brain regions (Buzsáki & Draguhn, 2004). At a systems electrophysiological level, these oscillations represent the firing of large groups of neurons, whose increases and decreases in firing rate are reflected in the amplitude (the summed electrical potential of the group of neurons, including their physical orientation) and phase angle (the group's timing in the oscillation, thought to represent precise synchrony of

neurons) of the signal (Buzsáki et al., 2012; Sauseng & Klimesch, 2008). The synchronous or sequential firing of groups of neurons in different brain regions, whether or not there is a causal effect or simply coincident activity, is reflected in the coordination of those signals, and is thought to play a role in simultaneous brain function across different regions that can thus be synchronized across temporal scales, timing, and space (Buzsáki & Draguhn, 2004; Wang, 2010). These signals can be measured with various electrophysiological methods, which directly record the electrical properties of the neurons, and functional magnetic resonance imaging (fMRI), which indirectly measures this underlying neural activity via the varying oxygenation of blood.

General resting state connectivity and the existence of discrete RSNs were initially defined in and are still commonly studied with fMRI, which enjoys the advantage of examining the entire brain non-invasively (Conner et al., 2011). However, fMRI can only measure oscillations at very slow frequencies <1Hz for both technological and physiological (blood flow dynamics are much slower than neuronal dynamics) reasons, and further blood flow does not perfectly correspond to neuronal activity, necessitating the use of electrophysiological methods (Winder et al., 2017). Oscillatory dynamics at faster frequencies up to 200Hz are meaningful in brain function, arising from the summed, synchronous activity of neuronal firing, with the specific frequency of the group of neurons deriving from the number of neurons, their cell types, their spatial orientation, and a variety of other factors (Buzsáki & Draguhn, 2004). These group oscillatory functions and properties may represent the physiological transition between individual neurons and complex behavior (Buzsáki & Draguhn, 2004).

Fast neuronal dynamics are far outside the temporal range of fMRI; measuring high frequencies with EEG and MEG is unreliable due to the degradation of high frequency/low amplitude signals as a function of both distance from the signal source and recording through skull and scalp, which also induces spatial smearing of the signals (Groppe et al., 2013; Hämäläinen et al., 1993). Therefore, the capacity to record neurophysiological potentials directly from the signal source allows for the investigation of broadband connectivity properties including precise timing information, the relationship between cytoarchitectural boundaries between brain regions and connectivity patterns, high frequency/highly localized patterns, and more (Buzsáki & Draguhn, 2004; Gordon et al., 2016).

Properties of connectivity may be different across the various frequency bands of interest (delta, 1-4Hz; theta, 4-8Hz; alpha, 8-12Hz; beta, 13-30Hz; gamma, 30-70Hz; high gamma, 70-200Hz). These bands not only have different oscillatory features, but also have been linked to specific functions. In particular, the high gamma range is associated with local

neuronal firing rates rather than a higher-level cognitive function (Buzsáki & Draguhn, 2004). Delta activity is linked to sleep, theta to memory, alpha to attention, beta to movement, and gamma to binding of multidimensional stimuli and to attention (Crone et al., 2011; Groppe et al., 2013; Wang, 2010). While EEG and MEG have the obvious advantages of being non-invasive and widely available, my contributions using electrocorticography (ECoG) to reach into and resolve higher frequencies and with superior spatial resolution (especially over EEG).

The combination of advantages and limitations associated with ECoG enable unique insights relative to other recording modalities (i.e., EEG, MEG, and fMRI). Recordings can be acquired over the course of up to 7 days, usually with 3-4 days available to conduct complex cognitive tasks. This collection period is critical in light of the fact that electrodes do not shift between sessions, so unlike other common recording modalities in humans we can assume that the electrodes are recording from the same neuronal population at every session over the course of different days. The electrode placement directly on surface of brain eliminates noise and attenuation from scalp and skull. This yields high fidelity recordings of high frequencies (whose low amplitudes are more attenuated by scalp and skull than are lower frequency signals with higher amplitudes), which is one of the most distinct advantages of the method in the context of behavioral studies. It also obviates the need to reconstruct signal space from sensor space, as each electrode only records the neurons directly below it (Buzsáki et al., 2012). The electrodes can be placed on any portion of the cortical surface, including medial and inferior regions relatively inaccessible with noninvasive electrophysiology, though they are constrained to the surface.

However, there are also significant limitations associated with the use of ECoG. It is currently only used in patients with medically intractable epilepsy, who often have comorbid neurological conditions and/or psychological limitations. Though I remove the electrodes identified as being located over the seizure focus from my analysis, we cannot completely rule out epilepsy-related or other effects on the rest of the cortex. The placement of the electrodes requires surgery. Relatively few individuals are eligible for the procedure and thus sample sizes are small (N=8 in each of my main projects). Though ECoG is increasingly used in conjunction with penetrating depth electrodes (which are not included in these studies), ECoG itself is restricted to the cortical surface. Electrodes are only placed over a portion of the cortical surface, which is determined exclusively by clinical need. This usually entails placement over a large portion of only one hemisphere, or smaller segments of both temporal lobes, or combinations such as inferior prefrontal cortex/superior parietal lobe or inferior temporal

lobe/supplementary motor area. However, despite the limitations of ECoG, it provides unique insights when taken in combination with other recording modalities with contrasting limitations.

The descriptions of resting state fMRI connectivity are confirmed by electrophysiological studies that offer insight into the neural source of the hemodynamic response. In the context of my work presented here, which is exclusively in electrophysiology, we are interested in the comparability of findings based on these two phenomena. For example, the default mode network, a key RSN which is usually investigated with fMRI in the <1Hz frequency range, can also be detected with high gamma (C. Keller et al., 2013) and with the slow amplitude envelope of high gamma in the same frequency range as the BOLD signal (Ko et al., 2011). Studies using simultaneous electrophysiological and fMRI measurement also observe correspondence between the BOLD signal and slow cortical potentials (both <1Hz) (Hiltunen et al., 2014; Kramer et al., 2011). Given these findings, and the role of high gamma as a marker of local neural activity (Crone et al., 2011), it is possible that the BOLD response phenomenon has its roots in firing rates of local neural populations, that is, the neural functions underlying the high gamma component of the electrophysiological signal (Hermes et al., 2012). This signal is a plausible candidate for the physiological source of the BOLD signal that is used to define the network, and provides a basis for extrapolating findings previous studies using fMRI into our electrophysiological investigations (Logothetis et al., 2001). I present exclusively electrophysiological work, but much of its rationale is grounded in fMRI findings that can be generally considered a guide, if not fully translatable between the modalities.

In the electrophysiological domain, there are a wide variety of computational methods that have been used to characterize connectivity across the cortex (Bastos & Schoffelen, 2016; Greenblatt et al., 2012). These measures characterize activity from regions that vary together, and include specific pairwise connections between regions of interest. As I noted above, basic oscillatory signals from which functional connectivity is commonly derived can be broken into two components: the amplitude of the signal (colloquially, its strength) and phase (the position in the oscillatory cycle), which can be interpreted as markers of co-incident activity and precise synchronization, respectively (Buzsáki et al., 2012; Sauseng & Klimesch, 2008). Consequently, connectivity calculated with different computational methods, on different features of the signals and temporal parameters, or on signals recorded with different recording modalities can result in substantially different views of connectivity patterns. These views are not truly contradictory, so much as revealing different embedded patterns; for example, two signals may have strong phase-related connectivity, indicating that the signals are precisely temporally aligned, but weak amplitude-related connectivity, indicating that their relative activity levels are not linked (Bastos

& Schoffelen, 2016; Greenblatt et al., 2012). In the context of my study of the resting state, this raises the question of what different connectivity measures and their values can be interpreted to mean, for the functions of spontaneous coupling, for its evolution over time, and for potential resting state functions linked to the prior engagement of specific tasks such as skill learning. (It is also important to note that for a given neural recording of a specified length, connectivity calculations will operate on a different number of oscillatory cycles for each frequency band, which may influence the estimated connectivity properties.)

In the course of these projects, I employed four connectivity measures, selected with a view to capturing diverse properties of the signals on the basis of phase, amplitude, or both. Phase locking value ($PLV= \sum e^{i*\Delta\theta}$, where $\Delta\theta$ is the instantaneous difference in phase between the signals) measures the consistency of the phases between the signals, and provides an estimate of sustained synchrony. Phase slope index ($PSI=(\sum_{f\in F} C_{ij}^{*}(f)C_{ij}(f+\delta f))$ *,* where $\Im$ is the imaginary part, C is the complex coherency, and f is each frequency to a maximum of F) measures consistency of the lag and direction of lag between two signals, and provides an estimate of sustained pseudo-causal flow between two signals. Amplitude correlation ($\rho_{X,Y}= \frac{E[(X-\mu X)(Y-\mu Y)]}{\sigma X*\sigma Y}$, where E is the expected value) is a simple Pearson's correlation performed on the amplitude component of each signal, and provides an estimate of co-activation between the pair of regions. Finally, magnitude-squared coherence ($C_{xy}(f)= \frac{|G_{xy}(f)|^2}{G_{xx}(f)*G_{yy}(f)}$, where $G_{xy}(f)$ is the cross-spectral density and G_{xx} and G_{yy} are the autospectra of x and y) is a measure of the variance shared between a pair of signals, and estimates shared content. Importantly, each of these four measures reveals complementary properties of the relationship between the pairs of neural signals (synchrony, lagged/directional synchrony, co-activation, and shared variance respectively), so using them in combination generates a more comprehensive understanding of the relationship between neural signals than any one measure alone. Each of these measures is also susceptible to different artifacts in the signals.

1.2 Research approach

These previous findings and open questions bring me to my primary interest: is resting state connectivity a stationary phenomenon? That is, does idle coupling between cortical regions change from one resting state period to another, and if so, what parameters that dictate this change? I investigate this in the context of both spontaneous fluctuations in connectivity and task-driven changes associated with learning a new skill.

Resting state connectivity varies not only as a result of external influences, but also spontaneously, driven by endogenous shifts in ongoing resting state processes. Spontaneous variation in resting state connectivity over the course of hours to days has been observed with fMRI (Allen et al., 2014; Baker et al., 2014; Fox & Raichle, 2007; Guo et al., 2012; Handwerker et al., 2012; Laumann et al., 2015; Zuo & Xing, 2014), EEG (Brookes et al., 2014; Espenhahn et al., 2016), and MEG (Garces et al., 2016; Z. Liu et al., 2010) studies. These changes encompass the whole cortex, and in electrophysiological studies, the full range of canonical frequency bands. In ECoG, the only prior studies of resting state variability have looked intra-session – that is, how connectivity shifts moment to moment (Hutchison, Womelsdorf, Gati, et al., 2013; Kramer et al., 2011). I am interested in the longer-term shifts between different sessions, which also serve as a baseline of variability onto which task-related changes are overlaid.

Broad longitudinal changes in the brain, and in resting state connectivity specifically, can be attributed to the learning or acquisition of a new motor skill. These changes range from the level of individual synapses to mesoscale white matter networks to whole-brain functional connectivity (Dayan & Cohen, 2011). The bulk of this work in humans has been performed in fMRI, and thanks to the superior spatial resolution of the modality it has been able to resolve learning-related connectivity changes associated with both cortical and deep subcortical regions (reviewed in Kelly & Garavan, 2005; Dayan and Cohen, 2011), as well as changes to the physical structure of gray and white matter after extended practice (Driemeyer et al., 2008; Scholz et al., 2009). In particular, changes to the connectivity patterns of the anterior cingulate, posterior and inferior parietal lobe, middle and inferior frontal gyri, medial frontal wall, dorsolateral prefrontal cortex, premotor and supplementary motor cortex, motor cortex, primary visual cortex, and various subcortical regions are robustly associated with acquisition of a new motor skill (Albert, Robertson, & Miall, 2009; Danielle Smith Bassett et al., 2011; Darainy et al., 2014; Kelly & Garavan, 2005; F. T. Sun et al., 2007; Taubert et al., 2011; Vahdat et al., 2011). As I am interested in electrophysiology, it is of note that the few studies measuring learning-related changes in connectivity with EEG or MEG identify overlapping regions of interest, though they are constrained to the cortical surface (De Vico Fallani et al., 2010; Houweling et al., 2008; Mehrkanoon et al., 2016; J. Wu et al., 2014).

However, two major gaps in the literature remain: the use of electrocorticography to assess learning-related effects on connectivity and differentiating task-related change from spontaneous patterns of variation in electrophysiology (which has been examined in fMRI studies). To my knowledge, no prior study of skill learning in any modality has used longitudinal

resting state ECoG recordings and intraclass correlation or similar measures of between session variability to establish a baseline of spontaneous variability to which post-task variability can be compared.

The resting state's putative role in memory consolidation links it to the learning of a new skill, as do prior studies investigating just that. However, working backwards from this endpoint: in order to identify changes in connectivity attributable specifically to the task, I first identified the changes that occur spontaneously, in the absence of any significant cognitive demands and especially in the absence of learning. Further, to identify changes in resting connectivity that are independent of behavior, I had to start with establishing patterns in a single time point, as well as methods of measuring and interpreting connectivity. These are the three steps of my investigations that I present in the main body of this work. (Two additional studies, both relevant to this theme but outside of this arc of logic, are presented at the end.)

1.3 Electrocorticography and resting state connectivity

The first section of my dissertation uses electrocorticography (ECoG) to reveal patterns of resting state connectivity at a single time point.

ECoG has been used to study other aspects of resting state connectivity before this (Crone et al., 2006; Groppe et al., 2013; C. Keller et al., 2013; Ko et al., 2011, 2013), but has not been used to addressed my specific question of the patterns of phase synchrony at rest, and especially the patterns of high frequency (high gamma, 70-200Hz, and amplitude modulations of high gamma, 0.1-1Hz envelope of high gamma). Previous work in on other aspects of resting state connectivity in ECoG, and other studies in simultaneous EEG-fMRI (Conner et al., 2011; D Mantini et al., 2007), MEG (Florin & Baillet, 2015; Hillebrand et al., 2012), and MRI (Gohel & Biswal, 2015) had identified other major properties and patterns of resting state connectivity that informed my hypotheses. This work probed at the synchrony of brain regions at rest, the properties of idling oscillations at different, canonical frequency bands, and the specificity of patterns within and between different regions.

This portion of my work has several novel contributions that built on prior work. My approach included general resting state connectivity bounded by cytoarchitecture rather than focusing on RSNs, I focused on phase-based (i.e., timing) connectivity measures, and I included inferred causality that was not based on cortical stimulation. Like other studies in ECoG, I was also able to take advantage of improved resolution of the high gamma band (and, consequently, the amplitude modulation of high gamma) that is not constrained as with non-invasive modalities. With this study I also established a set of connectivity measures and statistical

techniques, some of which go on to appear in my later work, and others of which informed the choice of others either as complements or replacements.

This and the second portions of the work focus on the synchronizing function of resting rhythms, coordinating brain regions with each other. In the following steps, I more closely examine the memory consolidation-related features as well.

1.4 Spontaneous change in resting state connectivity in electrocorticography

After establishing the patterns of resting state connectivity in a single resting state session, and the computational methods for that evaluation, my next step was to expand the scope to include multiple resting state sessions acquired longitudinally and within subjects. Variation in resting state functional connectivity over the course of hours to days is thought to exist, on the basis of observations in other imaging modalities, but it has never been previously examined in electrocorticography. In this chapter of my work, I explore how connectivity varies between resting state sessions acquired in a serial fashion across time and in different parts of the brain. How much does it vary in different aspects of the neural signals? These questions have been previously investigated in MEG (Garces et al., 2016; Z. Liu et al., 2010), EEG (Brookes et al., 2014; Espenhahn et al., 2016), and extensively in fMRI (e.g., Allen et al., 2014; Baker et al., 2014; Guo et al., 2012; Handwerker, Roopchansingh, Gonzalez-Castillo, & Bandettini, 2012; Laumann et al., 2015; Zuo & Xing, 2014; reviewed in Fox and Raichle, 2007), which identified long-term inter-session variability in resting state connectivity over hours, days, or longer. Inter-session variability has never been previously examined in ECoG: all of the prior studies looking at variation in connectivity in ECoG looked at variation in connectivity strength within a single session, where variability is occurring on the timescale of seconds. However, I looked over the course of multiple sessions in the span of a day. In addition to the insights into high frequencies not available to be analyzed in prior studies with other modalities, I also needed to establish the patterns of variation as measured by ECoG, so that they could be used as a baseline for the third, learning-specific step of my work.

The specific contributions of this portion of my work that distinguishes it from related work were: my range of multiple connectivity measures applied to a wide range of frequency bands of interest, a focus on connectivity bounded between cytoarchitecturally defined cortical fields, and my use of ICC for data reduction across individuals and brain regions. While each of these three components had been used individually to approach studies of change in resting state connectivity in other modalities, they had never been combined, and the question had

never been approached in ECoG at all (which means that the high end of the frequency range was expanded beyond all previous studies as well).

Following my first project, I also added two new connectivity measures (amplitude correlation and coherence) in this project. My aim was to capture the complementary aspects of phase and amplitude as discussed above. I also utilized a different longitudinal statistical method that detects significant changes in connectivity rather than significant connectivity values that stand out from noise.

Recognizing that inter-session variation in resting state connectivity exists independent of a task, what is the meaning and importance of such variation? Supposing that resting state activity and connectivity reflects defined cognitive functions, fluctuations in its strength or patterns may represent variation in the ongoing internal dialogue between sessions or background changes in cognitive state. It may also reflect long-term variation in the physiological properties underlying the observed connectivity. Regardless of the cognitive or physiological cause, these changes can be thought of as background or baseline inter-session fluctuations, which would continue to persist regardless of pre-resting state interaction and which task-linked changes in resting state overlay.

1.5 Task-linked changes in resting state connectivity in electrocorticography

In the third and final chapter of my work, I separate out changes in resting state connectivity that are attributable to learning a motor skill, from those that can be distinguished from random, spontaneous variation. The previous stage of my work determined the magnitude and distribution over cortex of spontaneous, between-session variation not specifically linked to an experimental task. I capitalized on this prior effort to identify significant changes different from spontaneous change. As far as I know, the use of spontaneous changes in resting connectivity to establish a baseline for changes specifically attributable to learning in electrophysiology is unique to my work. I was also interested in whether pre-task connectivity before any exposure to a task would predict ultimate performance, as has been found in previous studies in other modalities (J. Wu et al., 2014).

I divided those subjects who successfully learned the skill (i.e., showed significant improvement in performance) from those who did not improve, and examined connectivity changes (before and after task). These changes were then contrasted with spontaneous/baseline level of change and relative to each other. Previous studies had identified both pre-task connectivity (Sala-Llonch et al., 2012; J. Wu et al., 2014) and changes in connectivity after task learning associated with the level attained in the task (Albert, Robertson,

& Miall, 2009; Darainy et al., 2014; De Vico Fallani et al., 2010; Houweling et al., 2008; Kelly & Garavan, 2005; Mehrkanoon et al., 2016; F. T. Sun et al., 2007; Taubert et al., 2011; Vahdat et al., 2011), but neither pre-task connectivity nor post-task changes had never been previously examined in ECoG. This work examined both pre and post-task connectivity features that were predictive of or linked to successfully learning the skill, specifically focusing on distinguishing those who learned the skill from those who did not.

The specific contributions of this portion of my work that distinguishes it from related work include isolating and eliminating spontaneous or background variation that is not related to the task, examining the relationship of task performance to both pre-task connectivity and changes between pre and post-task connectivity, and my comparison of multiple connectivity measures that focus on different signal properties. This is to my knowledge the first study to investigate changes in connectivity following task learning while tracking broadband changes within the human cortical spectrum.

1.6 Overview of dissertation structure

Bearing in mind that the final goal of my work was to examine the changes in resting state connectivity that can be attributed to skill learning, I start by establishing a basic framework and methods. I begin by characterizing fundamental, single-session electrophysiological properties of resting state connectivity across functional cortical regions (Chapter 3). Having established these patterns, and equally importantly laying the foundations of standard methods for revealing functional connectivity, I then move to examining how those patterns change between multiple resting state sessions (Chapter 4). Finally, I use the baseline levels of spontaneous fluctuation in connectivity between sessions to identify changes that are outside that range and can be specifically attributable to learning (Chapter 5). I also identify pre-task connectivity patterns that correlate with improvements in task skill (Chapter 5). I conclude with an outlook on future research and the applications of my work (Chapter 6). After the conclusions, I include some tangents related to the broad theme of resting state connectivity properties but which did align with the central theme of the three main chapters. These studies illustrate changes in resting state connectivity associated with learning a brain-computer interface (which also serve to demonstrate the evolution of the methods I use to evaluate changes in connectivity), and differences in connectivity between humans and the non-human primates and sheep used in much fundamental research (Chapter 7).

2. Measuring brain connectivity at a single time point

2.1 Motivation and role in broader project

In this first chapter, I examined spatial patterns of resting state connectivity in a single session. The focus centered on pairwise connections across the whole cortex. I used a wide frequency range (1-200Hz) in order to establish a comprehensive portrait of resting state connectivity patterns. I also investigated the role and phase synchrony of the 0.1-1Hz amplitude modulations of high gamma, which had been previously demonstrated to closely match fMRI signal, and may be its physiological source (Hermes et al., 2012; Logothetis et al., 2001). I exclusively focused on phase-based connectivity measures, as I was interested in the precise synchrony and timing they represent, whereas later I would explore additional components of the complex signal including amplitude-only and coherence measures, in order to capture a wider range of the signal dynamics. I also filtered connections by whether they were statistically significant with a randomized permutation phase shuffle approach, whereas later I focused on whether connections changed significantly with respect to themselves.

I observed that longer distance connections largely depended on slow frequencies (delta, theta, and amplitude modulation of high gamma), while shorter distance connections ranged across the frequencies tested. Frontal regions included a larger fraction of significant low frequency connections, while parietal and temporal regions had more of the higher frequency connections. Inferred phase-based causality also indicated a frontal-to-parietal direction of influence. This paper laid the foundations for several specific connectivity properties in ECoG in preparation for examining changes to those properties later.

The following chapter was first published as:

Casimo, K., Darvas, F., Wander, J. D., Ko, A. L., Grabowski, T. J., Novotny, E. J., … Weaver, K. E. (2016). Regional Patterns of Cortical Phase Synchrony in the Resting State. Brain Connectivity, 6(6), 470–481. https://doi.org/10.1089/brain.2015.362

2.2 Introduction to section

The study of spontaneous interactions in the resting state has become a cornerstone of functional imaging research. Mean correlations of endogenous BOLD signal interactions in the 0.01 Hz range reflect to a good degree the brain's intrinsic functional architecture (Biswal et al., 2010; J S Damoiseaux et al., 2006), and frequently mimic patterns of task-based fMRI responses that support an array of complex cognitive functions (S. M. Smith et al., 2009). While resting state fMRI studies have substantially contributed to our understanding of the brain's intrinsic and dynamic functional connectivity, the electrophysiological interactions that define the resting state at a systems level are still far less understood (Cabral et al., 2014).

Correlation between resting state fMRI (rsfMRI) and electrophysiological recordings has been well characterized in the slow frequencies accessible with fMRI (<1Hz), especially with simultaneous EEG-fMRI. While the full range of regional band-limited power estimates has been shown to correlate with the spatial distribution of rsfMRI functional networks (e.g., Conner, et al., 2011; Hiltunen et al., 2014; Mantini, et al., 2007; Schölvinck et al., 2013), subdural recordings have shown that <1Hz electrophysiological fluctuations within the amplitude envelope of high frequency oscillations better reflect physiological processes related to the generation of resting state BOLD interactions (Keller et al., 2013; Ko et al., 2011; Ko et al., 2013; Nir et al., 2008).

We investigated oscillations across the full range of band-limited frequencies, including the high gamma (HG; 70-200Hz) band and low-frequency modulations of the HG amplitude. The contribution of this full broadband range to our understanding of resting state functional connectivity has not been previously reported. This is in part due to the greater susceptibility of EEG and MEG surface recordings to biological noise of non-cortical origin and volume conduction effects due to their distal recording position. The combined effect reduces spatial and amplitude resolution, particularly within the HG band (Buzsáki & Draguhn, 2004). Alternatively, electrocorticographic recordings, (ECoG), in which electrodes are placed directly on the human cortical surface, enables high fidelity sampling of high frequencies such as HG with high spatial and temporal precision, though the recording area is limited spatially (Olson et al., 2015). This access to the HG band is particularly important given its postulated role as a marker of evoked local activity (Crone et al., 2006; Miller et al., 2009). Further, ECoG potentials are only minimally impacted by cardiorespiratory artifacts (Kern et al., 2013), a critical factor for studies of slow and infraslow bands that overlap within the same frequency range (Deckers et al., 2006).

Coordinated phase relationships are thought to represent the synchronization of large populations of neurons across and between local and distant sites (Darvas et al., 2009;

Hillebrand, et al., 2012; Sauseng & Klimesch, 2008). Such coordination is postulated to represent a mechanism of information transfer, driving influence from one region to another, or synchronizing spatially separated regions of a network (Tognoli & Kelso, 2009). In the resting state, such connectivity encompasses a broad array of connections across the cortex, some of which form organized, intrinsic resting state networks (Biswal et al., 2010; Weaver et al., 2016). Here, we focus on anatomical variation in resting state functional connectivity across the whole cortex, rather than a specific network.

Most electrophysiological studies to date have characterized phase synchrony in event-related, behavioral settings (Lachaux et al. , 1999; Sauseng & Klimesch, 2008); few studies have investigated phase interactions in resting state, particularly those in the HG range. Hillebrand et al. (2012) found frequency-dependent spatial variation in phase coherence, but did not examine directional effects. In contrast to more commonly used correlation and coherence measures, such phase-based approaches of functional connectivity are minimally impacted variations in amplitude, and allow inference of directionality and thus causal relationships (Hillebrand et al., 2012; J. Lachaux et al., 1999; Nolte et al., 2008).

We examined a wide range of canonical band-limited frequency ranges. The amHG band correspond to the time scales observed in BOLD rsfMRI FC, and correlate with observed general network distribution of rsfMRI FC studies (C. Keller et al., 2013; Ko et al., 2013). Here, we examine the slow (0.1-1Hz) range of amHG. It has been proposed that the amHG band may reflect long-range patterns of connectivity and overlapping neural processes that generate spontaneous BOLD signals, providing a link between electrophysiology and well-documented rsfMRI connectivity (Gohel & Biswal, 2015; He et al., 2010; Keller et al., 2013, 2014; Ko et al., 2011, 2013). While one focus of this study is phase interactions within the slow amHG band, connectivity across more canonical frequency bands are consistent when measured by stable patterns of high pairwise correlation (Kramer et al., 2011). Therefore, we investigated properties of phase synchrony to assess band-limited functional connectivity (Greenblatt et al., 2012) across the broad neurophysiological spectrum.

Prior work on resting state functional state connectivity (e.g., Kramer et al., 2011) found network patterns that were specific to different frequency ranges. Typically, there was a higher degree of spatial variability in higher frequencies (i.e. ≥beta), and more consistent patterns at lower frequencies (see also Groppe et al., 2013; Sauseng & Klimesch, 2008). Causal relationships have also been explored to a limited extent (Sauseng & Klimesch, 2008), but overall, electrophysiological characteristics of resting networks are still poorly understood.

We employed the phase locking value (PLV; Lachaux et al., 1999), and for directional (pseudo-causal) relationships, the phase slope index (PSI; Nolte et al., 2008). PLV measures the consistency in phase difference between two linear oscillators, while PSI is a measure of pseudo-causal interaction between two signals based on a consistent phase lag across frequencies (Nolte et al., 2008). We used PSI as an alternative to other causality estimates because it is straightforward to interpret and non-parametric, unlike Granger causality, though it does not estimate flows in each direction separately and its estimates are pseudo-causal rather than causal in the Granger sense (Greenblatt et al., 2012).

The goal of this study was two-fold: 1) evaluate regional cortical phase relationships in resting state electrophysiology in subdural recordings, with specific focus on variation across frequency bands and anatomical regions; and 2) to test for the presence of directional interactions. Our subdural measurements are highly spatially specific and minimally impacted by volume conduction effects. We find that the dominant frequency of phase interactions increases as a function of anterior-posterior position along the lateral surface of the cortex. We also observed frontal-parietal directional connectivity across a wide range of frequencies. These effects may reflect the unique contributions of each region to the overall organization of spontaneous phase interaction within the resting state.

2.3 Methods

Subjects

Eight patients (5 female, age range 10-42) with intractable epilepsy were recruited from the surgical programs at Harborview Medical Center (HMC, n=5) Regional Epilepsy Center and Seattle Children's Hospital (SCH). Each subject had only left (n=6) or right (n=2) hemisphere electrodes (Figure 2.1). All patients were hospitalized for long-term electrophysiological monitoring with implanted electrocorticographic (ECoG) electrodes as part of surgical treatment. Placements of ECoG arrays were determined by clinical needs. All recordings were obtained with subdural platinum ECoG arrays (Ad-Tech, Racine, WI or Integra Lifesciences, Plainsboro, NJ; electrode surface diameter 2.3mm, 1cm inter-electrode spacing). All patients provided informed consent in accordance with University of Washington and Seattle Children's Hospital Institutional Review Boards.

Data Acquisition

MRIs were acquired pre-operatively at HMC on a Phillips 3T Achieva and at SCH on a Siemens 3T Magnetom MRI scanner, both with an 8-channel SENSE head coil. For all patients, an MPRAGE high-resolution T1 sequence (echo time 6.5ms, repetition time 3ms, flip angle 8**°,**

matrix size 265x265, 170 sagitally aligned 1mm slices) was acquired for anatomical volume registration and surface reconstructions.

ECoG recordings were obtained post-operatively at the patient's bedside. The patient was instructed to remain awake, silent, and still during an 8-minute block. Patients were verbally reminded to stay awake if they appeared to be falling asleep. Guger g.USBamps (GugerTech, Graz, Austria) amplifiers were used for all recordings. Recordings were controlled using BCI2000 software (Schalk et al., 2004). Data was recorded with DC coupling and with a hardware-imposed notch filter at 60Hz to remove line noise and low-pass filtered at 500Hz. A standard scalp reference electrode was used for HMC patients and a subgaleal reference for SCH patients (Olson et al., 2015).

Analysis

Anatomical labeling:

ECoG electrode positions were rendered and labeled as previously detailed ((Tim Blakely et al., 2009; Hermes et al., 2010). Electrode positions were identified on a high-resolution post-operative CT scan using BioImage Suite (Papademetris et al., 2006). The CT scan for each subject was then co-registered to native pre-operative 1mm isotropic MPRAGE using an affine registration. The transformation matrix was applied to all electrode position coordinates. A secondary transform warped the native MPRAGE and native electrode positions into 1mm MNI brain space. MNI coordinates were transformed to Talairach space using the MNI anatomical labeling atlas, and Brodmann area (BA) labels were estimated using the Talairach Daemon Client (Lancaster et al., 2000). For visualization, electrodes were placed on the MNI cortical surface. Surface renderings were generated by segmenting the MNI volume using FreeSurfer image analysis suite (freely available to download at http://surfer.nmr.mgh.harvard.edu/) and MATLAB. Center coordinates from each electrode at the group level were projected onto surface rendering, and electrodes separated by individual subject (Figure 2.1a), lobe (Figure 2.1b, Frontal, temporal, parietal), or Talairach functional anatomical labels (Figure 2.1c, Frontal: dorsolateral prefrontal cortex [dlPFC], inferior frontal, sensorimotor; temporal: inferior temporal, superior temporal; parietal: lateral parietal, superior parietal).

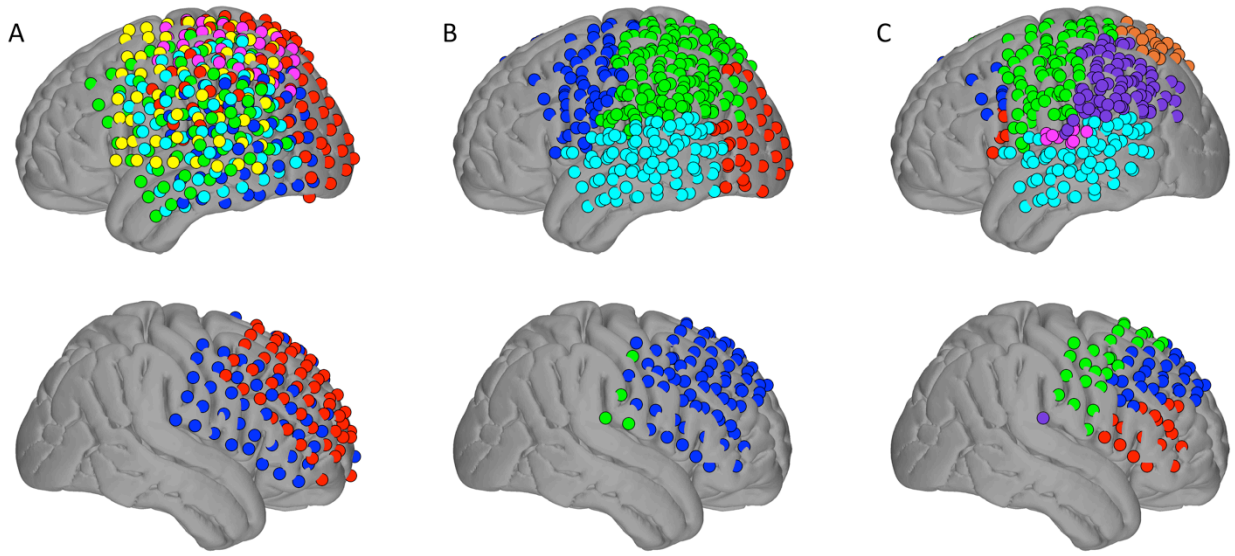

Figure 2.1
All electrode placements for all subjects (8 subjects, 480 electrodes). Interhemispheric and occipital lobe electrodes were excluded from analysis due to insufficient sample size. A) Color coded by subject. B) Color coded by lobe. C) Color coded by anatomical ROI. (Blue: dorsolateral prefrontal cortex (dlPFC), BA 8, 9, 46. Red: inferior prefrontal cortex (PFC), BA 44-47. Green: sensorimotor, BA 1-6. Light blue: inferior temporal, BA 20-22. Pink: superior temporal, BA 41-42. Purple: lateral parietal, BA 39-40. Orange: superior parietal, BA 7.

Pre-processing:

All data processing and analysis was performed in MATLAB (MathWorks, Natick, MA). Pre-processing included manual removal of non-physiological artifacts, interictal activity, or excessive noise. A common average reference across the ECoG grid was used to reject common mode noise.

Signals were band-pass filtered for each frequency band of interest (delta, 0-4Hz; theta, 4-8Hz; alpha, 8-12Hz; beta, 12-18Hz; high gamma, 70-200 Hz) using a zero-phase shift, fourth order Butterworth filter. Instantaneous phase and amplitude estimates across all channels for each frequency band were calculated using the Hilbert transform.

For the HG band only, amplitude modulation (amHG) was also calculated. Estimates were computed according to previously published methods (Keller et al., 2013; Ko et al., 2013) and consisted of a secondary, fourth order Butterworth filter applied to the HG amplitude envelope. Amplitude modulation was calculated for only the HG band in light of previous results suggesting that slow amHG oscillations uniquely represent resting state network properties (Ko et al., 2011, 2013).

Power spectra for each electrode across the full resting state time series were estimated using a fast Fourier Transform (FFT) averaged across 2sec of tapered Hanning windows. We sorted the electrodes according to two anatomical distinctions: I) lobe (Figure 2.1b) and 2)

cytoarchitecture regions of interest (Figure 2.1c). We then averaged the power spectra for all electrodes in each region (Figure 2.2a).

Connectivity measures:

Pairwise connectivity between electrodes for each individual subject was evaluated with the phase locking value (PLV) and the phase slope index (PSI). Both measures were calculated for all frequency bands of interest between all channel pairs for each subject, then pooled between subjects. Mean PLV was calculated across the full time series using the formula introduced in Lachaux et al., (1999), $PLV = \sum e^{i*\Delta\theta}$ where $\Delta\theta$ is the instantaneous difference in phase between the two signals. It provides a measure of the consistency in phase difference between two signals for the full resting state period.

PSI was calculated using the formula introduced in Nolte et al. (2008), $PS = \Im(\sum_{f \in F} C_{ij}^{*}(f) C_{ij}(f + \delta f))$, where $\Im$ is the imaginary part, C is the complex coherency, and f is each frequency to a maximum of F. It is a pseudo-causal measure of lag between two signals, indicating the consistency of the direction of lag, which Nolte et al. (2008) define as the primary direction of information transfer. A positive PSI between signals A and B reflects an overall greater influence of signal A on signal B, but does not exclude a lesser influence of signal B on signal A. PSI is an asymmetric measure: for a pair of signals, one will be assigned a positive value denoting pseudo-causal effect, while the other will be assigned the negative of that value denoting being caused. The mean PSI of a region will therefore be positive if the signals from that region provide more than they receive information. Since PSI is calculated using a sliding window over the full time-series, and too few delta cycles were included in the window for statistical significance to be established using the phase shuffling procedure, this band was excluded from PSI analysis.

Statistical analysis:

We corrected for multiple comparisons and established statistical significance thresholds for both the PLV and PSI interaction metrics using a non-parametric maximum value permutation test (Bullmore et al., 1999). Surrogate data were generated by randomly shuffling the phase component of the broadband signal, extracting frequency bands, and re-calculating PLV and PSI as above. Each iteration of this permutation procedure produced (n*n)/2−n interactions, where n is the number of electrodes (ranging from 48-64 per subject, yielding 1152-2048 surrogate interactions, respectively). To generate a null distribution of no significant phase interaction, we retained the resulting 50 highest values and then repeated the process twenty times. This approach produced a distribution of 1000 repetitions and served as basis for statistical comparison. Since we retain the maximum values across all channel pairs under the

null hypothesis, this approach minimizes the impact of artificially low estimates stemming random selection (i.e., selecting low synchrony estimates arising from, for example, large cortical distances between coupling electrodes)..

The 95th percentile of the empirical null distribution created with this procedure was set as the threshold for statistical significance (i.e., alpha level $p<0.05$ chance under the null hypothesis of random phase interactions). Values above this threshold were considered statistically significant and retained for further analysis, and values below were discarded. The total number of electrodes varied substantially between regions, as did the number of subjects contributing electrodes (Table 2.2A). This was due to the variation in clinical needs and placement of electrodes for individual patients. Importantly, only statistically significant comparisons were included in analysis. In addition, mean PLV or PSI values at the group level for any lobe or BA functional region were discarded if a regional contrast did not have at least 2 subjects contribute electrodes (Table 2.2A, B). Statistical significance in mean phase interactions across functional boundaries were evaluated using a one-way Analysis of Variance (ANOVA), investigating significant individual contrasts using Tukey-Kramer post hoc tests.

Age	Gender	Grid location (# electrodes)
11	F	L temporal (48)
13	F	L frontal (64)
13	M	L temporal (64)
20	M	L temporal/parietal (64)
35	F	L temporal (64)
37	F	L frontal/temporal (64)
37	M	R frontal, interhemispheric (64)
42	F	R frontal, interhemispheric (64)

Table 2.1
Subject demographics, electrode placement, and number of electrodes.

A

dlPFC	inf. frontal	sensorimotor	inf. temporal	sup. temporal	inf. parietal	sup. parietal
5	4	8	4	5	7	4

B

dlPFC	inf. frontal	sensorimotor	inf. temporal	sup. temporal	inf. parietal	sup. parietal
0	0	3	9	9	12	0
0	0	17	0	0	12	7
2	0	16	14	10	15	0
0	0	13	0	2	14	12
1	5	24	5	11	14	3
6	1	15	9	12	15	0
22	12	21	0	0	1	2
31	5	14	0	0	0	0

Table 2.2
Electrode location by functional region. A) Number of subjects contributing electrodes to each functional region. B) Number of electrodes contributed by each subject to each region. Subjects are in same order as Table 2.1.

2.4 Results

Electrode location

Electrode locations varied between patients, but were generally located in the lateral portions of the parietal, posterior frontal, and temporal regions (Table 2.1, Table 2.2, & Figure 2.1). The grid extended partially into the occipital lobe in one subject (<12 channels), and two subjects had interhemispheric electrode strips (both 1x8). However, due to the limited number of samples from these regions, these electrodes were excluded from analysis. Additionally, only four of eight subjects had electrodes in the superior parietal regions, two of which had 2 and 3 electrodes, substantially limiting the number of possible pairs involving this region (Figure 2.1a). All further analysis considers statistically significant pairwise analyses, drawn from an initial pool of 480 electrodes. We observed regional spatial patterns of resting power spectra consistent with previous studies (Figure 2.2a; Duff et al., 2008; Groppe et al., 2013).

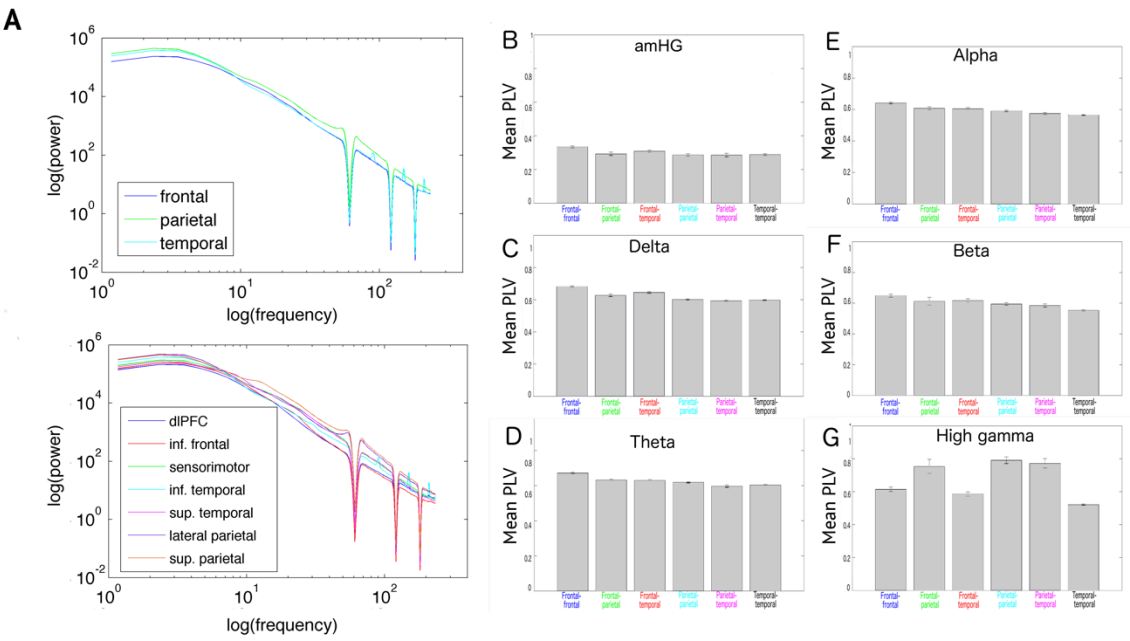

Figure 2.2
Power spectra for all subjects by lobe (A, top) and by functional anatomical region (A, bottom). The spatial distribution of spectra is consistent with prior studies (c.f. Groppe et al., 2013). Mean lobe-to-lobe PLV for B) 0.1-1Hz slow amplitude modulation of high gamma (amHG) oscillations, C) delta (0-4Hz), D) theta (4-8Hz), E) alpha (8-12Hz), F) beta (12-18Hz), G) high gamma (70-200Hz). Intra-frontal PLVs are consistently higher than intra-parietal or parietal-temporal PLVs across alpha and slower (including amHG) frequencies, while both intra-frontal and intra-parietal connectivity is significantly lower than other pairs' PLVs in HG.

