UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title
Attention and Disruption: Modulating the Interactions between Large-scale Brain Networks and the Properties of Local Cortical Regions

Permalink
https://escholarship.org/uc/item/5mr0654b

Author
Gratton, Caterina

Publication Date
2013

Peer reviewed|Thesis/dissertation

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

Attention and Disruption: Modulating the Interactions between Large-scale Brain Networks and the Properties of Local Cortical Regions

By

Caterina Gratton

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Neuroscience

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Mark D'Esposito, Chair
Professor Michael Silver
Professor Richard Ivry
Professor Stephen Palmer

Fall 2013

Abstract

Attention and Disruption: Modulating the Interactions between Large-scale Brain Networks and the Properties of Local Cortical Regions

By

Caterina Gratton

Doctor of Philosophy in Neuroscience

University of California, Berkeley

Professor Mark D'Esposito, Chair

An abundance of research has demonstrated that the brain is both organized into discrete, highly specialized, regions and that distributed regions work together in the form of large-scale networks. It remains unclear, however, how to reconcile these two views of brain function. In this thesis, I present several projects that investigated the specialized properties of individual regions, how these change depending on cognitive demands and external disruption, and how these changes interact with modulations across large-scale networks. Across these projects, I use a combination of methods including psychophysics, univariate and multivariate modeling of functional Magnetic Resonance Imaging (fMRI) data, Transcranial Magnetic Stimulation (TMS), and the study of patients with focal lesions.

The first section of this thesis focuses on visual attention to high-level face stimuli. Here, I demonstrate that highly specialized representations exist for individual faces, but that engaging in sustained top-down attention can cause selective modulations in these representations. This suggests that large-scale networks can interact with, and modify, the properties of local regions in an extremely selective manner.

The second section of the thesis examines how disrupting local regions, through focal brain lesions or TMS, impacts the structure of large-scale networks. The first two chapters in this section examine how acute and chronic disruptions modulate the magnitude of interactions throughout cognitive control networks. An additional project explores methodological considerations involving the variability in TMS effects and how these might be measured most effectively. The final project in this section examines how focal brain lesions also affect the organization of distributed networks. This section illustrates how the properties of local regions can modulate both the magnitude and organization of distributed networks.

In sum, these projects demonstrate that large-scale networks can both drive and be driven by the highly specialized properties of local regions. This research will hopefully provide some insight into how brain regions may concurrently demonstrate highly specialized brain activity and participate in distributed brain networks, and flexibly modify these interactions, in order to carry out complex behavioral functions.

Table of Contents

CHAPTER 1 ..... **1**

SECTION 1: ATTENTION

CHAPTER 2 ..... **5**
2.1 ABSTRACT 5
2.2 INTRODUCTION 5
2.3 MATERIALS AND METHODS 6
2.4 RESULTS 13
2.5 DISCUSSION 21

SECTION 2: DISRUPTION

CHAPTER 3 ..... **27**
3.1 ABSTRACT 27
3.2 INTRODUCTION 27
3.3 MATERIALS AND METHODS 29
3.4 RESULTS 33
3.5 DISCUSSION 39

CHAPTER 4 ..... **42**
4.1 ABSTRACT 42
4.2 INTRODUCTION 42
4.3 MATERIALS AND METHODS 44
4.4 RESULTS 47
4.5 DISCUSSION 57

CHAPTER 5 ..... **63**
5.1 ABSTRACT 63
5.2 INTRODUCTION 63
5.3 MATERIALS & METHODS 65
5.4 RESULTS 71
5.5 DISCUSSION 77

CHAPTER 6 ..... **83**
6.1 ABSTRACT 83
6.2 INTRODUCTION 83
6.3 MATERIALS AND METHODS 85
6.4 RESULTS 90
6.5 DISCUSSION 97

CHAPTER 7 ..... **100**

REFERENCES **102**

Acknowledgements

While I am the solitary author on this dissertation, in actual fact there is a large group of people who helped immeasurably in putting together this body of work. Although this short section is not sufficient, I would like to express all of the gratitude I feel for their help. First, I would like to thank my advisors, Mark D'Esposito and Michael Silver, for the support they have provided and the effort they have put into training me. They have provided me with freedom to explore but also always asked the right questions to guide my projects to much clearer and more interesting conclusions. And most of all, I would like to thank them for their constant optimism, which allowed me to believe that these projects might work after all.

Over my time here, I have benefited from the help and friendship of all of the members of the D'Esposito and Silver laboratories. I would like to particularly thank my officemates in Barker (Emi Nomura, Deanna Wallace, Jessica Cohen), Giannini (Courtney Gallen, and Maxwell Bertolero), and Minor (Rachel Denison and Ariel Rokem) who have very kindly shared their time and wisdom with me whenever I turned my chair to ask a question. In addition, I would like to provide a special thanks to Emi Nomura, who trained me in all things fMRI, effectively allowing me to do any science at all, and to Fernando Perez, who showed incredible patience and generosity in working with me to develop Python code for network analysis. Finally, I would like to acknowledge all of the help and hard work that my collaborators (Emi Nomura, Fernando Perez, Taraz Lee, Kartik Sreenivasan, Renee Visser, and Andy Kayser) have put into the projects in this dissertation. I would like to thank these amazing individuals for sharing their knowledge, taking the time to discuss these projects, training me in new methods, and, most of all, for being such great friends.

This research would not have been possible without the help of an industrious set of research assistants (Sahar Yousef, Antony LaBarbera, Jeff Defond, Mehrnaz Ahrar, and Jahlela Hasle), as well as the help of Miguel Perez, Josh Hoffman, and Drew Erickson in collecting patient scans and Bob Knight, Donatella Scabini, and Jessica Black in recruiting patients. I would also like to thank the BIC staff, especially Ben Inglis, Matthew Brett, and JB Poline, for all of their technical help on these projects, and Julie Ashworth for keeping such a well maintained computing cluster.

In addition, I would like to thank the other members of my thesis and qualifying exam committees (Rich Ivry, Steve Palmer, Jack Gallant, David Whitney, and Dan Feldman) for their probing questions, which pushed me to be a better and more careful scientist. I would like to thank the Helen Wills Neuroscience Institute for providing such a scientifically rich and friendly environment in which to work and Kati Markowitz in particular for always taking care of us graduate students. I would also like to thank the National Science Foundation and Department of Defense for their generous funding over the past five years.

Finally, I would like to thank my family: my sister, Cristina, my parents Monica and Gabriele, and my husband Terence. Thank you for being there every day, for listening to my triumphs and frustrations, for brainstorming new ideas with me, for teaching me statistics, signal processing, aging, chemistry, and genetics, for distracting me when needed, and for showing me that research and family-life can work together. Thank you all for always reminding me of why science is so neat.

Dedicated to my husband, Terence Crofts,
whose encouragement, support, insightful comments, and motionless brain have made this research infinitely better and a lot more fun to carry out.

Chapter 1

INTRODUCTION

Throughout the history of neuroscience there have often been two ways of viewing functional specialization in the brain. At one extreme is the localized processing perspective, which posits that each brain region is highly specialized for carrying out particular functions. At the other extreme is the holistic processing perspective, where behavior is seen as arising from the interactions among many (or all, in the case of equipotentiality) brain areas (Catani et al., 2012).

Compelling evidence exists that supports both of these viewpoints. Studies of lesions (in humans and animals) as well as electrophysiological recordings have suggested that different cortical regions can have highly specialized functions. However, observations of non-specific deficits after some lesions, disconnection syndromes, and diaschisis, as well as the tracing of a multitude of white matter pathways has suggested that brain regions operate in concert to carry out behavioral functions. Therefore, the localized and holistic processing perspectives have oscillated back and forth in popularity over the course of the past two centuries (Catani et al., 2012).

The advent of noninvasive techniques for functional imaging of the human brain has added evidence to both sides of this debate. The earliest studies using functional Magnetic Resonance Imaging (fMRI) and Positron Emission Tomography (PET) demonstrated that a number of brain regions show specialized responses to the presentation of particular types of stimuli or tasks (e.g., (Corbetta et al., 1991; Engel et al., 1994; Kanwisher et al., 1997; Petersen et al., 1988; Zeki et al., 1991)). Over time, using more sophisticated psychophysical and analysis techniques, many studies have exquisitely delineated the tuning properties of early sensory and motor regions to show that they have very specific functional properties (e.g., (Naselaris et al., 2009; Serences et al., 2009; Silver and Kastner, 2009)). However, studies of correlated temporal variability (functional connectivity) in fMRI and characterization of white matter pathways using Diffusion Imaging (DI) have also made it clear that brain areas work together in the form of large-scale networks (Fox and Raichle, 2007; Honey et al., 2009). These networks are evident as robust entities during rest (Damoiseaux et al., 2006; Shehzad et al., 2009) and, most importantly, are related to task performance and disease states (Greicius, 2008; Sadaghiani et al., 2010; Smith et al., 2009; Zhang and Raichle, 2010).