Non-directional connectivity between lobes

Connectivity within and across lobes

Average PLVs across channel pairs varied significantly across both lobe pairs and frequency bands, though absolute variation was small. The mean PLV of the amHG band (Figure 2.2b) across all lobe pairs was much lower than mean PLV in other frequency bands (Figure 2.2c-g). This may reflect the lower total number of cycles available for computing synchrony in this range compared to other canonical bands. Despite the relatively low PLV estimates in the amHG band, the pattern of increased frontal-frontal and frontal-temporal PLV is comparable to the other slow frequency bands (e.g., delta/theta, Figure 2.2b-d). Notably, mean amHG PLV was significantly higher within the frontal lobe than between temporal and parietal lobes and within the parietal lobe; while these differences were small, they were statistically significant (one-way ANOVA, Tukey-Kramer post-hoc tests, $p<0.001$).

We observed little absolute difference between delta, theta, alpha, and beta in the distributions of mean PLV between lobes (Figure 2.2c-f), but as above, these differences were statistically significant. In these bands, the intra-frontal mean PLV is consistently higher, and

intra-parietal PLV lower, than any other intra- or inter-lobe relationship, with the difference in mean inter-lobe PLV decreasing as frequency increases (one-way ANOVA, Tukey-Kramer post-hoc tests, $p<0.001$). In contrast, in the HG band, intra-frontal mean PLVs were lower than all relationships except frontal-parietal and intra-parietal mean PLVs (one-way ANOVA, Tukey-Kramer post-hoc tests, $p<0.001$) (Figure 2.2g).

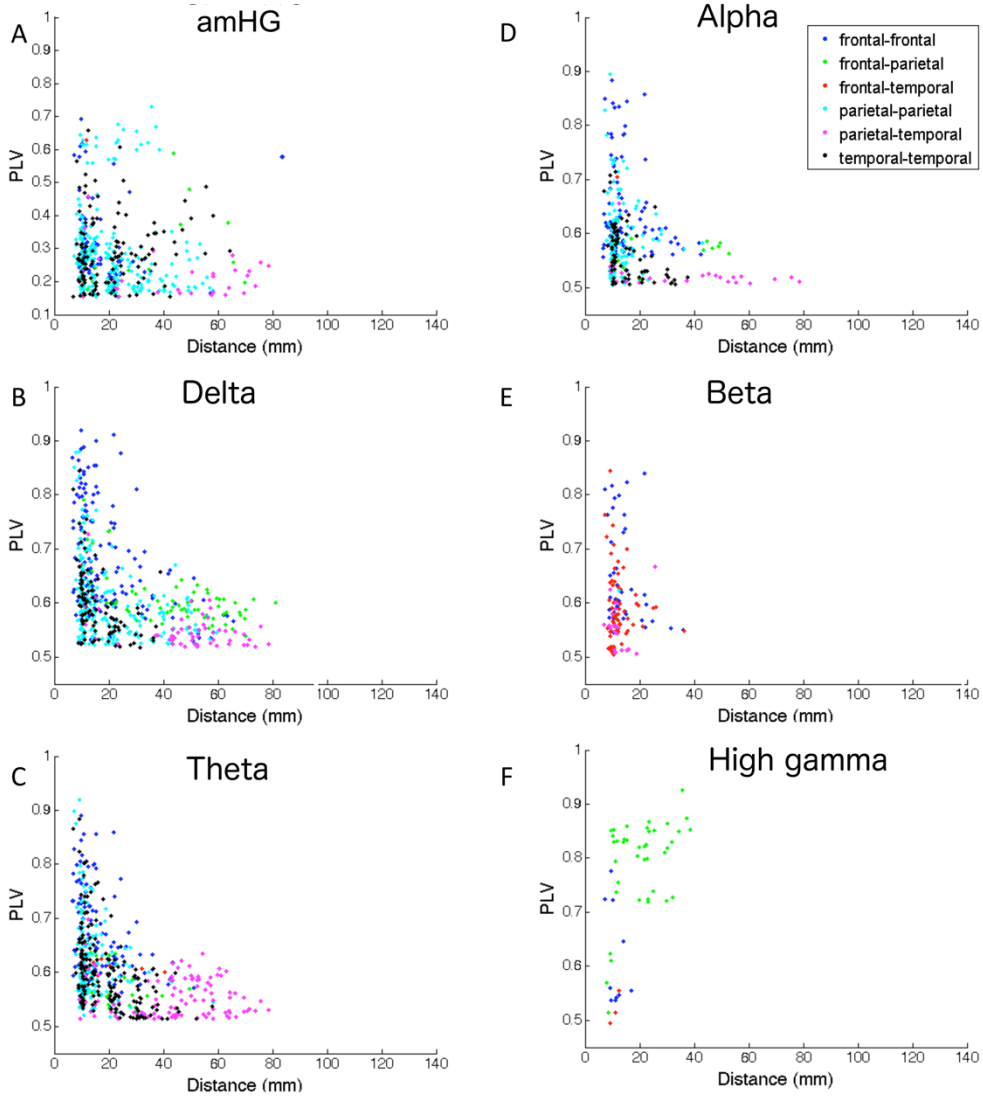

Figure 2.3
A-F, PLV as a function of distance between electrode pairs, color-coded by lobe pair. Inter-electrode spacing is 10mm. A) 0.1-1Hz slow amplitude modulation of high gamma (amHG) oscillations, B) delta (0-4Hz), C) theta (4-8Hz), D) alpha (8-12Hz), E) beta (12-18Hz), F) high gamma (70-200Hz). Colors: blue, frontal-frontal. Green: frontal-parietal. Red: frontal-temporal. Light blue: parietal-parietal. Pink: parietal-temporal. Black: Temporal-temporal. Note, vertical scale of A is different from all others. G) Mean distance between pairs of electrodes with significant PLV as a function of frequency band. Unsurprisingly, the highest PLVs were typically found between nearest neighbor electrodes, and long-range connectivity was found almost exclusively in slow (≤alpha, including amHG) frequencies.

Connectivity over distance

We examined PLV as a function of distance between each electrode pair (Figure 2.3). PLV was lower between more distant electrodes irrespective of frequency, and there were fewer phase relationships between widely spaced sites. Generally, more distant electrode pairs showed significant PLV only in slower frequencies (≤alpha), while nearer neighbor electrode pairs had significant PLV in a much wider range of frequencies. The highest PLVs, indicating very consistent resting state phase relationships, were overwhelmingly nearest neighbor electrodes (inter-electrode distance ≤10mm) located within the same lobe. Distant parietal-temporal pairs in particular exhibited a narrower range of low PLVs relative to other lobe pair relationships (Figure 2.3, pink markers); in contrast, frontal-parietal pairs had somewhat higher PLVs, and shorter inter-electrode distance, than parietal-temporal pairs (Figure 2.3, green markers).

Due to its proximity to cortex, ECoG is less affected by volume conduction than non-invasive methods are, but we cannot exclude potential effects on linear measures of synchrony. Low frequency oscillations (≤alpha) in particular, known to spread over large areas (Groppe et al., 2013), can cause spuriously high synchrony measures due to the influence of a distant generator that impacts multiple sites. Synchrony values, measured with PLV, will be a mixture of true synchrony and volume conduction.

Our permutation test does not account for this type of spurious synchrony, as our resampling methods destroy any linear correlations between signals and thus do not preserve volume conduction effects. If there were substantial volume conduction, we would expect significant PLV for all electrode pairs in a band, and that nearby sources will always appear more synchronized than distant ones. However, importantly, most electrode pairs did not exhibit significant synchronization as measured by this test, and particularly for HG did not follow the distance-PLV relationship predicted by volume conduction. We conclude that the effects observed are more likely to come from true neural synchrony rather than primarily from volume conduction.

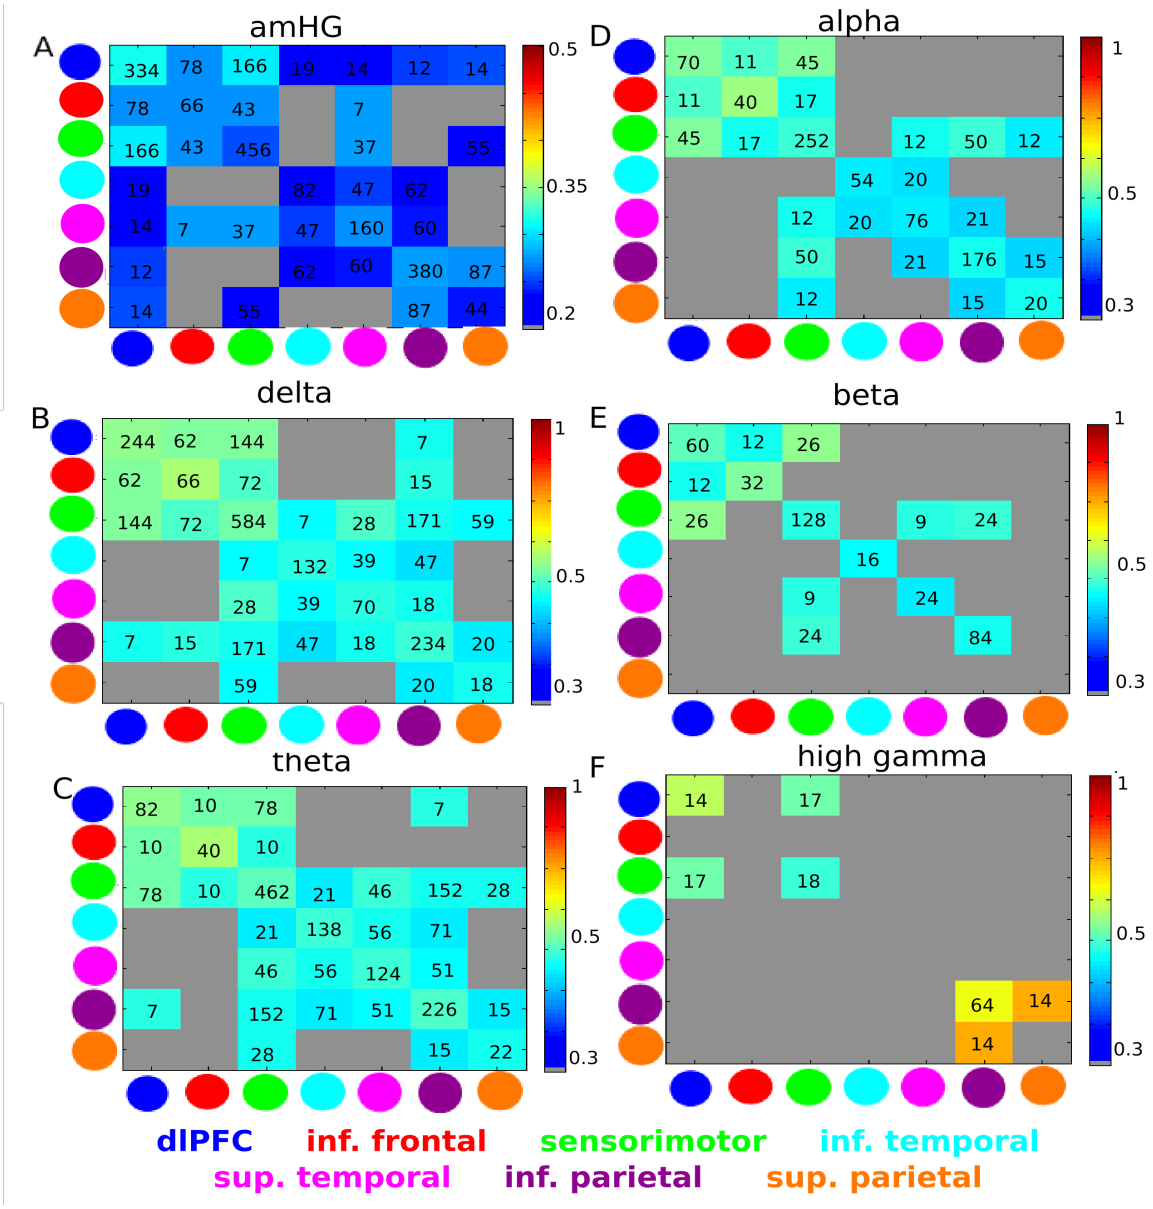

Figure 2.4
Mean region-to-region PLV for A) 0.1-1Hz slow amplitude modulation of high gamma (amHG) oscillations, B) delta (0-4Hz), C) theta (4-8Hz), D) alpha (8-12Hz), E) beta (12-18Hz), F) high gamma (70-200Hz). Along the axes of the connectivity matrix, functional anatomical regions are ordered roughly anterior to posterior. Value in each box indicates number of pairs included in each comparison; comparisons with fewer than seven pairs were excluded. Gray indicates pairs of regions with no significant interactions. We observed an increase in dominant frequency with significant PLV along an anterior-to-posterior gradient, such that significant HG interactions (F) were found more in parietal regions, while amHG interactions were more prevalent in frontal regions (A).

Non-directional connectivity between functional anatomical regions

We examined mean PLV between functional anatomical regions (Figure 2.1c). The lower PLVs seen in the amHG band relative to all other bands were again evident (Figure 2.4a). Again, regions near each other (along diagonal in plots of Figure 2.4) had generally higher PLVs

than more distant regions (near corners). For the reasons discussed above, this is more likely reflective of true synchrony and not volume conduction. Electrode grid size is limited to 7-8 cm. Due in part to the prevalence of the sensorimotor cortex as a clinical target for grid placement, coverage of the anterior frontal and superior parietal sites was limited relative to other regions. Thus, there are few opportunities to observe connectivity between very distant sites.

Overall, we observed a gradient of preferred PLV frequency across a lateral anterior-to-posterior cortical surface. Specifically, in lower (≤alpha) frequency ranges, mean PLVs were higher in the anterior relative to the posterior regional pairs (Figure 2.4a-d). Conversely, in higher (≥beta) frequency ranges, mean PLVs were higher in posterior relative to anterior pairs (Figure 2.4e-f). Consistent with the results of PLV over distance, resting state HG connectivity overwhelmingly serves nearest neighbor regions (Figure 2.4f). As noted above (Figures 2-3), generally speaking, lower frequencies are more dominant relative to higher frequencies in long-range connections and over a wider range of anatomical areas. Given this, low frequencies have higher frontal/temporal synchrony, and high frequencies have higher PLV temporal/parietal synchrony. Unfortunately, the small number of sampling electrodes in the superior parietal region, dlPFC, and inferior frontal regions also limits our interpretations.

Directional connectivity between functional anatomical regions

Although Figure 2.4 clearly illustrates the trend of the dominant phase locking relationship shifting to higher frequencies while moving along the lateral posterior surface, it does not indicate a causal or temporal order of these relationships. We used PSI, a pseudocausal measure of connectivity. Net flow as measured by PSI (Nolte et al., 2008) is a metric of the overall direction of information transfer between signals, and is applied in this context here for the first time.

We focused on PSI between functional anatomical regions (Figure 2.5). Due to the limitations outlined in the methods, we excluded the delta band. Generally, PSIs are consistent with PLVs, including PSIs being higher in ≤alpha frequencies (Figure 2.4b-d, Figure 2.5a-b), and PSIs within a lobe (directional flow from one part of a single lobe to another) being typically higher than between two different lobes (Figure 2.5).

PSI was especially strong within and between frontal regions, from frontal to parietal regions, and within and between parietal regions (Figure 2.5). Interestingly, PSI is lower in all bands between temporal regions relative to all other pairs of regions, which may reflect the lack of lateral temporal involvement in known RSNs (He et al., 2008). Overall, while there was widespread, multi-frequency frontal-to-parietal flow, there was virtually no parietal-to-frontal flow (Figure 2.5). There appears to be a potential dlPFC-sensorimotor-parietal chain of flows in lower

frequencies as this information moves anterior to posterior. This finding is notable as it demonstrates asymmetry in resting state interactions that cannot be studied with sufficient resolution and precision with fMRI due to its low temporal resolution and range of frequencies, and can only be investigated in a limited way with EEG given the far lower signal quality.

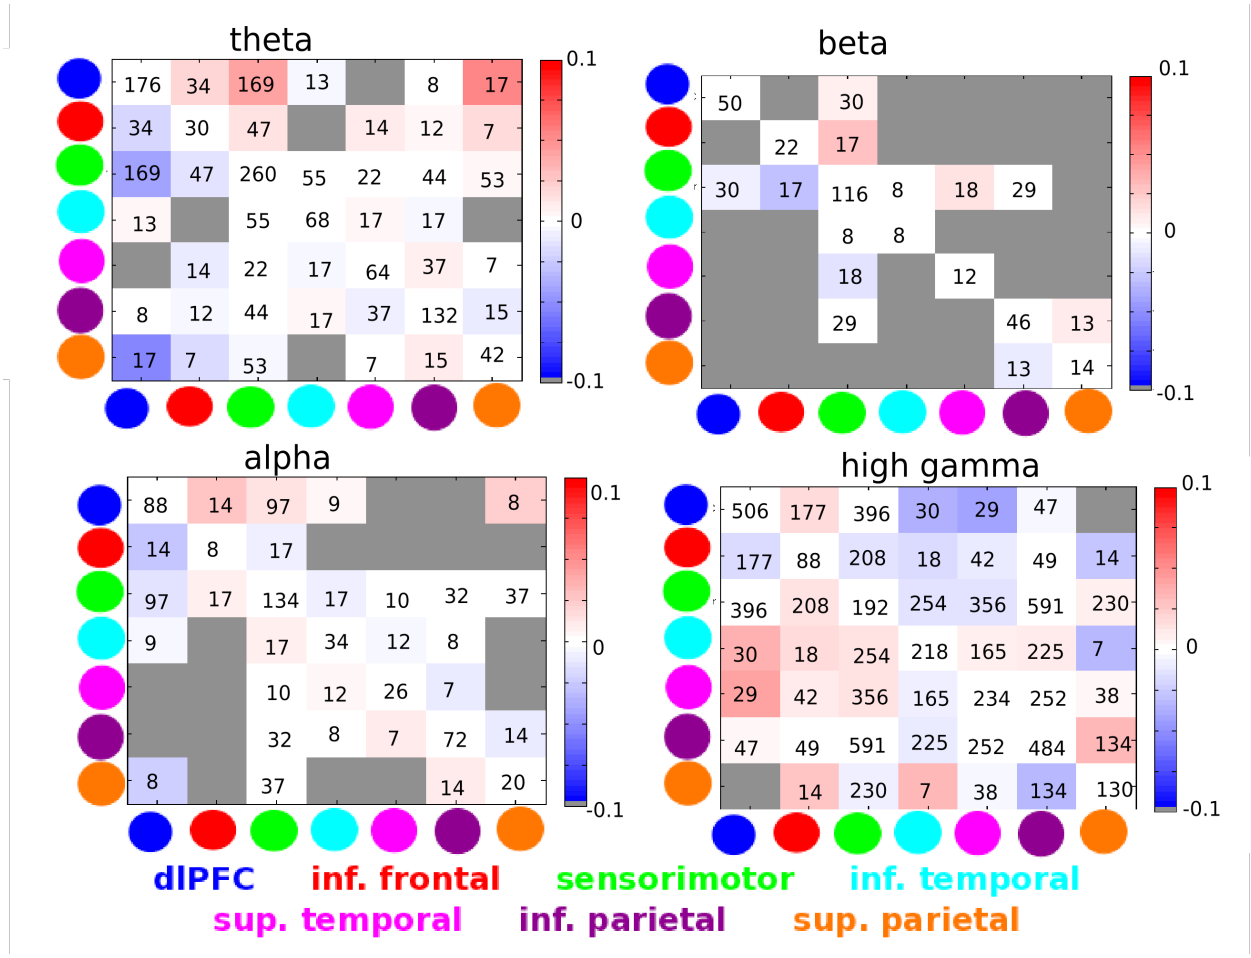

Figure 2.5
Mean region-to-region PSI for A) theta (4-8Hz), B) alpha (8-12Hz), C) beta (12-18Hz), D) high gamma (70-200Hz). Note that amHG and delta are not included due to technical limitations (see methods for details). Value in each box indicates number of pairs included in each comparison; comparisons with fewer than seven pairs were excluded. Rows are pseudo-causal and columns are pseudo-caused; red indicates flow (from region in row to region in column), blue indicates reverse flow (from region in column to region in row). Gray indicates pairs of regions with no significant interactions. Note the frontal-parietal, but not parietal-frontal direction of overall flow across multiple frequency bands. Plot thresholds set at ±0.1 based on the mean maximum and minimum observed PSI values.

2.5 Discussion

We used phase synchrony to identify patterns of linear, band-limited electrophysiological functional connectivity in the resting state. We highlight regional phase synchrony across the full human cortical frequency spectrum within the resting state ranging from delta through HG, and

including the amHG band. Our approach reveals differences between cortical regions in the frequencies with strongest phase-based connectivity, as well as local patterns of HG connectivity that occurred in multiple discrete near-neighbor groups. We also demonstrated that properties of resting state amHG PLVs are more similar to those of slow frequencies. This is in contrast with HG band PLVs, which show fewer significant relationships that were generally isolated to local functional boundaries, and to a lesser extent within lobes (Figure 2.3, 2.4). Further of note is the broad connectivity of the sensorimotor region in the lower frequency ranges (e.g. delta and theta; Figure 2.4). This pattern is inconsistent with the previously observed role of beta in sensorimotor regions during the resting state using MEG (Hillebrand et al., 2012) and the well-established beta idling rhythm that dominates the resting sensorimotor power spectrum (Groppe et al., 2013).

We also observed a frontal to parietal directional shift spanning multiple frequency bands of resting state phase synchrony estimates. While the present study does not measure the direct influence of frontal activity on parietal activity, our PSI analyses imply that frontal regions are consistently influencing the spontaneous activities of more posterior, and specifically parietal, zones. While phase coherence in the electrophysiology of resting state has been studied (Hillebrand et al., 2012), this is the first examination of causal properties of resting state electrophysiological phase synchrony.

The range of oscillatory frequencies generated by the brain allows multiple processes to be simultaneously communicated across multiple regions. The precise type of information carried by each frequency may vary, and is dependent on the regions involved and the distance and type of connection (Weiss & Mueller, 2012). Regions within a network, though spatially distant, must have some degree of interaction in order to communicate during the resting state. Phase synchrony, as examined here, is one physiological mechanism likely contributing to this level of interaction.

The phase locking value and non-causal interactions

Consistent with previous studies of resting state cross-frequency coupling (Florin & Baillet, 2015; Weaver et al., 2016), long distance connectivity appears to be largely dependent on low frequencies (Figure 2.3a-d). Across all regions, lower frequency (alpha and slower) phase synchrony was associated to a much greater degree with both short and long-range connectivity. In contrast, with the exception of amHG band, higher frequencies (i.e., beta, HG) appeared to be almost entirely linked to local cortical regions (Figure 2.3e-f).

Phase synchrony was noted between sensorimotor electrodes and electrodes from each other region across multiple frequency bands. However, the greatest magnitude effect and most

consistent phase relationships noted occurred in the beta band, particularly between sensorimotor cortical electrodes and inferior temporal electrodes (Figure 2.3e). This is consistent with results from a recent MEG study demonstrating a consistent pattern of resting state phase synchrony in beta and low gamma bands between sensorimotor regions and parietal and temporal regions (Hillebrand et al., 2012). This suggests that the classic sensorimotor beta "idling" rhythm may serve to coordinate resting state coupling of spontaneous sensorimotor activity with other functionally homogeneous neuronal assemblies spanning the lateral surface of the cortex.

We also noted frequency dependence of phase synchrony along roughly the anterior-posterior axis of the lateral neocortical surface. Generally, lowest (≤delta) frequencies had higher synchrony in frontal, especially prefrontal, regions, as well as the greatest anatomical extent of connectivity. In intermediate frequencies (theta through beta), preferential synchrony shifts towards posterior frontal and temporal areas and becomes sparser, while the highest frequencies (HG) have only very limited local synchrony, almost exclusively in parietal areas. Unfortunately, our ability to draw conclusions about the superior parietal region in particular is limited by the low number of electrodes located there.

This regional variation in functional connectivity patterns across frequencies, or alternatively, in preferred frequency of synchrony between pairs of regions, may be cytoarchitectural, anatomical, or functional in origin. Parietal and prefrontal association areas are cytoarchitecturally distinct from other primary cortical areas, such as sensorimotor or temporal regions (Amunts et al., 2007). These regions may support a neurophysiological environment more capable of high frequency resting state synchrony. For example, parvalbumin-containing inhibitory interneurons found in the association cortex are involved in coordinating gamma oscillations (Sohal et al., 2009). It is conceivable that the unique cytoarchitecture or anatomy in frontal and parietal association areas may enable resting state high frequency synchrony to a much greater degree. Further, the known specific cognitive and functional roles of each classic idling rhythm in the resting state, such as beta in motor, may here be reflected in the variance across cortical regions that do not equally contribute to those roles (Groppe et al., 2013).

amHG oscillations showed phase synchrony spanning much longer distances than synchrony of the HG band. The lower overall phase consistency (i.e., PLVs) between the amHG and all other bands, noted above, is likely a result of fewer total cycles available for computing phase interactions across our 8 minute sampling period. Consequently, desynchronization lasting even a few cycles would represent a larger fraction of the total cycles in the amHG

sample than in any of the faster frequencies, artificially lowering phase locking. Because cutoffs for statistical significance were calculated for each frequency band individually, these values are still significant despite the discrepancy with the other bands that arises from the desynchronization effect. Despite this overall drop in absolute value, the spatial patterns of connectivity indicate amHG is more analogous to low frequency bands rather than to the HG band itself. Consistent with previous studies (Keller et al., 2013; Ko et al., 2011), these results support the hypothesis that amHG cycling may act as a long-range carrier for entraining HG that is otherwise restricted to local neighborhoods. Potentially, phase coordination of the amHG band could synchronize and modulate local HG phase and amplitude activities across long-distances (Chawla et al., 1999).

The phase slope index and causal interactions

This study is the first to investigate causal phase relationships in the electrophysiology of resting state signals sampled directly from the cortical surface. Generally, connectivity patterns observed with PSI were consistent with PLV. PSI values appear, intuitively, quite low. Values close to zero may be due to the averaging of positive and negative numbers (indicating inconsistent flow direction), or because these networks tend to have low overall net effect on each other relative to all the inputs to a region. There are no established expected ranges for PSI in brain data. We found that net flow of spontaneous connectivity generally moves within frontal regions, most strongly from dlPFC to sensorimotor; from frontal regions to parietal regions; and from inferior to superior parietal. The lack of parietal-to-frontal flow cannot be attributed solely to the low number of electrode frontal-parietal electrode pairs, as unidirectional frontal-to-parietal flow was observed, as were significant PLVs between the regions. Notably, temporal regions have overall very weak causal relationships, even relative to the low PSI values overall; however, in the amHG band only, there is higher net flow from inferior to superior temporal and from inferior temporal to inferior parietal. The overall anterior-to-posterior direction of flow is consistent with findings from previous studies (Nolte et al., 2008), but is demonstrated more robustly here with greatly sensitivity to the high frequency signals.

The cognitive or neural mechanisms that mediate this directional flow of information are not clear, though the direction of connectivity is evident. It is likely that the same cytoarchitecture and regional anatomical differences driving PLV interactions discussed previously also contribute to causal PSI estimates, though we can derive a higher-level model of mean causal directionality from frontal and parietal regions within spontaneous resting state recordings. Several ubiquitous cortical networks, such as the default mode network, the frontoparietal network, and the dorsal attention network have localized hubs in frontal and

parietal neocortex, likely due to regional differences between primary, secondary and association cytoarchitecture. Hub membership may dictate directionality as measured by PSI, potentially by regional differences in intrinsic properties of idling rhythms. Additional studies are needed to investigate regional differences in resting state phase selectivity and dominant idling power distributions.

As discussed above, the frontal-parietal flow of information may reflect known differences in the contribution of those regions to overall cognitive or anatomical variability. In this work, this need to synchronize cortical areas is indicated by PLV, and a net flow or influence is indicated by PSI.

Advantages of ECoG for resting state recordings, limitations, and future directions

An advantage of this study was our use of ECoG, rather than scalp EEG or MEG. Relative to surface-based recordings, ECoG has minimal recording noise and high spatial resolution due to avoiding scalp, bone, and fluids. This results in a relatively low influence of volume conduction and greater signal-to-noise ratio, especially in the high frequency end of the spectrum. Consequently, we were able to leverage ECoG's advantages in both very high and very low frequency ranges to investigate the phase properties of the amHG band for the first time.

Despite the inherent advantages that ECoG provides relative to surface level measurements of resting state synchrony, ECoG suffers from incomplete spatial coverage of the full cortical space. Consequently, we did not generally have spatial coverage over multiple network hubs, and were therefore not able to investigate phase dynamics across a true network structure (e.g. default mode network, dorsal attention network). This limits the number and patterns of neural populations that are available for investigation and restricts a complete interpretation of network synchrony. Additionally, samples can only be recorded from individuals presenting with intractable focal epilepsy. Due to the low number of cycles included in an eight-minute period, we were unable to examine the 0.01-0.1Hz amHG range that more closely corresponds to frequencies observed in fMRI. These limitations should be considered with these results.

This study included three adolescent (ages 10, 13, and 13) subjects. We lack sufficient sample size here to investigate development of networks in detail. Prior studies indicate that patterns of phase coherence (Thatcher et al., 2008), default mode network connectivity (Ko et al., 2011) and imaging markers of resting state functional connectivity within the areas sampled here (de Bie et al., 2012) are established well before age 10, our youngest subject. Future work will include younger subjects to specifically study early childhood developmental effects.

While PLV and PSI are both effective at capturing linear relationships inside of a single frequency band, they are not sensitive to non-linear or cross-frequency interactions. Other methods such as bi-phase locking (Darvas, Ojemann, & Sorensen, 2009) and phase-amplitude coupling (Miller et al., 2014; Miller, Foster, & Honey, 2012; Penny et al., 2008) have been used to examine such patterns in non-resting state data. In the future, they may be used to examine other features of resting state data that the present analysis does not capture (e.g., Weaver et al., 2016).

Hillebrand et al. (2012) demonstrated a positive relationship between regional spectral power and phase synchrony. Although power effects are not the focus of the current set of analyses, their observations may have been influenced to some degree by volume conduction concerns inherent to MEG measurements. The relationship between the phase synchrony of two distinct neural populations and the amplitude or power of a region is not simple, despite the intuition that low signal power would underlie diminished population synchrony (Chawla et al., 1999; Daffertshofer & van Wijk, 2011).

Unlike Granger causality, which provides two numbers for a pair of signals (causal effect in each direction), PSI provides only one number for net flow, and does not separately estimate potential flow in the "reverse causal" direction (Granger, 1969; Nolte et al., 2008). However, despite these limitations, PSI has a much lower rate of false positives and also specifically is calculated from phase relationships between signals, which were of primary interest in this study.

Conclusion

Overall, our results indicate a strong consistent regional phase-related connectivity across frequency bands within the resting state. We conclude that: 1) the dominant frequencies that drive phase interactions increase across the anterior-to-posterior position along the lateral surface of the cortex, 2) the phase of amHG signal follows connectivity patterns similar to slow frequency bands (delta, theta) than HG, and may act as a carrier for HG connectivity that is otherwise restricted to local neighborhoods, and 3) there appears to be an overall anterior to posterior directional flow across frequencies as measured by PSI. The observed spontaneous, consistent interactions may imply the existence of underlying oscillatory generators that drive intrinsic networks, while the variation across frequency and space hints at the differences across the brain that maybe dictated by cytoarchitectural or functional neuroanatomy. The findings contribute to understanding the variation and flow of spontaneous interactions at multiple time scales in an unprovoked fashion.

3. Estimating spontaneous change in brain connectivity over time

3.1 Motivation and role in broader project

At this point, I transitioned from looking at resting connectivity at a single time point to examining how this observation varies over time – first, in its spontaneous, non-task-linked form. At this point, I also acknowledged and remedied some limitations of the original set of methods I used. I excluded PSI, as its values had relatively little spread, which limited its interpretability relative to PLV; this problem would only compound as the number of regions of interest involved grew from 7 in the previous paper to 10 in the following paper (i.e., since PSI is a directional measure, it produces twice as many values as non-directional measures). I added amplitude correlation in order to probe phase (synchrony) and amplitude (activity level) effects, and coherence, to probe the combination of the two features. To capture spontaneous, longitudinal variation, I applied intra-class correlations (ICC).

In the following paper, I used three different connectivity measures and assessed variation in resting state connectivity. I examined the differences between the three connectivity measures, the relationship between connectivity strength and stability, and differences in variation across the cortex. I found that coherence was overall lower and less variable than both PLV and amplitude correlation. In phase locking and amplitude correlation, strong connectivity predicted high stability (high ICC value), but high stability did not predict strong connectivity – that is, there were many weakly connected but stable links. This relationship between connectivity and stability was lost in coherence. I also found substantial variation in the connectivity of individual regions, though connectivity was overall most stable within single regions (regardless of strength). As examples: the parahippocampal and entorhinal cortex had very stable but weak connectivity except within itself, where connectivity was stable and strong. Dorsolateral prefrontal cortex connectivity was both weak and unstable in all of its non-self connections. Inferior parietal lobule had low stability, though its connectivity strength varied little.

The following chapter has been submitted as:

Casimo, K., Madhyastha, T.M., Ko, A.L., Grassia, F., Ojemann, J.G., Weaver, K.E. (in preparation). Spontaneous variation in electrocorticographic resting state connectivity. Brain Connectivity.

3.2 Introduction to section

Many previous studies have observed synchronized or coordinated activity in the brain at rest (reviewed in Biswal et al., 2010). This spontaneous resting state connectivity represents organization within and between functionally related brain areas, and may provide an anatomical mechanism facilitating information transfer within the brain (Buckner & Krienen, 2013; Garces et al., 2016). Such resting state connectivity has historically been most commonly evaluated using fMRI (Cabral et al., 2014), but can also be investigated with electrophysiological methods (EEG: Astolfi et al., 2009; Betzel et al., 2012; Chu et al., 2012; MEG: Baker et al., 2014; Brookes et al., 2014; ECoG: Casimo et al., 2016; Ko, Weaver, Hakimian, & Ojemann, 2013; review, Schölvinck, Leopold, Brookes, & Khader, 2013).

There are multiple computational approaches that have come to define functional connectivity, all of which are independent of and yet directly linked with the underlying modality. In fMRI, functional connectivity is most often defined as the correlation of the BOLD signal between regions of interest. In contrast to the hemodynamic response delay that makes precise timing information inaccessible with fMRI, electrophysiological properties respond in real time, enabling a wider range of connectivity methods, which each probe a different set of features of the signals (Bastos & Schoffelen, 2016; Greenblatt et al., 2012). Further, because electrophysiological signals reflect brain activity in real time, while the BOLD signal is subject to a physiological delay in response, connectivity calculated from electrophysiological signals can be used to probe precise timing-related relationship between brain regions (Sauseng & Klimesch, 2008). Such phase synchrony is thought to represent precisely timed communication and synchronization between brain regions (Greenblatt et al., 2012; Sauseng & Klimesch, 2008). In contrast, amplitude-based measures are less specific with regard to fast timing information, and represent more general, sustained co-variation in activity level, rather than synchronization (Sauseng & Klimesch, 2008). Finally, connectivity measures that incorporate both phase and amplitude components of the signal variously reflect the overall interdependence of two signals (Bastos & Schoffelen, 2016). Evaluating connectivity with multiple mathematical measures provides multifaceted insight into the complex dynamics linking brain regions.

Spontaneous inter-session variation in resting state connectivity

We focus here on intra-individual variation in resting state connectivity between different resting state sessions (as opposed to within a single resting state session, i.e., dynamic connectivity). Extensive studies in fMRI (Allen et al., 2014; Baker et al., 2014; Handwerker et al., 2012; Laumann et al., 2015; Zuo & Xing, 2014) have identified spontaneous fluctuations in

resting state connectivity strength that occur across the brain, but which vary between different regions and which seem especially flexible in inferior parietal and cingulate cortex (Allen et al., 2014), and especially stable in sensory and motor cortex (Espenhahn et al., 2016). Equivalent studies in EEG (Brookes et al., 2014; Espenhahn et al., 2016) and MEG (de Pasquale et al., 2010; Garces et al., 2016) have also found frequency-related patterns of spontaneous intra-individual, inter-session variation in resting state connectivity.

The most common method for evaluating the reliability or the level of variation in connectivity over the course of multiple experimental sessions is the intraclass correlation (Danielle S Bassett et al., 2006; Guo et al., 2012; Madhyastha et al., 2014) This method is important in that it enables comparisons of multiple subjects at different time points, and can easily accommodate missing data or datasets of different sizes, improving robustness (McGraw & Wong, 1996). Using this method to establish a baseline of spontaneous variability in inter-session connectivity contributes to more robust studies of resting state connectivity under a variety of conditions. For example, this would enable studies investigating change as a result of some intervention or a disease state to attribute any observed shifts in connectivity specifically to that intervention or disease, rather than to spontaneous change that would occur in the absence of any experimental task.

However, to our knowledge, inter-session variation in resting state connectivity has not been previously quantified in electrocorticography (ECoG). Electrocorticography is used for localizing seizure focus and surgical planning. Its in-dwelling, subdural nature provides time-frequency and spatial resolution superior to any other macro-scale modality available in humans. However, its drawbacks include limited spatial sampling within each individual, sampled regions are not fully overlapping across subjects, and the subject pool is by definition not typical, usually restricted to patients suffering from intractable focal epilepsy. Despite these limitations, prior studies of resting state in ECoG have identified patterns of connectivity similar to those identified with other modalities (Weaver et al., 2016), but with vastly superior fidelity of the high frequency range (i.e., high gamma, 70-20 Hz) and a greater capacity to estimate the amplitude modulation of high frequency signals not easily accessible with other surface-level recording technologies (Casimo et al., 2016; Groppe et al., 2013; C. Keller et al., 2013; Ko et al., 2011, 2013). Overall, ECoG studies investigating resting state connectivity in a single session are consistent with studies using fMRI and MEG (de Pasquale et al., 2010; Ko et al., 2011; Weaver et al., 2016). We build on this previously established congruence between modalities in a single session, and prior studies with fMRI and MEG in inter-session resting state variability, to here investigate inter-session changes in resting state connectivity in ECoG for the first time.

We investigated spontaneous inter-session, intra-individual variation in resting state connectivity recorded with ECoG over the course of three resting state sessions on two separate days. We calculated three different pairwise connectivity measures to capture different signal features between all electrode pairs in each resting state sessions for each of our eight subjects. We then identified the brain anatomical regions for those electrodes. Finally, we calculated intraclass correlation using all subjects' pairs of electrodes binned across anatomical regions. We also assessed the relationship between average connectivity strength, distribution of connectivity values, and the stability of the connectivity over time. Our results contribute to an understanding of a spontaneous, "baseline" level of inter-session variability in resting state connectivity. We find that high values of two of our three connectivity measures predict high stability, but high stability does not predict strong connectivity. Self-connections and interactions with primary sensorimotor regions are generally stable. Interactions with parahippocampal and entorhinal medial temporal regions and fronto-parietal connections are notably unstable. Our use of ECoG enables unique examination of high frequencies and precision in temporal relationships.

3.3 Methods

Subjects and data acquisition

Eight subjects with medically intractable epilepsy were recruited from the surgical programs at Harborview Medical Center Regional Epilepsy Center. Each subject had only left (n = 5) or right (n = 3) hemisphere electrodes (Figure 3.1). All patients were hospitalized for long-term seizure monitoring with implanted ECoG electrodes as part of pre-surgical planning. Placements of ECoG arrays were determined entirely by clinical needs. All recordings were obtained with subdural platinum ECoG arrays (Ad-Tech, Racine, WI or Integra Lifesciences, Plainsboro, NJ; electrode surface diameter 2.3mm, 1 cm inter-electrode spacing) with a standard scalp reference. All patients provided informed consent in accordance with University of Washington Institutional Review Board.

We obtained clinical recordings (XLtech, Natus Medical Inc., Pleasanton, CA) and video monitoring from postoperative days 4 and 5 and manually identified, and extracted three ten-minute sessions where the patient was resting (i.e. not engaged in any particular task, not moving, and awake). Recordings were collected at 1000Hz sampling rate with a high-pass filter at 0.1Hz. We identified two sessions from one day and one session from the other day, with some subjects having two sessions on postoperative day 4 and some having two on day 5. All sessions were at least 90 minutes before or after any clinically notated sub-clinical or clinically

manifest ictal responses. The subjects did not participate in any other research studies between these resting state sessions.

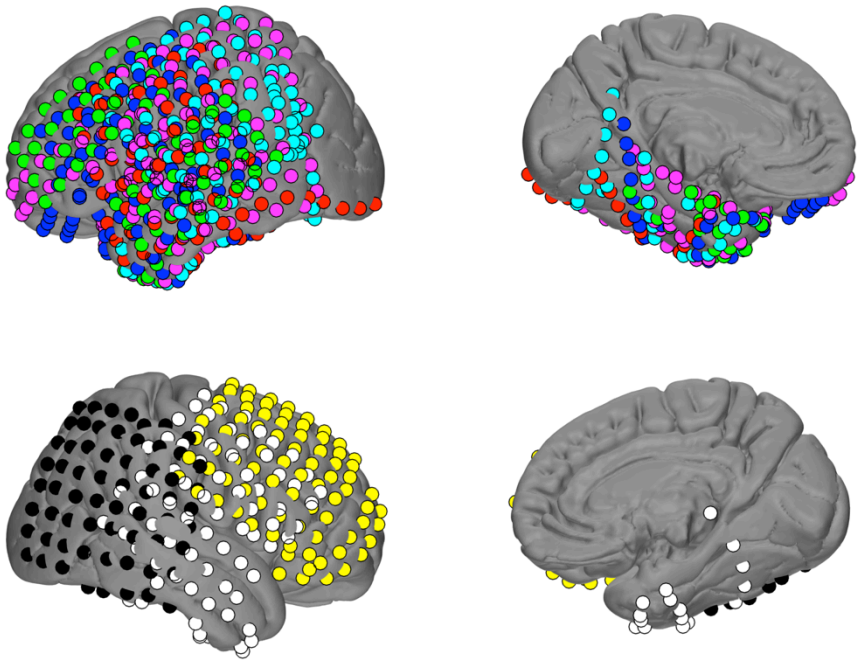

Figure 3.1
Electrode positions for eight subjects. Region pairs were excluded from analysis if they had ≤2 subjects or ≤10 electrodes, but are shown here.

Preprocessing and electrode localization

All ECoG data processing and analyses were performed in Matlab (Mathworks, Natick, MA). Preprocessing included selecting time segments that minimized artifacts, and calculating a common average reference across each grid and strip electrodes to reject common noise. We eliminated penetrating depth electrodes (present in 4 of 8 subjects) and electrodes clinically identified as being over the seizure focus and from further analysis. We also ensured that the rest periods were at least 90 minutes before or after a seizure and included no interictal activity.

For the duration of each resting state session, we extracted amplitude and phase angle using a non-analytic Morlet wavelet with ¼ octave resolution over pseudofrequencies 1-200Hz. We then averaged across phases and amplitudes for frequencies falling into our bands of interest (delta, 0.1-4Hz; theta, 4-8Hz; alpha, 8-12Hz; beta, 13-30Hz; gamma, 30-70Hz; HG, 70-200Hz). All following estimates of connectivity were calculated on these time series.

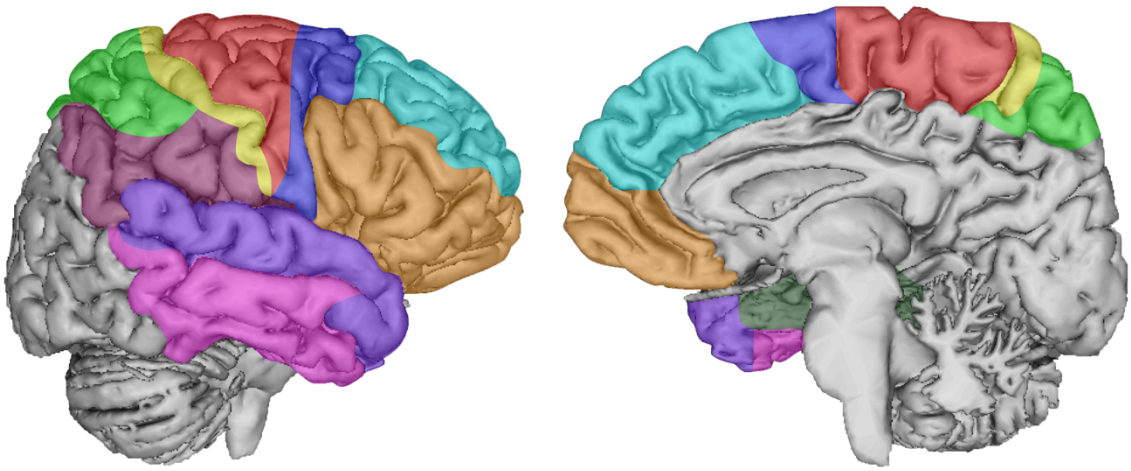

Figure 3.2
Outlines of functional brain regions used for regional pairwise connectivity analyses. List of Brodmann areas included in each functional region are shown in Table 3.1. Colors: red, primary somatosensory cortex; yellow, primary motor cortex; lime green, superior parietal lobule; blue, premotor and supplementary motor cortex; teal, dorsolateral prefrontal cortex; pink, inferior and middle temporal gyri; dark green, parahippocampal and entorhinal gyri; burgundy, inferior parietal lobule; purple, superior and transverse temporal gyri; orange, inferior frontal gyrus.

ECoG electrode locations were rendered and labeled as previously described (Blakely et al., 2009; Casimo et al., 2016). We first localized electrodes in native space based on preoperative MRI scan, then transformed each into standardized space (MNI 152 brain), and Brodmann area (BA) labels assigned to each electrode using FSL atlas tools and confirmed by Talairach daemon atlas (Figure 3.2). We then grouped these BA labels into larger functional zones (Table 3.1) based on electrode coverage and previous published criteria (Casimo et al., 2016). For each participant, we then sorted all electrodes in a pairwise manner, and binned each electrode pair into its equivalent pair of functional zones. Finally, we aggregated electrode pairs across subjects into their functional zones. Each subject only had electrodes over the right or left hemisphere, and we aggregated all electrodes across both hemispheres at this stage. Any pair of functional regions with fewer than ten electrode pairs, or fewer than three subjects, was excluded from further analysis. We grouped all electrode pairs according to their functional regions (including pairs of electrodes within a single functional region).

Region name	Abbreviation	BAs included
Primary somatosensory cortex	S1	1, 2, 3
Primary motor cortex	M1	4
Superior parietal lobule	SPL	5, 7
Premotor, supplementary motor cortex	PM	6, 8
Dorsolateral prefrontal cortex	dlPFC	9, 10, 46
Inferior and middle temporal gyri	I/MTG	20, 21
Parahippocampal and entorhinal gyri	PH-ER	27, 28, 34, 35, 36, 37
Inferior parietal lobule	IPL	39, 40, 43
Superior, transverse temporal gyri	STG	22, 38, 41, 42
Inferior frontal gyrus	IFG	11, 44, 45, 47

Table 3.1
Functional brain regions. Electrode coordinates were identified and automatically assigned to a Brodmann area. We then grouped Brodmann areas into ten broader functional regions. Brodmann areas not included (e.g. visual cortex) had two or fewer subjects with electrodes in that region.

Functional connectivity calculation

We calculated pairwise functional connectivity in each of the six frequency bands of interest, evaluating connectivity between all electrode pairs within each subject. We included 1) phase locking value (PLV), 2) amplitude correlation, and 3) coherence (Garces et al., 2016; Greenblatt et al., 2012; J. Lachaux et al., 1999). Mean pair-wise connectivity was estimated across the full time series. Because eligible electrodes across subjects ranged from 64-116, the number of pairwise electrode connections (N*N/2 - N) ranged from 1984-6612.

No single connectivity measure can capture the full extent of temporal and frequency dynamics of neural signals. As such, we present results from three selected linear connectivity measures. We include PLV, correlation of the signal amplitudes, and coherence. Together, these three measures target the specific interactions of phase only, amplitude only, and the two features combined. More specifically, PLV is particularly sensitive to variation in temporal phase synchrony, while amplitude correlation captures co-variation in signal envelope between two-time series, and coherence captures the widest set of signal features but at the expense of losing specificity. For all three measures, we obtained a single mean estimate of connectivity over the entire time series.

Statistics: intraclass correlations

We estimate the level of variation in functional connectivity using intraclass correlation (ICC) (McGraw & Wong, 1996). ICC is a widely used method of estimating the reliability, or level of variation, in a group of subjects at different measurement points (here, time points of resting state sessions). We employed the A-1 ICC as defined by McGraw and Wong, which is a one-way random effects model. The model accepts data in two dimensions: subjects (electrode

pairs' functional connectivity from all subjects) and measurements (the three resting state sessions). We also estimated the 95% confidence interval of the ICC.

We calculated the ICC in each region pair to evaluate the variation of all the electrode pairs' connectivity across the three resting state sessions. We repeated the ICC calculation for each frequency band of interest, in each pair of regions, for each of the three connectivity measures.

We aggregated these results across all eight subjects for all functional pairs.
Note that ICC is calculated on the aggregate of all electrode pairs between that pair of regions, so there is not a one-to-one correspondence between ICC values and each electrode pair's connectivity value (because each ICC value represents tens or hundreds of electrode pairs). Further, because ICC acts to combine the level of variation in each contributing electrode pair into an aggregated estimate of variability, it is entirely possible to have wide variation between the absolute connectivity values for each electrode pair and have a high value of ICC for stability, if each individual pair was relatively stable over time. Conversely, it is possible to have a narrower range of connectivity values and still have a low ICC, if each individual pair was relatively unstable within that narrow range.

Finally, we also calculated the average connectivity for each pair of functional regions by taking the mean of all electrode pairs' connectivity values in all three resting state sessions. We repeated this process over the six frequency bands of interest and three connectivity measures. We compared different connectivity values and ICC values across frequency bands using one-way ANOVAs. We evaluated the relationship between connectivity and ICC using Spearman's correlations.

3.4 Results

Reliability across the whole brain

We first examined global properties of pairwise connectivity and longitudinal stability across the sampled portions of the brain. Figure 3.3 shows the distribution of mean pairwise connectivity values and the ICCs for each of the three connectivity measures and six canonical frequency bands. Though displayed together, the pair boxes for each connectivity measures each represent different data: the spread of the connectivity distributions represents the consistency of connectivity between and within all measured functional brain regions, and the ICC represents consistency of connectivity over time.

First, we observe the connectivity values themselves between the three measures. We observe that in each frequency band, PLV is generally quite tightly clustered, especially in high gamma (Figure 3.3F). Coherence also falls into a very narrow range of values in each frequency

band and varies little between frequencies, but it is highest in gamma and high gamma (Figure 3.3E, F). In contrast, amplitude correlation values fall in a relatively wider distribution in frequencies slower than beta (Figure 3.3A-D), with an especially wide distribution of connectivity values in delta (Figure 3.3A). Within each connectivity measure (PLV, amplitude correlation, coherence), the connectivity values across different frequencies were all significantly different (one-way ANOVAs, all $p<0.01$). Within each frequency band, the three connectivity measures were significantly different (one-way ANOVAs, all $p<0.01$), driven largely by PLV being stronger than the other two measures.

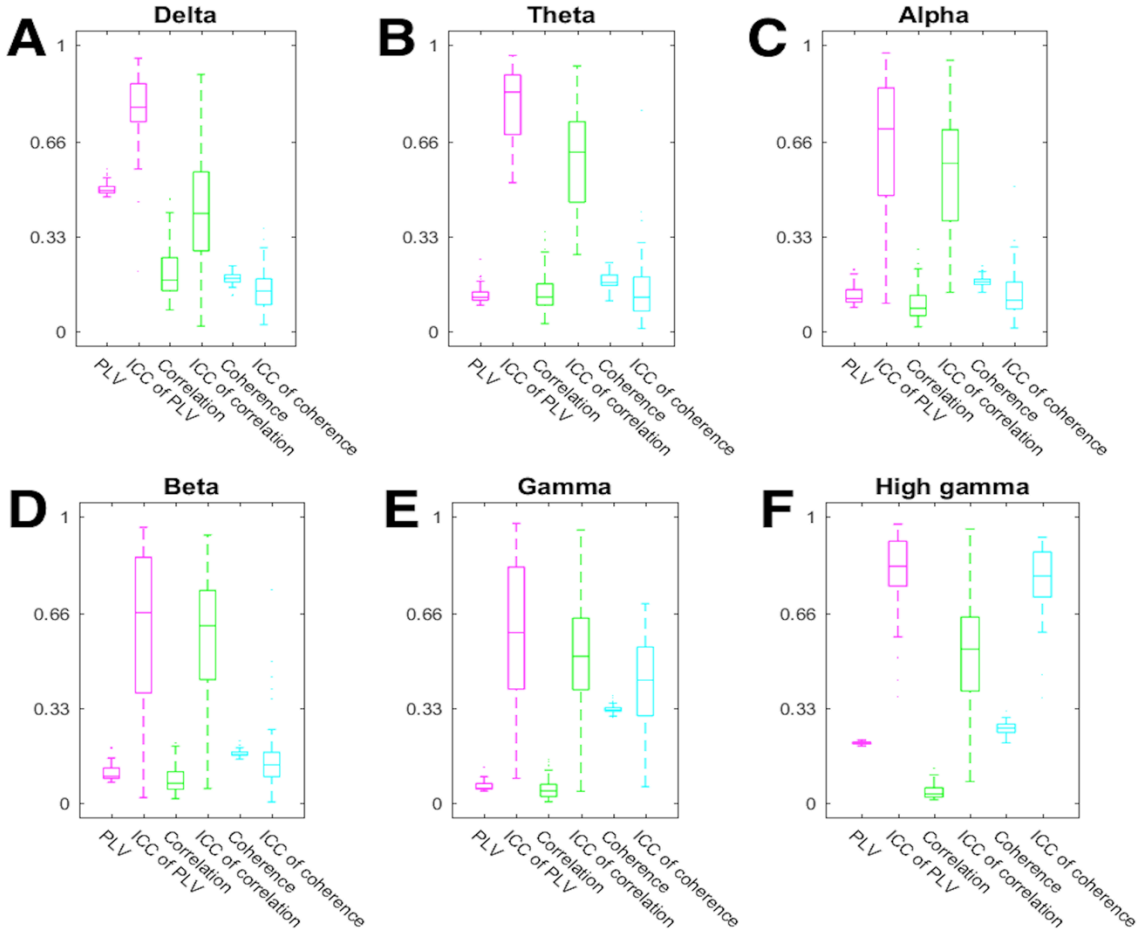

Figure 3.3
ICC and actual connectivity values for all three connectivity measures, aggregating all functional regions of interest. Each subplot represents one frequency band of interest. Box plots show median, interquartile range, and additional values outside interquartile range. Within each subplot, from left to right: (a, magenta) phase locking values, (b, magenta) ICC of PLV, (c, green) amplitude correlation values, (d, green) ICC of amplitude correlation, (e, cyan) coherence values, and (f, cyan) ICC of coherence. Individual values in boxes a, c, and e represent mean pairwise connectivity values between all pairs of electrodes in that pair of regions (from all eight subjects). Individual values in boxes b, d, and f show the ICC stability estimates between each pair of functional regions, each of which represents the aggregate of all subjects.

Next, we assessed longitudinal stability of connectivity by computing the distributions of the ICC measures of stability in each frequency band. There appears to be a frequency-dependent function influencing variation in the distribution of PLV ICCs between regions, with relatively narrow distributions of ICCs in delta and high gamma (Figure 3.3A, F), somewhat wider distribution in theta (Figure 3.3B), and very wide in alpha, beta, and gamma (Figure 3.3C-E). The ICCs of amplitude correlation are consistently widely distributed across all frequency bands, but the individual frequencies have similar mean values. In stark contrast to the patterns seen in PLV and amplitude correlation ICCs, coherence is generally unstable (low ICCs), with the exception of the high gamma band (Figure 3.3F). The distributions of coherence ICCs in all frequencies are very narrow, which is not true for the other connectivity measures (Figure 3.3). Within each connectivity measure's ICCs (PLV, amplitude correlation, coherence), the ICCs across frequencies were all significantly different (one-way ANOVAs, all $p<0.01$). Within each frequency band, the three connectivity measures' ICCs were also all significantly different (one-way ANOVAs, all $p<0.01$).

However, it is difficult to determine the specificity of the relationship between ICC stability and each connectivity measure using this method of comparison. We next investigate the properties of this relationship.

Relationship between reliability and strength of connectivity

We were interested in the relationship between each pair of regions' connectivity strength and ICC of that connection over the three resting state sessions, rather than the independent distributions of the two values. This provides insights into the interaction between the magnitude of connectivity and the consistency of connectivity over time in a regionally specific context. The relationship between connectivity and ICC stability for a given region pair varies both with frequency and connectivity measure.

To probe this relationship, we plotted the ICC stability for a specific functional region pair (Figure 3.2) against the mean PLV (Figure 3.4A-F, top row), amplitude correlation (Figure 3.4G-L, middle row), and coherence (Figure 3.4M-R, bottom row) of that region pair, across all electrode pairs and all three resting state sessions. As described, based on ECoG coverage across subjects and previously published criteria, we had binned all electrodes into one of 10 functional brain regions. In all subplots of Figure 3.4, each individual marker represents the overall stability and mean of connectivity strength between a given pair of functional brain regions (functional regions shown in Figure 3.2). Correlation coefficients are shown at the top of each subplot.

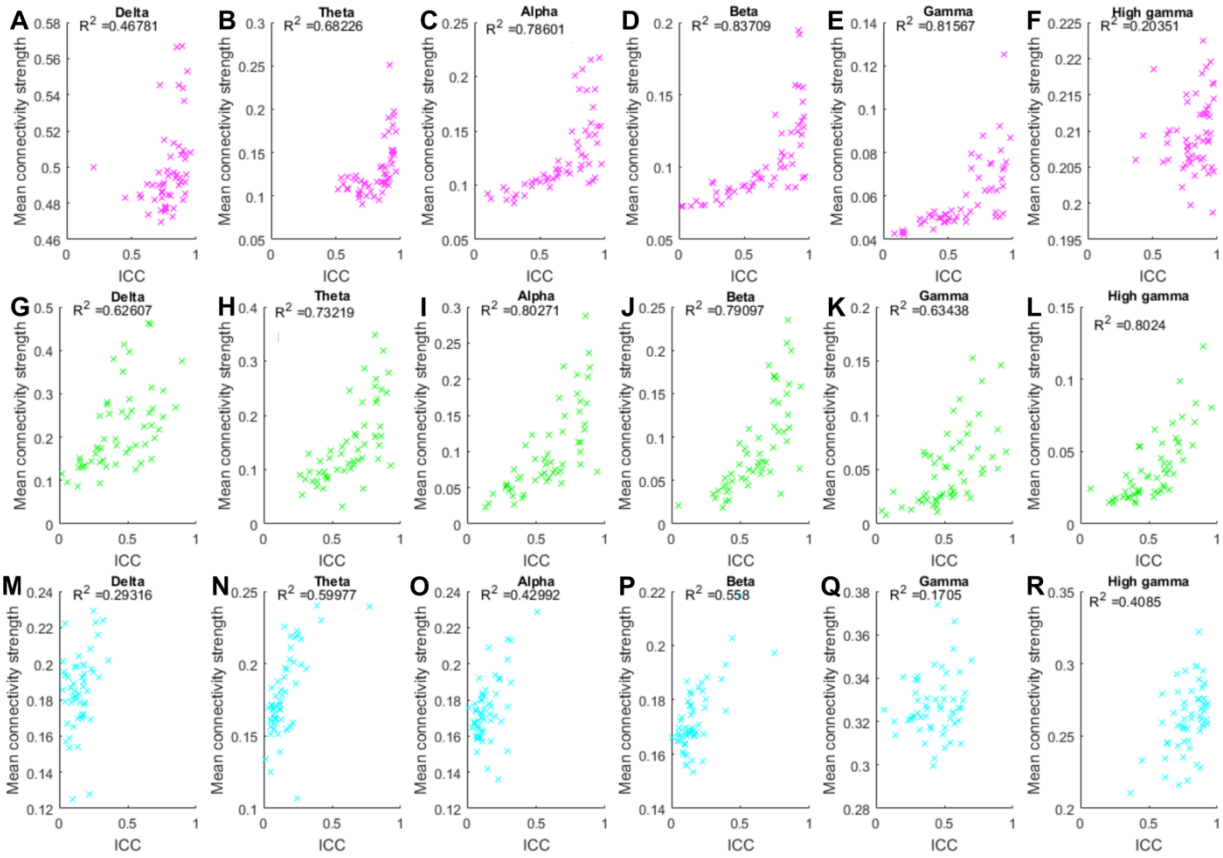

Figure 3.4
Scatter plots of the relationship between ICC (X axis) and mean connectivity (Y axis, each point represents mean connectivity between all pairs of electrodes in that pair of functional regions at all three time points). Top row (A-F), magenta: connectivity measure shown with its own ICC calculations is PLV. Strong PLV is overwhelmingly associated with stability (high ICC). However, high ICC is not always associated with strong PLV (i.e., many weak but stable connections). Middle row (G-L), green: amplitude correlation. The relationship between ICC and amplitude correlation is qualitatively similar to PLV, but trend is less defined, and there are high ICC/low amplitude correlation values. Bottom row (M-R), cyan: coherence. Note the minimal relationship between ICC and coherence.

We first examined the relationship between ICC and PLV. The R^2 Spearman's correlation between ICC and PLV is especially strong in theta, alpha, beta, and gamma (Figure 3.4B-E, R^2 range 0.68-0.83), moderate in delta (Figure 3.4A, R^2=0.46), and weak in HG (Figure 4F, R^2=0.20). In bands with lower R^2 correlation coefficient (delta, HG) this effect appears to be driven at least in part by the limited range of ICCs. Regardless of the R^2, strong PLV (Figure 3.4A-F) uniformly predicts a more stable connection (i.e., ICC near 1), as represented by the cluster of values in the top right of the plot and lack of values in the top left (i.e., no low ICCs associated with high PLVs). This relationship is maintained across all frequency bands despite the differences in the range of PLVs between bands. However, high ICCs are not always associated with high PLVs: there are many connections between functional regions occupying the bottom right of the graphs, indicating a stable and weak connection. In delta, theta, and high

gamma (Figure 3.4A, B, F), there are few connections with low ICC, but in alpha, beta, and gamma (Figure 3.4C-E), low ICCs are all associated with weak PLVs.

We examined the same relationship for amplitude correlation (Figure 3.4G-L, middle row) and coherence (Figure 3.4M-R, bottom row). In delta and gamma in particular, high amplitude correlation is not uniformly associated with high ICC as it is with PLV, though the trend is still visible (i.e., there are some high ICCs associated with lower amplitude correlation values, Figure 3.4G and K). There is little variation between frequency bands in the relationship between ICC and amplitude correlation (R^2 range 0.62-0.80 across all frequencies). The absolute range of amplitude correlation values varies more and is wider within each frequency than in PLV. Unlike PLV, which is associated with temporally specific phase synchronization, amplitude correlation captures co-variation that is less temporally specific. Consequently, high amplitude correlation may arise from a wider range of neural circumstances than the specific conditions leading to synchronization.

In contrast to both PLV and amplitude correlation, high coherence does not predict high ICC (Figure 3.4M-R). The R^2 values only rise to the middle of the range seen with the other connectivity measures (maximum R^2=0.59, theta). Where high ICCs do appear, far less frequently than they do in PLV and coherence and mainly in delta and high gamma, the strongest coherence values are not linked to the highest ICCs (Figure 3.4M and R). The weak predictive relationship may be partially due to the limited range of ICCs in coherence relative to the other connectivity measures and partially due to basic physiology: coherence is sensitive to both the phase and amplitude variation of the complex signal components, and not only one or the other.

Linking reliability and strength to brain anatomy

We illustrate regional variation of functional connectivity and ICC stability between each of the same ten individual regions shown in the scatter plots above. Now we highlight specific anatomical differences, rather than the relationship between functional connectivity and ICC directly. The prior sections detail broad trends that apply to these results, but these spatial patterns are the key to the overall connectivity.

We focus further in-depth analysis on PLV (Figure 3.5), because the relationship between connectivity and ICC shows most variation between anatomical regions of the three connectivity measures (results from amplitude correlation and coherence are shown in Supplementary Figures 3.1 and 3.2). We calculated pairwise PLV between all of the electrodes in each subject for the ten functional brain regions shown, and repeated this in each of the six

frequency bands of interest. All electrode pairs' PLVs (red) and ICC with 95% confidence interval (blue) are shown (Figure 3.5).

Because there are 53 regional pairwise connections (with sufficient number of subjects and electrodes for analysis, of 55 possible), we will focus our discussion on selected regions of interest (Figure 3.5). We observe disparate connectivity and stability patterns of parahippocampal/entorhinal gyri (PH-ER, highlighted in dark green), inferior parietal lobule (IPL, highlighted in burgundy), and dorsolateral prefrontal cortex (dlPFC, highlighted in teal). We selected these regions both for their roles in attention and memory consolidation, functions associated with resting state, and for their notable connectivity properties. For each brain region pair and frequency range combination, we can observe the ICC stability estimate, median of PLVs, and distribution of PLVs.

Two broad patterns and one main exception emerge in the region pairs' ICC estimates (Figure 3.5 inset). In Pattern A, representing 18 region pairs, there is relatively high ICC throughout the frequency range (Figure 3.5, e.g., dlPFC/dlPFC, PH-ER/PH-ER, S1/IPL). Notably, despite Pattern A describing only 34% of all region pairs, fully half of these are all nine connections within a single region (the tenth, M1/M1, was excluded due to too few electrode pairs). In Pattern B, representing 34 region pairs, there is wide variation in ICC across the frequency range that is strongest in delta, and often also stronger in theta or HG (Figure 3.5, e.g., PH-ER/STG, IPL/IFG). The principal exception is PH-ER and IPL (Figure 3.5, neither pattern), where delta ICC is unusually low. This may be influenced at least in part by the small number of electrode pairs between this pair of regions (N=14).

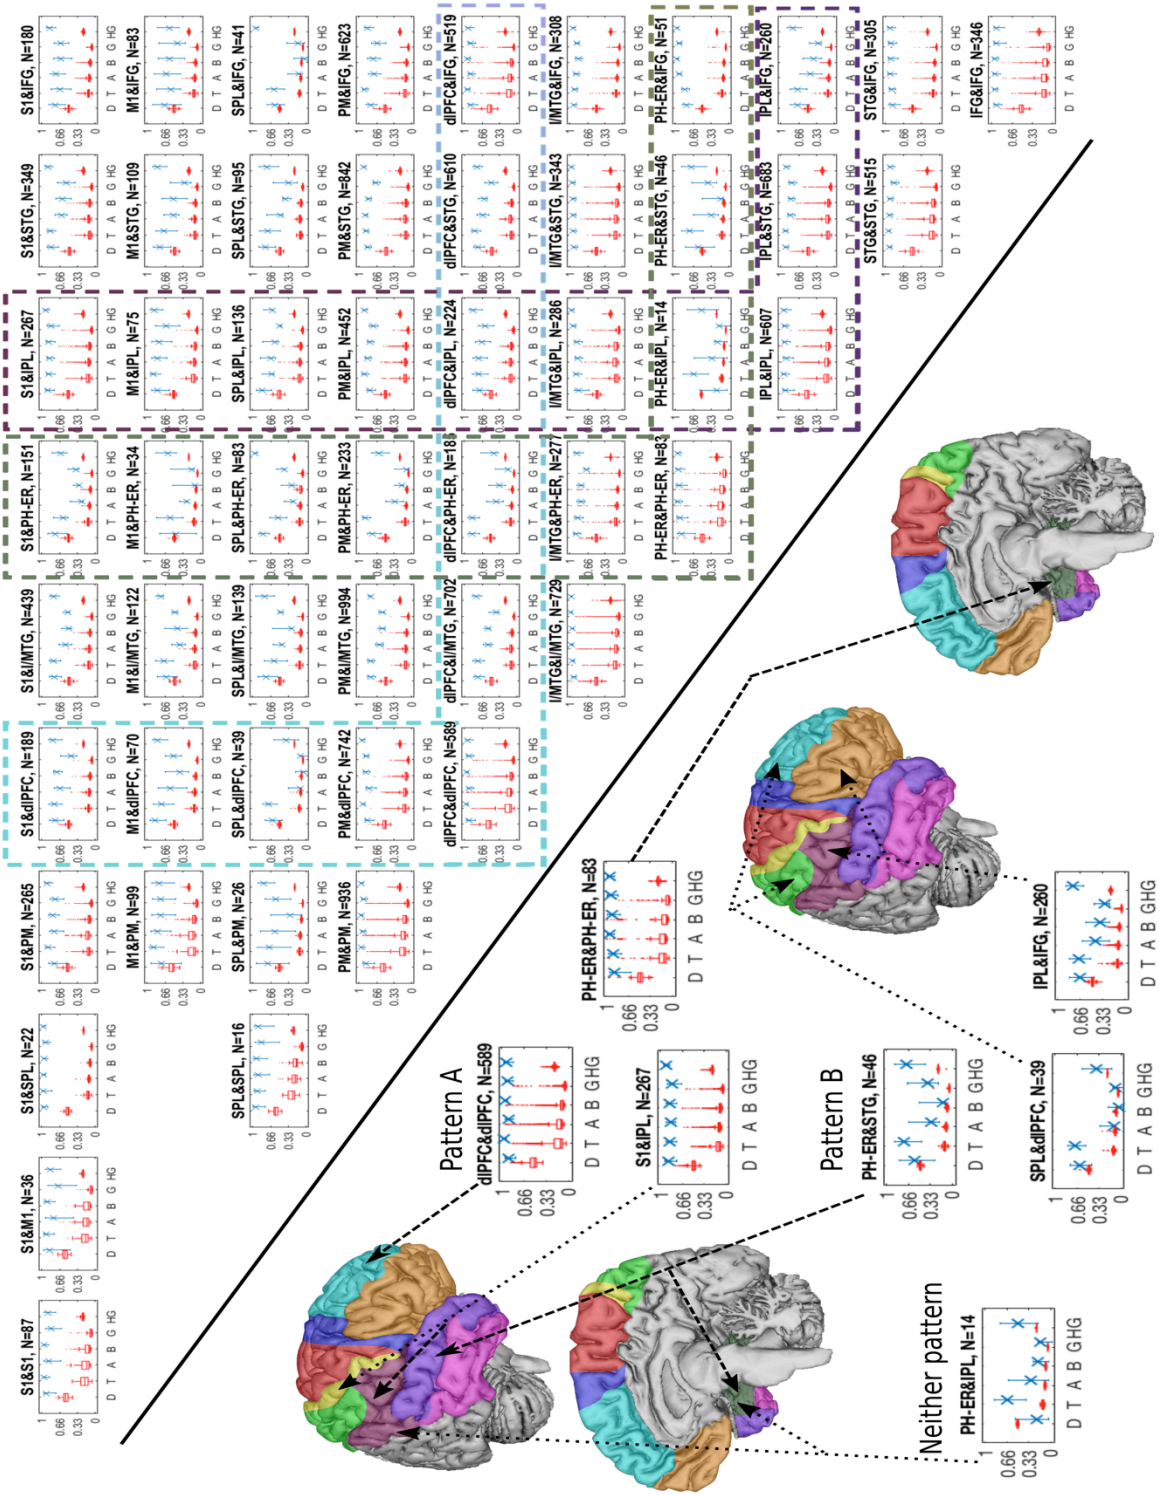

Figure 3.5
Pairwise PLV between specific pairs of functional brain regions. Across the entire figure, each subplot represents the PLV between the pair of regions listed in the subplot title (abbreviations in Table 3.1). Within each subplot, from left to right are the six frequency bands of interest (delta, theta, alpha, beta, gamma, and high gamma) marked along the horizontal axis. For each frequency band, the red box plot represents the real PLV of all pairs of electrodes (from all subjects) within that pair of functional regions,

including all three time points. Number of electrode pairs is value N listed in subplot title. The blue X and bars represent the ICC and 95% confidence interval of the ICC, calculated from the data shown in the box plot. Because PLV is a symmetric and non-directional connectivity measure, only the upper diagonal is shown; lower diagonal represents identical data. We highlight connectivity patterns associated with three regions of interest: dorsolateral prefrontal cortex (dlPFC, teal box), parahippocampal-entorhinal gyrus (PH-ER, dark green box), and inferior parietal lobule (IPL, burgundy box). These regions demonstrate properties of stable functional isolation, disconnection of the intuitive relationship between ICC and connectivity strength, and unstable weak connectivity, respectively. The inset (left) shows two distinct patterns of stability that appear in different brain regions. In some regions (Pattern A) stability is high across all frequencies. In others (Pattern B) it is high in the delta and HG range but dips in the middle frequencies. The exception is PH-ER/IPL connectivity, which is relatively unstable across all frequencies.

The connectivity patterns of parahippocampal/entorhinal gyri (PH-ER) may indicate relative resting state functional isolation and/or potentially increased local functional heterogeneity (Figure 3.5, dark green boxes; Walhovd et al., 2016). Mean resting state connectivity is generally fairly low. However, across the different anatomical regions, the ICC of these weak connections varies substantially (as noted with 4.), from low stability/weak PLV in all frequencies between PH-ER and inferior parietal lobe, to high stability/weak PLV in all frequencies between PH-ER and inferior frontal gyrus and within PH-ER. In contrast, there is widespread elevated phase synchrony from PH-ER to other regions, and correspondingly elevated stability, in delta. PLV in all other frequency bands only rises above 0.25 in self-synchrony and synchrony with inferior/middle temporal gyrus (which it borders). Interestingly, the distribution of PLVs of PH-ER to itself and to inferior/middle temporal gyrus is very wide compared to PH-ER's connectivity to other regions as well as being very stable. This indicates that there is substantial heterogeneity in the pairwise electrode connectivity within these region pairs, but that each electrode pair is relatively stable. This may reflect analogous heterogeneity in function within PH-ER, but stability in that diverse coupling over time.

The connectivity of inferior parietal lobule (IPL) illustrates how ICC stability and PLV connectivity values may be non-intuitively linked (Figure 3.5, burgundy boxes). While it may seem counterintuitive that the regional pairs with low stability can also have narrow PLV ranges, this indicates that within that narrow range, a given electrode pair will experience significant variation between the three time points. In contrast, for those region pairs with high stability but a wide range of connectivity values, any single electrode pair's PLV varies little, though there is much variation between different pairs. The distribution of PLVs here is notably narrow despite its instability – as ICC represents the collective stability as an aggregate of each individual value's variability, this is possible, if unintuitive. Notably, PH-ER and inferior frontal gyrus also have extremely low ICC relative to all other regions' connections to IPL, and ICC of the connection to superior parietal lobule is somewhat reduced relative to other regions. Curiously,

much as delta stands out from the other frequency bands for its elevated PLVs, its range is also relatively compressed relative to the other bands in these regions.

Finally, the connectivity of dorsolateral prefrontal cortex (dlPFC) provides an example of a region with both generally weak connectivity and low stability (Figure 3.5, teal boxes). Across all frequencies, connectivity of dlPFC is only highly stable to itself and to neighboring premotor cortex and inferior frontal gyrus. As described above, the relationship between the height of PLV distribution and ICC can vary dramatically; among these three regions of high ICC, PLV to IFG and within dlPFC are very wide with low centers, while PLV to premotor is somewhat narrower, but also with low centers. The narrowest distributions, however, are associated in this case with regions with unstable/low ICC connectivity to dlPFC, especially PH-ER and superior parietal lobule. The relatively low ICCs between dlPFC and most other regions across multiple frequencies may reflects the diverse range of cognitive functions associated with the region; that is, dynamic changes in functional connectivity patterns in the resting state (Madhyastha et al., 2014).

3.5 Discussion

In this study, we investigated intra-individual, inter-session variability in spontaneous resting state functional connectivity patterns across ten different functional cortical regions over a wide frequency range (up to 200Hz). To our knowledge, this is the first such investigation of spontaneous intra-individual, inter-session change in resting state connectivity of electrocorticographic cortical potentials.

Overall, the data illustrated in Figure 3.5, taken in conjunction with the alternate views of the same data presented in Figures 3.3 and 3.4, demonstrate several key properties of the relationship between the stability of connectivity measured with ICC, the overall strength of connectivity, and the distribution of individual electrode pairs' connectivity within a broader context.

Most importantly, we examined the relationship between connectivity strength, distribution of connectivity values, and ICC specifically for PLV. The relationship between the width of distribution and stability is complex and sometimes non-intuitive, but taken together the two properties together illuminate a broader scope of the fluctuating connectivity of the region (Figure 3.5). We especially wish to highlight properties such as stable functional isolation (illustrated by weak but stable connectivity of parahippocampal-entorhinal cortex), the unintuitive link between connectivity strength and stability (illustrated by the narrow range but instability of inferior parietal lobule), and the intersection of weak connectivity and stability (illustrated by the consistently weak connections of dorsolateral PFC). In combining connectivity and stability with

anatomy, low ICCs may be linked to frequently-changing allegiances, both between regions and locally, and a wider range of connectivity values within a single region pair may indicate a wide range of shared and/or overlapping cognitive functionality and therefore functional connectivity needs.

If one of the functions of synchrony is to transfer information within and between brain regions, a more heterogeneous brain region would potentially have lower stability in its connectivity patterns. Across the 53 region pairs, ICC almost exclusively follows either a pattern of being consistently high in the six frequency bands (Pattern A), or of being highest in delta and often also theta and HG, with a dip in the middle frequency bands (Pattern B, Figure 3.5 inset). In particular, though Pattern A describes just 34% off all connections, it describes all nine within-region connections, and the majority of the rest involve primary motor or sensory cortex. In contrast, the lower-stability Pattern B encompasses the majority of connections, and includes all the included connections associated with the fronto-parietal network and default mode network. The one exception was the relationship between parahippocampal and entorhinal cortex (PH-ER) and inferior parietal lobule (IPL), where delta ICC was unusually low, though as the number of included electrodes was small (N=14) we cannot draw strong conclusions.

In our broader examinations of the data (Figure 3.3 and 3.4), differences between the connectivity measures may be linked to the various characteristic properties underlying each of the three: PLV with temporal synchronization (Greenblatt et al., 2012; Sauseng & Klimesch, 2008), zero-lag amplitude correlation with sustained co-activation (Sauseng & Klimesch, 2008), and coherence with a combination or interdependence of the two (Bastos & Schoffelen, 2016).

We observed that PLVs were generally stronger and more stable than amplitude correlation and coherence values, which is consistent with other electrophysiological investigations of both single-session and variation in resting state connectivity (Casimo et al., 2016; Garces et al., 2016). ECoG is less impacted by volume conduction than other methods because of its immediate proximity to the cortical signal source (Casimo et al., 2016). As such, the higher PLVs we observed relative to amplitude correlation and coherence are not likely attributable entirely to volume conduction, and may reflect specific temporal synchronization between electrodes (measurable with PLV) rather than sustained co-activation of the regions in question (measurable with correlation) (Sauseng & Klimesch, 2008). Connectivity strength in all three measures is frequency dependent, with the highest connectivity in all three in delta, but also surprisingly high in HG, given that the increased number of cycles relative to slower frequency bands likely augments the variability in HG synchrony.

We also observed that coherence in general was notably unstable, and also relatively low, across the entire brain. We postulate that this may be because PLV and amplitude correlation each only directly depend on one component of the complex-valued representation of the neural signal (phase and amplitude, respectively) but coherence is calculated on the complete signal, and therefore is sensitive to potentially related linear variation in phase and amplitude when taken together. This may also contribute to the suppressed range of both coherence and ICCs of coherence, which are on the whole narrower in range than either amplitude correlation or PLV at any frequency.

We found a strong pattern across all frequencies of high PLV being exclusively associated with high ICC, but high ICCs (high stability) could be associated with a wide range of PLVs. This is consistent with Garces's findings in MEG (Garces et al., 2016). By extension, low ICCs were also exclusively associated with low PLVs. However, this relationship was weakened in amplitude correlation, and nonexistent in coherence. We again postulate that for variation in coherence this may be related to the connectivity measure depending on both the phase and amplitude components of the signals. The presence of some low ICC but high amplitude correlation region pairs may be related to more generalized co-variation in activity patterns that does not depend on the specific timing synchronization underlying with PLV.

Overall, combining the distribution-based view (Figure 3.3) and the relationship between ICC and connectivity (Figure 3.4) with anatomy (Figure 3.5) provides a complex, comprehensive view of spontaneously fluctuating connectivity over time. Our observations of the relationship between functional connectivity and ICC across functional regions are consistent with, and greatly extend in the frequency domain, prior work in EEG and MEG (Brookes et al., 2014; de Pasquale et al., 2010; Espenhahn et al., 2016; Garces et al., 2016). Of these, our methodology is most similar to Garces and colleagues, who tracked spontaneous variation in resting state measurements using MEG (Garces et al., 2016). Our findings parallel these previous efforts, including observations of differences between connectivity measures and the differences in spontaneous variation across anatomical regions, and extend their work, with our unique insight into parahippocampal/entorhinal regions not easily accessible with MEG and into higher frequencies.

This work extends findings on the stability of resting state connectivity previously investigated with fMRI and MEG. ECoG studies are novel as electrode positions are fixed and are rarely moved once implanted for the relatively long-term monitoring period (one week). Consequently, ECoG provides greater spatial resolution and higher fidelity of the full human cortical frequency spectrum than is available to other modalities while ensuring reliable and

stable recordings from a consistent area of cortex for the whole recording period. Although ECoG has limitations in its application and subject pool that do not similarly constrain other methodologies, our results contribute unique insight to that broader context of research, especially with regard to higher frequencies and phase synchrony. In particular, we were able to demonstrate a different pattern of connectivity and stability in PH regions compared to the more stable IPL patterns and widely varying dlPFC connectivity.

Limitations and further work

The principal limitations of this study are related to the nature of ECoG studies. Because ECoG is currently used in humans exclusively for clinical purposes, our subjects are affected by medically intractable epilepsy. Though we eliminated the electrodes over the seizure focus to mitigate the effect of abnormal cortex, we cannot rule out all effects. Further, because the placement and extent of the ECoG coverage is determined entirely by clinical needs, there are some brain regions (notably, the medially located cingulate cortex) with no coverage at all, and others (visual cortex) with too little coverage for the data to be usable (only one subject out of eight, with <20 electrode pairs). The adoption of stereoEEG may increase the ability to study some of these regions. Finally, we have a relatively small pool of eight subjects, which is still typical for ECoG studies (Casimo et al., 2016; Foster & Parvizi, 2012; C. Keller et al., 2013). We did not identify any relationship between the number of electrodes in a region pair and the patterns of ICC and/or connectivity strength, though we cannot completely rule out a sample size effect.

Further research could include linking the electrophysiological patterns we demonstrate here to the known variation in resting state connectivity as measured by fMRI within a single individual, that is, to qualitatively determine whether the variability in connectivity we observe here is the same in both modalities. To our knowledge, while both analyses have been performed separately, electrocorticographic and BOLD variability have never been assessed in the same subject pool. Further work could also investigate the relationship between intra-individual and inter-session variability.

In addition to further research on spontaneous variation in resting state connectivity, our findings contribute to describing a baseline or background level of variation relative to task-related variation. Specifically, the estimated stability of resting state connectivity in our data can be used to investigate the effect of task-related perturbations on resting state connectivity that are above or below spontaneous or “normal” levels, or to investigate differences in connectivity patterns between individuals.