How, then, can the localized and holistic processing viewpoints be reconciled? An associationist (Catani et al., 2012) or modular (Bullmore and Sporns, 2009; Friston and Price, 2011) perspective on brain function has been proposed, that includes both functional specialization and integration. Within this framework, the brain operates as a complex system with multiple large-scale networks (or modules), each of which are comprised of a number of regions with different functional specializations and roles within these networks.

This allows for regions to show specialization at a local level but also allows for complex brain function to arise from the flexible coordinated interaction among many different regions and modules.

Despite the theoretical appeal of this framework, relatively few studies have examined the interactions between local regional specialization and network integration. In this dissertation, I have characterized changes in local functional properties and how these functional properties interact with distributed large-scale networks after two different types of perturbations on the system: engagement of sustained visual attention and focal brain disruptions from lesions or transcranial magnetic stimulation (TMS).

The first project from my dissertation has focused on visual attention as a model cognitive process that requires changes on both local and network spatial scales. I found that attention can modify local functional properties (tuning) for high-level stimuli, in much the same way as it has been demonstrated to act for lower-level stimuli (Gratton et al., 2013). This study provides evidence that local functional properties for high-level features are not immutable, but instead can vary depending on current task demands, perhaps through interactions with large-scale fronto-parietal attention networks that have been shown in the past to influence visual cortical regions (Corbetta and Shulman, 2002; Lee and D'Esposito, 2012; Ruff et al., 2008; Ruff et al., 2006; Zanto et al., 2011).

In the second line of research, I have instead examined the effects of disrupting local properties, either through focal brain lesions or TMS, on the interactions among large-scale networks. The first project discussed in this section examined the effects of chronic lesions to one or the other of two putatively independent cognitive control networks. Here, we demonstrated that focal lesions affect the magnitude of interactions throughout a network, even in regions remote from the damaged location (but tend to leave other, independent, cognitive control networks intact; (Nomura et al., 2010)). The next chapter (**Ch. 3**) describes a follow-up combined TMS-fMRI project, which showed that acute disruptions caused by TMS show widespread impact, well beyond the damaged region's network (Gratton et al., submitted-a). The large amount of variability seen in measurements from this study also prompted a follow-up methodological study (Gratton et al., submitted-b) (**Ch. 4**), which showed that variability in the network effects of TMS is related to variability in local changes in perfusion under the site of stimulation. The final project from this section (**Ch. 5**) examined how focal lesions affect network organization, rather than just the magnitude of interactions within pre-defined networks. This project demonstrated that focal lesions can cause widespread reorganization of large-scale networks, but only when those lesions impact brain locations critical for inter-network communication (Gratton et al., 2012). This section suggests that changes in the functional properties of local regions (i.e., through disruption) may also influence the properties of large-scale networks.

Together, these projects demonstrate that local regional specialization and network integration are closely intertwined with one another, with each level of functional organization able to influence the other.

Dissertation Outline

- Section 1 - Attention
  - Ch. 2: Effects of attention on local tuning for faces (Gratton et al., 2013)
- Section 2 - Disruption
  - Ch. 3: Effects of lesions on cognitive control networks (Nomura et al., 2010)
  - Ch. 4: Effects of TMS on cognitive control networks (Gratton et al., submitted-a)
  - Ch. 5: Measuring the impact of TMS using perfusion (Gratton et al., submitted-b)
  - Ch. 6: Effects of lesions on whole-brain modular organization (Gratton et al., 2012)

- SECTION 1 -
ATTENTION

Chapter 2

ATTENTION SELECTIVELY MODIFIES THE REPRESENTATION OF INDIVIDUAL FACES IN THE HUMAN BRAIN

2.1 Abstract

Attention modifies neural tuning for low-level features, but it is unclear how attention influences tuning for complex stimuli. We investigated this question in humans using fMRI and face stimuli. Participants were shown six faces (F1–F6) along a morph continuum, and selectivity was quantified by constructing tuning curves for individual voxels. Face-selective voxels exhibited greater responses to their preferred face than to non-preferred faces, particularly in posterior face areas. Anterior face areas instead displayed tuning for face categories: voxels in these areas preferred either the first (F1-F3) or second (F4-F6) half of the morph continuum. Next, we examined the effects of attention on voxel tuning by having subjects direct attention to one of the superimposed images of F1 and F6. We found that attention selectively enhanced responses in voxels preferring the attended face. Taken together, our results demonstrate that single voxels carry information about individual faces and that the nature of this information varies across cortical face areas. Additionally, we found that attention selectively enhances these representations. Our findings suggest that attention may act via a unitary principle of selective enhancement of responses to both simple and complex stimuli across multiple stages of the visual hierarchy.

2.2 Introduction

Attention powerfully modifies the way we see the world, allowing us to filter vast amounts of information and to preferentially process those components that are most relevant to our current behavioral goals. However, despite the complexity of natural visual scenes, most studies examining modulation of neural representations by attention have focused on simple stimuli with low-level features. In these studies, sustained top-down attention selectively modifies representations in early visual areas by enhancing responses to stimuli at attended locations or containing attended features (e.g., (Gandhi et al., 1999; McAdams and Maunsell, 1999; Serences et al., 2009; Treue and Maunsell, 1999)). In contrast, less is known about the effects of attention on representations of more complex stimuli that are represented in higher visual areas. While attention can influence the processing of complex categories such as faces and scenes by enhancing responses across entire category-selective areas of cortex (Corbetta et al., 1991; Gazzaley et al., 2005; Wojciulik et al., 1998), the mechanism by which attention modulates the tuning of responses to individual stimuli within these categories is unknown.

Faces are extremely prevalent and salient stimuli in our environment. Individual neurons (reviewed in Gross, 2011; Gross et al., 1972) and entire brain regions (Kanwisher et al., 1997; McCarthy et al., 1997) exhibit specialized responses to faces compared with other stimuli. These regions comprise a distributed network (Fairhall and Ishai, 2007; Moeller et al., 2008; Zhang et al., 2009) that includes the Occipital Face Area (OFA), the Fusiform Face Area (FFA), the face-selective portion of the Superior Temporal Sulcus (fSTS), and more anterior parts of the inferior temporal lobe (Haxby et al., 2000). Both face features (Freiwald et al., 2009) and sets of face morphs (Leopold et al., 2006) are represented in a continuous fashion by neurons within these areas, perhaps within the context of an orderly face space (Valentine, 1991). Thus, because of their systematic representations and their high salience, face stimuli are ideal for examining whether the effects of attention on complex feature tuning mirror the effects that have previously been reported for low-level feature representations.

However, due to the relatively coarse spatial resolution of functional MRI (fMRI), in which the smallest volumes of measurement (voxels) are typically several cubic millimeters and contain thousands of neurons with different stimulus preferences, it has proven difficult to detect signals that are selective for individual faces. Nevertheless, recent evidence from both fMRI adaptation (fMR-A) paradigms (Gilaie-Dotan and Malach, 2007; Loffler et al., 2005; Rotshtein et al., 2005) and multi-voxel pattern classifiers (MVPA) (Kriegeskorte et al., 2007; Natu et al., 2010; Nestor et al., 2011) suggests that it may be possible to detect tuning for individual faces within face-selective regions. Here, we adapted a method recently developed by Serences and colleagues (Serences et al., 2009) for constructing feature tuning curves for single voxels and demonstrate that single voxels carry information about individual faces. In addition, we found that directing attention to one of a pair of superimposed faces selectively enhanced responses in voxels representing the attended face.

2.3 Materials and Methods

2.3.1 Participants

Ten healthy participants (7 female, mean age = 25, age range 19 – 32) each completed two separate 2-hour scanner sessions that were each accompanied by a 1-hour practice session within the preceding week. Informed consent was obtained from subjects in accordance with procedures approved by the Committee for the Protection of Human Subjects at the University of California, Berkeley.

2.3.2 Behavioral Tasks

Face localizer task. In each fMRI session, participants completed a face localizer run of 336 sec either once or twice. This task consisted of sequential 16-second blocks of faces, scenes, and fixation. During face and scene blocks, participants performed a 1-back task, pressing a

response button whenever an image was identical to the previous image. Faces and scenes were presented at a rate of one per second.

Face mapping task (Experiment 1). In each fMRI session, participants completed a 5-minute face mapping run five or six times. A block consisted of either continuous presentation of a central fixation cross or repeated presentation of a single face, during which participants counted the number of contrast decrements in the face that occurred during the block. Each run consisted of 21 12-second blocks (3 blocks for each of seven conditions (6 face morphs and fixation), **Figure 2.1A**). In each block (except fixation), a single face was flashed on the screen for 10 seconds at a rate of 2 Hz (300 ms on, 200 ms off), followed by a 2-second response period during which the fixation cross was red. Fixation blocks consisted of 12 seconds of a white fixation cross. Block order was pseudo-randomized, with the constraint that the same condition was never repeated in two successive blocks. Participants were instructed to count the number of times that the contrast of the entire face was reduced for the duration of a single 300 ms presentation (0-3 contrast decrements possible per block, 25% probability of a given number of contrast decrements for each block). The amount of contrast decrement in the face was adjusted across runs to maintain approximately 75% correct trials (chance performance = 25%).