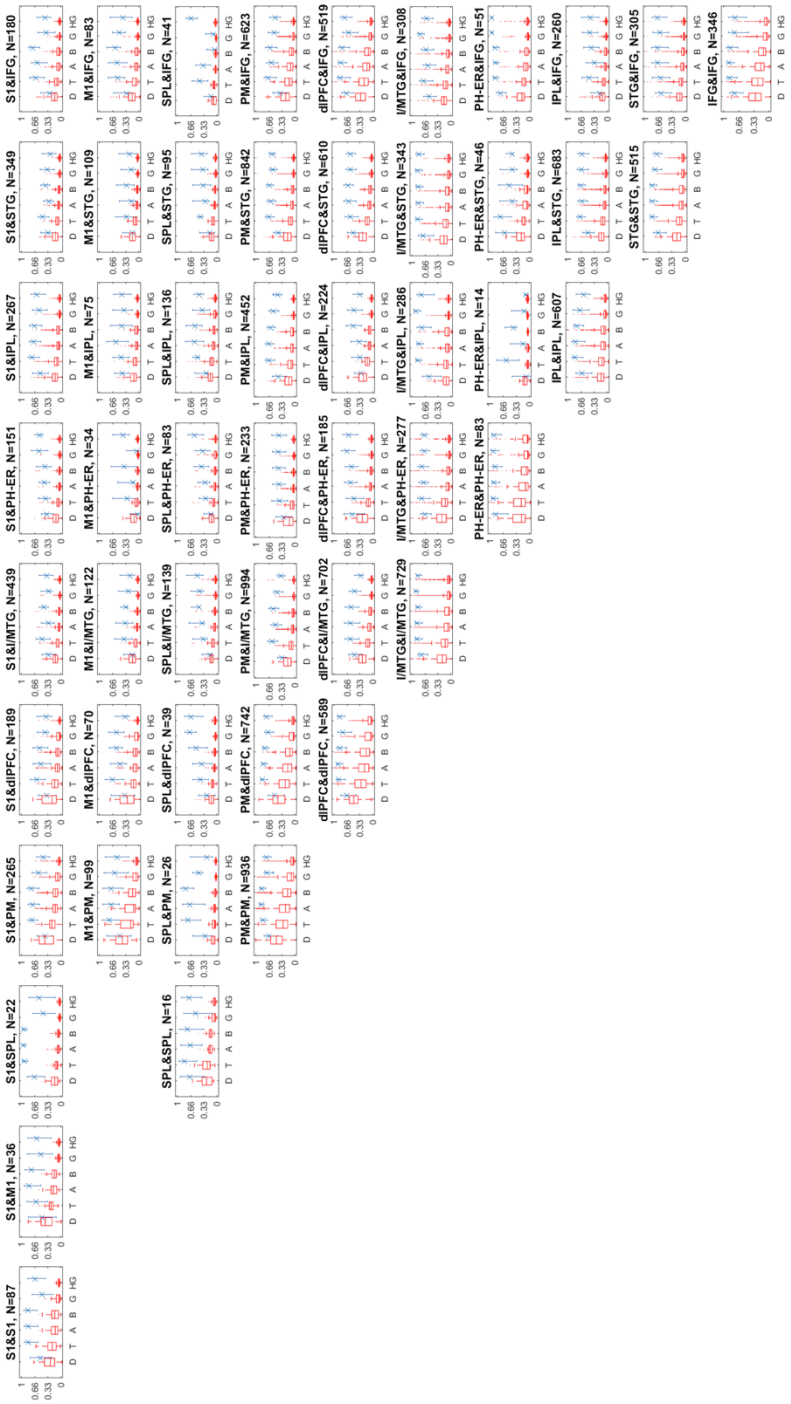

Supplementary Figure 2.1 Pairwise amplitude correlation between specific pairs of functional brain regions. Across the entire figure, each subplot represents the amplitude correlation between the pair of regions listed in the subplot title (abbreviations in Table 1). Within each subplot, from left to right are the six frequency bands of interest (delta, theta, alpha, beta, gamma, and high gamma) marked along the horizontal axis. For each frequency band, the red box plot represents the real amplitude correlation of all pairs of electrodes (from all subjects) within that pair of functional regions, including all three time points. Number of electrode pairs is value N listed in subplot title. The blue X and bars represent the ICC and 95% confidence interval of the ICC, calculated from the data shown in the box plot. Because correlation is a symmetric and non-directional connectivity measure, only the upper diagonal is shown; lower diagonal represents identical data.

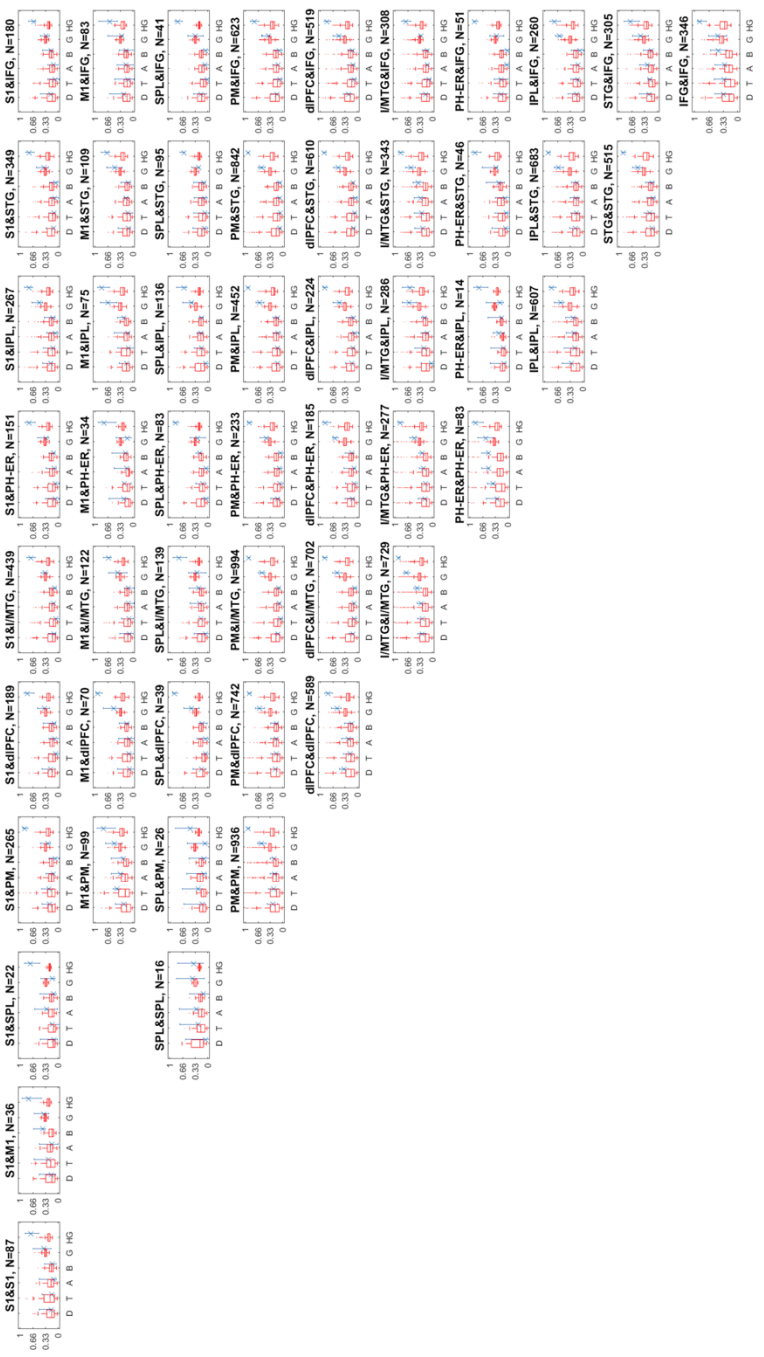

Supplementary Figure 3.2
Pairwise coherence between specific pairs of functional brain regions. Across the entire figure, each subplot represents the coherence between the pair of regions listed in the subplot title (abbreviations in Table 1). Within each subplot, from left to right are the six frequency bands of interest (delta, theta, alpha, beta, gamma, and high gamma) marked along the horizontal axis. For each frequency band, the red box plot represents the real coherence of all pairs of electrodes (from all subjects) within that pair of functional regions, including all three time points. Number of electrode pairs is value N listed in subplot title. The blue X and bars represent the ICC and 95% confidence interval of the ICC, calculated from the data shown in the box plot. Because coherence is a symmetric and non-directional connectivity measure, only the upper diagonal is shown; lower diagonal represents identical data.

4. Estimating change in brain connectivity attributable to learning

4.1 Motivation and role in broader project

In this third paper I addressed my original hypothesis of learning-related change in resting state connectivity, as I had built up the methods and underlying properties of the resting state in isolation necessary to do so. I used the same connectivity methods in the following paper as I did in the previous paper, and largely the same statistics. The major conceptual addition was using the spontaneous change in resting state data set to establish a baseline distribution of variation as a statistical contrast to compare the motor task data set against, thereby determining what level of variation exceeded random variation occurring in the absence of explicit procedural learning. I also examined the relationship between pre-task connectivity (that is, the connectivity state when the subjects were first approaching the task) and degree of improvement on the task, to identify whether there were any pairs of regions (or within-region connections) whose strength was predictive of how well they would learn.

I separated the subjects by their skill level on the task (improved vs. did not improve) and identified changes in connectivity strength between different cortical regions that were significantly different from the spontaneous levels of change I identified in my previous work. I found changes in connectivity that could be specifically attributable to the skill across multiple cortical regions. The connectivity changes attributed to attempting to learn the skill were especially differentiated between the skill level groups in inferior parietal lobule across all three connectivity measures, and were also different between the groups in parahippocampal/entorhinal gyrus in phase locking value, dorsolateral prefrontal cortex in amplitude correlation, and primary motor cortex in coherence. This paper built directly on the findings and methodologies of my prior work and established a method for robustly isolating learning-relating connectivity changes from those associated with mere exposure to the task and from those that arise spontaneously.

The following chapter will be submitted for publication as:

Casimo, K., Wu, J., Madhyastha, T., Rao, R.P.N., Ojemann, J.G., Weaver, K.E. Persistent impacts on resting state functional connectivity following learning of a skill. Neuroimage.

4.2 Introduction to section

The brain changes as part of the process of learning a new skill. This learning-related plasticity is encoded at the neuronal level as changes to synaptic strength or even the presence of new synapses (Anderson et al., 1996; Dayan & Cohen, 2011; Xu et al., 2009). Such alterations are likely associated with density changes in gray (Driemeyer et al., 2008; Taubert et al., 2011) and white matter (Scholz et al., 2009) that have been observed following extended practice of a new skill. These changes to brain structure underlie changes to patterns of activity and functional connectivity, including resting state functional connectivity.

The behavioral resting state and resting state functional connectivity are associated with memory consolidation processes, general attentiveness or readiness to respond to stimuli, and self-referential thought (Schölvinck et al., 2013). It varies both spontaneously (Allen et al., 2014; Casimo et al., n.d.; Garces et al., 2016) and in association with various tasks, including learning a new skill. Changes in brain connectivity have been observed during the initial acquisition phase of a new motor task (Bassett et al., 2011). Further, using fMRI, changes in resting state functional (Albert et al., 2009; Kelly & Garavan, 2005; Vahdat et al., 2011) and effective (Buchel et al., 1999) connectivity linked to learning a motor task encompass functionally related regions across the cortex. The magnitude of change in connectivity is correlated with the degree of task improvement (Buchel et al., 1999; Vahdat et al., 2011). These changes were specific to learning the skill and the associated improvement in performance, and not with merely performing the task at a consistent, learned level. Further, TMS disruption of M1 resting state activity after procedural learning reduces subsequent performance and disrupts of connectivity of M1 (Censor et al., 2014; Kraeutner et al., 2016). Collectively, these results suggest that specific connectivity patterns are essential for successful task learning and execution.

Neuroimaging investigations tracking the impact of procedural learning on brain activation and functional connectivity converge on a set of regions of particular interest: inferior frontal gyrus, premotor and primary motor areas, superior parietal lobule, inferior parietal lobule including angular gyrus, S1, and various subcortical regions, especially hippocampus (Albert et al., 2009; Buchel et al., 1999; Hardwick et al., 2013; Vahdat et al., 2011). The inferior parietal lobule has been particularly and consistently highlighted for its involvement in visuospatial learning and consolidation (Crottaz-Herbette et al., 2014; Igelström & Graziano, 2017; Kraeutner et al., 2016; Shum et al., 2011). In addition to these regions' association with general visuospatial tasks, mental rotation, which is necessary in our novel motor task, has been specifically linked to M1, angular gyrus (part of inferior parietal lobule), subiculum (near

parahippocampal and entorhinal cortex) and various subcortical regions (Kim et al., 2012; Luaute et al., 2009; J. Zacks et al., 1999; J. M. Zacks et al., 2002).

The majority of prior functional connectivity studies tracking the effects of procedural learning on resting state connectivity have been performed in fMRI. While contributing valuable insights, fMRI is limited as an indirect measure of neural activity, and so can indicate regions of interest but cannot resolve specific frequency-dependent effects. Studies using non-invasive scalp-based electrophysiological methods identify effects in the same regions identified in fMRI studies, but further observe frequency-dependent patterns of learning-related change (Houweling et al., 2008; Mehrkanoon et al., 2016; Wu et al., 2014). Connectivity of M1 has been specifically identified as not only affected by a motor learning task, but pre-task connectivity patterns have been identified as predictive of skill learning (Wu et al., 2014). Non-invasive electrophysiology methods, however, are limited by relatively poor spatial resolution restricted generally to the cortical surface, but do provide high temporal resolution in the low end of the human cortical frequency spectrum.

Here, we report results tracking resting connectivity before and after procedural learning using electrocorticography (ECoG). Subdural ECoG recordings provide the same temporal advantages as surface electrophysiology, but additionally provide superior fidelity in the faster gamma and high gamma frequency ranges (>70 Hz). ECoG also affords spatial resolution superior to other electrophysiology methods, but at the cost of limited spatial coverage of only partial sampling of cortex as well as and the requirement of sampling from clinical subjects. Despite these restrictions, ECoG provides an important adjunct to fMRI and other electrophysiology from its temporal precision and spatial accuracy.

We examined persistent changes in resting state functional connectivity between various pairs of brain regions following learning a skill, results which were significantly different from spontaneous resting state variation. Eight subjects completed a novel motor task that also required mental rotation, four of whom successfully improved on the task. We identified differences in the changes in pairwise functional connectivity following task execution between subjects who learned the task and subjects who did not. We also examined whether pre-task connectivity strength across various pairwise connections predicted their eventual skill level on the task.

4.3 Methods

We recruited subjects with medically intractable epilepsy who were undergoing surgery at the Harborview Medical Center Regional Epilepsy Center. Each subject had only left (N = 4) or right (N = 4) hemisphere electrodes (Figure 4.1A and B). The patients were hospitalized for long-term seizure monitoring with implanted ECoG electrodes for pre-surgical planning. They provided informed consent in accordance with the University of Washington Institutional Review Board. The placement of the patients' electrodes was determined exclusively by clinical need. Each subject was implanted with between 64-128 subdural platinum ECoG electrodes (Ad-Tech, Racine, WI, or Integra Lifesciences, Plainsboro, NJ; electrode surface diameter 2.3mm, inter-electrode spacing 1cm).

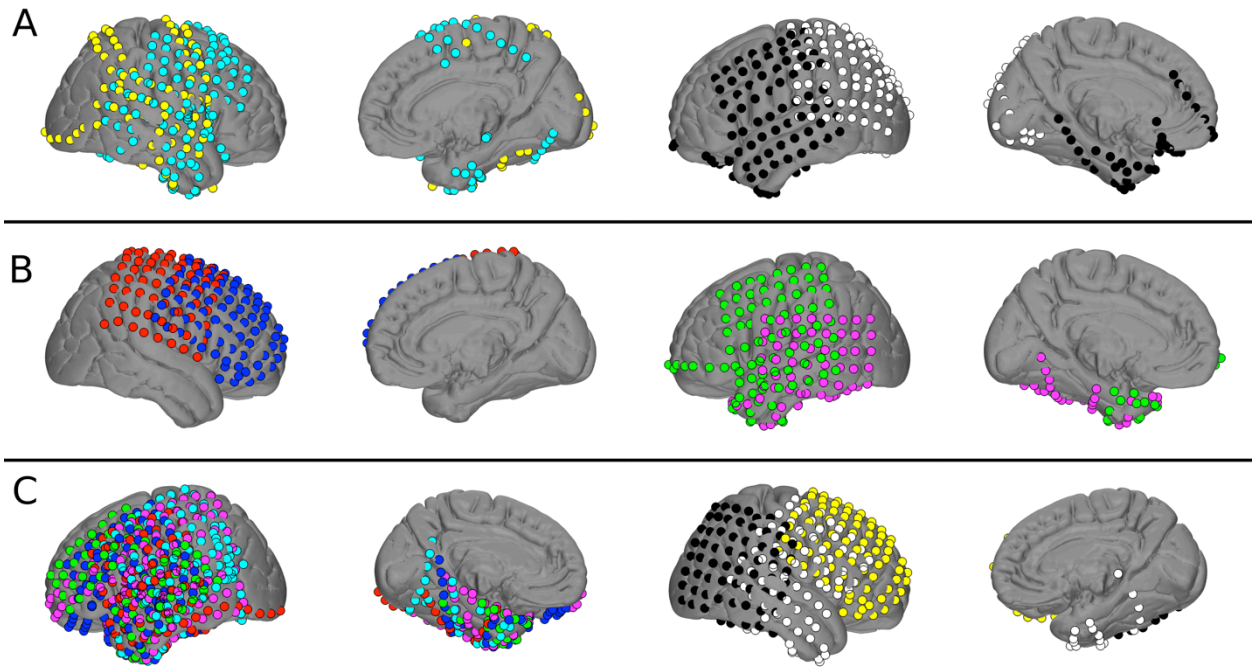

Figure 4.1
A: electrode positions for four subjects who successfully learned task, defined as a reduction in path length traveled to complete the task. B: electrode positions for four subjects who did not successfully learn task, defined as no change or increase in path length traveled to complete the task. C: electrode positions for 8 control subjects who did not complete the task.

Behavior

Participants completed a novel, computerized motor learning task. Subjects were asked to navigate a cursor through a simple maze using a track pad and collect twelve targets, where each target is randomly generated within a bounded area along the maze (Figure 4.2A). The maze was then virtually rotated so the perceived controls were offset by 90˚ (Figure 4.2B). The

original, starting side of the maze was always indicated by a green label independent of rotation angle offset. Completion of the rotated maze requires the subjects to comprehend and execute the rotated controls. The subjects were asked to complete the task using the hand contralateral to their ECoG electrodes (which were always on only one hemisphere), regardless of dominant hand.

Subjects completed between 3-7 blocks of the 90˚-rotated version of the task consisting of 12 targets each. Behavioral contrasts from the original, training orientation and the 90˚ rotation condition served as the primary metric of learning. Additional block conditions where the maze was rotated 45˚ or 180˚ were excluded from analysis, as all subjects were able to successfully perform the task with these conditions.

We calculated task performance for each block as the difference between the optimal path length and the executed path length between the twelve targets. The overall task performance is represented as the slope of the task performance over all runs of the 90˚ task. Thus, improvement on the task is represented by a negative performance score (that is, the total deviation from the optimal path length decreased with practice). All further analyses were performed separately on the four subjects whose task performance improved (Figure 4.1A, “learners”) and the four subjects whose task performance did not improve (Figure 4.1B, “non-learners”). The four non-learners all indicated that they understood what they were being asked to do even though they could not successfully execute it.

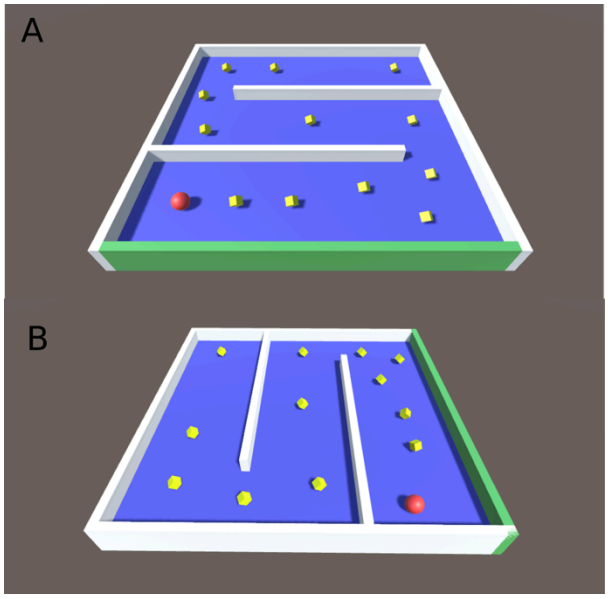

Figure 4.2
A: basic version of the behavioral task. Subjects navigate the red ball using a computer track pad to collect the yellow targets. B: rotated version of the behavioral task, which requires subjects to learn the rotated controls. The controls are oriented as though the subject is still located on the side of the maze with the green wall (as they were in the basic version), such that “up” on the screen is actually “right” on

the controls. Accuracy is scored as the total path length traveled minus the minimum path length. Improvement score is calculated as the slope of the accuracy scores over all runs through the maze, so a negative score represents greater improvement.

Resting state data was acquired before and after motor task performance. Each session was six minutes in length. Subjects were asked to lie quietly in bed and avoid moving, talking, or falling asleep. Patients were told to allow their minds to wander for the duration of each resting state period.

Data acquisition and processing

All ECoG recordings were obtained with a Tucker-Davis Technology system (Alachua, FL) at 1000Hz. All sessions were at least 90 minutes before or after any clinically notated subclinical or clinically manifest ictal activity.

All ECoG and behavioral data analyses were performed in Matlab (Mathworks, Natick, MA). Preprocessing included eliminating electrodes clinically identified as being over the seizure focus, applying a common average re-reference across each grid and strip of electrodes to reject common noise, applying a notch filter to remove line noise at 60Hz and its harmonics, and visually confirming the absence of artifacts. For the full duration of each resting state session recording, we extracted amplitude and phase angle of the signal using a non-analytic Morlet wavelet, with a ¼ octave resolution for pseudo-frequencies 1-200Hz. We then averaged the phases and amplitudes for the frequencies falling into the six frequency bands of interest (delta, 0.1-4Hz; theta, 4-8Hz; alpha, 8-12Hz; beta, 13-30Hz; gamma, 30-70Hz; HG, 70-200Hz). The following connectivity estimates were calculated from these data.

Functional connectivity estimates

Numerous electrophysiological approaches exist to quantify functional connectivity. Because no one connectivity measure can capture all of the temporal and frequency dynamics inherent to complex neural signals, we applied three, classic linear connectivity measures. These were phase locking value (PLV), amplitude correlation, and coherence (Garces et al., 2016; Greenblatt et al., 2012; J. Lachaux et al., 1999). These measures are based on the specific interactions of the phase component only, the amplitude component only, and the two features combined, respectively. PLV provides an estimate of phase-based temporal synchronization. Amplitude correlation (Pearson's correlation on the amplitude component of the signals only) provides an estimate of co-variation in the overall signal envelope, that is, coincident increases and decreases in activity level. Coherence captures the content shared between the two signals in both phase and amplitude, but loses the feature specificity in the other two measures.

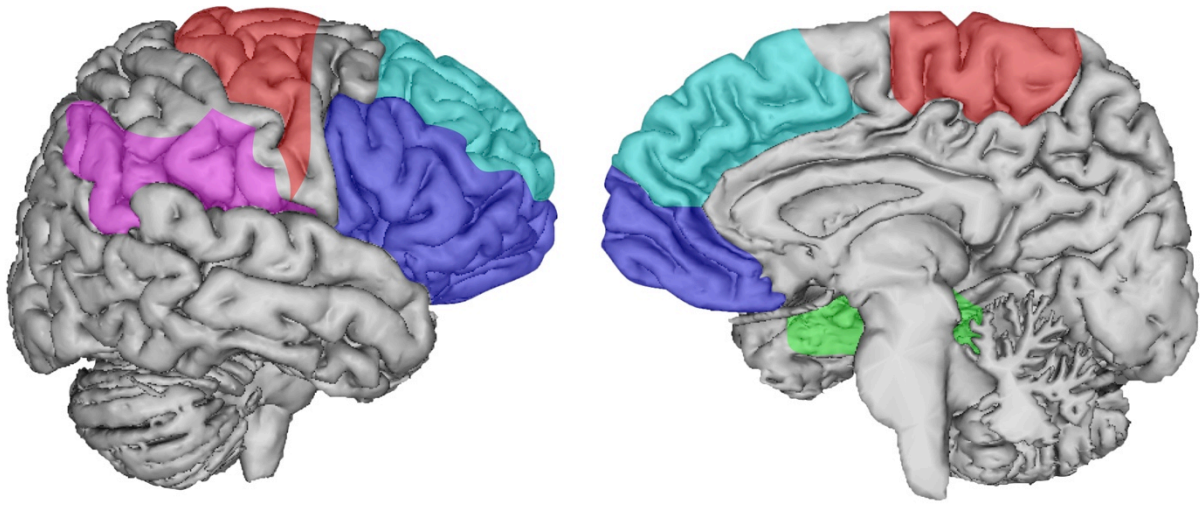

Figure 4.3
Outlines of functional brain regions used for regional pairwise connectivity analyses. List of Brodmann areas included in each functional region are shown in Table 1. Colors: teal, dorsolateral prefrontal cortex; pink, inferior parietal lobule; green, parahippocampal and entorhinal gyri (PH-ER); blue, inferior parietal lobule (IPL).

We calculated pairwise functional connectivity between all electrode pairs in each subject. We estimated pairwise connectivity across the full resting state period for all of the pairs of electrodes located in the four functional zones of interest (i.e., 10 possible pairwise connections between functional zones). After aggregating across subjects in each learning category, this ranged between 16-775 pairs of electrodes in a given pair of regions.

Region name	Abbreviation	BAs included
Primary motor cortex	M1	4
Dorsolateral prefrontal cortex	dlPFC	9, 10, 46
Parahippocampal and entorhinal gyri	PH-ER	27, 28, 34, 35, 36, 37
Inferior parietal lobule	IPL	39, 40, 43
Inferior frontal gyrus	IFG	11, 44, 45, 47

Table 4.1
Functional brain zones. We identified and automatically sorted all electrodes to Brodmann areas. We then grouped those Brodmann areas into four functional zones of interest. All functional zones had electrodes from at least two subjects in each learning category.

Electrode localization and aggregation

We rendered and labeled the positions of ECoG electrodes as previously described (Blakely, Miller, Zanos, Rao, & Ojemann, 2009; Casimo et al., 2016). We first identified the location of electrodes in native space based on a preoperative MRI and postoperative CT scans, and then transformed those locations into standardized space (1 mm MNI 152 brain). We then assigned Brodmann area (BA) labels to each electrode with FSL atlas tools and Talairach daemon atlas. We focused on five functional zones derived from Brodmann areas and previously published criteria for electrode region assignment (Casimo et al., 2016). These

regions were selected for their involvement in cognitive functions related to skill learning: executive (dorsolateral prefrontal, inferior frontal gyrus), sensory integration and mental rotation (inferior parietal lobule), motor execution (M1), and memory (parahippocampal/entorhinal gyri). These regions both 1) had data from at least 2 out of 4 subjects in each behavioral category (i.e., improved at task vs. did not improve at task), and 2) were relevant to the execution of the task and/or consolidation of procedural memories (Table 4.1, Figure 4.3).

For each participant, we used functional zone labels to sort electrodes in a pairwise manner (e.g., for a pair with one electrode in inferior frontal gyrus and one in dorsolateral prefrontal cortex, the pair would be labeled IFG/dlPFC). We then used these labels to aggregate electrodes across all subjects in each learning level (learners vs. non-learners) and between left and right hemispheres.

Statistics: intraclass correlations

To characterize stability of functional connectivity across resting state sessions, we computed the intraclass correlation (ICC; McGraw & Wong, 1996). ICC is used to estimate the level of variability in a group of subjects at different measurement times. In this case, the variability is in connectivity strength between pairs of electrodes, and the time points are before and after the maze task. We used the A-1, or one-way random effects, model of ICC. The calculation is based on two-dimensional data, in which are the electrode pairs' functional connectivity and the two resting state sessions. We calculated the 95% confidence interval of the ICC value for each pair of regions of interest, in each of the six frequency bands of interest, and in each of the three connectivity measures. Because ICC is calculated across all electrode pairs for a given region pair/frequency/connectivity measure combination, each ICC value represents tens or hundreds of electrode pairs.

To estimate statistical significance of task-related ICC values, we applied the same methodology to calculate the three connectivity estimates on the same five regions of interest pairs, and their corresponding ICCs, on a separate group of eight subjects (Table 4.2). These control subjects were undergoing the same clinical monitoring procedures (ECoG), but did not perform the task (Figure 4.1C; Casimo et al., n.d.). For each of these control subjects, three resting state recordings were acquired as described above. We used the longitudinal resting state recordings and ICC estimates of difference in connectivity in this 'control' subject pool to characterize baseline levels of spontaneous variation that are expected in the absence of any task (Casimo et al., n.d.). That is, we generated a distribution of baseline variation of ECoG connectivity, and compared task-related variation to this baseline. Any change in resting state

connectivity that arose after task engagement and that fell outside of this range was considered statistically significant and specifically task-linked.

	Learned task	Did not learn task	Control
N	4	4	8
Age	Mean = 32; range 21-42	Mean = 28; range 20-39	
Left hemisphere electrodes	2	2	5
N female	2	3	4

Table 4.2
Comparison between subjects who performed learning task and control subjects.

For each connectivity measure, region pair, and frequency combination, we compared the ICC value associated with learning-related variation to the confidence interval of the spontaneous variation and retained only the learning-related ICCs (and corresponding real connectivity values) whose confidence intervals did not overlap with the confidence interval of spontaneous change.

We did not use resting state sessions from the pool of subjects who performed the task as a control simply because we were unable to collect a sufficient number of resting state sessions from all eight of them during the period before they were first exposed to the task (Table 4.2). However, the baseline longitudinal variability in the control group was established over a longer time period (two days) and more sessions (three) than the two resting state sessions used in our learning set (two sessions, immediately before and immediately after task on the same day). This makes the statistical test more conservative, as the potential range of spontaneous change is broader than it would be in the shorter, fewer-samples learning set. The ICCs and connectivity estimates shown in Figures 4.4-6 are only those where the ICC associated with learning-related change was outside of the confidence interval of ICC associated with spontaneous change (values shown in Casimo et al., n.d.).

In the learners, we also directly compared pre-task and post-task connectivity for those connections that became more unstable than the control group's baseline level of variation (Figure 4.7). This comparison illustrates not only how our estimate of variation (ICC) in resting state connectivity is affected by learning, but also directly compares the actual connectivity values. For these selected region pair/frequency combinations, we performed a Wilcoxon signed rank test (α=0.01, FDR corrected) on the pre-task and post-task resting connectivity values.

Finally, to investigate the link between pre-task connectivity and task performance, we compared task improvement score to pre-task connectivity for each region pair/frequency combination using both a linear regression and a group comparison approach. We first

calculated the R^2 value on the relationship between task improvement score (0 for subjects who did not improve) and pre-task connectivity and its associated p value. We performed this analysis on just the four subjects who successfully learned the task, and also on all eight subjects. We also performed a Wilcoxon signed rank test (α=0.01, FDR corrected) to compare connectivity scores of the four learner subjects and four non-learner subjects, again in each region pair/frequency combination for each of the three connectivity measures.

4.4 Results

We identified variation in resting state connectivity following learning of our novel motor task that was significantly more stable or significantly more unstable than spontaneous levels of change. We compare differences in this task-linked variation between the four subjects who improved at (i.e., "learners") the task vs. the four subjects who did not improve at the task ("non-learners") (Table 4.3). We examine differences in ICC and connectivity between the learner and non-learner groups for each of the three connectivity measures separately. All of the following results are based on changes in connectivity that fall outside of the range of spontaneous inter-session variation in connectivity, and thus can be attributed specifically to the execution (successful or not) of the task.

Subject number	Electrode side	Performance score	Learning category
1	L	>0	Non-learner
2	R	>0	Non-learner
3	R	>0	Non-learner
4	L	>0	Non-learner
5	L	-343	Learner
6	R	-81	Learner
7	R	-696	Learner
8	L	-414	Learner

Table 4.3
Behavioral outcomes for each subject, the number of twelve-target blocks of the 90° task condition (where learning occurred), and the improvement score (slope of the total deviation from the ideal path length in each task block; larger numbers indicate greater improvement). Subjects were sorted into "learner" or "non-learner" on the basis of improvement score. All subjects used the hand contralateral to their electrode placement to complete the task.

Connections that become more unstable after the task had greater variance from before to after the task than the typical non-task-linked levels of variability. They changed more than would be expected in the absence of the task attempt. They can be described as reflecting critical post-task consolidation functions that changed substantially as a result of attempting the task (Albert, Robertson, & Miall, 2009).

Connections that became more stable after the task are less intuitive to conceptualize. They had less variance from before to after the task than typical non-task-linked levels of variability. Statistically significant increases in ICC suggest connections that are still impacted by the task, but they stayed closer to their pre-task connectivity levels than would have been expected in the absence of any intervention. They can be thought of as locked into pre-task connectivity patterns, such that their variation does not even rise to the level of variation seen in the absence of an outside stimulus.

One region pair (PH-ER/IPL) was excluded from analysis because there were fewer than 10 electrode pairs between the two regions in the learning subject group.

Task-related impacts on phase locking value

We first examined impacts on resting state phase locking value (PLV) following task execution between learner and non-learner groups (Figure 4.4). PLV measures consistency of phase relationships between signals from different regions and can be interpreted as the level of synchrony between regions (J. Lachaux et al., 1999).

Counting both across region pairs and across frequencies (total of 90 connections), in the learning subjects 8 connections were more unstable (i.e. statistically significant reduction in ICC), 8 were more stable (i.e., significant increase in ICC), 68 were not different from spontaneous levels of variability, and 6 were excluded due to <10 electrodes in the region pair. In the non-learning subjects, 15 were more unstable, 8 were more stable, and 67 were not different from spontaneous levels of variation. The only region pair where the impacts on connectivity attributable to the task were similar between the learner and non-learner groups was parahippocampal-entorhinal gyri/inferior frontal gyrus (PH-ER/IFG).

Overall, the changes in PLV from pre-task to post-task that are significant relative to the control group occur largely within the same region pairs in the learners and non-learners, but in different frequency bands. In particular, connectivity involving IPL is particularly different between the learners and non-learners.

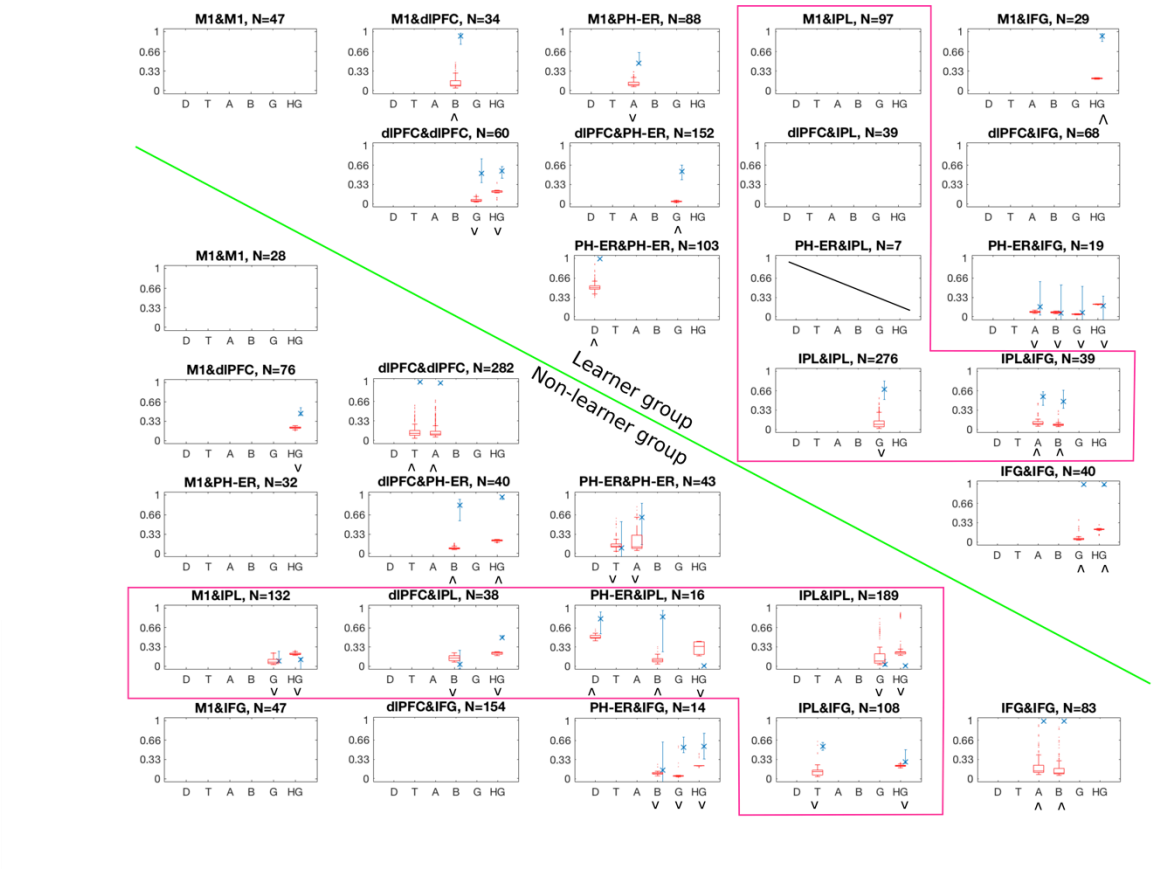

Figure 4.4
Pairwise ICCs and phase locking values (PLV) between the four functional regions of interest. Upper diagonal, four subjects who successfully learned tasks (electrode positions shown in Figure 4.1A). Lower diagonal, four subjects who did not learn task (electrode positions shown in Figure 4.1B). Across the figure, each subplot represents the connectivity between the pair of regions listed in the subplot title (abbreviations in Table 4.1). In each subplot, frequency bands are listed from left to right (delta, theta, alpha, beta, gamma, high gamma). The red box plot shows both pre and post connectivity values for all electrode pairs in that pair of regions, including values before and after the task. The blue x and bars represent the ICC and 95% confidence interval of ICC calculated from the data in the red box plot. Values shown are those that are significantly different from spontaneous or baseline levels of variation. Marks under each bar/box plot represent whether that connection was more stable (∧, higher ICC) or more unstable (∨, lower ICC; see Figure 4.7 for more details) than spontaneous ICC value.

Connectivity involving the inferior parietal lobule particularly distinguished the learner and non-learner groups. Connections involving IPL that varied as a result of the task appeared in both learners and non-learners, but did not overlap between the groups in in which frequencies' PLV was affected by learning and whether they became more stable or unstable. No particular region's connectivity with IPL emerged as especially different between learners and non-learners; the pattern that emerged was diffuse but distinctive on the basis of learning level.

Notably, we observed only a few, scattered resting state connectivity values involving M1 were impacted by task execution or successful learning. Moreover, no dominant frequency response within M1 emerged, despite the fact that M1 involvement is critical to executing this motor task and motor learning (Albert, Robertson, & Miall, 2009). Connections involving dlPFC and IFG were impacted in both groups, especially in lower frequencies for dlPFC and higher frequencies in IFG, but these effects were few enough that no systematic pattern distinguishing the learning groups emerged.

Finally, in the learner group, for the 8 region pair/frequency combinations that became more unstable than spontaneous change, we separated the pre-task and post-task connectivity values to more precisely identify the changes to PLV associated with learning (Figure 4.7A). Of these 8 comparisons, IPL/IPL PLV in gamma and M1/PH-ER in alpha were significantly different when comparing pre-task and post-task connectivity (Wilcoxon signed rank test, $\alpha=0.01$, FDR corrected). This indicates that not only were the levels of change in these two region pairs/frequencies (i.e., ICCs) significantly different from the control group's spontaneous level of change, the difference itself was also statistically significant ($p<0.01$). The other 6 comparisons that were more unstable than spontaneous levels of change, but not significantly different from pre-task to post-task when compared directly, still reflect statistically significant task-related variability that is different from spontaneous variation

Task-related impacts on amplitude correlation

We examined impacts on amplitude correlation following task execution between learner and non-learner groups (Figure 4.5). Amplitude correlation measures co-activation of pairs of signals from different regions and can be interpreted as co-varying activity but not specifically temporal synchrony (Greenblatt et al., 2012). Again, connectivity involving IPL particularly distinguishes the learners and non-learners. Overall, the task-related effects on resting state connectivity are far more widespread in space and frequency bands in non-learners than in learners.

Counting both across region pairs and across frequencies (total of 90 connections), in the learning subjects 7 resting connections were more unstable, 3 were more stable, 64 were not different from spontaneous levels of variability, and 6 were excluded due to <10 electrodes in the region pair. In the non-learning subjects, 19 were more unstable, 0 were more stable, and 61 were not different from spontaneous levels of change. Amplitude correlation varied far less beyond baseline change than PLV. Of note, no affected connections in non-learners became more stable, and few did in the learners. Notably, almost all of the resting state connections

impacted by attempting the task (in both learners or non-learners) became more unstable in the process, regardless of skill improvement level.

The impacts on connectivity attributable to the task were similar between the learner and non-learner groups in dlPFC/dlPFC, dlPFC/IPL, and dlPFC/IFG. In all three cases, connectivity in only one frequency was affected (in these three region pairs, the single frequency affected was HG, gamma, and HG respectively), indicating that local neural firing rates in these region pairs varied beyond baseline as a result of attempting the task at all. Additionally, the task-changes in connectivity IFG/PH-ER in the learners represented a subset of those in the non-learners.

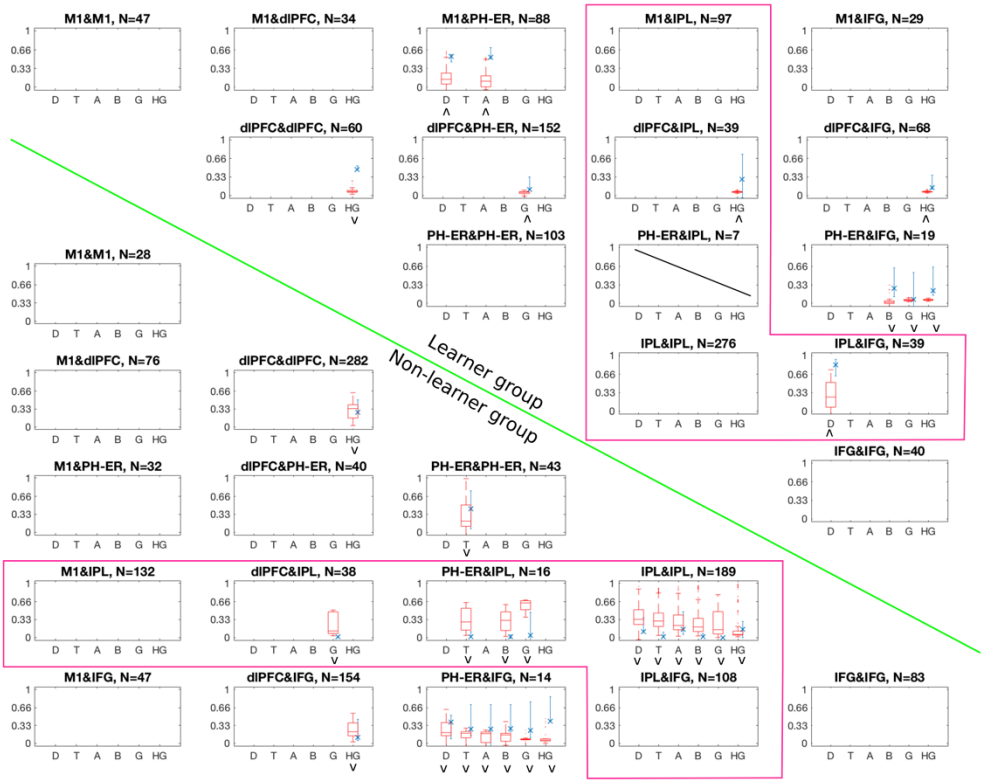

Figure 4.5
Pairwise ICCs and amplitude correlation values between the four functional regions of interest. Upper diagonal, four subjects who successfully learned tasks (electrode positions shown in Figure 4.1A). Lower diagonal, four subjects who did not learn task (electrode positions shown in Figure 4.1B). Subplot layout, colors, values, and labels are same as Figure 4.4.