Stimuli consisted of grayscale images of an old male face (F1), a young female face (F6), and four morphs along a continuum between F1 and F6 (F2-F5; **Figure 2.1A**). Images were obtained from the Center for Vital Longevity's face database (Minear and Park, 2004) and were morphed using Abrasoft's FantaMorph4 software after selecting a number of common image points in F1 and F6. The two faces on either end of the morph continuum were chosen to span a large range of facial features, including both gender and age. Intermediate faces were evenly spaced along the morph continuum (e.g., F2 was 80% F1 and 20% F6). After morphing, all faces were matched for average luminance and contrast (standard deviation of the luminance) across the entire face and placed on a grayscale background, leaving in ears and hair. Face images were 12 degrees of visual angle along the longer vertical axis (forehead to chin).

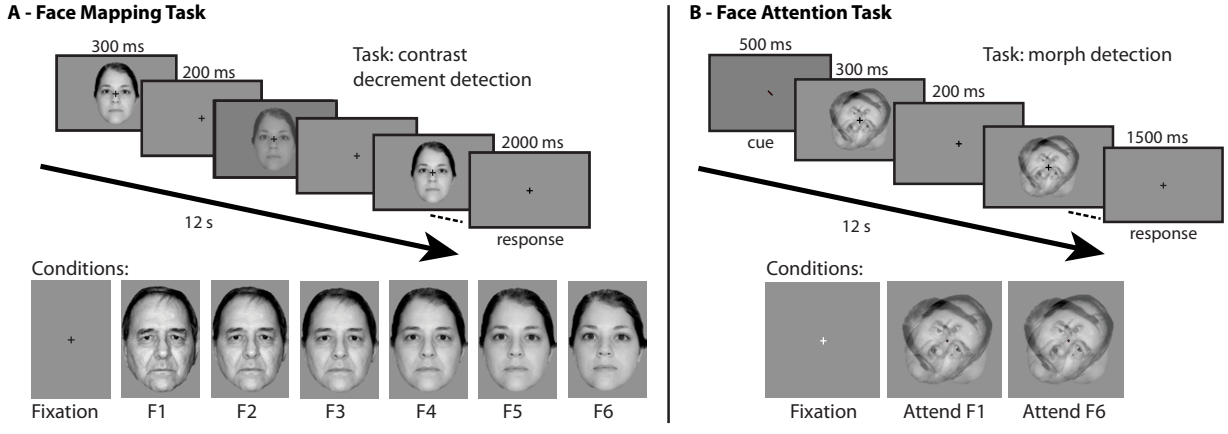

Figure 2.1: Trial schematic for face mapping (Expt. 1) and face attention (Expt. 2) tasks. (A) Participants detected brief contrast decrements over the entire face during presentation of a series of flashing faces.

Blocks were either fixation or one of six faces along a morph continuum (F1–F6). (B) Attention blocks were either fixation or flashing superimposed images of F1 and F6. Participants attended to either F1 or F6 and detected brief morphs in the attended face towards a third face while ignoring morphs in the unattended face. The face to be attended was cued by the color (red or green) of the fixation cross. For illustrative purposes, the fixation crosses and contrast decrement displayed here are larger than those actually presented to participants, and fixation crosses are shown here in black.

Face attention task (Experiment 2). In each session, participants completed 6-7 runs of the five-minute attention task in addition to the face mapping task (Expt. 1) described above. Participants viewed a superimposed image of F1 and F6, with each face tilted 45 degrees in opposite directions, and were instructed to attend to just one of the two faces and to detect brief morphs in that face (**Figure 2.1B**). Each run consisted of 24 12-second blocks, with 8 blocks in each of three conditions: attend F1, attend F6, or fixation. In attend F1 and attend F6 conditions, blocks began with a cue consisting of a red or green oriented bar indicating the identity of the face that was to be attended. The color associated with a particular face attention condition (e.g., green for attend F1) was consistent for a given subject but counterbalanced across subjects. The orientation of the cue indicated the orientation of the face to be attended. The cue remained on the screen for 500ms, followed by the onset of a flashing composite image of F1 and F6 in the center of the screen (**Figure 2.1B**). Face orientations were counterbalanced and pseudorandomly assigned across blocks in a scan. The composite face flashed at a rate of 2 Hz (300 ms on, 200 ms off) for 10 sec. The fixation cross (red or green) remained on the screen throughout the block. Each block ended with a response period (1500 ms) during which the fixation cross was black.

The participants' task was to count the number of morphs in the attended face towards a third face (0-3 morphs per block, number of morphs randomly chosen and occurring at random times within the trial after the first flash). Morphs, when present, were displayed continuously for an entire flash (300 ms). The percent morph to the third face was titrated for each subject and each attention condition (attend F1 or attend F6) to maintain approximately 75% correct trials (chance performance = 25%). The third face (i.e., 100% morph) was never shown to the subjects in any condition. The unattended face also morphed with the same probability as the attended face, but participants were instructed to ignore these changes. Block order was pseudorandom.

2.3.3 MRI Procedures

MRI data acquisition. T2*-weighted echo planar images (EPI) were collected on a whole-body 3 Tesla Siemens MAGNETOM Trio MR scanner using a 32-channel radiofrequency head coil. Each localizer run consisted of 168 volumes (TR = 2 s, TE = 25 ms, FA = 70 deg, voxel size = 3x3x3.3 mm, with 34 slices approximately 12˚ from axial, no GRAPPA acceleration). Each face mapping run (Expt. 1) was 132 volumes (TR = 2 s, TE = 20 ms, FA = 70 deg, voxel size = 2.22 x 2.22 x 2.3 mm, with 46 slices approximately 12 degrees relative to axial, GRAPPA factor = 4).

Each face attention run (Expt. 2) consisted of 150 volumes with the same acquisition parameters as the mapping runs. Due to scanner malfunction, face mapping runs were

acquired with the localizer run parameters in three sessions, and face attention runs were acquired with the localizer run parameters in two sessions and with slightly altered parameters (matching typical face attention runs but with voxel size = 2.22 x 2.22 x 3.3 mm, 34 slices) in one session. The acquisition volume covered approximately the whole brain, although it was positioned to cover all of anterior temporal cortex at the expense of some regions in the parietal cortex. Localizer runs had larger voxels in order to obtain a higher signal-to-noise ratio, whereas face mapping and attention runs had smaller voxel sizes in order to more densely sample the face regions for voxel tuning analyses. Structural images were acquired using a high-resolution MP-RAGE T1-weighted sequence.

MRI data analysis**.** Image preprocessing was conducted in AFNI (http://afni.nimh.nih.gov/afni/; (Cox, 1996)) and consisted of the removal of non-brain voxels from the EPI volumes, despiking, motion correction (within and between runs of the same type), and rigid-body alignment to the structural image. In addition, spatial smoothing (4 mm FWHM) was applied to localizer runs.

General linear models (GLM) were computed for each type of run (localizer, face mapping, and face attention) to estimate response amplitude for each condition. All GLMs included head motion and detrending (third-degree polynomial) regressors. The GLM for the localizer task contained separate regressors for face and scene blocks. For the face mapping task (Expt. 1), separate GLMs were computed for each run (in order to generate independent beta estimates for tuning analyses), and each GLM included regressors for each of the six face conditions and fixation, a response period regressor, and a single parametric "adaptation regressor".

The motivation for including the adaptation regressor was to mitigate possible effects of face adaptation on our estimates of response amplitude. For example, because F1 and F6 are on the ends of the face morph continuum and therefore less likely than F2-F5 to be preceded by a stimulus nearby in face space, they would, on average, be the least adapted of the face stimuli. To account for this, we included a parametric regressor in our GLM estimation of the response weights that was scaled according to the relative distance, in face space, between the current block and the previous block. Thus, if the current block was very similar to the previous block in face space (e.g., the current block was F1 and the previous block was F2), this regressor would be assigned a value of 1, indicating the high potential contribution of face adaptation. If, instead, the current block was very different from the previous block in face space (e.g., the current block was F1 and the previous block was F6), this regressor would be assigned a value of 0.2, corresponding to the potentially small contribution of adaptation. A value of zero was assigned when the current block was preceded by a fixation block. All results were largely the same whether or not this regressor was included.

For the face attention task (Expt. 2), a single GLM was computed for all runs, with separate regressors for each of the three conditions (attend F1, attend F6, and fixation) and for the response period. In two sessions, the response period was not in a separate volume from the stimulus block and was therefore not included as a separate regressor.

Region of Interest (ROI) selection**.** For each session, face areas were defined based on a contrast of faces vs scenes from the face localizer task ($p<0.001$, uncorrected; this somewhat liberal threshold was used since a second voxel selection step was included in computing tuning curves). Voxels in temporal and occipital cortex that responded significantly more to faces than to scenes were included in a single "all face areas" region for primary analyses. This region included clusters of voxels within the OFA, FFA, and fSTS, as well as several small clusters or single voxels, mainly in the anterior temporal lobe, that were combined and labeled as anterior temporal areas (fAT; **Figure 2.2**). Collectively, these areas correspond to regions classically identified as part of the face processing network (Haxby et al., 2001).

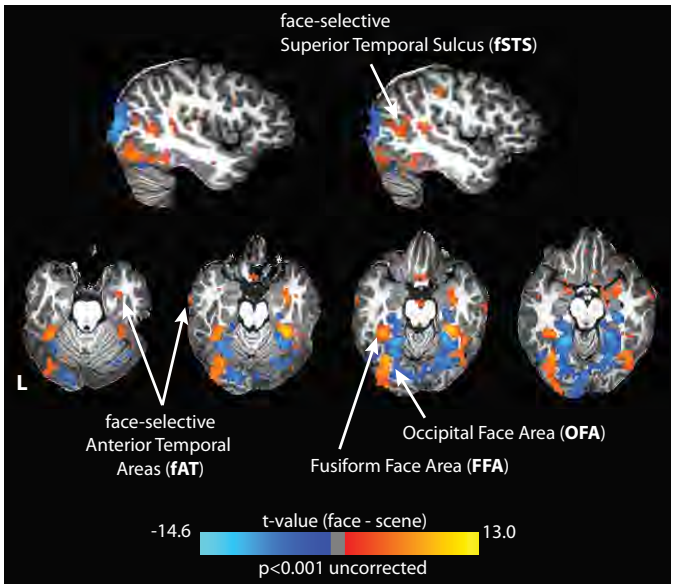

Figure 2.2 Identification of face areas. Face-selective regions in temporal and occipital cortex were defined in each session using a localizer (faces > scenes; displayed in warm colors). An example from one subject is shown here.

Voxel tuning curve procedure. We adapted a method described by Serences et al. (Serences et al., 2009) to classify the face preference and construct face tuning curves for individual voxels. For each face-selective (preferring faces to scenes) voxel identified with the face localizer, separate beta values were extracted for each of the six faces in the morph continuum for each face mapping run, using the contrast of each face condition (F1 through F6 blocks) vs fixation. These beta values were normalized within each voxel by subtracting the voxel's mean response across all six face conditions. This was done in order to remove differences across voxels in their overall responsiveness, thereby facilitating measurement of differences in the pattern of responses to different faces within a voxel. We then used all but one face mapping run to classify each voxel as preferring one of the six face morphs, based on the face that elicited the highest average beta value across these classification runs. The left-out run that was not used in classification of the preferred face for each voxel was then used to measure responses to each of the six face types, creating a tuning curve

for each voxel. This leave-one-out procedure was repeated with all possible combinations of classification and analysis runs for cross-validation.

A further selection was applied during this procedure: only those voxels that had the same preferred face in more than half of the classification runs were considered to be "reliable" voxels and included in further analysis (on average, 22 voxels per session for all face areas combined). Due to the small size of the face-selective anterior temporal area, this reliability criterion was relaxed to a threshold of more than 20% agreement in preferred face across classification runs, thereby resulting in similar numbers of reliable voxels for posterior and anterior face areas.

Tuning curves were averaged across all voxels with the same face preference (F1 through F6) and across all runs of the cross-validation procedure. Because the number of reliable voxels varied widely across individual sessions, group analyses were weighted by the number of voxels that each session contributed to the analysis using the procedure described in (Bland and Kerry, 1998). In addition, in order to examine whether some areas had tuning that was more categorical in nature rather than specific for individual faces along the morph continuum, the same procedure described above for classifying voxels based on individual face preference was used to classify voxels as preferring either the first half (F1-F3, the "male" half) or the second half (F4-F6, the "female" half) of the face morph continuum. For the categorical classification, beta values were averaged across the three conditions within each half of the face continuum (i.e., averaged responses to either F1 through F3 or F4 through F6) for each voxel.

Statistical Analyses of Tuning Curves**.** To assess how reliably our tuning curve procedure classified the preferred face type for individual voxels, weighted beta values from each session were entered into a repeated-measures ANOVA with two factors: voxel type (F1-F6) and condition (preferred vs all non-preferred face types, e.g., for F1 voxels, F1 condition responses vs the average of F2-F6 condition responses). This and all other ANOVAs were conducted using R's (R Development Core Team, 2011) ezANOVA package. The p-values for the main effect of condition in the ANOVA are the results of one-tailed tests, reflecting the a priori directional hypothesis that preferred conditions should elicit higher responses than non-preferred conditions. These results were further examined with post-hoc one-tailed t-tests comparing responses of preferred to non-preferred conditions for each of the classified voxel types. These same methods and statistical procedures were also applied to secondary analyses examining the tuning separately in posterior and anterior face areas and in motor cortex and early visual cortex control regions for tuning to all six individual faces as well as for tuning to face categories (F1-F3 vs F4-F6).

The weighted beta values from each session were also used to calculate the intraclass correlation coefficient (ICC) between the two fMRI sessions from the same participant. A separate ICC value was computed for each of the six voxel types, and these ICC values were averaged. Mean ICC values were -0.07, indicating a low correspondence at the individual voxel level between fMRI sessions for the same subject. We therefore analyzed all data separately for each session (N=20). It is unsurprising that ICC values would be low, since face preference was determined separately for each voxel, and the exact population of

neurons in a voxel is likely to be strongly affected by the specific head position and slice prescription in each session.

In addition, in order to probe whether tuning curves changed systematically along the morph continuum, we conducted a session-specific linear regression analysis on the difference between F1 and F6 tuning curves (F1 minus F6). These correlation (r) values were normalized by Fisher transformation, and a one-tailed paired t-test was then used to test for significant differences from zero (reflecting the a priori hypothesis that a negative relationship should exist; i.e., the difference between F1 and F6 responses should be more positive when F1 is presented and more negative when F6 is presented). Note that here, unlike in the tuning curve analysis, the values were not weighted by the number of reliable voxels in each session, as multiple voxel types were combined for this analysis.

Additional analyses quantified variability in the number of voxels representing different faces. The percentages of reliable voxels classified as preferring each face were entered into a one-factor repeated-measures ANOVA. Here, unlike the other ANOVAs, more than two groups were compared, so statistical results were corrected for sphericity using the Huyn-Feldt method in R's ezANOVA package (R Development Core Team, 2011).

Analysis of attention modulation**.** Beta values from the face attention task were computed for each reliable voxel and examined according to each voxel's preferred face type (defined by the classification procedure on the independent face mapping data set from Expt. 1, but using all mapping runs, rather than one run left out). The differences in beta values between attend-F1 (compared with fixation) and attend-F6 (compared with fixation) contrasts were then calculated for each of the six voxel types. We tested for a U-shaped response pattern with two-tailed paired t-tests that compared responses in voxels preferring one end of the morph continuum to responses in voxels preferring intermediate faces (F1 vs F2-5 and F6 vs F2-5).

In order to directly compare the two attention conditions, beta values from the attend F6 condition were subtracted from those from the attend F1 condition after subtracting the average response to each condition (the attend F6 condition elicited overall higher responses, perhaps contributing to the slightly higher morph detection rate observed in this condition). We tested for attentional modulation of responses with paired two-tailed t-tests on the normalized beta values within each voxel type. These p-values were corrected for multiple comparisons using the False Discovery Rate (FDR) procedure, implemented in R's p.adjust program. An analogous procedure was used to analyze attention affects in voxels classified categorically (F1-F3 vs F4-F6).

Control Analyses**.** In order to determine whether voxel-based face tuning curves were only present in face-selective (i.e., preferring faces to scenes) regions, we conducted tuning curve and attentional modulation analyses in a control area comprised of bilateral motor cortical regions. This area was selected from the FSL Harvard Oxford Probability atlas (thresholded at 25% probability, (Desikan et al., 2006)) and reverse normalized into each individual subject's native space. The control region was then masked to remove non-brain voxels (leaving an average of 7218 voxels, SEM=457; **Table 2.1**). All analyses described

above (including selection of reliable voxels, classification, tuning, and attention analyses) were then conducted for this control region. This constitutes a strong test of possible bias in our procedures and of the specificity of our findings to visual regions, since using a large area with many voxels increases the probability of finding some voxels with face tuning.

We also conducted a second control analysis using an area centered over the calcarine sulcus and covering early visual cortex. This region was selected from the FSL Harvard Oxford Probability atlas (*intracalcarine*, thresholded at 25% probability (Desikan et al., 2006)). In addition, because this region did not extend to the occipital pole (and might, therefore, have excluded foveal visual field representations within the stimulus), we manually drew an extension for this region to the occipital pole. We then reverse normalized this region into each individual subject's native space and masked it to remove any non-brain voxels (leaving an average of 2022 voxels per session, SEM=154; **Table 2.1**). Note that this region is similar in size to the combined posterior and anterior face-selective area used for our main experimental analyses. As before, all previously described analyses were conducted for this early visual control region, constituting a strong test of the specificity of our findings to higher-level visual regions.