As in PLV, changes to connectivity involving IPL particularly distinguished learners and non-learners. In particular, IPL/IPL connectivity became more unstable than spontaneous variation in all frequencies in non-learners, and did not vary relative to spontaneous variation in

learners. This suggests that the non-learners' IPL/IPL amplitude correlation changes fail to support learning the task.

In contrast to the limited change in connectivity beyond baseline in the learner group, resting connectivity in the non-learner group varied more broadly across region pairs and frequencies. In the non-learners, connectivity also exclusively became more unstable than baseline variation. This suggests inappropriately high levels of change in co-varying regional activation as measured by amplitude correlation are linked to poorer task performance.

Finally, in the learner group, for the 3 region pair/frequency combinations that became more unstable than spontaneous change, we separated the pre-task and post-task connectivity values to more precisely identify the changes to amplitude correlation associated with learning (Figure 4.7B). None of these comparisons were significantly different (Wilcoxon signed rank test, α=0.01, FDR corrected). However, these levels of change (ICC values) were still statistically significantly different from spontaneous levels of change.

Task-related impacts on coherence

Finally, we examined impacts on resting state coherence following task execution between learner and non-learner groups (Figure 4.6). Coherence is calculated from both the phase and amplitude of a signal and represents the shared content between the signals (Sakkalis, 2011). Overall both groups are broadly impacted by learning, but the spatial and frequency patterns of those impacts in M1 connections are especially different between groups.

Counting both across region pairs and across frequencies (total of 90 connections), in the learning subjects 5 connections were more unstable, 27 were more stable, 52 were not different from spontaneous levels of variability, and 6 were excluded due to <10 electrodes in the region pair. In the non-learning subjects, 8 connections were more unstable, 15 were more stable, and 67 were not different from spontaneous levels of change. In both the learning and non-learning groups, a much higher proportion of those connections that were impacted by the task became more stable in coherence than in the other measures.

There were far more task-related changes in connectivity in both the learner and non-learner groups in coherence than in the other two measures, making it more difficult to distinguish between them. However, the only region pair where the impacts on connectivity attributable to the task were similar between the learner and non-learner groups was dlPFC/PH-ER. This executive-memory link may be related to consolidation of the experience of the task regardless of performance.

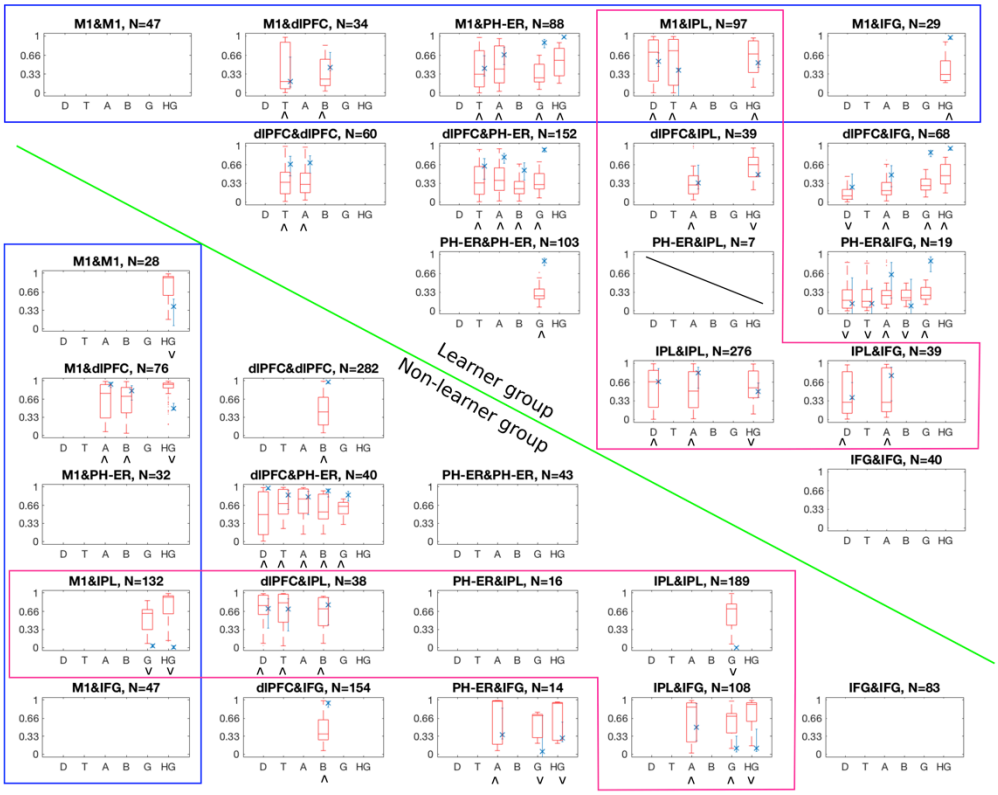

Figure 4.6
Pairwise ICCs and coherence values between the four functional regions of interest. Upper diagonal, four subjects who successfully learned tasks (electrode positions shown in Figure 4.1A). Lower diagonal, four subjects who did not learn task (electrode positions shown in Figure 4.1B). Subplot layout, colors, values, and labels are same as Figure 4.4.

In the learner group, four region pairs stood out as having connections with impacts spread over the majority of frequency bands: M1/PH-ER, dlPFC/PH-ER, dlPFC/IFG, and PH-ER/IFG. This suggests task-related changes to executive function (dlPFC/IFG) and executive function/memory interactions (other three pairs) require broadband involvement to successfully learning the task.

While M1 connectivity was only minimally impacted by task exposure in the other two measures, it stands out here as broadly, and differently, impacted in both learning groups. Especially notable are the broadband impacts on M1/PH-ER connectivity (i.e., motor and memory interactions) in the learners not seen in the non-learners. Non-learner impacts were restricted to high frequencies, while statistically significant effects extended into low frequencies in most of the M1 interactions in learners.

Finally, there were substantial inconsistencies between the three connectivity measures in the regions and frequency bands affected by the task. In the learner group, all three connectivity measures were affected by learning the task in alpha of M1/PH-ER and gamma of dlPFC/PH-ER. In the non-learners, all three measures were affected in gamma of IPL/IPL. For PH-ER/IFG connectivity in all three measures, effects were in beta and gamma in learners, and gamma and HG in non-learners.

Finally, in the learner group, for the 6 region pair/frequency combinations that became more unstable than spontaneous change, we separated the pre-task and post-task connectivity values to more precisely identify the changes to amplitude correlation associated with learning (Figure 4.7C). Of these 6 region pair/frequency combinations, only IPL/IPL in HG was significantly different when comparing pre-task and post-task connectivity (Wilcoxon signed rank test, $\alpha=0.01$, FDR corrected). The other 5 comparisons where the direct comparison was not significant still reflect statistically significant task-related variability that is different from spontaneous variation.

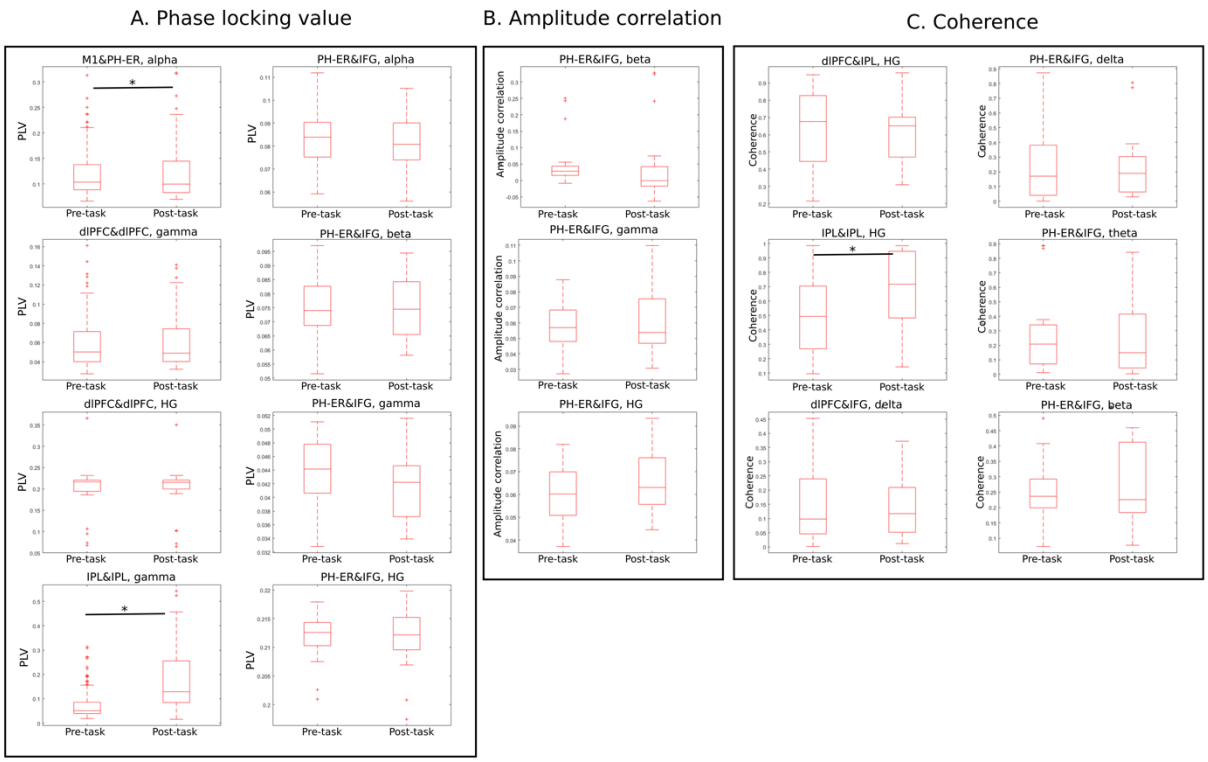

Figure 4.7
Box plots separate pre-task and post-task connectivity values in learners for those connections that became more unstable relative to spontaneous levels of variability (i.e., there was a greater difference between pre- and post- task connectivity than there would have been in the absence of the task). Each section corresponds one connectivity measure (A: PLV, B: amplitude correlation, C: coherence), and each subplot to one region pair and frequency that was marked with a ∨ in Figure 4, 5, or 6. Connections whose connectivity values became more stable are not shown, because their higher ICC values mean that there is little difference between pre- and post-task connectivity values. Distributions of pre-task to post-task connectivity values of all electrodes in each region pair shown were compared using a Wilcoxon signed rank test (α=0.01, FDR corrected); significantly different comparisons are marked with a star. Where comparisons are not significantly different, note that the ICC value indicating level of stability was still significantly different from the control group.

Task performance and pre-task connectivity

Finally, we examined the relationship between level of task improvement and pre-task connectivity, to identify whether connectivity patterns before task exposure mediate the level of improvement at the task. We examined the relationship between pre-task connectivity both for just the four subjects who improved at the task, as well as for all eight subjects. No pre-task connectivity values in any region pair or frequency were predictive of task improvement, across either the four learner subjects or all eight subjects (all $p>0.05$).

In addition to this treatment of task improvement as a continuous value, we also used task improvement scores to divide the pre-task connectivity values into the groups of learners

and non-learners as we did above. We compared pre-task connectivity between the learners and non-learners for all pairs of regions, in each frequency band, in each of the three connectivity measures using Wilcoxon signed rank tests; none were significantly different (all $p>0.05$).

4.5 Discussion

We investigated persistent electrocorticographic changes in resting state functional connectivity after a motor learning task, and those signal features that distinguished task improvement from simply engaging in the task but not improving. We used three connectivity measures to capture broad features of the connectivity patterns, which differ between phase locking value (timing and synchrony), amplitude correlation (co-activation), and coherence (shared content). Prior studies have indicated that resting state functional connectivity varies spontaneously in the absence of an outside stimulus structured task engagement, and that these stability patterns vary across brain regions (Casimo et al., n.d.). We compared the stability of selected functional connections before and after attempting our motor task to the stability of multiple sessions in the absence of any regular goal-driven task.

We report systematic differences of task learning on resting state functional connectivity variation patterns that are specifically linked to a measurable degree of task improvement. In all three connectivity measures we investigated, we were able to identify specific region pairs and frequency bands that distinguished task learners from non-learners and which were not otherwise attributable to endogenous inter-session variability. Our approach identified connections involving the IPL in all three connectivity measures, and M1 in coherence. In agreement with previous studies (Crottaz-Herbette et al., 2014; Igelström & Graziano, 2017; Kraeutner et al., 2016; Shum et al., 2011) the IPL played a particularly prominent role in the resting state motor learning consolidation process (Figures 4.4-6).

In PLV and amplitude correlation, we observed that a greater proportion of functional connections became more unstable than became more stable after task execution, but in coherence a greater proportion of connections became more stable than more unstable. Task-related changes in resting connectivity were not consistent between the three measures and differed between the learners and non-learners in each of the three connectivity measures. This can be at least partly attributed to the fact that the three measures reveal connectivity properties based on different combinations of phase (synchrony), amplitude (co-activation), or both.

We interpret connections that became more unstable (i.e., functional connectivity value changed more following the task than it would have spontaneously) as representing critical post-task functions within relevant neural networks operating differently than they did before

incorporating the task. We interpret connections that became more stable (i.e., connectivity value after task execution was more similar to pre-task connectivity than the level of spontaneous change would have been) as being locked into pre-task connectivity patterns. That is, engagement with the task caused the persistence of pre-task connectivity patterns that otherwise would have varied.

A functional connection becoming more or less stable is not necessarily an estimate of its adaptive value critical for task improvement. It is conceivable that under our statistical approach increased stability after task execution (that is, changing less relative to normal change in the absence of a task) is a signature of task learning (or execution) impediment, reflecting a failure of resting state consolidation. We interpret such connections (that become more stable) as representing solidification of resting state connectivity or an active consolidation process not present before the task. Functional connections becoming more stable were especially prevalent in coherence, and the discrepancy between measures may be related to the relationship between phase and amplitude captured by coherence that is not detected in the other two measures.

Differences between the groups in the connectivity impacts of successful or unsuccessful task learning may reflect the cognitive state associated with their skill level. Changes in connectivity patterns relative to baseline variability seen only in the learning group, such as broadband changes in of M1/PH-ER in coherence, may reflect consolidation, cognitive strategies or involvement associated with understanding and improving at the task. Changes seen only in the non-learning group, such as all frequencies of IPL/IPL in amplitude correlation, may reflect the consolidation of the ineffective or counterproductive strategies that hindered task performance, reflecting possible involvement of different brain regions and/or functional connectivity patterns than the successful strategies. Changes seen in both groups reflect engagement with or attention to the task, the consolidation the simple motor demands of using the track pad (separate from the choice of movements), or other cognitive functions necessary for the consolidation of the task experience regardless of skill level.

The role of inferior parietal lobule in resting state consolidation of motor learning

The neuroanatomy of resting state consolidation is reflected in these patterns, especially in prevalence of IPL-related connectivity as a point distinguishing learners from non-learners. IPL has been previously linked to the execution and consolidation of visuospatial and visuomotor tasks similar to this maze task (Albert et al., 2009; Buchel et al., 1999; Crottaz-Herbette et al., 2014; Hardwick et al., 2013; Igelström & Graziano, 2017; Kraeutner et al., 2016; Luaute et al., 2009; Shum et al., 2011). Specifically, it is thought to be involved in visual and

motor integration required for procedural learning, as indicated by electrophysiology (Houweling et al., 2008; Wu et al., 2014) and imaging (Chapman et al., 2010; Crottaz-Herbette et al., 2014; Luaute et al., 2009) studies in humans, as well as by the effects of stimulation interfering with IPL function reducing consolidation of skills (Censor et al., 2014; Kraeutner et al., 2016; Shum et al., 2011). Our findings suggest that the learning-related changes in resting state functional connectivity we observed reflect offline consolidation of this newly acquired complex visual and motor skill. Indeed, a future experiment could evaluate the correspondence between this plastic consolidation period and offline behavioral gains measurable in a second session of the task.

High frequencies

We further indicate that high frequency connectivity patterns (gamma and high gamma), especially in interactions involving IPL, were a key feature of the task-related resting state connectivity changes we observed. In PLV and amplitude correlation, over half of the ICC values that were significantly different from control variation were associated with gamma or high gamma connectivity patterns; in coherence, just under half of values were associated with those bands. As high frequency interactions are generally associated with local function, the high prevalence significant synchrony or co-activation in those bands across longer distances suggests a level of local-to-local functional coordination in the consolidation of the task.

Further, when we compared pre-task and post-task connectivity values for the connections in learners that became more unstable with learning (i.e., learning-induced variation), two of the three comparisons that were significantly different were in gamma and high gamma, and both were also in IPL/IPL connectivity (Figure 4.7). Our finding that intra-IPL connectivity not only differs significantly from spontaneous change as assessed by ICC, but also differs when directly compared between pre-task and post-task values, further underscores its importance. Our unique insights derived from the use of ECoG demonstrate a robust involvement of these high frequencies in the consolidation of a learned skill.

Predicting skill level from pre-task connectivity

In contrast to prior work (Wu et al., 2014), we observed no significant relationship between pre-task connectivity and task improvement in any of the three connectivity measures, in any region pair, or in any frequency, when task improvement score was treated as a continuous variable. Further, when the subjects were categorically grouped based on their task improvement score, there were still no differences between the groups as delineated by any of those categories. This lack of a response may be due to poor sensitivity of the task outcome measures to neurophysiological changes, or insufficient electrode coverage. Further, the observed null relationships may reflect differences in post-task connectivity shifts that

distinguish the learners and non-learners, effects that cannot be attributed to differences between the groups in their starting connectivity. Instead, the post-task changes relative to baseline result from their learning, or not learning, the task. The pre-task connectivity patterns do not explain a predisposition to learning or not learning the task, though other non-connectivity brain dynamics may.

Limitations and further studies

This study, as in all studies using ECoG in humans, was limited by the spatial placement and extent of coverage of electrodes. The electrode placements were determined entirely by the patients' clinical needs. Our coverage of the parahippocampal and entorhinal regions was the most limited of the regions we closely examined, though over half of the subjects had electrodes there, restricting the generalizability of those findings. We examined pairwise resting state connectivity outside of the context of canonical resting state networks. Due to insufficient electrode coverage in critical regions (e.g., only one subject had any cingulate electrodes, which is a key node of the default mode network), we were not able to examine large scale canonical RSNs. We also had no data from non-cortical regions inaccessible with surface ECoG electrodes. Though moving a computer cursor is a familiar simple motor skill to the participants, our experimental perturbation reflects their novel attempts (successful or not) at learning a new mechanism of using a track pad.

This work was performed in conjunction with another study examining changes in activity in selected brain regions during the execution of the task (Figure 4.2). Between these two studies, we have developed a broader understanding of the effects of learning a new motor/rotation task on brain activity and connectivity patterns and the relationship of those patterns to skill level. This study contributes to understanding the corresponding persistent changes in functional connectivity associated with learning a novel motor skill, which bolsters understanding of changes that appear during the execution of the task itself.

Future studies may focus on expanding these electrophysiological findings to larger areas of cortex already examined in studies using fMRI. The addition of more subjects would increase spatial coverage. Combining depth electrodes placed in the hippocampus and other subcortical regions with the electrodes over surface cortical regions we examined would enable a broader view of offline functional connectivity alterations that occur after learning. Though the time restrictions on research with ECoG subjects make such lengthy collection periods challenging, further studies may attempt to collect the baseline resting state sessions to determine spontaneous change in connectivity from the same subjects as performed the task.

5. General conclusions

Resting state connectivity contributes to a variety of ongoing endogenous cognitive and physiological functions of the brain. It specifically plays a critical role in memory consolidation. Modeling the patterns of spontaneous and task-related variation in brain connectivity contributes to our understanding of what affects the dynamics of various properties of electrophysiological connectivity. In this work, I (1) described resting state connectivity in a single session using techniques not previously applied to electrocorticographic data, (2) quantified change in that connectivity between multiple resting state sessions, which had never been previously investigated with ECoG, and (3) identified changes in resting state connectivity specifically attributable to learning a new task, which had also never been previously investigated with ECoG. Though it has limitations, ECoG technology was critical for my research questions because, in contrast to other human electrophysiology or fMRI imaging, it has uniquely high resolution at fast frequencies thought to represent local neural activity, can record from cortical regions relatively inaccessible with EEG and MEG, and the electrode placement has guaranteed consistency over the course of multiple days. My work provides unique insights into spontaneous and task-related changes in resting state connectivity that provide substantial insight beyond prior studies in other modalities.

5.1 Review of findings

I first demonstrated general properties of resting state connectivity in electrocorticography in a single session (Chapter 2). Prior studies using ECoG had examined the congruence with patterns of connectivity observed in fMRI studies, but none of those studies had broadly characterized the properties of electrophysiological connectivity on the basis of cytoarchitecture (i.e., Brodmann areas) as I did here (Conner et al., 2011; Hiltunen et al., 2014; C. Keller et al., 2013; Ko et al., 2011, 2013; D Mantini et al., 2007; Nir et al., 2008; Schölvinck et al., 2013). Most prior electrophysiological studies of phase synchrony were in the context of events or behavior (Groppe et al., 2013; Hillebrand et al., 2012; Kramer et al., 2011; J. Lachaux et al., 1999; Nolte et al., 2008; Sauseng & Klimesch, 2008). I found that the dominant frequency of resting connectivity as measured with phase locking value varied across the lateral cortical surface, with lower frequencies more dominant (i.e., stronger connectivity; larger number of pairs with connectivity statistically significantly different from random interactions) in frontal regions, and higher frequencies in temporal and parietal regions. I also found that pseudo-causal directional influences flowed primarily from frontal to parietal regions. My findings reinforce the importance of electrophysiological studies of the resting state given the presence

of frequency-specific relationships between brain regions, and especially the use of ECoG to resolve higher frequencies.

I then investigated patterns of spontaneous fluctuations in resting state connectivity between different sessions (Chapter 3). Prior studies had investigated inter-session variation in connectivity in EEG (Brookes et al., 2014; Espenhahn et al., 2016), MEG (Garces et al., 2016; Z. Liu et al., 2010), and fMRI (Allen et al., 2014; Baker et al., 2014; Handwerker et al., 2012; Laumann et al., 2015; Zuo & Xing, 2014). Prior studies of spontaneous change in resting state connectivity in ECoG had only investigated these fluctuations within a session, whereas I looked at inter-session variability (Hutchison, Womelsdorf, Gati, et al., 2013; Kramer et al., 2011). I also employed intraclass correlations, which provide a method for tracking variation across a group of values that each change over time. I found inconsistency in the degree of variation between the three connectivity measures I used (phase locking value, amplitude correlation, and coherence). In PLV and amplitude correlation (i.e., measures that are calculated on only the phase or amplitude component of the signal), there was a significant predictive relationship that stronger connectivity was predictive of high stability, but not vice versa – there were also many stable but weak connections. This relationship between connectivity strength and stability did not exist with coherence. When I examined individual regions' connectivity patterns, I found that within-region connections were far more stable over time than connectivity between different regions. In highlighting three exemplar regions, I illustrated inconsistencies in the relationship between connectivity and stability: parahippocampal-entorhinal cortex was notable for stably weak connectivity except within itself, dorsolateral prefrontal cortex connectivity was both weak and unstable, and inferior parietal lobule had little stability within narrow connectivity bounds. My findings contribute to a broader understanding of changes in resting state connectivity in the absence of a task, especially with regard to high frequencies and phase synchrony.

Finally, having established a baseline of endogenous variation that occurs in the absence of a task, I investigated changes in resting state connectivity from before to after learning a new skill that could be specifically attributed to the learning process (Chapter 4). Prior studies of learning-related changes in resting state connectivity in other modalities had identified changes in connectivity across the cortex that were associated with learning a skill, changes which were not associated only with regions involved in task execution (Albert, Robertson, & Miall, 2009; Darainy et al., 2014; De Vico Fallani et al., 2010; Houweling et al., 2008; Kelly & Garavan, 2005; Mehrkanoon et al., 2016; F. T. Sun et al., 2007; Taubert et al., 2011; Vahdat et al., 2011). The execution of mental rotation tasks was especially associated with inferior parietal

lobule, among other regions (J. M. Zacks, 2008). This question had never been previously investigated with electrocorticography.

I again employed intraclass correlations, this time to determine whether the level of variability associated with learning the task was significantly different from the spontaneous change seen in the absence of any learning (defined as non-overlap of 95% confidence intervals of spontaneous and task-related ICCs). This statistical approach had never been used in ECoG before and is highly conservative. I found systematic differences in the changes in resting state connectivity from before to after task execution between those who became adept at the task and those who did not, which were not consistent between the three connectivity measures I used, but in such a way that revealed underlying properties of the independence of synchrony and co-activation. Specifically, I found that connections involving inferior parietal lobule were different between learners and non-learners in all three connectivity measures, and further, that these effects were concentrated in high frequencies and in intra-IPL connectivity. I also found patterns distinguishing learners and non-learners in M1 coherence. Surprisingly, I found that no pre-task connectivity patterns were predictive of the eventual skill level at the task. The patterns of changes in resting state connectivity I identified involve regions associated both with sensorimotor task execution (M1, IPL) and resting state memory consolidation functions (PH-ER, prefrontal regions). My findings provide insight into the differences between learning, attempting, and failing at a task, and form a foundation for further studies looking to isolate changes in resting state connectivity associated with learning.

5.2 Conclusions

My findings suggest that not only are electrophysiological patterns of resting state connectivity affected by behavioral tasks such as learning a new skill, they cannot be treated as stationary or uniform across the cortex even in the absence of external stimuli. When assessing the properties of connectivity associated with a task, we must compare task-related changes against endogenous variation that appears spontaneously (Chapter 4). Spontaneous variation, which I treat as a baseline that post-task changes must exceed in order to be deemed specifically task-related, is highly differentiated across brain regions (Chapter 3). Not only that, but the features of connectivity within a single resting state session are also variable across the cortex, and across different frequency bands and in different properties of the signals (i.e., different connectivity measures) (Chapter 2).

Previous studies had identified that spontaneous, endogenous variation in resting state connectivity exists, and also that connectivity changes after learning a new task. My specific contributions include expanding the frequency range of prior studies, investigating the

electrophysiology of regions not accessible with non-invasive methods, and investigating the properties of electrophysiological connectivity over the course of multiple days because the electrode positions are fixed in place.

In the broader context of previously unanswered questions and prior results on related questions, I begin to address the importance and function of variation in resting state connectivity. The ubiquity of meaningful resting state synchrony and co-activation across the cortex indicates its importance, as it is metabolically expensive. Variation across the cortex in the context of a single session indicates that different brain regions are not contributing equally to resting state function, which implies their connectivity patterns reflect ongoing resting state-related function that is dependent on different specific regions. Spontaneous changes in connectivity patterns in the absence of a task likely reflect endogenous dynamic variation in resting state function, and that non-stationary brain function reflects ongoing variable cognitive functions. Finally, those changes in connectivity that I identified as being specifically task-related likely possibly reflect consolidation of task engagement, performance, and/or (successful or unsuccessful) task strategies.

5.3 Limitations

Electrocorticography, as I have discussed, has unique advantages in the realm of human electrophysiology, but that high-quality data comes with significant limitations from the nature of the technology, the patient population, and the clinical setting. ECoG provides unique insights relative to other methods used in humans, particularly with regard to the properties of high frequency oscillations thought to represent local cortical activity (Buzsáki & Draguhn, 2004). However, the subjects in my work all had medically intractable epilepsy, which is often accompanied by various neurological comorbidities and/or cognitive limitations. I removed electrodes from analysis that were clinically identified as being located over seizure foci and other regions highly impacted by ictal activity, but I cannot rule out the effect of their epilepsy on the remaining tissue. Furthermore, ECoG also requires multiple surgeries to place and then remove the electrodes.

Each individual subject contributed data from a limited cortical area, which restricted the inferences I could make about spatially distant regions (e.g., prefrontal and occipital regions) or interhemispheric connections. However, every area where I drew inferences had coverage from more than half of the subjects included in that project. While the spatial resolution of ECoG is superior to EEG and MEG, and unlike those methods ECoG can resolve cortical signals with high confidence from subtemporal and interhemispheric regions (when clinically indicated), it is

still lower resolution than fMRI and is restricted to the cortical surface (except when used in conjunction with penetrating depth electrodes, which are not included here).

I had relatively few subjects – a maximum of 8 in any experiment – and their electrode placements were determined entirely by clinical needs. This meant I had no control over which regions had a sufficient number of electrodes and subjects to complete analyses. Though it is tempting to point to the parahippocampal-entorhinal gyrus, inferior parietal lobule, dorsolateral prefrontal cortex, and primary motor cortex, among others, as the critical regions of interest in the study of endogenous/spontaneous and learning-related change in connectivity, this may be an artifact of the regions where I could dependably collect data. These brain areas had many electrodes from a majority of the subjects, while key regions like the cingulate and primary visual cortex were excluded not because of a lack of interesting connectivity patterns, but because of a lack of data. I was also not able to incorporate existing, though limited, data from penetrating depth electrodes in both cortical and subcortical regions. As the use of penetrating depth electrodes becomes more common, sometimes in conjunction with a surface grid and sometimes not, future work will need to start to include these electrodes both for practical reasons of number of subjects, and to expand coverage to more regions than possible with surface electrodes alone.

Once the data have been collected, I selected from a vast array of methods for calculating both connectivity between time signals and variation over time (Bastos & Schoffelen, 2016; Greenblatt et al., 2012; McGraw & Wong, 1996). In various portions of the work I presented, I used four measures of connectivity: phase locking value, phase slope index, amplitude correlation, and (magnitude squared) coherence. All four of the measures I used are linear, which quantify connectivity within one frequency band, and do not measure interactions between multiple frequencies or with a time delay (except phase locking value, which does account for lag). There are at least 15 other linear measures in common use in electrophysiological brain data, each of which provides a different type and focus of information about the relationship between the signals (Bastos & Schoffelen, 2016; Greenblatt et al., 2012). My choice in measures was far from comprehensive with regard to the nature of the connectivity patterns in question, but provided robust information about fundamental properties of the signals that are also clearly interpretable in terms of neural and cognitive function.

Also available are an array of pseudo-causal measures that capture directional (and putatively causal) relationships and non-linear measures to investigate interactions between different frequencies (Bastos & Schoffelen, 2016; Canolty & Knight, 2010; F Darvas et al., 2009; Greenblatt et al., 2012). However, pseudo-causal measures such as phase slope index

generate two values per interaction (signal A to B, and B to A), which was not clearly interpretable (Chapter 3). Non-linear measures such as phase-amplitude coupling can generate up to hundreds of connectivity values per signal pair depending on the size of the frequency space investigated and the level of granularity. In both cases, the increased number of connectivity values dramatically increases the complexity of interpretation.

5.4 Future work

Further studies may investigate additional properties of spontaneous or endogenous variation in resting state connectivity. My results were congruent with previous studies in other modalities investigating inter-session resting state connectivity, but did not directly link the patterns I observed in ECoG to the similar ones seen with other methodologies. Future studies could quantitatively compare connectivity patterns in resting state fMRI and ECoG in the same subjects (taken pre- and inter-operatively, respectively). The cingulate is an important region in resting state connectivity, both in the context of the default mode network and in general resting state connectivity (Allen et al., 2014; Jerbi et al., 2010), and is inaccessible with EEG and MEG, as those methods can only record from the lateral surface of the cortex. Unfortunately, cingulate coverage is relatively uncommon in ECoG and I was unable to acquire data from enough subjects to include it in my analysis. Further studies could employ ECoG to quantify the connectivity of the cingulate and compare it to prior findings from fMRI. I pooled intra-individual changes in resting state connectivity, but did not examine differences in variation patterns between different individuals. With a larger number of subjects, a future study could quantify differences in variation in connectivity between individuals, and potentially identify predictive factors between individuals that explain differences in spontaneous variation between them. Finally, prior studies of spontaneous variation in resting state connectivity have been a mix of inter-session and intra-session studies. Intra-session change in resting state connectivity has previously been examined in ECoG (Hutchison, Womelsdorf, Gati, et al., 2013; Kramer et al., 2011), and I have now examined inter-session variability, but a future study could link inter-session and intra-session variation in the same individuals. Besides all of these additional experimental parameters, future work may also employ alternative linear, pseudo-causal, or non-linear (cross-frequency) connectivity measures.

In the realm of learning, this study was the first to use ECoG to investigate resting state connectivity changes specifically attributable to attempting a task, though those properties had been previously examined in other modalities. The nature of ECoG is such that I had limited coverage of brain regions, especially as I then further divided my already small subject pool into those who successfully improved at the task and those who did not. Future work can start with

expanding coverage to more brain regions with a larger subject pool. The task I used was appropriate for this investigation because there was a clear division between subjects who improved at the task and those who did not, but the range of performance within those who improved was relatively narrow. A more challenging task, and correspondingly broader performance gradient, would improve the estimation process for whether any pre-task brain state corresponded to ultimate performance. I examined resting state connectivity before and after a single session of task practice; further work can examine connectivity during the task execution itself to examine task execution-related connectivity patterns (as opposed to the consolidation patterns I examined), or add additional practice sessions in order to assess the relationship between resting state connectivity and offline gains in skill level. Finally, because of restrictions in data collections, I used data from a different pool of subjects to establish the baseline levels of spontaneous connectivity for the various regions of interest. Further work in ECoG needs to replicate my work using each subject's own data to establish the baseline for spontaneous variability against which task-related variability can be compared. A sufficient number of pre-task resting state sessions will be challenging to acquire, which is why I was unable to include this comparison, but is essential to establish the robustness of the patterns I recorded.

6 and 7. Side projects

As an addendum, I have included two side projects that investigate properties of resting state connectivity, but did not fit into the linear theme connecting the three main papers included in the body of my thesis. Though they were tangents to that research, I still had interesting and relevant findings in each, and thought them worth including as representative of my overall interest in the resting state.

6. Changes in connectivity attributable to learning of a brain-computer interface device

The following is my investigation of changes in resting state connectivity related to brain-computer interface learning. I completed this project between the first two papers of my main work, and it uses mostly the same methods as the first (in particular, it lacks the use of spontaneous change in resting state connectivity as a baseline to identify changes specifically associated with learning the BCI). I did not include the paper in the main body because I did no other work on BCIs, and because I did not apply these methods or findings again in later studies. I found changes that roughly corresponded to those seen in prior studies using typical motor tasks; in particular, motor cortex connectivity increased, inferior parietal connectivity decreased, and shorter-distance connectivity was more impacted than distant connections. Some of the regions of interest that came up as significantly different across multiple frequency bands I later identified as regions of interest in the more robust rotation task. These findings suggest that learning a motor cortex BCI is not unlike learning a typical motor task in terms of its connectivity impact.

The following section was first published as:

Casimo, K., Weaver, K. E., Wander, J., & Ojemann, J. G. (2017). BCI Use and Its Relation to Adaptation in Cortical Networks. IEEE Transactions on Neural Systems and Rehabilitation Engineering, 25(10). https://doi.org/10.1109/TNSRE.2017.2681963

6.1 Introduction to section

The goal of this review is to discuss prior research in adaptive mechanisms related to how the brain processes and subsequently gains control of brain-computer interfaces (BCI). We emphasize and highlight three relevant fields: 1) a comparison of neural processes in procedural learning during motor tasks and in an electrocorticography-based BCI task, 2) long-term brain changes as a result of learning motor tasks generally and BCIs specifically and 3) the role of cortical networks in all of the above. Finally, we will discuss current work on changes in cortical networks following BCI use, and potential implications for future BCI design.

Brain-computer interfaces represent a promising clinical application for individuals who have been paralyzed and as part of prosthetic controls. While BCI designs have been refined to the point where they are reasonably accurate and reliable, the broad, long-term brain adaptations related to their use are not yet known. Within the domain of BCI design, we argue that it is not sufficient to study motor cortex alone, and we will review how studies of motor learning predicted, and prior studies of BCI learning found, that effects extend across the cortical surface.

We highlight cortical networks and connectivity, which have been less studied than evoked response patterns. Evoked changes in activity in the motor cortex, and other motor regions such as premotor cortex and basal ganglia, are well described for procedural skills, including BCIs (Dayan & Cohen, 2011; Wander et al., 2013). For motor and BCI tasks, changes in activity across the whole brain are increasingly characterized (Wander et al., 2013). In contrast, the patterns of connectivity, information flow, and long-term adaptation in the neural networks that support skill acquisition are not as well defined for BCI tasks. In particular, learning-related change in connectivity between regions other than the motor cortex over the course of development of BCI proficiency has been the subject of little examination. Modern tools of systems neuroscience, including MRI and functional MRI (fMRI), electroencephalography (EEG), electrocorticography (ECoG), and others have advanced contemporary views of brain organization centered on understanding the networked brain as a complex system of interacting units driven by functional connectivity, spanning from local "microcircuit" levels to macro-scale networks across large swaths of cortical area (Cabral et al., 2014). In light of this, we will discuss recent results demonstrating heterogeneous changes in connectivity during BCI use (Wander et al., 2016).

Our review, and our current research, illustrates changes in connectivity resulting from using a BCI and acquiring BCI proficiency. Although a universally agreed-upon definition of functional connectivity in systems electrophysiology has yet to emerge, multiple studies have

demonstrated statistical dependencies between disparate signals using classic measures of coherence, linear measures of phase and amplitude interactions, or non-linear metrics of cross-frequency coupling (reviewed in (Greenblatt et al., 2012; Schölvinck et al., 2013)). For example, early applications of phase-amplitude coupling (PAC) indicated that bursts of high gamma (HG) band activity were preferentially entrained around the trough of the low-frequency (theta to beta range) phase cycles of local field potentials (Canolty & Knight, 2010). This signal model posits that the excitability cycles of local populations need to be synchronized to facilitate information transfer between disparate populations with no phase delay (Onslow et al., 2011). Single unit recordings in animals and humans demonstrated that the slower entraining oscillation provides a gating mechanism for high-frequency activities within local circuits or small-world networks. This provides a theoretical, orchestrating mechanism capable of coordinating local post-synaptic integration and spiking activity amongst different network elements (Canolty & Knight, 2010; Jacobs et al., 2007).

Starting from this broad definition, we find changes in connectivity across a wide area of cortex immediately after BCI skill acquisition. Increased connectivity within the frontal cortex, including motor cortex and prefrontal regions, is contrasted with uniform decreases in connectivity in parietal cortex. We observe frequency-dependent and connectivity measure-dependent variation in temporal cortex. The results we present here from our analyses clearly demonstrate broad changes in network connectivity directly as a result of BCI use that occur between sites other than the BCI control site. These results are consistent with the limited prior work in BCI learning and with work in other motor learning, and are the first examination of adaptive connectivity in BCI learning in humans outside of the BCI control site.

When considering impacts of BCI use on cortical networks, it is important to recall that while the BCI itself is controlled by only one (or occasionally, a few) cortical site, the broader extent of the brain is not disengaged from the task. Measuring the activity and connectivity across a wide region of cortex is necessary for a full understanding of BCI use and, critically, BCI learning. This is therefore by definition essential for the study of cortical networks, and specifically how various aspects of learning BCI shape interactions across relevant functional networks. It is well established that motor learning generally, and more recently BCI learning, does involve distributed cortical resources (Dayan & Cohen, 2011). Accordingly, BCI studies that ignore cortex beyond the control site may not fully capture the dynamic state of the system. Understanding the adaptations and alterations in and across large cortical networks during BCI learning and subsequent to the development of BCI proficiency will be directly translatable to the long-term efficacy of BCI systems.

6.2 Review of literature

BCI Design and Capabilities

A wide variety of BCI design models have been developed; for a comprehensive review, see (Bensmaia & Miller, 2014; Hiremath et al., 2015). A popular and traditional design is to record from motor cortex or other motor-related regions such as premotor or parietal cortex. They use neuronal firing rates or variation in amplitude or power in a specific frequency range to control a cursor, robotic arm, or other assistive device. The BCIs used in this study employ such a power-based design. Other BCIs control schemes include connectivity (Billinger et al., 2013; Daly et al., 2012; Hamner et al., 2011; Koush et al., 2015) and event-related potentials (Buttfield et al., 2006; Sellers et al., 2010). Besides the direct applications in developing improved BCIs, these systems provide a unique opportunity to 1) study motor skill acquisition in the absence of actual motor movement, 2) enable significantly greater control over experimental parameters than in standard motor tasks, 3) guarantee complete novelty to the subject (Shenoy & Carmena, 2014).

In non-human primates, studies have employed various BCIs controlling, for example, a cursor (Ganguly & Carmena, 2009), robotic arm (Carmena et al., 2003; Ganguly et al., 2011), or the monkey's own arm (Moritz et al., 2008). Over the course of weeks to months of training, non-human primates are able to execute complex, dexterous tasks of the types necessary for BCIs to be a viable clinical option. In contrast to the studies of humans discussed below, non-human primate BCIs typically utilize single-units or small ensembles of <100 neurons (Carmena et al., 2003; Ganguly et al., 2011; Ganguly & Carmena, 2009; Moritz et al., 2008) as drivers of control signals. The benefits of single-unit and small ensemble recordings include, but are not limited to, the animal having very precise control over the signals generated, high spatial resolution, and the ability to match the desired BCI outcome to a specific pre-existing neural preference to aid learning (Carmena et al., 2003; Ganguly et al., 2011; Ganguly & Carmena, 2009; Moritz et al., 2008).

In humans, invasive methods of acquiring control signals include penetrating electrodes such as microgrid arrays (Gilja et al., 2015; Hochberg et al., 2006) and macroscale ECoG (Leuthardt et al., 2004; Schalk et al., 2008; Wander et al., 2013, 2016). Penetrating electrodes (e.g. Utah arrays) have been used in paralyzed individuals who have volunteered for long-term studies specifically investigating BCIs such as (Gilja et al., 2015; Hochberg et al., 2006). Penetrating electrodes provide very high spatial resolution and have been shown to be very effective for fine control of a robot or cursor, but the spatial distribution of these studies remains limited to date (Gilja et al., 2015; Hochberg et al., 2006). Less invasive methods include EEG

(Billinger et al., 2013; Daly et al., 2012) and fMRI (Koush et al., 2015), but suffer from an inability to reliably track signals over long periods of time. However, because it is noninvasive, provides superior temporal resolution relative to fMRI, and is readily usable in many patients, EEG in particular possessed distinct advantages in the BCI framework (Birbaumer et al., 1999; Collinger et al., 2013; Leeb et al., 2015; Muller-Putz et al., 2015; Pichiorri et al., 2015; Ramos-Murguialday et al., 2013; Schalk et al., 2008)

In contrast, ECoG can be used in experimental settings at various levels of moderate spatial resolution depending on grid type (typically 1-10mm) and provide high frequency resolution including up to 200Hz (Leuthardt et al., 2004; Schalk et al., 2008; Wander et al., 2013, 2016). Additionally, these electrodes are rarely shifted after their initial placement, unlike EEG electrodes. This provides the capacity to track changes in connectivity from a consistent neural population over a long period.

ECoG studies are conducted most frequently in individuals with epilepsy undergoing monitoring for one week as part of surgical epilepsy treatment, with placement decided entirely by clinical needs. Like penetrating electrodes, ECoG is subject to limited spatial coverage and requires surgery for indwelling placement. Relative to penetrating electrode studies, surface-based ECoG studies are limited by short experimental time and access only to cortical surface, although limitations are changing (McIntyre et al.). However, these compromises enable access to a much larger area of brain, which when considering the role of networks and distributed cortical involvement in the BCI task, this spatial coverage is by definition necessary.

While penetrating electrodes record from relatively small numbers of neurons, similar to the electrodes used in non-human primate studies, macro-scale ECoG electrodes record from over 100,000 neurons at a time (Chang, 2015). In BCIs that use ECoG, for example, the user must learn how to voluntarily modulate the signal amplitude generated by this large population collectively, which may contribute to impacts on other involved cortical regions and networks.

Motor Learning and BCIs as a Motor Skill

Motor, or procedural, learning is qualitatively different from other forms of memory, and depends on a distinct set of brain regions and neural processes relative to, for example, declarative memory. A learned motor skill can be divided into the selection of a movement from a pool of options, and the accurate execution of that movement (Chein & Schneider, 2005; Diedrichsen & Kornysheva, 2015). These patterns were first observed in straightforward motor tasks such as sequence learning, but studies of BCI skill acquisition to date have observed the same patterns (Sadtler et al., 2014; Wander et al., 2016). Execution of a BCI is more dependent on the execution of an abstract skill that cannot be specifically drawn from existing resources.

Nearly all, and the most clinically relevant, BCIs are motor-driven with the BCI control site located in the motor cortex, often in the hand or tongue region. In both human and non-human primate studies, studies started from the essential motor basis of the BCI task. As we will review, this approach has demonstrated that BCI learning resembles other motor learning both at the motor cortex and across the brain and in both activity and cortical network patterns.

A disperse group of brain regions has been consistently linked to motor learning. Primary motor cortex (M1), and other association motor regions such as the premotor cortex that exhibit high levels of connectivity with M1, are the most obvious candidates. However, changes in electrophysiological power or fMRI BOLD signal are also consistently found in human and/or non-human primates in the left dorsolateral prefrontal cortex (Chein & Schneider, 2005; Wander et al., 2013), left (Chein & Schneider, 2005) and right (Rota et al., 2011) ventrolateral PFC, bilateral anterior cingulate (Chein & Schneider, 2005), bilateral posterior parietal (Chein & Schneider, 2005; Wander et al., 2013), bilateral occipito-temporal junction (Chein & Schneider, 2005), inferior parietal (Albert, Robertson, & Miall, 2009), parts of cerebellum (Albert, Robertson, & Miall, 2009), and premotor cortex (Diedrichsen & Kornysheva, 2015; Wander et al., 2013). These regions are identified through changes in signal amplitude in one or more frequency bands, including beta (13-30Hz), which is commonly linked to motor activity, and HG (70-200Hz), generally believed to be a marker of local neural broadband population firing rate (Miller et al., 2014). Each of these areas and frequencies are assumed to represent multiple functions (such as thalamic influence), including non-motor functions, and their involvement may depend on the complexity of specific task demands (Buzsáki & Draguhn, 2004).

Changes in evoked activity patterns across the learning cycle, as discussed above, reliably follow a set of principles: 1) activity generally decreases in non-M1 regions as proficiency increases, 2) the brain regions recruited are the same for any motor task and 3) a region that was not involved in the initial stages of learning will not become involved later (Chein & Schneider, 2005). Importantly, only new learning generates major alterations in functional activity, while execution of previously learned skills or random movements does not (Monfils et al., 2005). Changes in both behavior and neural activity emerge over the course of hours to days, which may be accelerated if the new skill is closely related to or an extension of a previously mastered skill, possibly as a result temporarily repurposing or restructuring existing architecture and then storing a modified alternate version (Sadtler et al., 2014).

Notably, activity changes after BCI learning generally tends to mimic these rules in both humans and monkeys, and are concentrated in the same regions associated with learning of

typical motor tasks (Sadtler et al., 2014; Wander et al., 2013, 2016). The regions exhibiting the most notable activity decreases during task execution after proficiency is attained include, as predicted by the results listed above, prefrontal, premotor, and parietal cortex (Wander et al., 2013). This work by Wander and colleagues is critical to understanding BCI learning in the context of networks, as it was the first study to establish in humans that development of BCI proficiency depends on a distributed set of cortical regions. Understanding the longitudinal activity changes related to BCI and other motor learning may aid in the development of permanent BCIs that are capable of compensating for such change, as current systems either assume that the neural signals to be decoded are roughly static or, more often, do not attempt to re-adjust the classifiers (Gilja et al., 2015; Hochberg et al., 2006).

Evoked activity is only one way to characterize dynamic functional neuroanatomy outside of the primary motor circuit during motor learning. We can also measure connectivity, which investigates how regions couple, synchronize and communicate with each other, and can include both connectivity of regions to the BCI control site (i.e., a seed-point analysis) or of all regions to all other regions. In this work, we employ the phase locking value (PLV) (J. Lachaux et al., 1999) and simple correlation; in our other work, we have also used phase slope index (PSI) (Casimo et al., 2016; Nolte et al., 2008). PLV uses only the phase of the signal and provides an estimate of synchrony between two signals (J. Lachaux et al., 1999); its principal advantage derives from not using amplitude to provide a cleaner estimate of synchrony specifically (Greenblatt et al., 2012), though it is not scaled to the different number of cycles at different frequencies. In contrast, correlation includes both phase and amplitude components, rendering it susceptible to effects from variations in amplitude between channels, but capturing a larger set of signal features. We avoided Granger causality due to its high number of assumptions and *a priori* conditions (Greenblatt et al., 2012; Nolte et al., 2010). Although phase slope index does not measure true causality, restricted instead to pseudo-causal or directional inferences, it does not require setting a model *a priori* (Nolte et al., 2008). Additional measures not employed in this work, such as phase-amplitude coupling (PAC), amplitude-amplitude cross-frequency coupling (amp-amp CFC), and short time-windowed covariance (STWC) describe temporally complex relationships between signals at differing frequencies (PAC, amp-amp CFC) or time scales (STWC) (Timothy Blakely et al., 2014; Canolty & Knight, 2010). For further detail and additional connectivity measures, see (Greenblatt et al., 2012).

In other, non-BCI motor tasks, many studies have observed transient increases (Wander et al., 2016) followed by a substantial decreases in connectivity between the motor region executing the task and support regions that appear to only be involved in the acquisition phase

(Albert, Robertson, & Miall, 2009; Ganguly et al., 2011; Rota et al., 2011; Wander et al., 2016). However, some reports have indicated a distinction between increases in connectivity in some functional networks related to spatial and visual processing, and decreases specific to motor networks (Buchel et al., 1999; Doyon & Benali, 2005; Vahdat et al., 2011; Vincent et al., 2008). These changes also usually correlate with behavioral performance (Albert, Robertson, & Miall, 2009; Buchel et al., 1999; Ganguly et al., 2011; Vincent et al., 2008; Wander et al., 2016). In humans, training increases white matter fractional anisotropy in diffusion tensor imaging and grey matter density in and around the intraparietal sulcus, indicating that changes in functional connectivity may also be more than transient (Scholz et al., 2009).

Ultimately, however, the connectivity patterns associated with BCIs specifically have only begun to be investigated. Previous work from our lab has shown significant changes in short time-windowed covariance (STWC), a measure of within-frequency amplitude-amplitude connectivity (Timothy Blakely et al., 2014; H. Sun et al., 2015), of HG (70-200 Hz) band was restricted almost exclusively to connections between the BCI control electrode and motor, premotor, and prefrontal regions (Figure 6.1, A-C), and progressively decreased during the acquisition and use of a BCI (Figure 6.1, D) (Wander et al., 2016). It is important to note that this decrease is consistent with predictions based on other motor learning tasks, suggesting that not only activity patterns but also connectivity patterns of BCI learning and use follow those of motor tasks (Hardwick et al., 2013). However, while pure amplitude-based STWC experienced decreases, nonlinear cross-frequency phase-phase connectivity as assessed by biphase locking (bPLV) (Darvas et al., 2009; Felix Darvas et al., 2009) remained strikingly consistent throughout the period of task execution, and generally was observed between sites farther apart from each other than those with significant STWC (Wander et al., 2016). The heterogeneity of connectivity changes across the brain, across frequency ranges, and between amplitude and phase measures require further investigation. Further, this work does not investigate the persistence of these changes beyond the short experimental ECoG windows of a week or less with the BCI, with each dataset described here collected within a single session in one day, or interactions besides those seeded to the BCI control electrode.

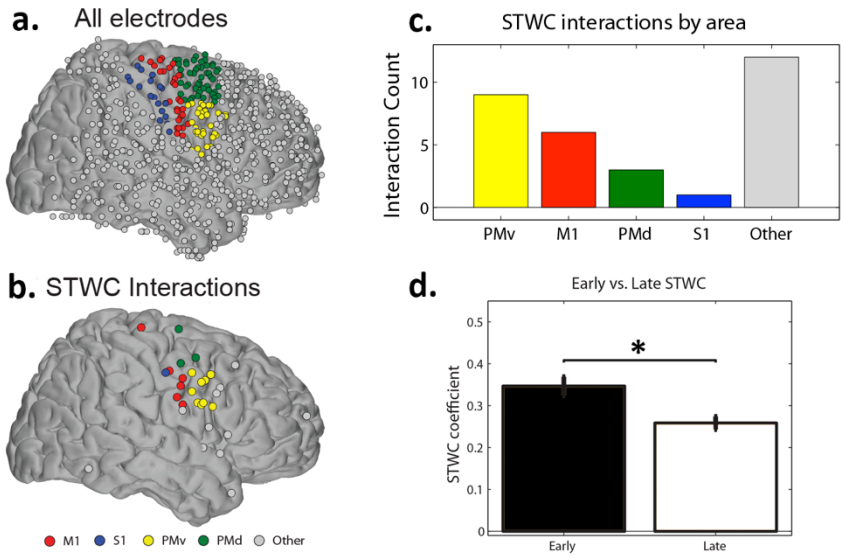

Figure 6.1
Evoked connectivity of control electrode across cortex during BCI execution, and adaptations with BCI learning. A: electrodes of n=9 subjects, with electrodes in regions of particular interest colored: red, M1; blue, S1; yellow, ventral premotor; green, dorsal premotor; gray, all others. Electrodes are shown on a Talairach brain, with anatomical labels derived from the HMAT standard atlas. B: electrodes with significant HG short time-windowed covariance (STWC) interactions with the BCI control electrode during the execution of the task. Colors as in A. Virtually all electrodes with significant HG amplitude-amplitude interactions were in frontal lobe motor circuits or farther anterior within prefrontal cortex proper. C: electrode count of all significant STWC in premotor, motor and sensory cortex. D: change in STWC coefficient of electrodes shown in B and C to BCI control electrode decreases after development of BCI proficiency (*p<0.05). Adapted with permission from (Wander et al., 2016).

The Resting State: Window Into Consolidation?

There are two candidate experimental windows for assessing connectivity related to execution of BCI or any other task: while the subject is executing the task, and after the task in complete. Task-driven connectivity of M1 and other areas of the cortex during BCI use can be assessed like it would be during any other motor or procedural task, and research to date has indicated similarities between them (Ganguly et al., 2011; Wander et al., 2016). This, however, still leaves open the questions of persistent connectivity changes to these networks beyond the task performance period.

Though it may seem counterintuitive as a tool for studying task-related changes, the resting state is ideal for studying connectivity in intrinsic networks. The resting state is defined as a behavioral state of relative inactivity and lack of directed engagement with an overt task. Prior studies of the resting state have revealed spontaneous clustering patterns that fall into consistent, intrinsic large-scale and small-world networks (Cabral et al., 2014). Such endogenous, stimulus-independent background activity is widespread and has been observed in nearly every vertebrate system examined to date (Zhang et al., 2010). Despite its

spontaneous nature, resting state activity is postulated to reflect a ready state of diffuse attentiveness, internally self-directed thought, consolidation of learning, and more (Marcus E Raichle, 2011; S. M. Smith et al., 2009). The resting state is particularly important for studying task-negative networks such as the default mode and frontal-parietal networks (Buckner et al., 2008), which are postulated to play a role in memory, skill consolidation (Albert, Robertson, & Miall, 2009), self-directed thought, and response readiness (Marcus E Raichle, 2011).

Resting state functional connectivity patterns undermine assumptions of relative stability in the brain. Changes in resting state connectivity occurring during the related learning process are related to consolidation (Albert, Robertson, & Miall, 2009). These changes are persistent and consistent across individual humans and non-human primates in “traditional” motor tasks (Albert, Robertson, & Miall, 2009; Chein & Schneider, 2005; Sadtler et al., 2014), though the effects of BCIs specifically have not been previously studied. Additionally, by definition, data collected during post-task resting state windows allow measurement of changes that persist for minutes to hours after task performance has ended (Albert, Robertson, & Miall, 2009). As any long-term clinical applications of BCI would in all likelihood target users with chronic needs and thus persistent neural adaptations cannot be ignored. In a chronically implanted BCI, understanding and compensating for long-term learning effects on the brain become highly relevant to maintaining consistent, low-maintenance use.

6.3 Current research

Considering the activity at the BCI control site and across the brain, connectivity both seeded to and independent of the BCI control site, and the resting state itself are all important to our goal of characterizing the dynamic relationship between BCI use and development of proficiency. Here, we present results illustrating changes in connectivity across cortex after four subjects were trained on a motor-based BCI. We find significant and spatially disparate disperse changes in connectivity immediately following a single session of BCI use.

In this ECoG-based BCI, the HG amplitude in hand motor region was used as the control signal. Of particular note, ECoG provides high sensitivity to the low-amplitude signals of the HG range, which is not easily accessible by surface electrodes. HG (70-200Hz) is considered the best proxy for local broadband neural activity (Miller et al., 2014) and has high temporal specificity, and thus is an optimal control signal for brain-computer interfaces (Ray et al., 2008; Wander et al., 2013).

Data Collection

All data were collected with surgically implanted, subdural ECoG electrode arrays as previously described (Wander et al., 2013; Weaver et al., 2016). Three adults (ages 31, 33, 44)

and one adolescent (age 13) successfully completed the SimulBCI task. In this BCI, subjects learn to control the vertical position of a cursor by volitional modulation of the amplitude of the HG signal in a selected area of motor cortex, while simultaneously controlling horizontal position with a keyboard with the hand corresponding to that motor cortex. They were asked to guide the cursor to one of eight targets arranged in a circle around a central starting point. In all of these individuals, the BCI was controlled by the signal from an electrode over the hand motor region. Before exposure to BCI, and immediately after the BCI session, we recorded data during an eight-minute resting state session. These two resting state data sets are used for all analyses. To visualize responses across the cortical surface, electrodes were registered first into native space on a pre-operative MRI scan, then into standardized space (the 152 MNI brain), from which Brodmann area labels were assigned to each electrode. Finally, all electrodes from all subjects were overlaid onto a single common MNI template brain (Figure 6.2).

Signal processing and statistics were also performed as previously described (Wander et al., 2013; Weaver et al., 2016). In brief, signals were first notch filtered, common average re-referenced, and manually screened for artifacts. Number of channels ranged from 64-82 across subjects. Instantaneous amplitude and phase were calculated for each frequency band of interest (delta: 0-4Hz; theta: 4-8Hz; alpha: 8-12Hz; beta: 13-30Hz; HG: 70-200Hz) using a Hilbert transform. Connectivity was calculated for each pair of electrodes in each frequency range for a given individual. This process was repeated for the pre-BCI and post-BCI resting state sessions, and the difference between the two was taken using paired t-tests. This process was repeated across subjects and served to normalize between-subject differences.

There are a variety of methods to assess connectivity, generally focusing on various properties related to phase and amplitude components, ability to make causal inferences, linearity or non-linearity, and more (Greenblatt et al., 2012). Here, we show results from two connectivity measures: traditional Pearson's correlation and phase locking value (PLV), both estimated at zero phase-lag. We employ both measures of connectivity as no single measure captures all of the temporal and frequency interplay of a large set of interacting signals; further measures could be added as well (Digiovanna et al., 2009; Greenblatt et al., 2012; J. Lachaux et al., 1999; Tort et al., 2010; Weaver et al., 2016). We highlight PLV, a solely phase-based measure, in order to assess synchrony between signals, including persistent high-frequency responses, though at the cost of excluding some signal features. In contrast, correlation is influenced by both phase and amplitude co-variation between signals, but captures a wider set of signal features but at the cost of potentially washing out smaller but still significant variations. All differences are presented as t-scores of the post-BCI minus pre-BCI connectivity value.

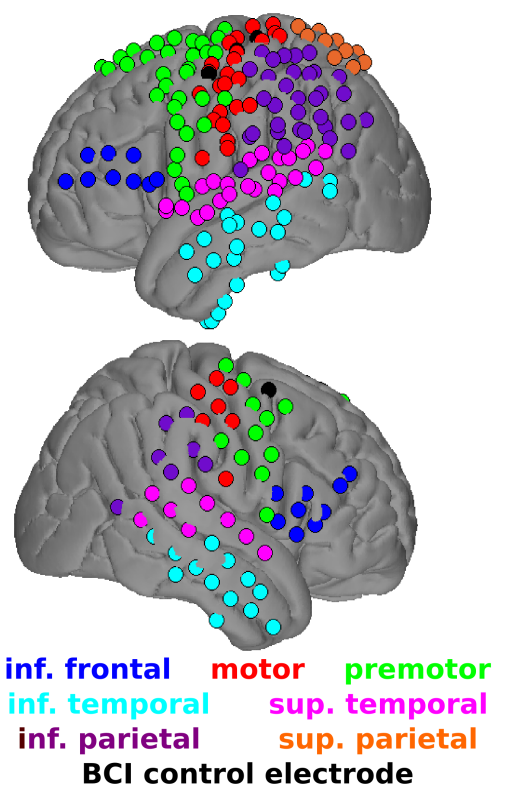

Figure 6.2
Electrode placement in four subjects. Three subjects had left-sided electrode implants; none had bilateral implants. Electrode that was used to control BCI is colored black. All other electrodes were placed on the MNI152 template brain, labeled by Brodmann area and color-coded: blue, inferior frontal; red, sensorimotor; green: premotor; cyan: inferior temporal; magenta: superior temporal; purple: inferior parietal; orange: superior parietal.

Statistical significance was determined based on a permutation test; a null distribution of connectivity values was created by separating the signal into ten-second bins, separating the phase and amplitude components of the signal, and randomly permuting the phases from each electrode between all the time bins. We then recalculated the connectivity measure on the permuted time series as previously described, retaining the top 5% of maximum values, and this repeated 20 times. Values greater than the 95th percentile of this randomly generated null distribution, reflecting the maximum output values from each permutation, were retained and are shown in the Figures below. This null distribution represents a high level of random chance phase locking, and is therefore more conservative than a truly random distribution of phases. Because this method controls for the maximum values across the entire set of electrodes, no further correction for multiple comparisons was required. We then pooled the comparisons across the four subjects and averaged all of the electrode pairs between pairs of regions (Figure 6.2).

Results

Behaviorally, all four subjects were able to successfully control the BCI device, which again had eight possible target zones for the cursor. After 20-30 minutes of practice, they successfully reached the target zone an average of 56.9% of their final session of 20 trials.

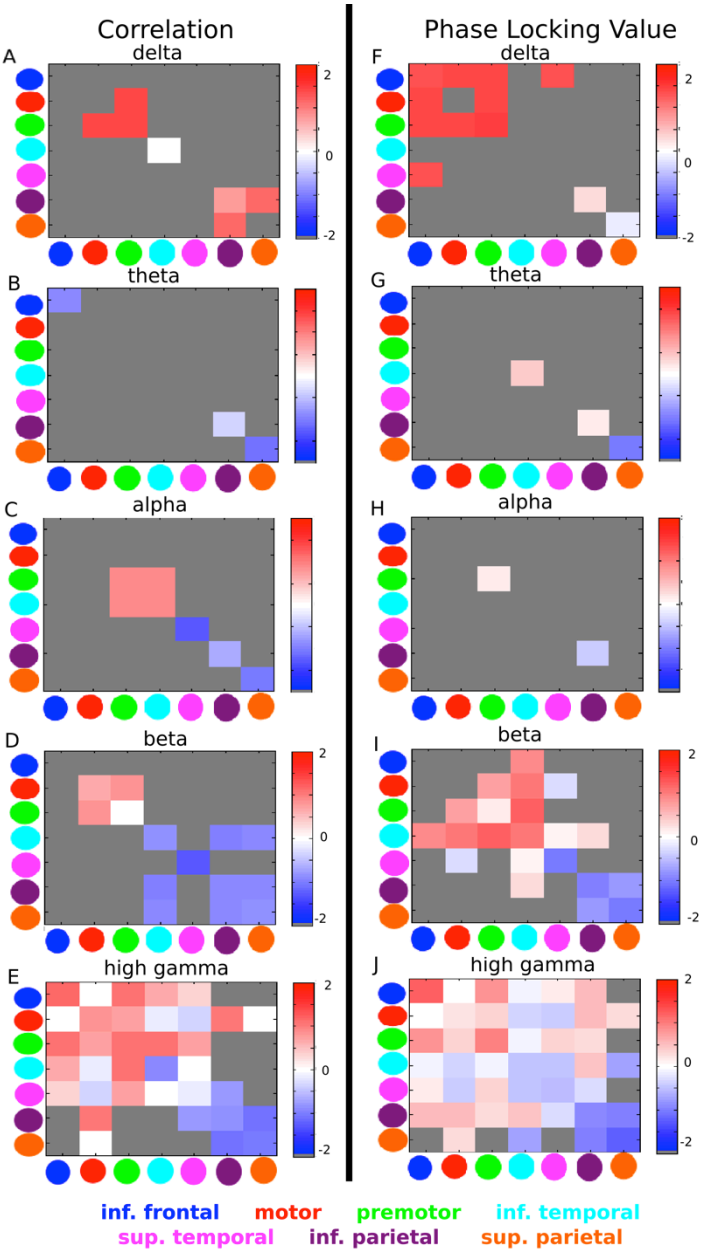

Figure 6.3
Change in correlation (A-E) and PLV (F-J) connectivity in resting state before and after first use of a BCI. Connectivity was calculated in delta (0-4Hz; A, correlation; D, PLV), theta (4-8Hz; B, correlation; G, PLV), alpha (8-12Hz; C, correlation; H, PLV), beta (12-18Hz; D, correlation; I, PLV), and high gamma (70-200Hz; E, correlation; J, PLV). Values are organized into region-by-region pair as shown on the axes, with the mean t-score binned across all electrode pairs and all subjects shown for each box. Differences in t-score reflect post-BCI minus pre-BCI; red indicates an increase in connectivity, blue indicates a decrease. Gray boxes indicate a lack of statistically significant interactions between those pairs of regions. In both correlation and PLV results, increases in connectivity post-task (red boxes) appear in frontal regions, almost exclusively in or anterior to motor cortex. These increases appear in both the frequency band acting as the BCI control signal (HG, E and J) and in other bands (delta, A and F; beta, D and I). Connectivity to and within the parietal cortex notably and consistently decreased (A-J). The majority of PLVs did not significantly change, but those that did were distributed across the whole cortex. Connectivity values in other canonical bands were non-significant outside of a few sporadic regions.

We found that connectivity results from correlation and PLV followed similar general patterns, though there were some notable differences between the two. In both measures, the overall trend was of increased connectivity between frontal regions and decreases between parietal regions, with more inconsistencies across frequencies and connectivity measures in temporal lobe interactions (Figure 6.3). In particular, correlation (Figure 6.3, A-E) and PLV

(Figure 6.3, F-J) increased strongly within and between inferior frontal, premotor, and motor regions across almost all frequencies of interest in the post-BCI resting state session relative to the pre-BCI session, with the notable exception of theta (Figure 6.3, B and G), which showed little change in either measure. Parietal connectivity, in contrast, decreased strongly in beta and HG in both correlation (Figure 6.3, D and E) and PLV (Figure 6.3, I and J). The greatest magnitude of changes, both in increases and decreases, were in the lowest (delta, Figure 6.3, A and F) and highest (beta, Figure 6.3, D and I; high gamma, Figure 6.3, E and J) frequencies, while middle frequency bands (theta, Figure 6.3, B and G; alpha, Figure 6.3, C and H) had substantially fewer connections with change in connectivity across time.

The most notable disagreement between our current connectivity measures, and at first glance with our previously published results, is in temporal lobe connectivity. The disparities are strongest in beta (Figure 6.3, D and I) and to a lesser extent in HG (Figure 6.3, E and J). Frontal-temporal PLV increased in both frequencies, while frontal-temporal correlation decreased. This may reflect similar underlying amplitude decoupling between regions driving the decrease in STWC, and maintenance in phase synchrony measured by bPLV, discussed above (Figure 6.2).

In addition to the specific interactions observed here, these results also clearly highlight the importance of assessing connectivity across the whole cortical surface as part of characterizing the effects of BCI learning and use. Expanding on prior work examining cortical interactions associated with BCI use (Wander et al., 2016), we show here for the first time significant changes in connectivity in the post-task resting state period as a result of BCI use between sites not directly involved in BCI control.

6.4 Conclusions

Changes in activity across cortex, in seed-point connectivity to motor cortex from across the rest of the brain, and to a lesser extent, in evoked connectivity all across the cortex have been amply demonstrated in the past in general motor task learning and execution (Albert, Robertson, & Miall, 2009; Chein & Schneider, 2005; Diedrichsen & Kornysheva, 2015; Monfils et al., 2005; Sadtler et al., 2014). As BCIs have emerged as a scientific tool and a viable clinical possibility, the activity patterns associated with their learning and use have been a staple of research, and indicated involvement of a variety of regions of cortex consistent with other motor skills (Carmena et al., 2003; Ganguly & Carmena, 2009; Schalk et al., 2008; Wander et al., 2013).

More recently, connectivity of the BCI control site to other areas of the brain has indicated that learning and execution of a BCI has heterogeneous impacts on interactions of

different types and to different regions (Figure 6.1) (Ganguly et al., 2011; Wander et al., 2016). That BCI learning would be analogous to other motor and, more generally, procedural learning was strongly predicted, but not previously demonstrated.

Here, we demonstrate for the first time post-task changes in connectivity across cortex including, but not limited to, the BCI control site (Figure 6.2, 6.3). We find changes in connectivity that are regionally distinct, with primarily increases between frontal regions including motor cortex, primarily decreases between parietal regions, and inconsistent in temporal cortex depending on frequency band, connectivity measure, and sub-region. Though increases in resting state functional connectivity may seem irregular at first glance, we note that these changes appear consistent with the amplitude-phase interaction differences previously observed (Wander et al., 2016). However, as a mechanism driving all of these changes in interactions has not been thoroughly established, and because our experimental design cannot definitively attribute the entirety of the changes in connectivity to the BCI learning itself, further research into connectivity changes that encompasses the full array of interactions is recommended. We also recommend that further studies include multiple training sessions on the BCI to capture a larger proportion of the learning process, in order to begin probing long-term changes in connectivity as proficiency further increases and task novelty is lost.

The brain's natural plasticity has already emerged as a design consideration for BCIs (Digiovanna et al., 2009; Shenoy & Carmena, 2014). Plasticity from the cellular (Monfils et al., 2005) to the macro-network (Nudo et al., 1996) level in electrophysiology, and in fMRI (Albert, Robertson, & Miall, 2009; Danielle Smith Bassett et al., 2011; Diedrichsen & Kornysheva, 2015), has been consistently observed in studies of non-BCI motor learning. As studies of BCI learning at the macro- and meso-levels find changes in activity and connectivity consistent with those of traditional motor tasks (Ganguly & Carmena, 2009; Wander et al., 2013, 2016), we posit that BCI-linked changes in connectivity at smaller scales (e.g., entrainment efficiency of local micro-circuits) will eventually also prove similar to those seen with other motor tasks. Given the demonstrated involvement of spatially dispersed regions of cortex in learning and execution of motor tasks, and more recently of motor-driven BCIs specifically, we hope that improved understanding of the role of networks in BCIs will expand to this trend to include more areas of the brain beyond motor cortex. We further anticipate expanded use of BCIs as a tool for basic research into cortical activity and connectivity in motor tasks, and eventual translation of this understanding into therapeutic applications of implanted BCIs in chronic rehabilitative medicine.

7. Comparing brain connectivity in humans, monkeys, and sheep

The following is my comparison of resting state between humans, monkeys, and sheep. I wrote this paper between the first two papers and during the second paper of my main work. This project was the first to include amplitude correlation but still used the permutation test for statistical significance, making it an intermediate step between my methods for examining a single resting session and those for examining spontaneous change. I did not include this paper in the main body because it is the only one to use animals, and because its findings did not contribute to the design of later studies. I found that humans and monkeys largely differed in their connectivity within individual regions, and that the long-range connections linking different regions were similar in the two species. Sheep, with data from only a single region, also differed from humans in somatosensory cortex's intra-regional connectivity. These findings suggest that studies comparing connectivity between regions in monkeys (or depending on multi-regional functions that require such inter-regional connectivity) are more likely to be generalizable than those that compare connectivity within a single region.

The following section was first published as:
Casimo, K., Levinson, L. H., Zanos, S., Gkogkidis, C. A., Ball, T., Fetz, E., … Ojemann, J. G. (2017). An interspecies comparative study of invasive electrophysiological functional connectivity. Brain and Behavior, 7(10). https://doi.org/10.1002/brb3.863

7.1 Introduction to section

Functional connectivity (FC) tracks dynamic, time-varying statistical interactions reflective of structured multi-synaptic connections, and is believed to indicate transient interactions between neuroanatomical hubs (Schölvinck et al., 2013). fMRI estimates of FC have proven particularly popular, due to the completely non-invasive nature of the technology. Many fMRI studies comparing connectivity across species have done so in the resting state, a period devoid of structured behavioral demands (Fox & Raichle, 2007; Zhang et al., 2010). Spontaneous, behaviorally independent FC generally conforms to an underlying anatomical framework and varies as a function of time (Jessica S Damoiseaux & Greicius, 2009; Deco et al., 2011). However, fMRI is limited in the temporal domain to oscillatory interactions occurring on the order of 0.01 Hz (Fox & Raichle, 2007).

In contrast, electrophysiological measurements of FC can resolve these oscillatory dynamics at millisecond time scale (Schölvinck et al., 2013). Neurophysiological studies on these oscillations have highlighted how canonical frequency band dynamics reflect different levels and modes of neural connectivity, such as thalamocortical and cortico-cortical interactions (Wang, 2010). Our examination of connectivity specifically in resting state enables comparison between species, independent of cross-species behavioral or performance discrepancies (Goulas et al., 2014; He et al., 2008). In addition to the obvious relevance of nonhuman primate research to human neurophysiology, nonhuman primates and other large animal models are frequently used as analogues for humans, especially for electrode development (Gierthmuehlen et al., 2014; Kohler et al., 2017).

Here, we aim to characterize and compare properties of resting state connectivity in the human and macaque brain using invasive electrophysiology, supporting our findings with comparable data from sheep. Our primary focus is comparison of human and macaque monkey connectivity between brain regions.

Comparative anatomy between humans and nonhuman primates

The brains of nonhuman primates resemble those of humans in cytoarchitecture (Hackett et al., 2001; Petrides & Pandya, 1999) and functional organization (Koyama et al., 2004; Rees et al., 2000). Nonhuman primates and humans share similarities in cortical neuronal density and projection patterns, both between layers and between different regions. Humans and non-human primates are broadly more similar in the relative size of various functional brain regions than humans are to more distantly related species (Krubitzer, 2007). Divergence is more in the relative sizes of the regions (Hackett et al., 2001; Krubitzer, 2007; Petrides & Pandya, 1999), in particular in the prefrontal cortex (Saleem et al., 2014).

For example, comparison of macaque and human visual system fMRI data reveals that relative to macaques, the human parietal, temporal, and frontal cortices have functionally expanded more than the occipital cortex (Buckner & Krienen, 2013; Orban et al., 2004). Overall, though, the relative level of similarity between macaque and human brains makes them a particularly useful model organism for comparative study.

Comparative functional connectivity

Macaques have often been compared directly to humans in both task-based (Grefkes et al., 2002; Joly et al., 2012) and resting state (Hutchison et al., 2013; Mantini et al., 2007; Margulies et al., 2009) fMRI studies. These extensive studies have broadly found that the species share many similarities in connectivity, especially inter-regional connectivity (Hutchison et al., 2012; Hutchison, Womelsdorf, Gati, et al., 2013; Margulies et al., 2009; Zhang et al., 2010). Patterned resting state connectivity has been observed in resting fMRI data in such diverse species as mice (Sforazzini et al., 2014), rats (Liang et al., 2011; Zhang et al., 2010), macaques (Vincent et al., 2007), and chimpanzees (Rilling et al., 2007). Monkey species in particular, with their extensive homology to humans, offer specific opportunities for studies that contribute to understanding of human resting state function (Hutchison & Everling, 2012).

Given the inherent challenges in comparing performance differences between species, investigating functional connectivity patterns in the resting state provides distinct advantages (reviewed in Hutchison & Everling, 2012). However, a true resting state is difficult to induce and characterize in animals without the use of sedation, as animals cannot be instructed to be specifically inactive. Anesthesia has been found to be only roughly equivalent to resting state in humans (Breshears et al., 2012), pigs (Tanosaki et al., 2014), macaques and other monkey species (Vincent et al., 2007), and mice (Grandjean et al., 2014). Our understanding of differences in connectivity under anesthesia relative to waking state remains incomplete, and limits the application of animal anesthesia studies to understand human resting state. Here, we circumvent this common limitation by only using waking humans and animals, and specifically delineating resting state in our animals through video monitoring.

While intracranial electrodes have long been used in monkey studies, including specifically for resting state connectivity (e.g., Fukushima et al., 2012), they have not to our knowledge been used for direct comparison to human resting state connectivity. Monkey functional (Belcher et al., 2013) and structural (Goulas et al., 2014) connectivity is observed to be globally similar to humans, but more dissimilar locally.

Advantages of electrophysiology for comparative investigation

Though resting state fMRI remains the gold standard for functional connectivity studies,

a number of potential confounds, including non-neuronal noise and artifacts introduced by motion, may impact inferences drawn from comparative studies (reviewed in Murphy, Birn, and Bandettini, 2013). More importantly, the hemodynamic BOLD signal is an indirect measure of neural activity and the cascade of physiological events linking neural activity to hemodynamics may not be homologous across species, as it heavily relies on glial function, which has tremendous cross-species variability (Oberheim et al., 2009). Electrophysiological methods, particularly invasive direct cortical recording with electrocorticography (ECoG), bypass some of these issues, though they have limits of their own.

ECoG monitoring for seizure localization in patients with intractable epilepsy provides a unique opportunity to record electrophysiological data directly from the human cortex, producing robust differentiation of all major canonical frequency bands. This oscillatory activity is far beyond the temporal resolution of fMRI. The high frequencies of the cortical spectrum (high gamma [HG], 70-200 Hz), a range that is thought to best reflect local cortical computation (Miller et al., 2009), are particularly challenging to capture clearly from outside the skull using surface recording technology such as EEG and MEG (Crone et al., 2006). Low signal-to-noise ratio (Crone, et al., 1998) and high spatial and temporal resolution make ECoG an especially valuable tool for connectivity analysis, especially for synchrony measures such as those we use here, as timing of the neural signals is captured with high fidelity by ECoG. However, the invasive nature of ECoG surgery translates to few patients and limited time to perform research, and each patient only has electrode coverage over a portion of the brain. Similar electrophysiological data from other animals, combined with data from experiments that can be performed in animals but not in humans, provide support for ECoG findings in humans.

Resting state connectivity has been successfully characterized in humans from multiple band-limited frequencies of ECoG data using a variety of measures (Casimo et al., 2016; Ko et al., 2011; Weaver et al., 2016) including coherence, amplitude correlation, phase locking, and various cross-frequency coupling approaches (Foster et al., 2015; Schölvinck et al., 2013). Recent work has indicated correspondence between electrophysiological properties recorded with ECoG, especially in 0.01-0.1 Hz modulations of the HG band, correlate well with BOLD signals, indicating spatial consistency and a possible physiological link between the phenomena (Keller et al., 2013; Ko et al., 2011; Ko et al., 2013). Although we are exclusively examining electrophysiological data here, this prior work indicates some degree of transferability to similar fMRI comparative studies exists.

ECoG has been used with nonhuman primates for large scale connectivity studies (Hutchison & Everling, 2012; X. Liu, Yanagawa, Leopold, Fujii, & Duyn, 2015; Vincent et al.,

2007; Wu et al., 2016; Yanagawa et al., 2013). However, to our knowledge, no direct comparison between humans' and nonhuman primates' resting electrophysiological connectivity has been conducted.

Sheep are also strong candidates for chronic implantation of ECoG grids, as unlike mice and rats, they have large enough cranial capacity for approximately the same size ECoG grids used in human experiments (Gierthmuehlen et al., 2014), which limits variation resulting from recording modality. Though sheep have not to our knowledge been evaluated for resting state connectivity in ECoG or any other modality, somatosensory evoked potentials have been successfully recorded in sheep implanted chronically with a microECoG grid (Gierthmuehlen et al., 2014). These grids, also used here, have smaller electrode diameters and decreased inter-electrode spacing than typical macro-scale ECoG grids used in invasive human epilepsy studies. Although this provides higher spatial resolution with less coverage, microECoG is still sensitive to population-scale potentials.

This study compares electrophysiological resting state connectivity using invasive ECoG recordings in rhesus macaque monkeys and humans, extending previous comparisons (Hutchison et al., 2013; Mantini et al., 2011; Margulies et al., 2009; Vincent et al., 2008) in the fMRI literature (reviewed in Hutchison & Everling, 2012). We focus on connectivity within and between pairs of homologous brain regions in humans and macaque monkeys. We supplement our findings in monkeys with additional data from somatosensory cortex in sheep.

7.2 Methods

Animal and human subjects

This study includes data from humans, macaque monkeys, and sheep. Age and sex data on all subjects is described in Table 7.1. All experiments described here were approved by the appropriate ethics boards: University of Washington IRB (humans), University of Washington IACUC (monkeys), Animal Committee of the University of Freiburg, Regierungspräsidium Freiburg, Baden-Wuerttemberg, and EU directive 2010/63/EU (sheep).

Electrode implants and data collection

Human subjects (Table 7.1) were undergoing long-term electrocorticographic (ECoG) monitoring with video for epilepsy at Harborview Medical Center in Seattle, WA. Patients were implanted with subdural platinum ECoG arrays (Ad-Tech, Racine, WI, or Integra Lifesciences, Plainsboro, NJ) for the clinical purpose of localizing medically intractable epilepsy. Electrodes were placed in grid and strip configurations (Figure 7.1A), with a scalp reference, 2.3mm electrode surface diameter, and 1cm inter-electrode spacing. Electrode locations in human subjects were based entirely on clinical indications. From a group of nine ECoG patients who

consented to research, electrode placement from four subjects overlapped with homologous, gross cortical structures covered by the aggregate of the monkeys' electrode placement (Figure 7.1B). Time series data were extracted from these four patients' clinical recordings taken with standard clinical epilepsy monitoring equipment (Xltek EEG Systems, Natus Medical Incorporated) sampled at 1000Hz.

Resting state electrophysiological recordings and accompanying video were obtained from three rhesus macaque monkeys (Table 7.1), which were housed at the Washington National Primate Research Center (RRID:SCR_002761) and were part of other ongoing research. Monkeys were implanted with microECoG electrodes (custom-made; 0.2mm diameter) (Figure 7.1B) and a distant skull screw ground. Monkey implant locations were determined based on the needs of other studies. Electrophysiological data were recorded with a Tucker-Davis Technologies system (Alachua, FL) at a sampling rate of 1200 Hz. Penetrating electrode arrays present in two of the three monkeys were excluded from analysis.

Sheep subjects (Table 7.1) were implanted with chronic, wireless implants (CorTec GmbH, Freiburg, Germany) consisting of a intracorporeal processing unit, a custom microECoG electrode array (0.8 mm contact diameter, 4mm inter-contact spacing, 16 recording contacts) and an extracorporeal unit for communication with the implant (Kohler et al., 2017). Chronic implantations were performed for the purpose of testing electrode design and long-term performance. The µECoG electrode arrays were placed over the somatosensory cortex and adjacent sensory areas (Figure 7.1C), which were localized by landmarks derived from prior studies using somatosensory evoked potentials (Gierthmuehlen et al., 2014) and anatomy (Dinopoulos et al., 1985; Johnson, Rubel, & Hatton, 1974). During resting state (unmoving, unrestrained) data were recorded at a sampling frequency of 1000Hz.

In all three species, video recordings (humans, monkeys) or live observations (sheep) were used to confirm that subjects were resting for an uninterrupted period with their eyes open or mostly open. In both humans and monkeys, three 10-minute periods in which the subject was quiet, unmoving, and awake were identified and labeled as rest. In sheep, equivalent 3-minute periods were identified. The identified time segments were then extracted from the larger ECoG data streams.