	Number of sessions (N)	Voxel number (mean +/- SEM)	Voxel count range
All face areas	20	2000 (+/- 341)	113 - 5541
Posterior face areas (OFA, FFA, fSTS)	18	1636 (+/- 318)	77 - 5400
Anterior face areas (fAT)	17	267 (+/- 95)	13 - 1507
Motor cortex (control)	20	7218 (+/- 457)	2415 - 8199
Early visual cortex (control)	20	2022 (+/- 154)	636 - 3325

Table 2.1: Number and size of face areas examined. The number of sessions for which each face-selective region could be identified and the mean size of each region are listed. Regions examined included posterior face areas (OFA, FFA, and fSTS), anterior face areas (small clusters in the anterior temporal lobe, fAT), and two control sites (motor cortex and early visual cortex). In most cases, regions were bilateral.

2.4 Results

Subjects participated in two experiments within the same fMRI session (**Figure 2.1**), and each fMRI session was analyzed separately (***see Materials and Methods***). The aim of Experiment 1 was to generate voxel-based tuning curves for a set of six morphed faces (F1-F6), and the aim of Experiment 2 was to measure the effects of feature-based attention to one of the previously presented faces (F1 or F6) on these voxel-based tuning curves.

2.4.1 Identification of face selective regions

In each session, clusters of voxels in temporal and occipital cortex that responded more to faces than scenes during the localizer task were selected for further analysis (**Table 2.1**). These clusters included regions classically identified as part of a face processing network (Haxby et al., 2000), including the FFA, OFA, fSTS, and regions in the anterior temporal lobe (fAT; see **Figure 2.2** for an example from one participant and **Table 2.1** for the size and number of face areas that were identified across sessions).

2.4.2 Experiment 1: Face Tuning

Behavior. In Expt. 1, blocks of a single face morph (F1, F2, F3, F4, F5, or F6) or fixation were presented to participants while they counted the number of contrast decrements within each face morph block (**Figure 2.1**). This encouraged them to maintain spatial attention on the face without attending to specific facial identities or features. The amount of contrast decrement was adjusted for each run in order to maintain accuracy at approximately 75% correct across runs (chance performance = 25%). Mean accuracy on the task was 73.4% (sd=6.6%), and the mean contrast decrement required to maintain this accuracy was 0.122 (sd=0.04).

Voxel-based tuning across all face-selective areas. We constructed voxel-based tuning curves (**Figure 2.3**) of responses to each face morph in the continuum by employing a cross-validation procedure described by Serences et al. ((Serences et al., 2009)**, see *Materials and Methods***). This method classifies the preferences of individual voxels directly from the amplitude of the hemodynamic responses to presentation of face stimuli and is distinct from methods that classify voxels based on their attenuation to repetition, as in fMR-A, or on the weights they contribute to a model, as in MVPA studies. Critically, this technique allows us to determine response profiles of single voxels (rather than across a population, as with MVPA) that can then be compared across experimental conditions.

We assessed the reliability of our voxel-based tuning curves using a two-factor random effects ANOVA, with factors of voxel type (F1-F6) and condition (preferred or non-preferred; e.g., for a voxel classified as F1, the preferred condition was responses in the F1 condition, and the non-preferred condition was the average of responses in the F2-F6 conditions). As a group, tuning curves demonstrated selectivity for the preferred face compared to non-preferred faces ($F(1,19)=4.20$, one-tailed $p=0.027$). There was no significant main effect of voxel type or interaction between voxel type and condition. Post-hoc t-tests revealed that four of the six tuning curves had significant selectivity for the preferred face (voxels preferring F1: $p=0.024$; F4: $p=0.041$; F5: $p=0.031$; F6: $p=0.012$; **Figure 2.3A**).

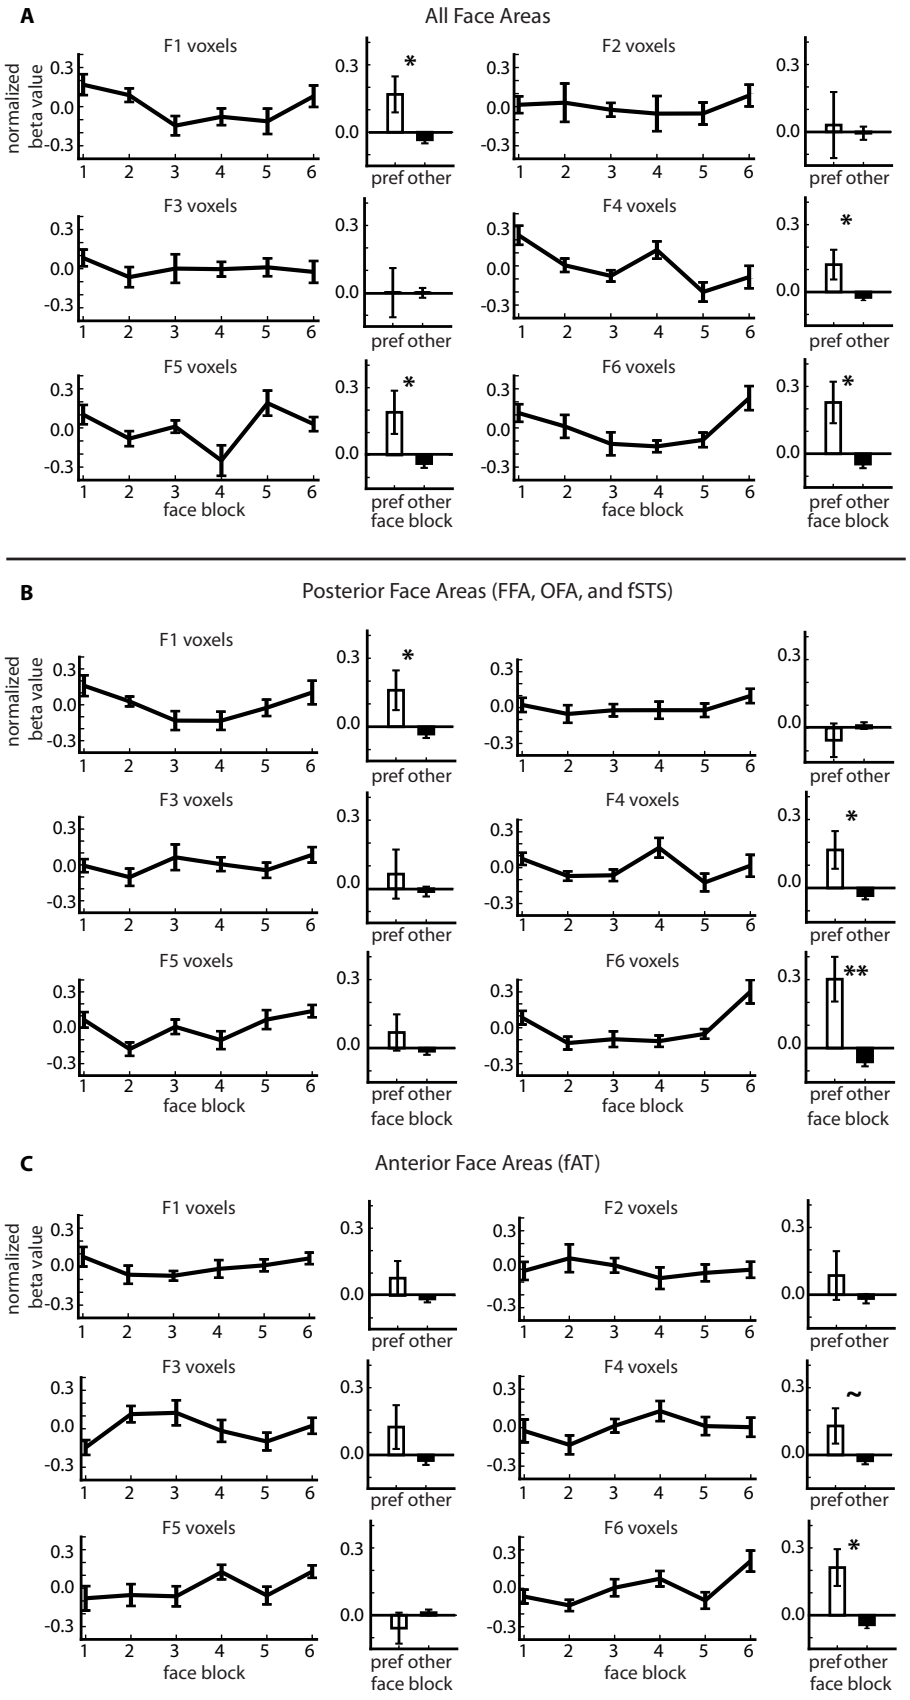
A
All Face Areas
F1 voxels
F2 voxels
F3 voxels
F4 voxels
F5 voxels
F6 voxels
normalized beta value
face block
pref other
B
Posterior Face Areas (FFA, OFA, and fSTS)
C
Anterior Face Areas (fAT)

Figure 2.3: Voxel tuning curves (Expt. 1). (A) Voxel-based tuning curves are shown for all reliable face-selective voxels in occipital and temporal cortex. Tuning curves for voxels preferring each face morph (F1-F6) are displayed in separate subplots across the six face conditions (error bars = SEM across fMRI sessions for this and all subsequent figures). The bar graphs to the right of each tuning curve plot the response to the preferred face compared to the average response to non-preferred faces for that voxel type (e.g., for F1 voxels, the response to the F1 condition vs the mean response to F2-F6 conditions). Data are also displayed for voxels in only posterior (B) and anterior (C) face areas. *, *p<0.05*; **, *p<0.01*; ∼, *p<0.10*

Tuning in posterior vs anterior face areas. Next, tuning curves were examined separately for voxels in more posterior (FFA, OFA, fSTS) (**Figure 2.3B**) and anterior (fAT) (**Figure 2.3C**) face areas. As a group, tuning curves in posterior areas demonstrated significant selectivity for the preferred face ($F(1,17) = 3.87$, one-tailed $p=0.033$). In anterior areas, tuning curves demonstrated only marginally significant selectivity for the preferred face ($F(1,15) = 2.86$, one-tailed $p=0.056$). Post-hoc t-tests showed that in posterior areas, three of six tuning curves showed significant selectivity (F1: $p=0.042$, F4: $p=0.030$; F6: $p=0.003$), while only one tuning curve showed significant selectivity in the anterior areas (F6: $p=0.010$; marginally significant in F5: $p=0.060$). There was no significant main effect of voxel type or interaction between condition and voxel type in either set of areas.