	Age	Sex	Electrodes (included, per area)
Human 1	43	M	108 (PFC: 31; M1: 2; PM: 10; PC: 3)
Human 2	44	M	90 (PFC: 8; M1: 1; PM: 12; PC: 5)
Human 3	20	M	64 (PFC: 0; M1: 0; PM: 0; PC: 31)
Human 4	31	F	82 (PFC: 9; M1: 1; PM: 12; PC: 9)
Monkey 1	6	M	32 (PFC: 3; PM: 6; M1: 18)
Monkey 2	5	M	20 (PFC: 4, M1: 11; PC: 2)
Monkey 3	4	M	35 (PFC: 9; PM: 10; M1: 1; PC: 8)
Sheep 1	Adult	M	16 (S1: 16)
Sheep 2	Adult	F	16 (S1: 16)

Table 7.1
Characteristics and electrodes for all subjects.

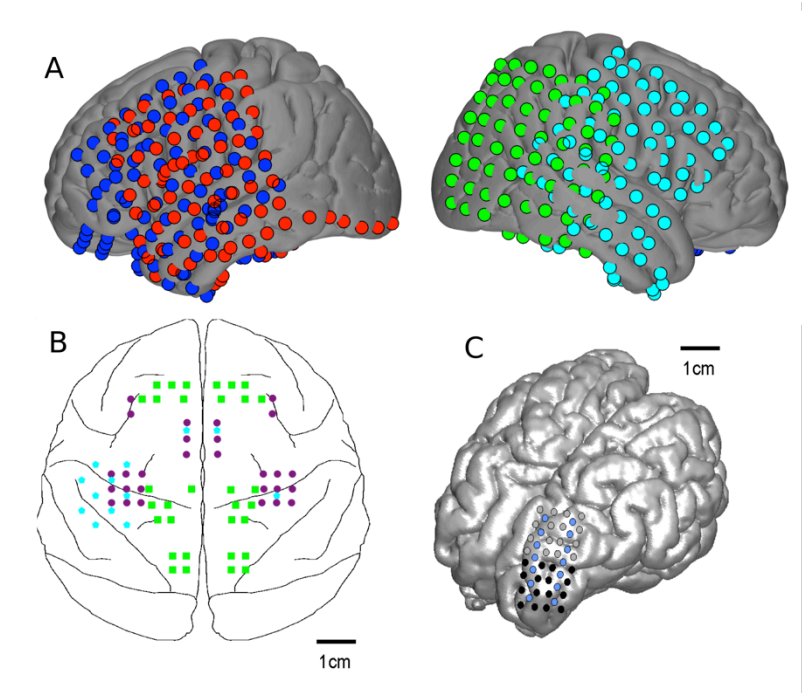

Figure 7.1
Locations of electrodes in human (A), macaque monkey (B), and sheep (C) subjects. Humans and monkeys are color coded by subject. Sheep electrodes are in the same locations for both sheep; black is recording, blue is reference, gray is unused.

Data analysis

Data processing

Data preprocessing for all species was conducted in MATLAB (MathWorks, Natick, MA; RRID:SCR_001622) computing environment as described in Casimo et al., (2016). This included common average re-referencing, notch filtering to remove line noise (humans, monkeys: 60Hz and harmonics; sheep: 50Hz and harmonics), and exclusion of interictal

artifacts in human subjects using clinical labels. Human and monkey recordings were resampled to 600Hz (Casimo et al., 2016). For all recordings, amplitude and phase angle were extracted using a non-analytic Morlet wavelet with ¼ octave resolution for pseudo-frequencies 1-200Hz. We then averaged across the phases and amplitudes for the frequency bins that fell within each canonical frequency band of interest (delta, 0-4Hz; theta, 4-8Hz; alpha, 8-12Hz; beta, 13-30Hz; gamma, 30-70Hz; HG, 70-200Hz).

Anatomical locations of humans' electrode locations were identified as previously described (Casimo et al., 2016). We aligned and registered each subject's preoperative MRI with a postoperative CT indicating electrode position using BioImage Suite (Papademetris et al., 2006; RRID:SCR_002986). We then registered each native T1 MRI to an MNI atlas in FreeSurfer (Reuter, Rosas, & Fischl, 2010; with MNI 305 atlas; RRID:SCR_001847), and applied the resulting transformation matrix to the electrode coordinates. Brodmann areas (BA) for each electrode were identified by registering the MNI-space electrode positions in Talairach space and labeled with the Talairach Daemon Client (Lancaster et al., 2000; RRID:SCR_000448). Monkey and sheep electrode locations were identified based on stereotaxic coordinates measured during the surgical placement of electrodes and compared to standard macaque (Saleem & Logothetis, 2012). Sheep electrode locations were estimated by landmarks (e.g., bregma) and previous somatosensory evoked potential studies, as MRI and stereotaxic surgical equipment were unavailable (Dinopoulos et al., 1985; Gierthmuehlen et al., 2014; Johnson et al., 1974).

Connectivity and individual-level statistical analysis

Pairwise connectivity across the entire resting state time series from each individual's recording was calculated for all electrode pairs within each individual human, monkey, and sheep. We calculated two connectivity measures: 1) amplitude-amplitude Pearson correlation, a measure of the consistency of variation in amplitude between two signals over time, and 2) phase-locking value (PLV; Lachaux et al., 1999), a measure of the consistency of the difference in phase between two oscillating signals over time (PLV = $\Sigma e^{i*\Delta\theta}$ where $\Delta\theta$ is the instantaneous difference in phase between the signals from any two electrodes). PLV is specifically sensitive to the degree of synchronization between two signals, separately from any offset or delay between the signals. Both connectivity measures specifically quantify synchrony between regions; phase synchrony implies a degree of functional integration between regions or parts of a region (Sauseng & Klimesch, 2008).

Statistical significance for pairwise connectivity was established using a nonparametric, maximum-value permutation test (Weaver et al., 2016). We generated surrogate data by

randomly shuffling the phase, for PLV, or amplitude, for amplitude correlation, of the broadband signal, re-extracting the frequency bands, and recalculating the connectivity measures as before. Each iteration of the permutation procedure produced (n x n)/2 – n interactions, where n is the number of electrodes for a given subject. In humans, this ranged from 64-108 and in monkeys from 20-35, and 32 in both sheep (Table 7.1). We generated a null distribution of values by retaining the top 50 highest values, and repeated the process 20 times, for a total of 1000 values in the statistical distribution. Because we retained the maximum value across all channel pairs, this approach is highly conservative, as it generates a null distribution of maximum connectivity values. We then identified the 95th percentile of this null distribution (i.e., alpha level of $p<0.05$ of random interactions) as the threshold for retaining values as statistically significant, and discarded connectivity values below this threshold. As the method for generating the null distribution includes the entire array of signals, no further correction for multiple comparisons is needed (Casimo et al., 2016).

Anatomical labeling and group-level statistical analysis

Using the electrodes' anatomical location as identified stereotactically in surgery, we grouped the monkey electrodes into four regions: prefrontal cortex, premotor and supplementary motor cortex, primary motor cortex, and parietal cortex (Thiebaut de Schotten et al., 2012). We then identified all electrodes in the human subjects clustered in these four areas for comparison, as the anatomical extent of the monkey electrodes was more limited. Electrodes were then pooled 1) between each pair of anatomical regions and 2) within a single anatomical region (self-pairing), and finally 3) channel-pair labels were binned across subjects within a species (Figure 7.2A).

Between species, we only quantitatively compared the mean connectivity values (within or between regions) that were statistically significant in both species (Figure 7.2B). There are multiple reasons that a connectivity value may not be statistically significant, most notably in this case if a pair of regions have relatively few or no electrode pairs linking them, so the lack of attendant connectivity value does not imply actual lack of connectivity (Laumann et al., 2015). Consequently, we restricted further analysis to comparisons where we could determine significant connectivity values for both species.

We performed a nonparametric Mann-Whitney U-test comparing humans and monkeys for each region-to-region or within-region connectivity strength. We corrected for multiple comparisons (total of ten comparisons, region-to-region or within-region, per band) with the Bonferroni method within each frequency band, and connections with statistically significant differences were retained (Figure 7.2C). Consequently, all results and figures (Figure 7.3) below

show connectivity between and within regions where humans and monkeys were both statistically significant separately, and were significantly different from each other.

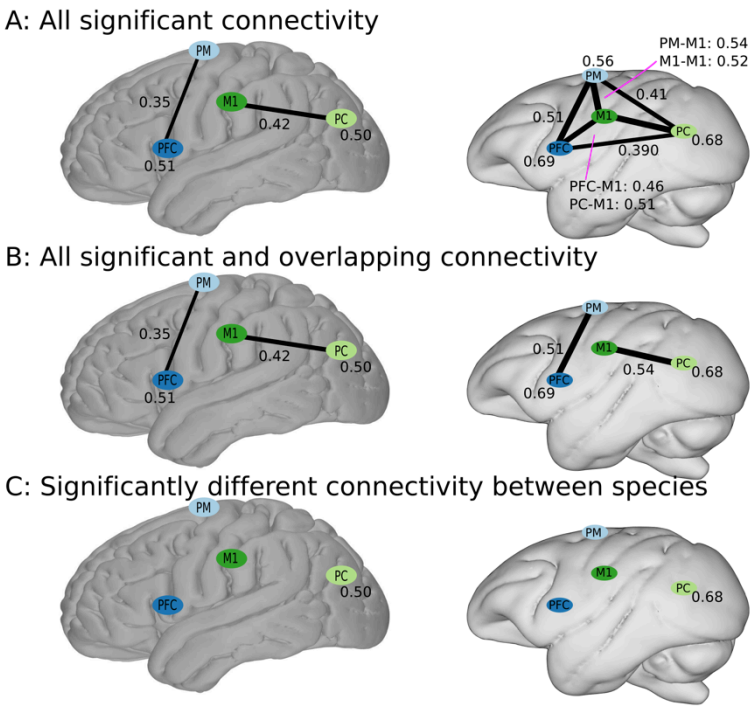

Figure 7.2
Demonstration of statistical analysis process in comparing humans and macaque monkeys. A: for each species, connectivity values for a given region or pair of regions are pooled within an individual, filtered for statistical significance using a permutation test, and significant values are retained and pooled among all individuals. B: connections between or within regions that had statistically significance in both species were retained. C: connectivity values for a given connection are compared between species with a Mann-Whitney U-test, and statistically significant comparisons' values are retained. Correction for multiple comparisons was with the Bonferroni method, correcting for the number of contrasts within each frequency band for each connectivity measure.

7.3 Results

Macaque monkeys and humans

The main focus of our approach is a comparative analysis of humans and macaque monkey connectivity estimates spanning homologous, gross cortical loci. We assessed differences in connectivity strength between the two species in 60 (four intra-region connections + six inter-region connections, at six frequency bands) potential connections, per connectivity measure, in each species. Overwhelmingly, the comparative cross-species analyses yielded similar results across frequency bands and connectivity approaches. However, we noticed several inter-species incongruities in the spatial and frequency distribution. Here we will focus on these differences.

Both connectivity measures are specifically measures of synchrony. They can be variously interpreted as the synchronized exchange of information between regions, combining information from multiple regions (such as in sensory integration), precise timing, or sequential activity (J.-P. Lachaux et al., 2012; Sauseng & Klimesch, 2008). From here forward, when we discuss "correlation strength" or "PLV strength," we refer to the average connectivity between the pair of regions or within the region mentioned. This specifically is the mean of the connectivity values between all of the pairs of electrodes contained within or between the defined regions. Consequently, connectivity values presented represent the mean synchrony observed between all the relevant regions.

Overall, inter-species connectivity differences were overwhelmingly observed in connectivity within a single region, far more often than in connectivity between two different regions. Further, two general trends appeared: first, correlation showed relatively larger differences in magnitude between species than PLV did between species. Second, the differences between species were more widespread across cortex in PLV relative to correlation. In both measures, the species differed more often in lower frequencies than in higher frequencies. However, it is important to note that the differences were the exception rather than the rule, and correlation and PLV strengths were not different between the species in 80% and 72% of comparisons, respectively.

Inter-species differences between correlation and phase locking

We first examined broad trends of differences between the two connectivity measures. Of the 60 comparisons of connectivity strength between species, 12 inter-species comparisons of correlation strength (20%) were significantly different (Figure 7.3A-F). All of these were within-region correlations, which represent 40% of the possible connections. In contrast, 17 comparisons of PLV strength were significantly different between species (28%) (Figure 7.3G-L), 42% more than there were differing correlations (Figure 7.3A-F). Twelve of these (71%) were within-region PLVs (Figure 7.3G-L). The differences between species are more widespread in phase-based than in amplitude-based synchrony, but the magnitude of the difference was generally larger in correlations than in PLVs.

Of the 12 inter-species differences in correlation strength, four of those connections' PLV strengths were not different between the species (33% of differing correlations). These were observed in intra-PFC connectivity in delta (Figure 7.3A, G) and theta (Figure 7.3B, H), and intra-premotor cortex connectivity in gamma (Figure 7.3E, K) and HG (Figure 7.3F, L).

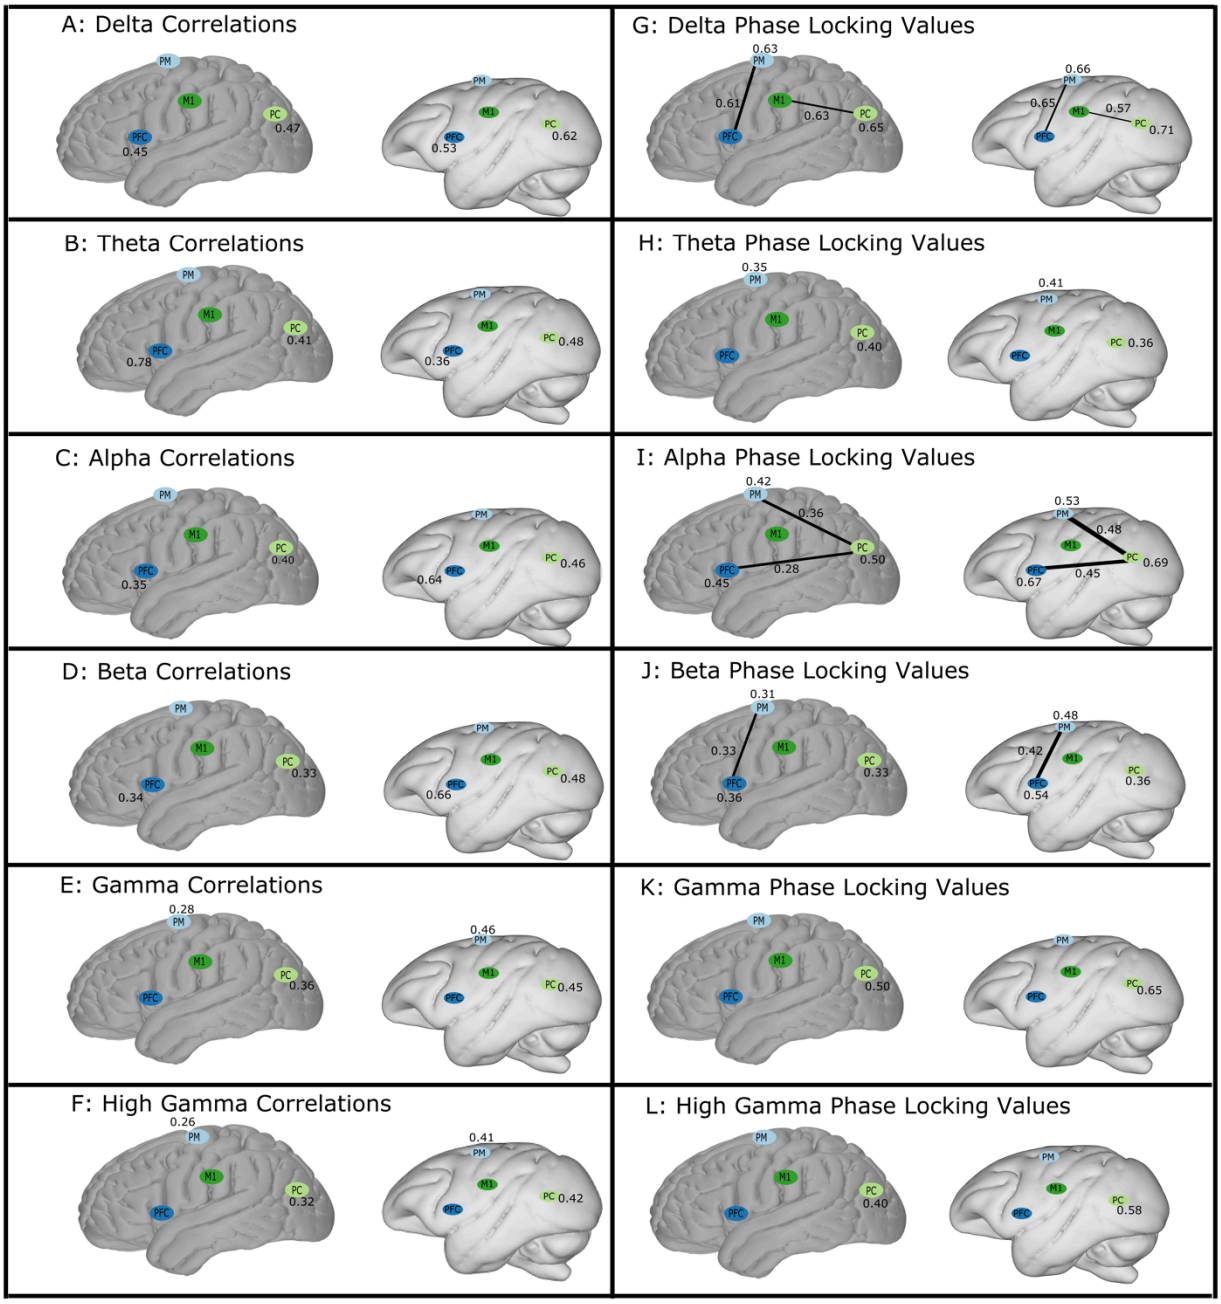

Figure 7.3
Comparison of strength of correlation (A-F) and phase locking value (G-L) in humans and macaque monkeys. Only connectivity values that were significantly different between the species are shown, as illustrated in Figure 7.2. Regions are prefrontal cortex (PFC), premotor cortex (PM), motor cortex (M1), and parietal cortex (PC). Connectivity values are shown for the six potential edges connecting different pairs of regions, and for within the four regions, for a total of ten possible connectivity values per connectivity measure and per frequency. Values shown next to each value or region represent connectivity strength.

Of the 17 inter-species differences in PLV strength, nine of those connections' correlation strengths were not different (53% of PLV inter-species differences). These were intra-premotor cortex connectivity in delta (Figure 7.3A, G), alpha (Figure 7.3C, I), and beta (Figure 7.3D, J); premotor-PFC connectivity in delta (Figure 7.3A, G) and beta (Figure 7.3D, J); M1-parietal cortex connectivity in delta (Figure 7.3A, G); intra-premotor cortex connectivity in theta (Figure 7.3B, H); premotor-parietal cortex connectivity in alpha (Figure 7.3C, I); and parietal-PFC connectivity in alpha (Figure 7.3C, I). Five of these nine inconsistencies between the connectivity measures were between two different regions, and four were connectivity within a region.

Inter-species differences between frequencies

We also observe differences between the species in both connectivity measures in the six frequency bands. As noted, there were 71% as many inter-species comparisons with significantly different correlation strengths as significantly different PLV strengths. Consequently, knowing there were disparities between the connectivity measures, we assessed patterns in both measures in the six frequency bands.

The frequency band with the least agreement between the connectivity measures is alpha, with five significant differences within and between regions in PLV strengths (Figure 7.3I), and two in correlation strengths (Figure 7.3C), both of which were also different in PLV strength. In contrast, in both gamma and HG, there were few differences between the species: two significantly different correlation strengths between species in each frequency (Figure 7.3E-F), and only one PLV strength was different between species in each frequency (Figure 7.3K-L), which overlapped with correlation differences in both. Finally, there were just two significant differences in theta in both correlation (Figure 7.3B) and PLV (Figure 7.3H) strength, but only one overlapped in space between measures. In no frequency band did either of the two connectivity measures completely agree.

Above we summarized aggregate differences between the species across all frequencies. Now we separate the differences by frequency: in delta, 20% of correlation strengths (Figure 7.3A) and 40% of PLV strengths (Figure 7.3G) were different between the species; in theta, 20% of correlation (Figure 7.3B) and 20% of PLV strengths (Figure 7.3H); in alpha, 20% of correlation (Figure 7.3C) and 50% of PLV strengths (Figure 7.3I); in beta, 20% of correlation (Figure 7.3D) and 30% of PLV strength (Figure 7.3J); in gamma, 20% of correlation (Figure 7.3E) and 10% of PLV strengths (Figure 7.3K); and in HG, 20% of correlation (Figure 7.3F) and 10% of PLV strengths (Figure 7.3L). All but one (within parietal cortex) of these differing connection strengths were more synchronous in monkeys than in humans.

Inter-species differences between regions

Finally, we examine differences within regions and between pairs of regions at each frequency band and in each connectivity measure. Again, as noted above, there were very few statistically significant differences in inter-regional connectivity strength. For connectivity between pairs of different regions, 14% of PLV strengths were different between species, but no correlation strengths were different. In contrast, for comparisons of connectivity strength within a single region, 50% of correlation strengths and 50% of PLV strengths were different between species.

Within prefrontal cortex, we observed differences between the species in alpha (Figure 7.3C, I) and beta (Figure 7.3D, J) correlation and PLV, as well as in delta correlation (Figure 7.3A). Within premotor cortex, the species were significantly different in delta through beta PLV strength (Figure 7.3G-J), and gamma and HG correlation strength (Figure 7.3E-F). Within motor cortex, the species were not significantly different in either measure at any frequency (Figure 7.3). Within parietal cortex, all comparisons between species were significantly different (Figure 7.3); human connectivity was stronger than monkey connectivity only in theta PLV, the only connection more synchronous in humans (Figure 7.3H).

We observed the following additional connectivity patterns across the six pairs of different regions, with monkey connectivity stronger in all connections. Between PFC and premotor cortex, the species differed in delta (Figure 7.3G) and beta (Figure 7.3J) PLV. Between PFC and parietal cortex, and also between premotor and parietal cortex, the species differed only in alpha PLV strength (Figure 7.3I). Finally, between motor and parietal cortex, the species differed in only delta PLV (Figure 7.3G). Between PFC and motor cortex, and between premotor and motor cortex, there were no significant differences in connectivity between the species in either measure at any frequency (Figure 7.3).

As noted above, of the 29 connections in both measures that were significantly different between the species, just one (theta PLV strength within parietal cortex) was stronger in humans.

Inter-species differences between humans and sheep

We also assessed intra-somatosensory (parietal) cortex connectivity in two sheep. As noted above, human and monkey connectivity within parietal cortex was significantly different in all contrasts but delta correlation (Figure 7.3). Where human and monkey intra-parietal (somatosensory) connectivity was significantly different, connectivity was stronger in monkeys in all except theta PLV, the only instance where human connectivity was stronger than monkey connectivity (Figure 7.3H). In contrast, sheep intra-somatosensory correlation in delta, gamma,

and high gamma, and PLV in all frequencies were significantly lower than in humans (Figure 7.4A, B). In no frequency, in either correlation or PLV, was sheep intra-somatosensory connectivity stronger than human connectivity (Figure 7.4A, B).

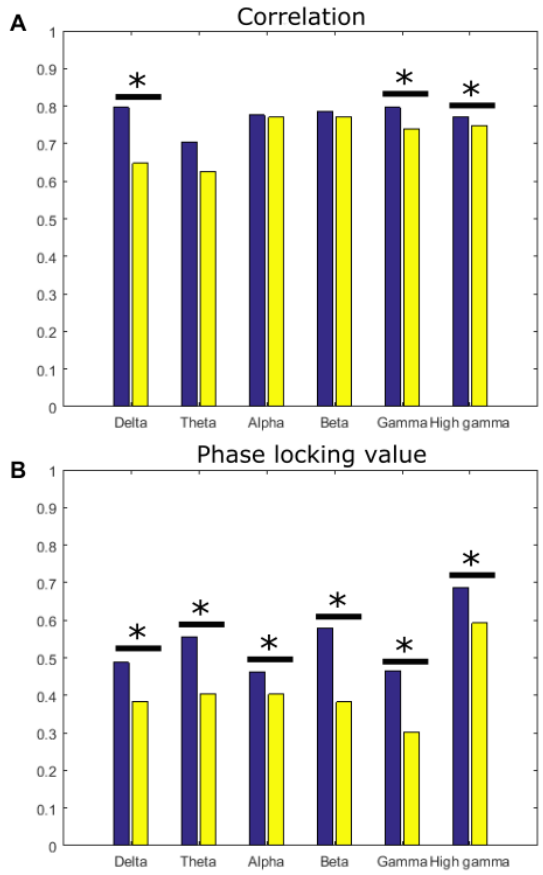

Figure 7.4
Comparison of strength of correlation (A) and phase locking value (B) between the same four humans used above (blue bars) and two sheep (yellow bars). Connectivity strength between species is compared with a Mann-Whitney U-test, and statistically significant differences are marked with a star.

7.4 Discussion

We compared resting state connectivity using subdural electrophysiological methods within and between cortical regions, focusing on humans and macaque monkeys, and between humans and sheep. We observe a general pattern of functional connectivity in a behaviorally unstructured state that was generally similar between different regions when comparing human and macaque monkeys, with greater discrepancies occurring within homologous regions than between regions.

We speculate that reduced local synchrony in humans relative to monkeys may reflect the balance between heterogeneity and integration in function within and between cortical

regions. Though we cannot rule out contributions from the differing electrode diameters, the similarity in the two species in connectivity between pairs of different regions implies that the increased heterogeneity in humans is local, and that functional interactions between different regions are largely preserved.

Among the differences between species, the greatest disparity was within-parietal cortex connectivity, which differed between humans and macaques in every frequency band in both connectivity measures (Table 7.1, Figure 7.3). Within-prefrontal cortex connectivity also differed in multiple frequency bands in both connectivity measures (Figure 7.3). In the minority of instances in which the species differ, human synchrony may be generally weaker than monkey connectivity potentially because of an increased functional granularity of functional parcellation or more fine-grained regional specialization in human brain (Saleem et al., 2014; Thiebaut de Schotten et al., 2012). The prefrontal cortex is notably larger and more cytoarchitecturally complex in humans than in other primate species (Buckner & Krienen, 2013; Thiebaut de Schotten et al., 2012). This might result in greater functional heterogeneity in the neuronal populations detected by each electrode, leading to decreased synchrony and therefore connectivity strength.

This conclusion is further supported by the observation that no differences between the species were identified in within- primary motor cortex connectivity, one of the most evolutionarily old and well-conserved parcels of neocortex (Arendt et al., 2015). In sheep and human parietal cortex, connectivity patterns also differed, as they did in human and monkey within-parietal comparisons, but human connectivity was stronger than sheep. Relative connectivity strength may be attributable to the number of different afferent sources to that cortical parcel, level of functional integration within the region, or other sources of heterogeneity.

Differences between connectivity measures

Collectively, the humans and monkeys are more similar in resting state connectivity patterns between different regions than within a single region. Further, sheep and human connectivity differences, evaluated only within one region, were dissimilar in both connectivity measures much as humans and monkey values were for a single region.

The two connectivity measures we evaluated, amplitude-amplitude correlation and PLV, were generally similar, with the bulk of the conflict between the two measures in monkey-human comparison appearing in between-region connectivity. This may be a result of the fact that signal phase becomes more desynchronized with respect to amplitude over greater distances (i.e. between regions) (von Stein & Sarnthein, 2000). Other potential factors that may account for these differences include that PLV is simply more sensitive to smaller inconsistencies in

phase difference than correlation is to small variation in amplitude, or that differences between species in the number of electrodes in any given region.

Differences between regions of the brain

The vast majority of comparisons of regions' connectivity in humans and macaques are not significantly different. This is consistent with prior fMRI comparisons between finding extensive homology between human and monkey structural and functional connectivity (Goulas et al., 2014; Hutchison et al., 2012; Hutchison & Everling, 2012; Margulies et al., 2009; Zhang et al., 2010). Our findings in electrophysiology concur with and support these fMRI observations, as the species' connectivity patterns, based on phase and amplitude interactions across a wide range of frequency bands, are generally similar.

The differences between species were overwhelmingly concentrated in within-region connectivity, rather than between regions. This suggests that specific local synchrony differs more substantively than long-distance, inter-regional functional connectivity. However, both amplitude-amplitude coupling and PLV are linear measures of interaction, and are not sensitive to cross-frequency interactions. Cross-frequency coupling (CFC) is now believed to underlie long-distance cortico-cortical entrainment to a greater degree possibly as a neurobiological mechanism facilitating information transfer (Canolty & Knight, 2010; Weaver et al., 2016). Further research using CFC may reveal features that are not accessible in linear measures.

We speculate that the strong similarities between species in motor cortex connectivity may indicate conserved function of the motor systems between humans and macaque monkeys, extending to multiple motor-related areas that are functionally homologous between these two species. This is consistent with prior work using tract-tracing methodology (Hackett et al., 2001; Thiebaut de Schotten et al., 2012).

In contrast, both within-premotor and within-parietal cortex connectivity was markedly different between the species across multiple frequency bands (Figure 7.3). For within-premotor cortex, this included motor-linked beta, potentially indicating localized differences in the functional architecture relating to motor planning behavior. For instance, humans may have a more complex or extensive motor repertoire than macaques, or there may be greater complexity in motor decision-making processes. Connectivity within parietal cortex had the most pervasive differences between the species, with statistically significant differences in all frequency bands and both connectivity measures. Human parietal association areas may incorporate more afferents, more sophisticated sensory integration, or more sensory processing function relative to monkeys.

Contributions from sheep recordings

We found the most pervasive disparity between humans and macaque monkeys was in within-parietal cortex connectivity (Figure 7.3). Our findings in sheep extend this result. The equivalent comparison of parietal cortex between humans and sheep differed in delta, theta, alpha, and beta frequencies in both connectivity measures (Figure 7.4). However, in contrast to the human-monkey comparisons where monkey connectivity was almost universally stronger, human connectivity in both measures was always stronger than in sheep. We suggest that low synchrony in sheep may be indicative of weaker sensory integration than in humans, as this observation specifically pertains to parietal cortex. We are not able to draw conclusions about overall sheep connectivity patterns due to lack of coverage. However, our findings indicate that sheep are suitable for investigations of connectivity, such as for stimulation-induced connectivity changes, that then can be translated to primates or humans for further investigation.

Limitations

The primary limitation of this study derives from the differing sensor properties used across the three species. Although all recordings were made with subdural electrodes, the precise composition and size of the electrodes and their spacing varied. In particular, the monkey and sheep electrodes were smaller than the humans', and thus each electrode samples from fewer neurons. We excluded the depth electrodes also present in two of the monkeys; a comparison of human and monkey depth electrode recordings would extend our findings.

This study, as with all ECoG studies in humans, and some degree in animals, was further limited by the spatial placement and coverage of electrodes. The macaque monkey and sheep electrode placements were determined by the needs of the other studies. The human electrode placement was determined by their clinical needs, though we were able to select human subjects whose coverage generally overlapped with that of the monkeys. We suggest further research include additional subjects, human and animal, to expand spatial coverage.

In addition, the resting state is not as well defined in animals with respect to human subjects. We utilized video of animal behavior to distinguish periods of inactivity from activity. However, there is less control over animal models when ensuring complete adherence to a true resting state. In sheep, we were only able to obtain 3-4 minute long recordings, in contrast to 10 minutes in humans and monkeys. This, as well as the limited spatial extent of the sheep electrode placement, leaves opportunity for further investigation.

Conclusions and further studies

In conclusion, we find that human and monkey electrophysiological functional connectivity in the resting state is largely similar, particularly in inter-regional connectivity. The minority of significant differences occur mostly in local, intra-regional connectivity, again

consistent with prior studies finding consistency between humans and nonhuman primates in inter-regional connectivity (Hutchison et al., 2012; Hutchison, Womelsdorf, Gati, et al., 2013; Margulies et al., 2009; Zhang et al., 2010). Further, in nearly all of the disparate regions, the strength of connectivity was greater in monkeys. The similarity in synchrony noted, especially between humans and monkeys, suggests a physiological link between the species that bridges the cognitive and behavioral gap.

While we cannot rule out possible differences in sensor characteristics, the variation in synchrony levels may reflect differences in the degree of functional segregation between species. We were not able to capture fMRI data in this investigation, but these results are consistent with those of extensive previous imaging-based investigations (Hutchison et al., 2012; Hutchison, Womelsdorf, Allen, et al., 2013; Margulies et al., 2009; Zhang et al., 2010).

We suggest that this similarity supports the application of PLV and amplitude correlation in monkeys as one means of indirectly probing the architecture of neural circuits in humans. If further validated, comparability between species can enable findings from an animal model such as a monkey or a sheep, where more invasive electrophysiology (e.g. tracing, invasive implantations, single-unit recordings) are readily available, to be used to supplement electrophysiological or imaging techniques usable in humans. This may be applicable in both studies of the origins of resting state connectivity or of other task-related connectivity.

While sheep and humans differ more sharply than monkeys and humans do, an understanding of sheep connectivity properties aids in the use of sheep to study basic principles of connectivity, particularly in instances such as cortical stimulation where the ability to do highly invasive studies in sheep can provide guidance for further studies. Experimental outcomes in sheep can then be validated in other species, as we have done here. Sheep are a also popular model for testing a variety of neurological diseases including stroke (Wells et al., 2012), Parkinson's disease (Baskin et al., 1994; Hammock et al., 1989), and focal epilepsy (Opdam et al., 2002), and our findings potentially contribute to the cross-species validation of these disease models.

A significant degree of cross-frequency coupling has been observed in resting state in both humans (Weaver et al., 2016) and monkeys (Schroeder & Lakatos, 2009); differences in this phenomenon may also related to the single-frequency effects observed here. Further work should evaluate cross-frequency coupling (including phase-amplitude coupling or other measures) to interrogate this possibility.

8. Abbreviations

amHG: amplitude modulations of the high gamma band
BCI: brain computer interface
BOLD: blood-oxygen-level dependent
dlPFC: dorsolateral prefrontal cortex
ECoG: electrocorticography
EEG: electroencephalography
fMRI: functional magnetic resonance imaging
HG: high gamma (70-200Hz)
ICC: intraclass correlation
IFG: inferior frontal gyrus
I/MTG: inferior and middle temporal gyri
IPL: inferior parietal lobule
M1: motor cortex
MEG: magnetoencephalography
microECoG: microelectrocorticography
PC: parietal cortex
PFC: prefrontal cortex
PH-ER: parahippocampal and entorhinal gyri
PLV: phase locking value
PM: premotor cortex
PSI: phase slope index
ROI: region of interest
rsfMRI: resting state functional magnetic resonance imaging
RSN: resting state network
S1: primary somatosensory cortex
STG: superior temporal gyrus
SPL: superior parietal lobule

9. Acknowledgements

So many people have contributed to my success in graduate school and all of the projects I have worked on, both in my research and in everything else I did. I would like to thank a few who were particularly helpful, though I will inevitably fall short both of thanking everyone who deserves recognition, and of adequately expressing my gratitude.

First and foremost, I want to thank my advisor Jeff Ojemann and mentor Kurt Weaver, for their guidance, support, and commitment to my development as a scientist. I would also like to thank the rest of my supervisory committee – Raj Rao, Adrienne Fairhall, Chet Moritz, and Ric Robinson.

I also want to thank the members of the lab who I have worked with over the last five years: David Caldwell, Kelly Collins, Jenny Cronin, Felix Darvas, Chao-hung Clark Kuo, Gabriel Obregon, Miah Wander, Nancy Wang, Nile Wilson, and James Wu.

Teaching and outreach have been integral to my time in graduate school. I would like to thank Ric Robinson, Eric Chudler, and Dave Wolczyk, as well as the entire Neuroscience Community Outreach Group, for providing me with opportunities to develop these skills.

Finally, I would like to thank my family and friends, who have supported and believed in me through the process of completing 21st grade.

My graduate training was supported by:
University of Washington Institute of Neuroengineering: WRF Innovation Graduate Fellowship in Neuroengineering
NIH-UW Computational Neuroscience Training Program
Center for Sensorimotor Neural Engineering

10. References

Albert, N. B. … Miall, R. C. (2009). Resting state networks and memory consolidation. *Communicative & Integrative Biology*, *2*(6), 530–532. https://doi.org/10.4161/cib.2.6.9612

Albert, N. B. … Miall, R. C. (2009). The Resting Human Brain and Motor Learning. *Current Biology*, *19*(12), 1023–1027. https://doi.org/10.1016/j.cub.2009.04.028

Allen, E. A. … Calhoun, V. D. (2014). Tracking whole-brain connectivity dynamics in the resting state. *Cerebral Cortex*, *24*(3), 663–676. https://doi.org/10.1093/cercor/bhs352

Amunts, K. … Zilles, K. (2007). Cytoarchitecture of the cerebral cortex - more than localization. *Neuroimage*, *37*(4), 1061–1065. https://doi.org/10.1016/j.neuroimage.2007.02.037

Anderson, B. J. … Greenough, W. T. (1996). Motor-skill learning: changes in synaptic organization of the rat cerebellar cortex. *Neurobiology of Learning and Memory*, *66*(2), 221–229. https://doi.org/10.1006/nlme.1996.0062

Arendt, D. … Marlow, H. (2015). From nerve net to nerve ring, nerve cord and brain — evolution of the nervous system. *Nature Reviews Neuroscience*, *17*(1), 61–72. https://doi.org/10.1038/nrn.2015.15

Astolfi, L. … Babiloni, F. (2009). Estimation of effective and functional cortical connectivity from neuroelectric and hemodynamic recordings. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *17*(3), 224–233. https://doi.org/10.1109/TNSRE.2008.2010472

Baker, A. P. … Woolrich, M. (2014). Fast transient networks in spontaneous human brain activity. *ELife*, *3*, e01867–e01867. https://doi.org/10.7554/eLife.01867

Baskin, D. S. … Grossman, R. G. (1994). Development of a Model for Parkinson's Disease in Sheep Using Unilateral Intracarotid Injection of MPTP via Slow Continuous Infusion. *Life Sciences*, *54*(7), 471–479.

Bassett, D. S. … Bullmore, E. (2006). Adaptive reconfiguration of fractal small-world human brain functional networks. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(51), 19518–19523. https://doi.org/10.1073/pnas.0606005103

Bassett, D. S. … Grafton, S. T. (2011). Dynamic reconfiguration of human brain networks during learning. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(18), 7641–7646. https://doi.org/10.1073/pnas.1018985108

Bastos, A. M., & Schoffelen, J.-M. (2016). A Tutorial Review of Functional Connectivity Analysis Methods and Their Interpretational Pitfalls. *Frontiers in Systems Neuroscience*, *9*(January), 1–23. https://doi.org/10.3389/fnsys.2015.00175

Belcher, A. M. … Stein, E. A. (2013). Large-scale brain networks in the awake, truly resting marmoset monkey. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *33*(42), 16796–16804. https://doi.org/10.1523/JNEUROSCI.3146-13.2013

Bensmaia, S. J., & Miller, L. E. (2014). Restoring sensorimotor function through intracortical interfaces: progress and looming challenges. *Nature Reviews Neuroscience*, *15*(5), 313–325. https://doi.org/10.1038/nrn3724

Betzel, R. F. … Sporns, O. (2012). Synchronization dynamics and evidence for a repertoire of network states in resting EEG. *Frontiers in Computational Neuroscience*, *6*(September), 1–13. https://doi.org/10.3389/fncom.2012.00074

Billinger, M. … Müller-Putz, G. R. (2013). Single-trial connectivity estimation for classification of motor imagery data. *Journal of Neural Engineering*, *10*(4), 046006. https://doi.org/10.1088/1741-2560/10/4/046006

Birbaumer, N. … Flor, H. (1999). A spelling device for the paralysed. *Nature*, *398*(6725), 297–298. https://doi.org/10.1038/18581

Biswal, B. B. … Milham, M. P. (2010). Toward discovery science of human brain function. *PNAS*, *107*(10), 4734–4739. https://doi.org/10.1073/pnas.0911855107

Blakely, T. … Ojemann, J. G. (2009). Robust, long-term control of an electrocorticographic

brain-computer interface with fixed parameters. *Neurosurgical Focus*, *27*(1), E13. https://doi.org/10.3171/2009.4.FOCUS0977
Blakely, T. … Rao, R. P. N. (2014). Short-time windowed covariance: a metric for identifying non-stationary, event-related covariant cortical sites. *Journal of Neuroscience Methods*, *222*, 24–33. https://doi.org/10.1016/j.jneumeth.2013.10.005
Breshears, J. D. … Leuthardt, E. C. (2012). Mapping sensorimotor cortex with slow cortical potential resting-state networks while awake and under anesthesia. *Neurosurgery*, *71*(2), 305–316. https://doi.org/10.1227/NEU.0b013e318258e5d1
Brookes, M. J. … Barnes, G. R. (2014). Measuring temporal, spectral and spatial changes in electrophysiological brain network connectivity. *NeuroImage*, *91*, 282–299. https://doi.org/10.1016/j.neuroimage.2013.12.066
Buchel, C. … Friston, K. (1999). The predictive value of changes in effective connectivity for human learning. *Science*, *283*(March), 1538–1542. https://doi.org/10.1126/science.283.5407.1538
Buckner, R. L., & Krienen, F. M. (2013). The evolution of distributed association networks in the human brain. *Trends in Cognitive Sciences*, *17*(12), 648–665. https://doi.org/10.1016/j.tics.2013.09.017
Buckner, R. L. … Schacter, D. L. (2008). The brain's default network: anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences*, *1124*, 1–38. https://doi.org/10.1196/annals.1440.011
Bullmore, E. T. … Brammer, M. J. (1999). Global, voxel, and cluster tests, by theory and permutation, for a difference between two groups of structural MR images of the brain. *IEEE Transactions on Medical Imaging*, *18*(1), 32–42. https://doi.org/10.1109/42.750253
Buttfield, A. … Millán, J. D. R. (2006). Towards a robust BCI: error potentials and online learning. *IEEE Transactions on Neural Systems and Rehabilitation Engineering : A Publication of the IEEE Engineering in Medicine and Biology Society*, *14*(2), 164–168. https://doi.org/10.1109/TNSRE.2006.875555
Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, *304*(5679), 1926–1929. https://doi.org/10.1126/science.1099745
Buzsáki, G. … Koch, C. (2012). The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, *13*(6), 407–420. https://doi.org/10.1038/nrn3241
Cabral, J. … Deco, G. (2014). Exploring the network dynamics underlying brain activity during rest. *Progress in Neurobiology*, *114C*, 102–131. https://doi.org/10.1016/j.pneurobio.2013.12.005
Canolty, R. T., & Knight, R. T. (2010). The functional role of cross-frequency coupling. *Trends in Cognitive Sciences*, *14*(11), 506–515. https://doi.org/10.1016/j.tics.2010.09.001
Carmena, J. M. … Nicolelis, M. A. L. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biology*, *1*(2), 193–208. https://doi.org/10.1371/journal.pbio.0000042
Casimo, K. … Weaver, K. E. (n.d.). Spontaneous variation in electrocorticographic resting state connectivity. *Brain Connectivity*.
Casimo, K. … Weaver, K. E. (2016). Regional Patterns of Cortical Phase Synchrony in the Resting State. *Brain Connectivity*, *6*(6), 470–481. https://doi.org/10.1089/brain.2015.362
Censor, N. … Cohen, L. G. (2014). Interference with existing memories alters offline intrinsic functional brain connectivity. *Neuron*, *81*(1), 69–76. https://doi.org/doi:10.1016/j.neuron.2013.10.042
Chan, J. C. K., & LaPaglia, J. A. (2013). Impairing existing declarative memory in humans by disrupting reconsolidation. *Proceedings of the National Academy of Sciences*, *110*(23), 9309–9313. https://doi.org/10.1073/pnas.1218472110
Chang, E. F. (2015). Towards large-scale, human-based, mesoscopic neurotechnologies.