Additional tuning analyses. To further assess whether tuning differed systematically across the morph continuum, we computed the difference between the F1 and F6 within-voxel tuning curves (F1 minus F6) and conducted a linear regression analysis on data from each session. If tuning curves change smoothly along the morph continuum, there should be a systematic progression from stronger responses in F1 voxels for faces similar to F1 to stronger responses in F6 voxels for faces similar to F6, resulting in a negative slope in the difference between the F1 and F6 tuning curves. In posterior areas, this negative linear relationship between the location of the presented face within the morph continuum and the F1/F6 response difference across sessions was observed ($R = -0.22$, $p=0.035$), but this was not the case for anterior areas ($R = -0.03$; **Figure 2.4A**). This indicates that in posterior areas, response amplitude varied systematically as a function of the six morphed faces, with higher responses in F1 voxels for faces closer to F1 within the face continuum and higher responses in F6 voxels for faces closer to F6.

It is possible that in some brain areas, faces are represented in a more categorical fashion (e.g., F1 through F3 vs F4 through F6) rather than being represented as individual exemplars of faces along a morph continuum. This sort of division might allow for representations based on different categories within this set of face morphs, such as male vs female or old vs young. We therefore used the same cross-validation procedure to recompute the voxel tuning curves for only two classes of voxels: "male" voxels (preferring F1 through F3 more than F4 through F6) and "female" voxels (preferring F4 through F6 more than F1 through F3; note that use of the terms "male" and "female" here is somewhat arbitrary since the two halves differed along multiple categorical dimensions). This binary classification revealed significant tuning in anterior ($F(1,15)=5.77$, $p=0.030$) but not in posterior ($F(1,17)=1.10$) face areas (**Figure 2.4B**). Post-hoc t-tests demonstrated that in anterior areas, both "male" and "female" voxels showed significantly higher responses to their preferred face class relative to their non-preferred face class (first half: $p=0.040$;

second half: p=0.012). This indicates that, although anterior face areas only show marginally significant tuning to individual face morphs, they exhibit significant tuning when comparing responses to the two halves of the face continuum. That is, while posterior areas primarily show tuning for individual faces along the entire morph continuum, anterior areas mainly show tuning for face categories within this continuum.

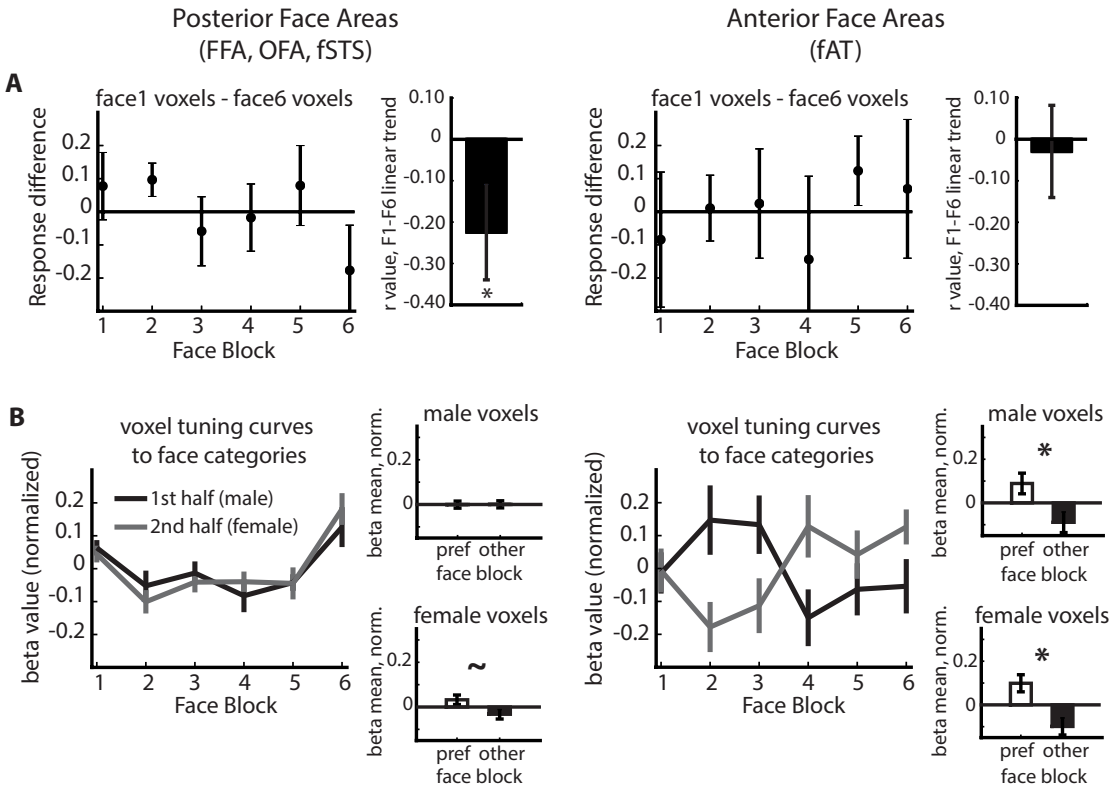

Figure 2.4: Posterior and anterior face areas have different tuning properties (Expt. 1). (A) A regression analysis was conducted on the difference between tuning curves for F1 and F6 voxels in individual fMRI sessions. A significant linear relationship was present in posterior but not anterior face areas. (B) Tuning curve plots show responses for voxels classified into two face categories (black = F1 through F3, gray = F4 through F6). Bar graphs compare responses to preferred and non-preferred conditions (e.g., for male voxels, preferred = F1, F2, and F3 blocks; non-preferred = F4, F5, and F6 blocks). Responses to the preferred half of the face morph continuum were significantly higher than responses to the non-preferred half in anterior (both halves $p<0.05$) but not in posterior face areas.

2.4.3 Experiment 2: Face Attention

In addition to the face mapping task, participants completed a feature-based attention task in the same session (**Figure 2.1B)**. This allowed for analysis of the effects of attention on voxel-based tuning for faces. Participants viewed superimposed images of F1 and F6 and were instructed to attend to just one of the two faces in each block and to count the number of morphs in the attended face towards a third face.

Behavior. The amount of morph towards the third face in each block was modified on each run for each attention condition in order to maintain morph detection performance at about 75% correct across runs (chance performance = 25%). Mean accuracy was 71.3% (sd=5.5%) for attending to F1 and 74.8% (sd=6.4%) for attending to F6. On average, a 34.7% (sd=7.4%) morph for F1 and 35.0% (sd=10.4%) morph for F6 were required to maintain this performance level.

Attention effects across all face areas. We used the voxel classifications from Expt. 1 to measure attentional modulation of reliable face-selective voxels to a composite display consisting of only the two faces at either end of the face morph continuum (F1 and F6), superimposed and tilted 45 degrees in opposite directions. Therefore, if the voxel classification procedure from Expt. 1 was accurate, responses to the composite display (relative to fixation) should be elevated in voxels preferring F1 and F6, compared to voxels preferring intermediate faces. Indeed, responses to the composite F1/F6 image during the attention task exhibited a U-shaped profile across voxel types in both attend-F1 and attend-F6 conditions, with highest responses in voxels classified as preferring F1 and F6 in Expt. 1, and lower responses in F2-F5 voxels (**Figure 2.5A**; F1 vs F2-F5: $t(19) = 2.84$, $p=0.011$; F6 vs F2-F5: $t(19) = 3.84$, $p=0.001$). Therefore, classification of voxels from Expt. 1 correctly predicted which voxels would exhibit the highest response in an independent data set from Expt. 2. This corroborates the findings from Expt. 1, providing further evidence that our voxel classification procedure correctly identified voxel preferences for individual faces. This concordance of the face mapping and attention data is notable, since it occurred despite differences between the two experiments in task, size of stimulus (in degrees of visual angle), and stimulus configuration.

To assess the effects of feature-based attention on responses to the pair of superimposed faces, we examined the difference between conditions (attend F1 minus attend F6, following subtraction of the average response to both attention conditions in Expt. 2, across all voxels that were classified as reliable in Expt. 1). Attending to F1, relative to attending to F6, selectively enhanced responses in F1 voxels (**Figure 2.5A**; $t(19)=2.97$, $p[FDR]=0.032$). Similarly, attending to F6, relative to attending to F1, selectively enhanced responses in F6 voxels ($t(19)=-2.84$, $p[FDR]=0.032$) and marginally enhanced responses in F5 voxels ($t(19)=-2.13$, $p[FDR]=0.092$).

Attention Effects in Posterior vs Anterior Face Areas: Posterior areas showed a similar pattern of attention effects to those seen for all face areas combined. Responses were U-shaped across voxel types, again suggesting that classification based on face mapping data correctly identified the six classes of voxels (**Figure 2.5B**; F1 vs F2-F5: $t(17)=2.75$, $p=0.014$; F6 vs F2-F5: $t(17)=4.30$, $p=0.0005$). Moreover, in posterior areas, attending to F1 selectively enhanced responses in F1 voxels ($t(17)=3.01$, $p[FDR]=0.017$), and attending to F6 selectively enhanced responses in F5 and F6 voxels (F5: $t(17)=-3.06$, $p[FDR]=0.017$; F6: $t(17)=-3.52$, $p[FDR]=0.016$) (**Figure 2.5B**).

In anterior face areas, the response profile was not U-shaped (**Figure 2.5C**; F1 vs F2-5: $t(15)=-0.351$, F6 vs F2-5: $t(15)=-0.429$), and there were no significant differences between attention conditions (**Figure 2.5C**; F1: $t(15)=-0.67$; F6: $t(15)=0.11$, all $p[FDR]>0.17$). In

addition, anterior face areas exhibited no significant attention effects even when responses were categorically grouped by voxels preferring each face class (F1 through F3 and F4 through F6) rather than by voxels preferring individual face morphs (F1 through 3: t(15)=-1.75, F4 through 6: t(15)=1.68). Note that numerically, these values in anterior face areas are in the opposite direction from that expected if attention to a feature enhanced responses in neurons selective for that feature.

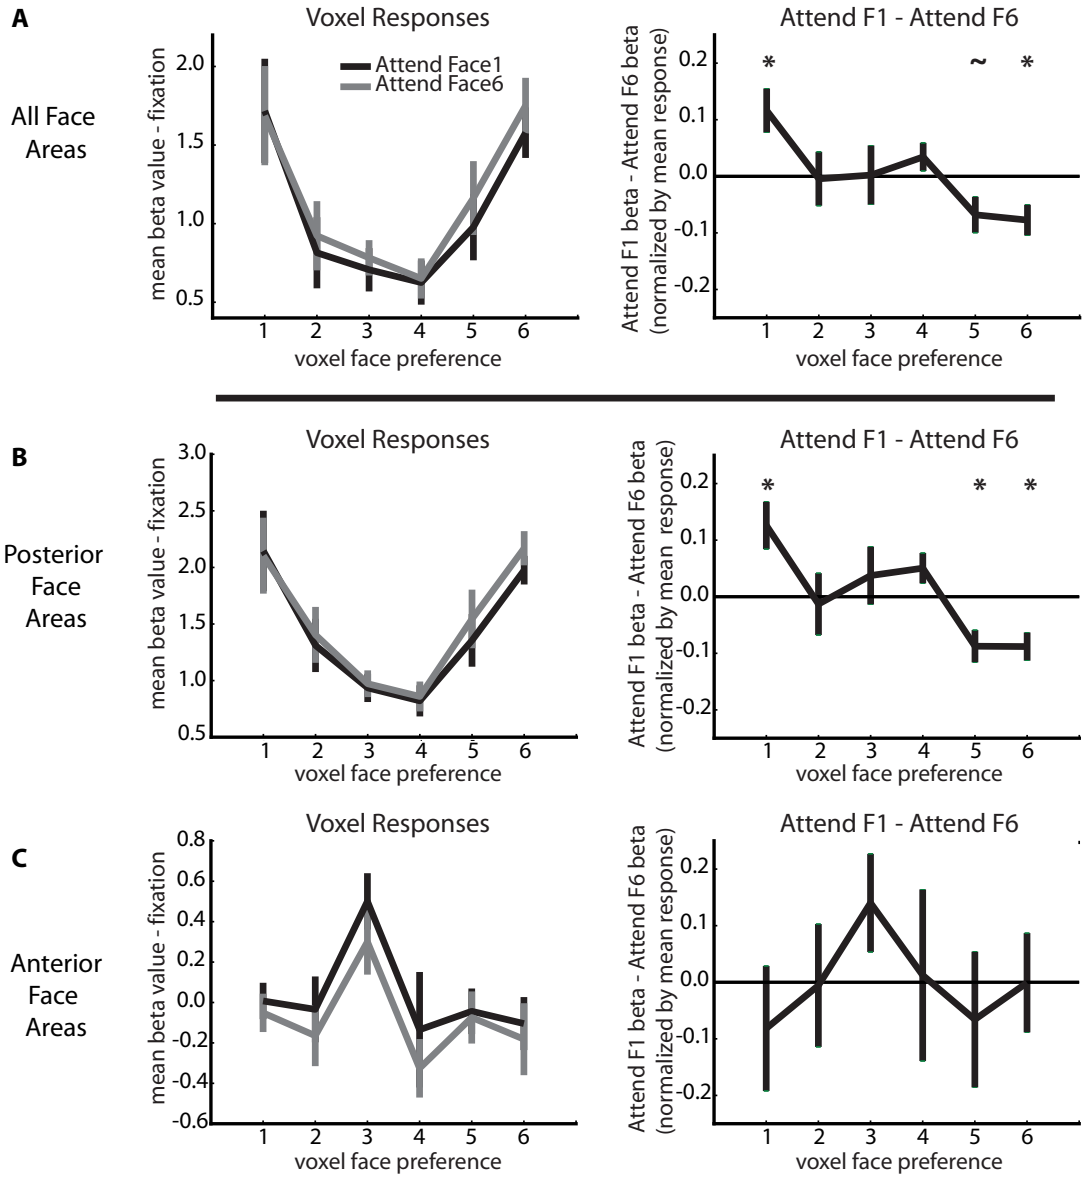

Figure 2.5: Responses during attention task (Expt. 2) across voxel types. (A, left) Beta values for each attention condition (compared to fixation) from voxels across all face areas are plotted according to the face preference those voxels displayed in Expt. 1. (A, right) The difference between attention conditions (normalized by subtracting the mean response to each condition across all voxel types) is shown for voxels across all face areas. The same information is plotted for voxels in only posterior (B) and anterior (C) face

areas. Note that in this figure (unlike Figs. 2.3 and 2.4), responses are plotted across voxel types on the x-axis rather than across stimulus conditions (within voxel types). *, *p<0.05*; ~, *p<0.10; FDR-corrected*