Neuron, *86*(1), 68–78. https://doi.org/10.1016/j.neuron.2015.03.037
Chawla, D. … Friston, K. (1999). The relationship between synchronization among neuronal populations and their mean activity levels. *Neural Computation*, *11*(6), 1389–1411.
Chein, J. M., & Schneider, W. (2005). Neuroimaging studies of practice-related change: fMRI and meta-analytic evidence of a domain-general control network for learning. *Cognitive Brain Research*, *25*(3), 607–623. https://doi.org/10.1016/j.cogbrainres.2005.08.013
Chu, C. J. … Cash, S. S. (2012). Emergence of Stable Functional Networks in Long-Term Human Electroencephalography. *Journal of Neuroscience*, *32*(8), 2703–2713. https://doi.org/10.1523/JNEUROSCI.5669-11.2012
Collinger, J. L. … Schwartz, A. B. (2013). High-performance neuroprosthetic control by an individual with tetraplegia. *Lancet*, *381*(9866), 557–564. https://doi.org/10.1016/S0140-6736(12)61816-9
Conner, C. R. … Tandon, N. (2011). Variability of the relationship between electrophysiology and BOLD-fMRI across cortical regions in humans. *Journal of Neuroscience*, *31*(36), 12855–12865. https://doi.org/10.1523/JNEUROSCI.1457-11.2011
Crone, N. E. … Franaszczuk, P. J. (2011). Cortical gamma responses: searching high and low. *International Journal of Psychophysiology*, *79*(1), 9–15. https://doi.org/10.1016/j.ijpsycho.2010.10.013
Crone, N. E. … Korzeniewska, A. (2006). High-frequency gamma oscillations and human brain mapping with electrocorticography. *Progress in Brain Research*, *159*, 275–295. https://doi.org/10.1016/S0079-6123(06)59019-3
Crone, N. E. … Lesser, R. P. (1998). Functional mapping of human sensorimotor cortex with electrocorticographic spectral analysis. II. Event-related synchronization in the gamma band. *Brain*, *121 (Pt 1*, 2301–2315. https://doi.org/10.1093/brain/121.12.2301
Daffertshofer, A., & van Wijk, B. (2011). On the Influence of Amplitude on the Connectivity between Phases. *Frontiers in Neuroinformatics*, *5*(6). https://doi.org/10.3389/fninf.2011.00006
Daly, I. … Warwick, K. (2012). Brain computer interface control via functional connectivity dynamics. *Pattern Recognition*, *45*(6), 2123–2136. https://doi.org/10.1016/j.patcog.2011.04.034
Damoiseaux, J. S. … Beckmann, C. F. (2006). Consistent resting-state networks across healthy subjects. *PNAS*, *103*(37), 13848–13853. https://doi.org/10.1073/pnas.0601417103
Damoiseaux, J. S., & Greicius, M. D. (2009). Greater than the sum of its parts: a review of studies combining structural connectivity and resting-state functional connectivity. *Brain Structure & Function*, *213*(6), 525–533. https://doi.org/10.1007/s00429-009-0208-6
Darainy, M. … Ostry, D. J. (2014). Plasticity in the Human Motor System Induced by Perceptual Learning. *IEEE*, 8–9.
Darvas, F. … Ojemann, J. G. (2009). Nonlinear phase-phase cross-frequency coupling mediates communication between distant sites in human neocortex. *Journal of Neuroscience*, *29*(2), 426–435. https://doi.org/10.1523/JNEUROSCI.3688-08.2009
Darvas, F. … Sorensen, L. B. (2009). Bi-phase locking - a tool for probing non-linear interaction in the human brain. *NeuroImage*, *46*(1), 123–132. https://doi.org/10.1016/j.neuroimage.2009.01.034
Dayan, E., & Cohen, L. G. (2011). Neuroplasticity subserving motor skill learning. *Neuron*, *72*(3), 443–454. https://doi.org/10.1016/j.neuron.2011.10.008
de Bie, H. M. a … Sanz-Arigita, E. J. (2012). Resting-state networks in awake five- to eight-year old children. *Human Brain Mapping*, *33*(5), 1189–1201. https://doi.org/10.1002/hbm.21280
de Pasquale, F. … Corbetta, M. (2010). Temporal dynamics of spontaneous MEG activity in brain networks. *PNAS*, *107*(13), 6040–6045. https://doi.org/10.1073/pnas.0913863107
De Vico Fallani, F. … Bablioni, F. (2010). Structural organization of functional networks from EEG signals during motor learning tasks. *International Journal of Bifurcation and Chaos*,

20(03), 905–912. https://doi.org/10.1142/S0218127410026198
Deckers, R. … de Zwart, J. (2006). An adaptive filter for suppression of cardiac and respiratory noise in MRI time series data. *Neuroimage*, *33*(4), 1072–1081. https://doi.org/10.1016/j.neuroimage.2006.08.006
Deco, G. … McIntosh, A. R. (2011). Emerging concepts for the dynamical organization of resting-state activity in the brain. *Nature Reviews. Neuroscience*, *12*(1), 43–56. https://doi.org/10.1038/nrn2961
Diedrichsen, J., & Kornysheva, K. (2015). Motor skill learning between selection and execution. *Trends in Cognitive Sciences*, *19*(4), 227–233. https://doi.org/10.1016/j.tics.2015.02.003
Digiovanna, J. … Sanchez, J. C. (2009). Coadaptive Brain – Machine Interface via reinforcement learning. *IEEE Transactions on Bio-Medical Engineering*, *56*(1), 54–64.
Dinopoulos, A. … Michaloudi, H. (1985). Thalamic projections to motor, prefrontal, and somatosensory cortex in the sheep studied by means of the horseradish peroxidase retrograde transport method. *Journal of Comparative Neurology*, *241*(1), 63–81. https://doi.org/10.1002/cne.902410106
Doyon, J., & Benali, H. (2005). Reorganization and plasticity in the adult brain during learning of motor skills. *Current Opinion in Neurobiology*, *15*(2), 161–167. https://doi.org/10.1016/j.conb.2005.03.004
Driemeyer, J. … May, A. (2008). Changes in Gray Matter Induced by Learning. *PLOS ONE*. https://doi.org/10.1371/ journal.pone.0002669
Duff, E. P. … Egan, G. F. (2008). The power of spectral density analysis for mapping endogenous BOLD signal fluctuations. *Human Brain Mapping*, *29*(7), 778–790. https://doi.org/10.1002/hbm.20601
Espenhahn, S. … Ward, N. S. (2016). Movement-related beta oscillations show high intra-individual reliability. *NeuroImage*, *147*(September 2016), 175–185. https://doi.org/10.1016/j.neuroimage.2016.12.025
Florin, E., & Baillet, S. (2015). The brain's resting-state activity is shaped by synchronized cross-frequency coupling of oscillatory neural activity. *NeuroImage*, *111*, 26–35. https://doi.org/10.1016/j.neuroimage.2015.01.054
Foster, B. L., & Parvizi, J. (2012). Resting oscillations and cross-frequency coupling in the human posteromedial cortex. *NeuroImage*, *60*(1), 384–391. https://doi.org/10.1016/j.neuroimage.2011.12.019
Foster, B. L. … Parvizi, J. (2015). Intrinsic and Task-Dependent Coupling of Neuronal Population Activity in Human Parietal Cortex. *Neuron*, *86*(2), 578–590. https://doi.org/10.1016/j.neuron.2015.03.018
Fox, M. D., & Raichle, M. E. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience*, *8*, 700–711. https://doi.org/10.1038/nrn2201
Fransson, P. (2005). Spontaneous low-frequency BOLD signal fluctuations: an fMRI investigation of the resting-state default mode of brain function hypothesis. *Human Brain Mapping*, *26*(1), 15–29. https://doi.org/10.1002/hbm.20113
Fukushima, M. … Averbeck, B. B. (2012). Spontaneous High-Gamma Band Activity Reflects Functional Organization of Auditory Cortex in the Awake Macaque. *Neuron*, *74*(5), 899–910. https://doi.org/10.1016/j.neuron.2012.04.014
Ganguly, K., & Carmena, J. M. (2009). Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biology*, *7*(7). https://doi.org/10.1371/journal.pbio.1000153
Ganguly, K. … Carmena, J. M. (2011). Reversible large-scale modification of cortical networks during neuroprosthetic control. *Nature Neuroscience*, *14*(5), 662–667. https://doi.org/10.1038/nn.2797
Garces, P. … Maestu, F. (2016). Quantifying the Test-Retest Reliability of Magnetoencephalography Resting-State Functional Connectivity. *Brain Connectivity*.

https://doi.org/10.1089/brain.2015.0416
Gierthmuehlen, M. … Ball, T. (2014). Mapping of sheep sensory cortex with a novel microelectrocorticography grid. *Journal of Comparative Neurology*, *522*(16), 3590–3608. https://doi.org/10.1002/cne.23631
Gilja, V. … Henderson, J. M. (2015). Clinical translation of a high-performance neural prosthesis. *Nature Medicine*, *21*(10), 6–8. https://doi.org/10.1038/nm.3953
Gohel, S., & Biswal, B. B. (2015). Functional Integration Between Brain Regions at Rest Occurs in Mutliple-Frequency Bands. *Brain Connectivity*, *5*(1). https://doi.org/10.1089/brain.2013.0210
Gordon, E. M. … Petersen, S. E. (2016). Generation and Evaluation of a Cortical Area Parcellation from Resting-State Correlations. *Cerebral Cortex*, *26*(1), 288–303. https://doi.org/10.1093/cercor/bhu239
Goulas, A. … Diesmann, M. (2014). Comparative Analysis of the Macroscale Structural Connectivity in the Macaque and Human Brain. *PLoS Computational Biology*, *10*(3), e1003529. https://doi.org/10.1371/journal.pcbi.1003529
Grandjean, J. … Rudin, M. (2014). Optimization of anesthesia protocol for resting-state fMRI in mice based on differential effects of anesthetics on functional connectivity patterns. *NeuroImage*, *102*(P2), 838–847. https://doi.org/10.1016/j.neuroimage.2014.08.043
Granger, C. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, *37*(3), 424–438.
Greenblatt, R. E. … Ossadtchi, A. E. (2012). Connectivity measures applied to human brain electrophysiological data. *Journal of Neuroscience Methods*, *207*(1), 1–16. https://doi.org/10.1016/j.jneumeth.2012.02.025
Grefkes, C. … Fink, G. R. (2002). Crossmodal processing of object features in human anterior intraparietal cortex: An fMRI study implies equivalencies between humans and monkeys. *Neuron*, *35*(1), 173–184. https://doi.org/10.1016/S0896-6273(02)00741-9
Groppe, D. M. … Mehta, A. D. (2013). Dominant frequencies of resting human brain activity as measured by the electrocorticogram. *NeuroImage*, *79*, 223–233. https://doi.org/10.1016/j.neuroimage.2013.04.044
Guo, C. C. … Seeley, W. W. (2012). One-year test–retest reliability of intrinsic connectivity network fMRI in older adults. *NeuroImage*, *61*(4), 1471–1483. https://doi.org/10.1016/j.neuroimage.2012.03.027
Hackett, T. A. … Kaas, J. H. (2001). Architectonic identification of the core region in auditory cortex of macaques, chimpanzees, and humans. *The Journal of Comparative Neurology*, *441*(3), 197–222. https://doi.org/10.1002/cne.1407
Hahn, B. … Stein, E. A. (2007). Cingulate activation increases dynamically with response speed under stimulus unpredictability. *Cerebral Cortex (New York, N.Y. : 1991)*, *17*(7), 1664–1671. https://doi.org/10.1093/cercor/bhl075
Hämäläinen, M. … Lounasmaa, O. V. (1993). Magnetoencephalography: theory, instrumentation, and applications to noninvasive studies of the working human brain. *Reviews of Modern Physics*, *65*(2), 413–497. https://doi.org/10.1103/RevModPhys.65.413
Hammock, B. D. … Catagnoli, N. (1989). A Sheep Model for MPTP Induced Parkinson-Like Symptoms. *Life Sciences*, *45*, 1601–1608. https://doi.org/10.1017/CBO9781107415324.004
Hamner, B. … del R Millan, J. (2011). Phase-based features for motor imagery brain-computer interfaces. In *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society* (pp. 2578–2581). IEEE. https://doi.org/10.1109/IEMBS.2011.6090712
Handwerker, D. a … Bandettini, P. a. (2012). Periodic changes in fMRI connectivity. *NeuroImage*, *63*(3), 1712–1719. https://doi.org/10.1016/j.neuroimage.2012.06.078
Hardwick, R. M. … Eickhoff, S. B. (2013). A quantitative meta-analysis and review of motor learning in the human brain. *NeuroImage*, *67*, 283–297.

https://doi.org/10.1016/j.neuroimage.2012.11.020
He, B. J. … Raichle, M. E. (2008). Electrophysiological correlates of the brain's intrinsic large-scale functional architecture. *PNAS*, *105*(41), 16039–16044. https://doi.org/10.1073/pnas.0807010105
He, B. J. … Raichle, M. E. (2010). The temporal structures and functional significance of scale-free brain activity. *Neuron*, *66*(3), 353–369. https://doi.org/10.1016/j.neuron.2010.04.020
Hermes, D. … Ramsey, N. F. (2010). Automated electrocorticographic electrode localization on individually rendered brain surfaces. *Journal of Neuroscience Methods*, *185*(2), 293–298. https://doi.org/10.1016/j.jneumeth.2009.10.005
Hermes, D. … Ramsey, N. F. (2012). Neurophysiologic correlates of fMRI in human motor cortex. *Human Brain Mapping*, *33*(7), 1689–1699. https://doi.org/10.1002/hbm.21314
Hillebrand, A. … Stam, C. J. (2012). Frequency-dependent functional connectivity within resting-state networks: an atlas-based MEG beamformer solution. *NeuroImage*, *59*(4), 3909–3921. https://doi.org/10.1016/j.neuroimage.2011.11.005
Hiltunen, T. … Palva, J. M. (2014). Infra-Slow EEG Fluctuations Are Correlated with Resting-State Network Dynamics in fMRI. *Journal of Neuroscience*, *34*(2), 356–362. https://doi.org/10.1523/JNEUROSCI.0276-13.2014
Hiremath, S. V … Boninger, M. L. (2015). Brain computer interface learning for systems based on electrocorticography and intracortical microelectrode arrays. *Frontiers in Integrative Neuroscience*, *9*, 40. https://doi.org/10.3389/fnint.2015.00040
Hochberg, L. R. … Donoghue, J. P. (2006). Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, *442*(7099), 164–171. https://doi.org/10.1038/nature04970
Houweling, S. … Beek, P. J. (2008). Neural changes induced by learning a challenging perceptual-motor task. *NeuroImage*, *41*(4), 1395–1407. https://doi.org/10.1016/j.neuroimage.2008.03.023
Hutchison, R. M. … Chang, C. (2013). Dynamic functional connectivity: Promise, issues, and interpretations. *NeuroImage*, *80*, 360–378. https://doi.org/10.1016/j.neuroimage.2013.05.079
Hutchison, R. M. … Everling, S. (2012). Functional connectivity of the frontal eye fields in humans and macaque monkeys investigated with resting-state fMRI. *Journal of Neurophysiology*, *107*(9), 2463–2474. https://doi.org/10.1152/jn.00891.2011
Hutchison, R. M., & Everling, S. (2012). Monkey in the middle: why non-human primates are needed to bridge the gap in resting-state investigations. *Frontiers in Neuroanatomy*, *6*(July), 29. https://doi.org/10.3389/fnana.2012.00029
Hutchison, R. M. … Menon, R. S. (2013). Resting-state networks show dynamic functional connectivity in awake humans and anesthetized macaques. *Human Brain Mapping*, *34*(9), 2154–2177. https://doi.org/10.1002/hbm.22058
Jacobs, J. … Fried, I. (2007). Brain Oscillations Control Timing of Single-Neuron Activity in Humans. *Journal of Neuroscience*. https://doi.org/10.1523/JNEUROSCI.4636-06.2007
Jerbi, K. … Lachaux, J.-P. (2010). Exploring the electrophysiological correlates of the default-mode network with intracerebral EEG. *Frontiers in Systems Neuroscience*, *4*(June), 27. https://doi.org/10.3389/fnsys.2010.00027
Johnson, J. … Hatton, G. (1974). Mechanosensory projections to cerebral cortex of sheep. *Journal of Comparative Neurology*, *158*(1), 81–107. https://doi.org/10.1002/cne.901580106
Joly, O. … Orban, G. A. (2012). Processing of vocalizations in humans and monkeys: A comparative fMRI study. *NeuroImage*, *62*(3), 1376–1389. https://doi.org/10.1016/j.neuroimage.2012.05.070
Keller, C. J. … Mehta, a. D. (2014). Corticocortical Evoked Potentials Reveal Projectors and Integrators in Human Brain Networks. *Journal of Neuroscience*, *34*(27), 9152–9163. https://doi.org/10.1523/JNEUROSCI.4289-13.2014

Keller, C. … Mehta, A. (2013). Neurophysiological investigation of spontaneous correlated and anticorrelated fluctuations of the BOLD signal. *Journal of Neuroscience*, *33*(15), 6333–6342. https://doi.org/10.1523/JNEUROSCI.4837-12.2013
Kelly, A. M. C., & Garavan, H. (2005). Human functional neuroimaging of brain changes associated with practice. *Cerebral Cortex*, *15*(8), 1089–1102. https://doi.org/10.1093/cercor/bhi005
Kern, M. … Ball, T. (2013). Heart cycle-related effects on event-related potentials, spectral power changes, and connectivity patterns in the human ECoG. *Neuroimage*, *81*(178–190). https://doi.org/10.1016/j.neuroimage.2013.05.042
Kim, S. M. … Frank, L. M. (2012). Spatial information outflow from the hippocampal circuit: distributed spatial coding and phase precession in the subiculum. *Journal of Neuroscience*, *32*(34), 11539–11558. https://doi.org/doi:10.1523/JNEUROSCI.5942-11.2012
Ko, A. L. … Ojemann, J. G. (2013). Identifying functional networks using endogenous connectivity in gamma band electrocorticography. *Brain Connectivity*, *3*(5), 491–502. https://doi.org/10.1089/brain.2013.0157
Ko, A. L. … Sorensen, L. B. (2011). Quasi-periodic fluctuations in default mode network electrophysiology. *Journal of Neuroscience*, *31*(32), 11728–11732. https://doi.org/10.1523/JNEUROSCI.5730-10.2011
Kohler, F. … Schuettler, M. (2017). Closed-loop interaction with the cerebral cortex: A review of wireless implant technology. *Brain-Computer Interfaces*, *4*(3), 146–154. https://doi.org/10.1080/236263X.2017.1338011
Koush, Y. … Scharnowski, F. (2015). Learning Control Over Emotion Networks Through Connectivity-Based Neurofeedback. *Cerebral Cortex*, bhv311. https://doi.org/10.1093/cercor/bhv311
Koyama, M. … Miyashita, Y. (2004). Functional magnetic resonance imaging of macaque monkeys performing visually guided saccade tasks: Comparison of cortical eye fields with humans. *Neuron*, *41*(5), 795–807. https://doi.org/10.1016/S0896-6273(04)00047-9
Kramer, M. a … Cash, S. S. (2011). Emergence of persistent networks in long-term intracranial EEG recordings. *Journal of Neuroscience*, *31*(44), 15757–15767. https://doi.org/10.1523/JNEUROSCI.2287-11.2011
Krubitzer, L. (2007). The Magnificent Compromise: Cortical Field Evolution in Mammals. *Neuron*, *56*, 201–208. https://doi.org/10.1016/j.neuron.2007.10.002
Lachaux, J.-P. … Crone, N. E. (2012). High-frequency neural activity and human cognition: past, present and possible future of intracranial EEG research. *Progress in Neurobiology*, *98*(3), 279–301. https://doi.org/10.1016/j.pneurobio.2012.06.008
Lachaux, J. … Varela, F. (1999). Measuring phase synchrony in brain signals. *Human Brain Mapping*, *8*(4), 194–208. https://doi.org/10.1002/(SICI)1097-0193(1999)8:4<194::AID-HBM4>3.0.CO;2-C
Lancaster, J. L. … Fox, P. T. (2000). Automated Talairach Atlas labels for functional brain mapping. *Human Brain Mapping*, *10*(3), 120–131. https://doi.org/10.1002/1097-0193(200007)10:3<120::AID-HBM30>3.0.CO;2-8
Laumann, T. O. … Petersen, S. E. (2015). Functional System and Areal Organization of a Highly Sampled Individual Human Brain. *Neuron*, *87*(3), 1–14. https://doi.org/10.1016/j.neuron.2015.06.037
Leeb, R. … Millán, J. D. R. (2015). Towards Independence: A BCI Telepresence Robot for People with Severe Motor Disabilities, *103*(6), 969–982. https://doi.org/10.1109/JPROC.2015.2419736
Leuthardt, E. C. … Moran, D. W. (2004). A brain-computer interface using electrocorticographic signals in humans. *Journal of Neural Engineering*, *1*(2), 63–71. https://doi.org/10.1088/1741-2560/1/2/001
Liang, Z. … Zhang, N. (2011). Uncovering Intrinsic Connectional Architecture of Functional

Networks in Awake Rat Brain. *Journal of Neuroscience*, *31*(10), 3776–3783. https://doi.org/10.1523/JNEUROSCI.4557-10.2011
Liu, X. … Duyn, J. H. (2015). Robust long-range coordination of spontaneous neural activity in waking, sleep and anesthesia. *Cerebral Cortex*, *25*(9), 2929–2938. https://doi.org/10.1093/cercor/bhu089
Liu, Z. … Duyn, J. H. (2010). Large-scale spontaneous fluctuations and correlations in brain electrical activity observed with magnetoencephalography. *NeuroImage*, *51*(1), 102–111. https://doi.org/10.1016/j.neuroimage.2010.01.092
Logothetis, N. K. … Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, *412*(6843), 150–157. https://doi.org/10.1038/35084005
Luaute, J. … Vuilleumier, P. (2009). Dynamic Changes in Brain Activity during Prism Adaptation. *Journal of Neuroscience*, *29*(1), 169–178. https://doi.org/10.1523/JNEUROSCI.3054-08.2009
Madhyastha, T. … Jäncke, L. (2014). Longitudinal reliability of tract-based spatial statistics in diffusion tensor imaging. *Human Brain Mapping*, *35*(9), 4544–4555. https://doi.org/10.1002/hbm.22493
Mantini, D. … Corbetta, M. (2007). Electrophysiological signatures of resting state networks in the human brain. *PNAS*, *104*(32), 13170–13175. https://doi.org/10.1073/pnas.0700668104
Mantini, D. … Vanduffel, W. (2011). Default Mode of Brain Function in Monkeys. *Journal of Neuroscience*, *31*(36).
Margulies, D. S. … Petrides, M. (2009). Precuneus shares intrinsic functional architecture in humans and monkeys. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(47), 20069–20074. https://doi.org/10.1073/pnas.0905314106
McGraw, K. O., & Wong, S. P. (1996). Forming inferences about some intraclass correlations coefficients. *Psychological Methods*, *1*(4), 390–390. https://doi.org/10.1037/1082-989X.1.4.390
McIntyre, C. C. … Lempka, S. F. Engineering the next generation of clinical deep brain stimulation technology. *Brain Stimulation*, *8*(1), 21–26. https://doi.org/10.1016/j.brs.2014.07.039
Mehrkanoon, S. … Summers, J. J. (2016). Upregulation of cortico-cerebellar functional connectivity after motor learning. *NeuroImage*, *128*, 252–263. https://doi.org/10.1016/j.neuroimage.2015.12.052
Miall, R. C., & Robertson, E. M. M. (2006). Functional Imaging: Is the Resting Brain Resting? *Current Biology*, *16*(23), 998–1000. https://doi.org/10.1016/j.cub.2006.10.041
Miller, K. J. … den Nijs, M. (2009). Power-law scaling in the brain surface electric potential. *PLoS Computational Biology*, *5*(12), e1000609. https://doi.org/10.1371/journal.pcbi.1000609
Miller, K. J. … Honey, C. J. (2012). Does rhythmic entrainment represent a generalized mechanism for organizing computation in the brain? *Frontiers in Computational Neuroscience*, *6*(October), 85. https://doi.org/10.3389/fncom.2012.00085
Miller, K. J. … Ojemann, J. G. (2014). Broadband changes in the cortical surface potential track activation of functionally diverse neuronal populations. *NeuroImage*, *85 Pt 2*, 711–720. https://doi.org/10.1016/j.neuroimage.2013.08.070
Monfils, M.-H. … Kleim, J. a. (2005). In search of the motor engram: motor map plasticity as a mechanism for encoding motor experience. *The Neuroscientist : A Review Journal Bringing Neurobiology, Neurology and Psychiatry*, *11*(5), 471–483. https://doi.org/10.1177/1073858405278015
Moritz, C. T. … Fetz, E. E. (2008). Direct control of paralysed muscles by cortical neurons. *Nature*, *456*(7222), 639–642. https://doi.org/10.1038/nature07418
Muller-Putz, G. … Millan, J. del R. (2015). Towards Noninvasive Hybrid Brain-Computer Interfaces: Framework, Practice, Clinical Application, and Beyond. *Proceedings of the*

IEEE, *103*(6), 926–943. https://doi.org/10.1109/JPROC.2015.2411333
Murphy, K. … Bandettini, P. A. (2013). Resting-state fMRI confounds and cleanup. *NeuroImage*, *80*, 349–359. https://doi.org/10.1016/j.neuroimage.2013.04.001
Nir, Y. … Malach, R. (2008). Interhemispheric correlations of slow spontaneous neuronal fluctuations revealed in human sensory cortex. *Nature Neuroscience*, *11*(9), 1100–1108. https://doi.org/10.1038/nn.2177
Nolte, G. … Muller, K.-R. (2010). Comparison of Granger Causality and Phase Slope Index. *Journal of Machine Learning Research*, 1–10.
Nolte, G. … Müller, K.-R. (2008). Robustly Estimating the Flow Direction of Information in Complex Physical Systems. *Physical Review Letters*, *100*(23), 234101. https://doi.org/10.1103/PhysRevLett.100.234101
Nudo, R. J. … Merzenich, M. M. (1996). Use-dependent alterations of movement representations in primary motor cortex of adult squirrel monkeys. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *16*(2), 785–807. https://doi.org/8551360
Oberheim, N. A. … Nedergaard, M. (2009). Uniquely Hominid Features of Adult Human Astrocytes. *Journal of Neuroscience*, *29*(10).
Olson, J. D. … Darvas, F. (2015). Comparison of subdural and subgaleal recordings of cortical high-gamma activity in humans. *Clinical Neurophysiology*, *127*(1), 277–284. https://doi.org/10.1016/j.clinph.2015.03.014
Onslow, A. C. E. … Jones, M. W. (2011). Quantifying phase-amplitude coupling in neuronal network oscillations. *Progress in Biophysics and Molecular Biology*, *105*(1–2), 49–57. https://doi.org/10.1016/j.pbiomolbio.2010.09.007
Opdam, H. I. H. … Bellomo, R. (2002). A sheep model for the study of focal epilepsy with concurrent intracranial EEG and functional MRI. *Epilepsia*, *43*(8), 779–787. https://doi.org/10.1046/j.1528-1157.2002.04202.x
Orban, G. A. … Vanduffel, W. (2004). Comparative mapping of higher visual areas in monkeys and humans. *Trends in Cognitive Sciences*, *8*(7), 315–324. https://doi.org/10.1016/j.tics.2004.05.009
Papademetris, X. … Staib, L. H. (2006). BioImage Suite: An integrated medical image analysis suite: An update. *The Insight Journal*, *2006*, 209. https://doi.org/http://www.bioimagesuite.org.
Patel, R. … Turner, G. R. (2013). Functional brain changes following cognitive and motor skills training: A quantitative meta-analysis. *Neurorehabilitation and Neural Repair*, *27*(3), 187–199. https://doi.org/10.1177/1545968312461718
Penny, W. … Ojemann, J. (2008). Testing for nested oscillation. *Journal of Neuroscience Methods*, *174*(1), 50–61. https://doi.org/10.1016/j.jneumeth.2008.06.035
Petrides, M., & Pandya, D. N. (1999). Dorsolateral prefrontal cortex: comparative cytoarchitectonic analysis in the human and the macaque brain and corticocortical connection patterns. *The European Journal of Neuroscience*, *11*(3), 1011–1036. https://doi.org/10.1046/j.1460-9568.1999.00518.x
Pichiorri, F. … Mattia, D. (2015). Brain-computer interface boosts motor imagery practice during stroke recovery. *Annals of Neurology*, *77*(5), 851–865. https://doi.org/10.1002/ana.24390
Raichle, M. E. (2011). The restless brain. *Brain Connectivity*, *1*(1), 836–839. https://doi.org/10.1089/brain.2011.0019
Raichle, M. E. … Shulman, G. L. (2001). A default mode of brain function. *PNAS*, *98*(2), 676–682. https://doi.org/10.1073/pnas.98.2.676
Raichle, M. E., & Snyder, A. Z. (2007). A default mode of brain function: a brief history of an evolving idea. *NeuroImage*, *37*(4), 1083-90; discussion 1097-9. https://doi.org/10.1016/j.neuroimage.2007.02.041
Ramos-Murguialday, A. … Birbaumer, N. (2013). Brain-machine interface in chronic stroke

rehabilitation: a controlled study. *Annals of Neurology*, *74*(1), 100–108. https://doi.org/10.1002/ana.23879
Ray, S. … Hsiao, S. S. (2008). Neural correlates of high-gamma oscillations (60–200 Hz) in macaque local field potentials and their potential implications in electrocorticography. *Journal of Neuroscience*, *28*(45), 11526–11536. https://doi.org/10.1523/JNEUROSCI.2848-08.2008.Neural
Rees, G. … Koch, C. (2000). A direct quantitative relationship between the functional properties of human and macaque V5. *Nat.Neurosci.*, *3*(7), 716–723. https://doi.org/10.1038/76673
Reuter, M. … Fischl, B. (2010). Accurate Inverse Consistent Robust Registration. *NeuroImage*, *53*(4), 1181–1196. https://doi.org/10.1016/j.neuroimage.2010.07.020
Rilling, J. K. … Votaw, J. R. (2007). A comparison of resting-state brain activity in humans and chimpanzees. *Proceedings of the National Academy of Sciences of the United States of America*, *104*(43), 17146–17151. https://doi.org/10.1073/pnas.0705132104
Rota, G. … Dogil, G. (2011). Reorganization of functional and effective connectivity during real-time fMRI-BCI modulation of prosody processing. *Brain and Language*, *117*(3), 123–132. https://doi.org/10.1016/j.bandl.2010.07.008
Sadtler, P. T. … Batista, A. P. (2014). Neural constraints on learning. *Nature*, *512*(7515), 423–426. https://doi.org/10.1038/nature13665
Sakkalis, V. (2011). Review of advanced techniques for the estimation of brain connectivity measured with EEG/MEG. *Computers in Biology and Medicine*, *41*(12), 1110–1117. https://doi.org/10.1016/j.compbiomed.2011.06.020
Sala-Llonch, R. … Bartrés-Faz, D. (2012). Brain connectivity during resting state and subsequent working memory task predicts behavioural performance. *Cortex*, *48*(9), 1187–1196. https://doi.org/10.1016/j.cortex.2011.07.006
Saleem, K. S., & Logothetis, N. K. (2012). *A combined MRI and histology atlas of the Rhesus monkey brain in stereotaxic coordinates, 2nd edition* (2nd ed.). Academic Press.
Saleem, K. S. … Price, J. L. (2014). Subdivisions and connectional networks of the lateral prefrontal cortex in the macaque monkey. *Journal of Comparative Neurology*, *522*(7), 1641–1690. https://doi.org/10.1002/cne.23498
Sauseng, P., & Klimesch, W. (2008). What does phase information of oscillatory brain activity tell us about cognitive processes? *Neuroscience and Biobehavioral Reviews*, *32*(5), 1001–1013. https://doi.org/10.1016/j.neubiorev.2008.03.014
Schalk, G. … Leuthardt, E. C. (2008). Two-dimensional movement control using electrocorticographic signals in humans. *Journal of Neural Engineering*, *5*(1), 75–84. https://doi.org/10.1088/1741-2560/5/1/008
Schalk, G. … Wolpaw, J. R. (2004). BCI2000: a general-purpose brain-computer interface (BCI) system. *IEEE Transactions on Bio-Medical Engineering*, *51*(6), 1034–1043. https://doi.org/10.1109/TBME.2004.827072
Schölvinck, M. L. … Khader, P. H. (2013). The contribution of electrophysiology to functional connectivity mapping. *NeuroImage*, *80*, 297–306. https://doi.org/10.1016/j.neuroimage.2013.04.010
Scholz, J. … Johansen-Berg, H. (2009). Training induces changes in white-matter architecture. *Nature Neuroscience*, *12*(11), 1370–1371. https://doi.org/10.1038/nn.2412
Schroeder, C. E., & Lakatos, P. (2009). Low-frequency neuronal oscillations as instruments of sensory selection. *Trends in Neurosciences*, *32*(1), 9–18. https://doi.org/10.1016/j.tins.2008.09.012
Sellers, E. W. … Wolpaw, J. R. (2010). A brain-computer interface for long-term independent home use. *Amyotrophic Lateral Sclerosis : Official Publication of the World Federation of Neurology Research Group on Motor Neuron Diseases*, *11*(5), 449–455. https://doi.org/10.3109/17482961003777470
Sforazzini, F. … Gozzi, A. (2014). Distributed BOLD and CBV-weighted resting-state networks

in the mouse brain. *NeuroImage*, *87*, 403–415. https://doi.org/10.1016/j.neuroimage.2013.09.050
Shenoy, K. V., & Carmena, J. M. (2014). Combining Decoder Design and Neural Adaptation in Brain-Machine Interfaces. *Neuron*, *84*(4), 665–680. https://doi.org/10.1016/j.neuron.2014.08.038
Smith, K. (2012). Idle minds. *Nature*, *489*, 356–358.
Smith, S. M. … Beckmann, C. F. (2009). Correspondence of the brain's functional architecture during activation and rest. *PNAS*, *106*(31), 13040–13045. https://doi.org/10.1073/pnas.0905267106
Sohal, V. … Deisseroth, K. (2009). Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. *Nature*, *459*(7247), 698–702. https://doi.org/10.1038/nature07991
Sun, F. T. … D'Esposito, M. (2007). Functional connectivity of cortical networks involved in bimanual motor sequence learning. *Cerebral Cortex*, *17*(5), 1227–1234. https://doi.org/10.1093/cercor/bhl033
Sun, H. … Ojemann, J. G. (2015). Sequential activation of premotor, primary somatosensory and primary motor areas in humans during cued finger movements. *Clinical Neurophysiology : Official Journal of the International Federation of Clinical Neurophysiology*, *126*(11), 2150–2161. https://doi.org/10.1016/j.clinph.2015.01.005
Tanosaki, M. … Okada, Y. (2014). Effective connectivity maps in the swine somatosensory cortex estimated from electrocorticography and validated with intracortical local field potential measurements. *Brain Connectivity*, *4*(2), 100–111. https://doi.org/10.1089/brain.2013.0177
Taubert, M. … Ragert, P. (2011). Long-term effects of motor training on resting-state networks and underlying brain structure. *NeuroImage*, *57*(4), 1492–1498. https://doi.org/10.1016/j.neuroimage.2011.05.078
Thatcher, R. W. … Biver, C. J. (2008). Development of cortical connections as measured by EEG coherence and phase delays. *Human Brain Mapping*, *29*(12), 1400–1415. https://doi.org/10.1002/hbm.20474
Thiebaut de Schotten, M. … Catani, M. (2012). Monkey to human comparative anatomy of the frontal lobe association tracts. *Cortex*, *48*(1), 82–96. https://doi.org/10.1016/j.cortex.2011.10.001
Tognoli, E., & Kelso, J. a S. (2009). Brain coordination dynamics: True and false faces of phase synchrony and metastability. *Progress in Neurobiology*, *87*(1), 31–40. https://doi.org/10.1016/j.pneurobio.2008.09.014
Tort, A. B. L. … Kopell, N. (2010). Measuring phase-amplitude coupling between neuronal oscillations of different frequencies. *Journal of Neurophysiology*, *104*(2), 1195–1210. https://doi.org/10.1152/jn.00106.2010
Vahdat, S. … Ostry, D. J. (2011). Functionally Specific Changes in Resting-State Sensorimotor Networks after Motor Learning. *Journal of Neuroscience*, *31*(47), 16907–16915. https://doi.org/10.1523/JNEUROSCI.2737-11.2011
Vincent, J. L. (2009). Learning and Memory: While You Rest, Your Brain Keeps Working. *Current Biology*, *19*(12), R484–R486. https://doi.org/10.1016/j.cub.2009.05.024
Vincent, J. L. … Buckner, R. L. (2008). Evidence for a frontoparietal control system revealed by intrinsic functional connectivity. *Journal of Neurophysiology*, *100*(6), 3328–3342. https://doi.org/10.1152/jn.90355.2008
Vincent, J. L. … Raichle, M. E. (2007). Intrinsic functional architecture in the anaesthetized monkey brain. *Nature*, *447*(7140), 83–86. https://doi.org/10.1038/nature05758
von Stein, A., & Sarnthein, J. (2000). Different frequencies for different scales of cortical integration: from local gamma to long range alpha/theta synchronization. *International Journal of Psychophysiology : Official Journal of the International Organization of Psychophysiology*, *38*(3), 301–313.

Walhovd, K. B. … Fjell, A. M. (2016). Premises of plasticity - And the loneliness of the medial temporal lobe. *NeuroImage*, *131*, 48–54. https://doi.org/10.1016/j.neuroimage.2015.10.060
Wander, J. D. … Darvas, F. (2016). Cortico-Cortical Interactions during Acquisition and Use of a Neuroprosthetic Skill. *PLoS Computational Biology*.
Wander, J. D. … Ojemann, J. G. (2013). Distributed cortical adaptation during learning of a brain-computer interface task. *PNAS*, *110*(26), 10818–10823. https://doi.org/10.1073/pnas.1221127110
Wang, X.-J. (2010). Neurophysiological and Computational Principles of Cortical Rhythms in Cognition. *Physiological Reviews*, *90*(3), 1195–1268. https://doi.org/10.1152/physrev.00035.2008
Weaver, K. E. K. E. … Darvas, F. (2016). Directional patterns of cross frequency phase and amplitude coupling within the resting state mimic patterns of fMRI functional connectivity. *NeuroImage*, *128*, 238–251. https://doi.org/10.1016/j.neuroimage.2015.12.043
Weiss, S., & Mueller, H. M. (2012). “Too many betas do not spoil the broth”: The role of beta brain oscillations in language processing. *Frontiers in Psychology*, *3*(JUN), 1–15. https://doi.org/10.3389/fpsyg.2012.00201
Wells, A. J. … Turner, R. J. (2012). A surgical model of permanent and transient middle cerebral artery stroke in the sheep. *PLoS ONE*, *7*(7), 1–9. https://doi.org/10.1371/journal.pone.0042157
Winder, A. T. … Drew, P. J. (2017). Weak correlations between hemodynamic signals and ongoing neural activity during the resting state. *Nature Neuroscience*, *20*(12), 1761–1769. https://doi.org/10.1038/s41593-017-0007-y
Wu, J. … Cramer, S. C. (2014). Resting-state cortical connectivity predicts motor skill acquisition. *NeuroImage*, *91*, 84–90. https://doi.org/10.1016/j.neuroimage.2014.01.026
Wu, T.-L. … Gore, J. C. (2016). Effects of anesthesia on resting state BOLD signals in white matter of non-human primates. *Magnetic Resonance Imaging*, *34*(9), 1235–1241. https://doi.org/10.1016/j.mri.2016.07.001
Xu, T. … Zuo, Y. (2009). Rapid formation and selective stabilization of synapses for enduring motor memories. *Nature*, *462*(7275), 915–919. https://doi.org/10.1038/nature08389
Yanagawa, T. … Fujii, N. (2013). Large-scale information flow in conscious and unconscious states: An ECoG study in monkeys. *PLoS ONE*, *8*(11), 1–13. https://doi.org/10.1371/journal.pone.0080845
Zacks, J. … Glover, G. H. (1999). Imagined transformations of bodies: An fMRI investigation. *Neuropsychologia*, *37*(9), 1029–1040. https://doi.org/10.1016/S0028-3932(99)00012-3
Zacks, J. M. (2008). Neuroimaging Studies of Mental Rotation: A Meta-analysis and Review. *Journal of Cognitive Neuroscience*, *20*(1). https://doi.org/https://doi.org/10.1162/jocn.2008.20013
Zacks, J. M. … Tversky, B. (2002). A parametric study of mental spatial transformations of bodies. *NeuroImage*, *16*(4), 857–872. https://doi.org/10.1006/nimg.2002.1129
Zhang, N. … King, J. (2010). Mapping resting-state brain networks in conscious animals. *Journal of Neuroscience Methods*, *189*(2), 186–196. https://doi.org/10.1016/j.jneumeth.2010.04.001
Zuo, X. N., & Xing, X. X. (2014). Test-retest reliabilities of resting-state FMRI measurements in human brain functional connectomics: A systems neuroscience perspective. *Neuroscience and Biobehavioral Reviews*, *45*, 100–118. https://doi.org/10.1016/j.neubiorev.2014.05.009