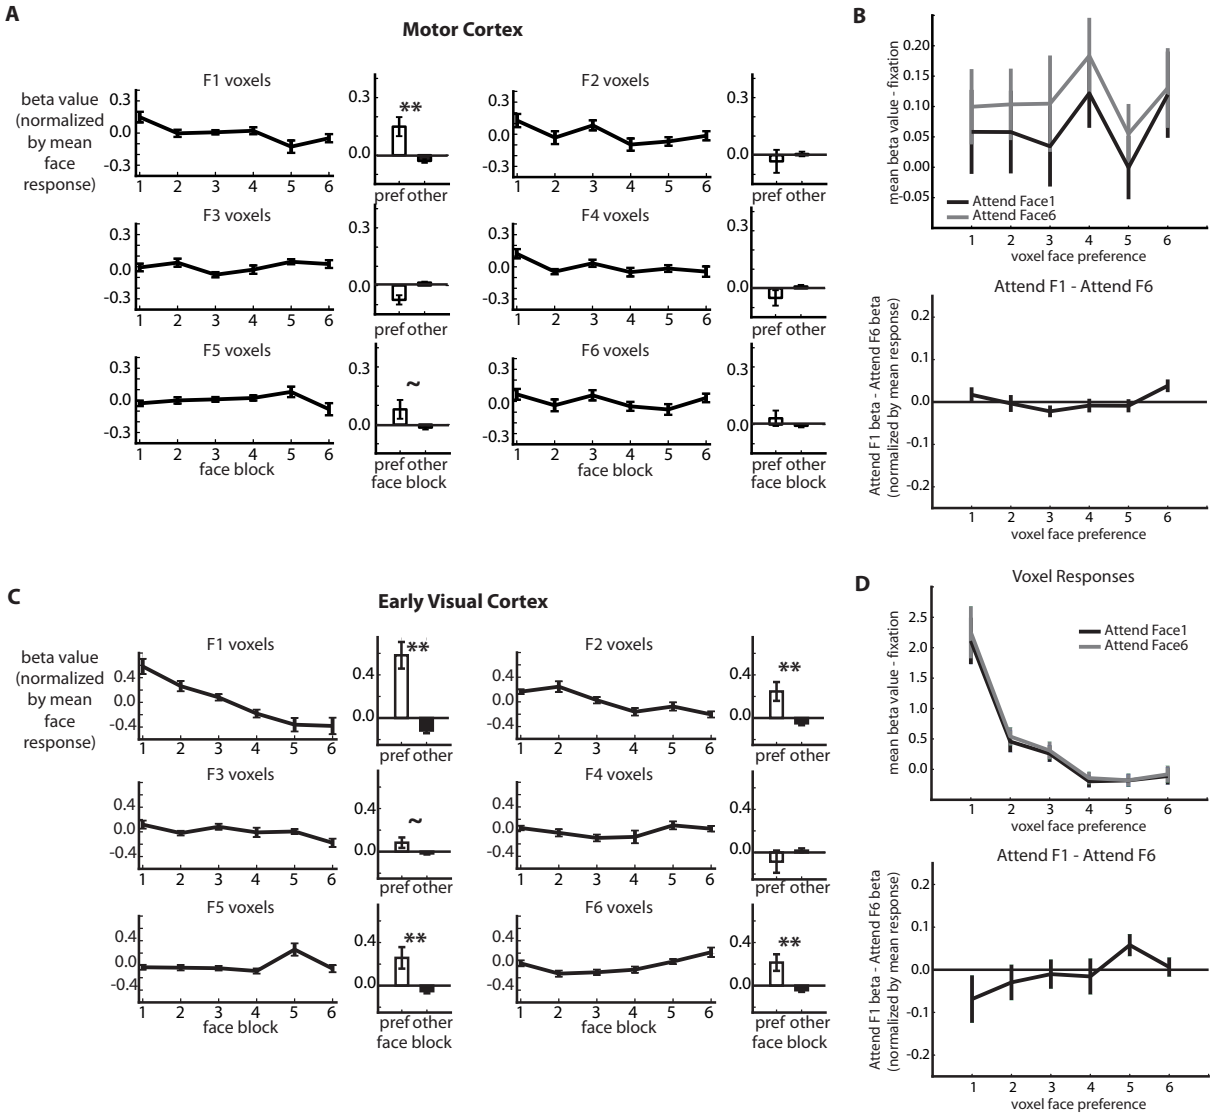

Figure 2.6: Tuning and attention results in control bilateral motor and early visual cortical regions. Group voxel-based tuning curves for bilateral motor (A) and early visual (C) regions, plotted using conventions from Figure 2.3 (**, *p<0.01*; ~, *p<0.10; note that the scale is larger in subplots for C to accommodate the greater range of beta values*). Responses in the two attention conditions for motor (B) and early visual (D) areas are displayed on the top subplot using the same conventions as Figure 2.5. The difference between attention conditions, normalized by the average response to each condition, is displayed in the bottom subplot for each area, as in Figure 2.5 (all n.s.).

2.4.4 Experiments 1 and 2: Control Area Tuning and Attention

In order to assess whether there was any bias in our methods for constructing voxel-based tuning for faces or for measuring attentional modulation of this tuning, we conducted all analyses on an anatomically-defined bilateral motor cortical region. Despite the large size of this region, which should increase the probability of finding a small subset of voxels with tuning, voxel tuning across the six faces was poor (**Figure 2.6A**). The main effect of condition (preferred versus non-preferred) was not significant ($F(1,19)=0.448$), and there was no detectable main effect of voxel type or interaction between voxel type and condition. Only one of the six tuning curves showed significantly higher responses for preferred, compared to non-preferred, faces in post-hoc t-tests (F1: $t(19)=3.0$, $p=0.004$). Furthermore, there was no significant attentional modulation in any voxel type (**Figure 2.6B, bottom**; F1: $t(19) = 0.47$, F6: $t(19) = 2.19$, note that the F6 value in motor cortex is in the opposite direction from that expected if attention to a feature enhanced responses in neurons selective for that feature).

In addition, we conducted a separate control analysis in an anatomically-defined early visual cortex region in order to assess the specificity of our results to higher-level visual regions. We observed strong voxel tuning to the six faces in Expt. 1 in this region (**Figure 2.6C**), and the main effect of condition was highly significant ($F(1,19)=27.52$, one-tailed $p=0.00002$). In addition, 4 of the 6 individual tuning curves had significantly higher responses to preferred, compared to non-preferred, faces, based on post-hoc t-tests (voxels preferring F1: $p=0.0001$; F2: $p=0.006$; F5: $p=0.019$; F6: $p=0.012$; F3: $p=0.052$). However, the U-shaped response profile was not present in Expt. 2 (**Figure 2.6D, top,** F1 vs F2-F5: $t(19) = 4.73$, $p=0.0001$; but F6 vs F2-F5: $t(19) = 0.32$), suggesting limited invariance to the differences in stimulus properties between Expt. 1 and 2, including contrast and orientation of the faces. Furthermore, there was no significant attentional modulation in any voxel type (**Figure 2.6D, bottom**, F1: $t(19)=-1.23$, F6: $t(19)=0.28$).

2.5 Discussion

We employed a method previously used to construct voxel-based tuning curves for low-level features (Serences et al., 2009) to identify individual voxels that were selective for faces along a morph continuum (F1-F6). Different cortical areas showed distinct patterns of tuning: posterior face areas (FFA, OFA, fSTS) were primarily selective for individual face morphs, while more anterior face areas mainly exhibited categorical tuning for faces (F1-F3 vs F4-F6, e.g., "male" vs "female" halves). Data from a separate task in which subjects attended one of a pair of superimposed faces (F1 and F6) validated these tuning results, showing that F1 and F6 voxels exhibited significantly larger responses to a stimulus containing both faces, compared to voxels preferring intermediate face types. Critically, directing attention to a particular face selectively enhanced responses in voxels previously defined as preferring that face. Our results demonstrate that fMRI can be used to classify individual face preferences in single voxels and provide direct evidence that the basic principles of selectivity of attentional effects that have previously been described for lower level features apply equally well to attentional modification of the representations of complex stimuli.

2.5.1 Relationship to previous studies of neural face representations

Our findings complement previous studies that examined neural selectivity for faces across different spatial scales. Single neuron responses can be selective for face stimuli (see review by Gross, 2011), face feature dimensions and feature combinations (Freiwald et al., 2009), face morph vectors (Leopold et al., 2006), and even individual identities (Quiroga et al., 2005). The selectivity for individual face stimuli that we observed may depend on a combination of signals that are tuned for multiple distinct elements (individual faces, face features, and face morphs).

In addition, fMR-A has been used to show that responses in face-selective areas are sensitive to differences between individual faces (Kadosh et al., 2010; Pourtois et al., 2005; Winston et al., 2004), face parts (Andrews et al., 2010; Harris and Aguirre, 2008), and face morphs (Gilaie-Dotan and Malach, 2007; Loffler et al., 2005; Rotshtein et al., 2005). However, fMR-A measures may not accurately reflect neuronal selectivity (due to a number of factors, including possible feedback/feedforward influences (Bartels et al., 2008; Kriegeskorte et al., 2007; Mur et al., 2010) and a lack of correspondence between neuronal response selectivity and neuronal adaptation selectivity (Sawamura et al., 2006)). In contrast, our study used a method that is not susceptible to these issues and directly demonstrates voxel-level selectivity for individual faces.

It has also been suggested that multi-voxel pattern classification methods can be used to reveal a distributed representation of high-level features (Haxby et al., 2001). These methods have recently identified fMRI responses differentiating individual faces (Kriegeskorte et al., 2007; Natu et al., 2010; Nestor et al., 2011). Our results similarly demonstrate that voxels from multiple areas contain signals that are selective for individual faces. However, we found that only a small proportion of voxels within these areas were both reliably selective for the presented face morphs and modified by attention. This may not be surprising, as the signals that drive classification presumably result from small differences in the relative proportions of neurons tuned to different faces within each voxel, and we only tested a small number of faces in our study. Our results extend previous MVPA studies by demonstrating that single voxels are tuned for facial identity.

As noted above, our results agree with previous electrophysiological and fMRI studies investigating neural tuning for faces. However, since we only tested a small set of individual face stimuli, it is difficult to determine the extent to which the tuning properties described here generalize to a larger set of stimuli. Future studies with larger stimulus sets can add to these findings by more precisely determining the stimulus parameters to which individual voxels are tuned.

2.5.2 Tuning curve properties in different face areas

Our results demonstrate that face areas vary in the relative levels of continuous featural codes versus categorical codes used to represent face stimuli. Specifically, posterior face areas showed selectivity for individual face morphs. Furthermore, the difference in tuning

between F1 and F6 voxels indicated a significant monotonic relationship, suggesting that tuning varied systematically along the morph continuum. In contrast, anterior face areas showed selectivity for face categories (i.e., F1 through F3 vs F4 through F6, corresponding to male vs female halves of the continuum).

Our results support existing models of differences in representations across face areas. Specifically, anterior temporal face areas have been proposed to connect physical face information with more semantic, biographical detail, with posterior regions (OFA, FFA, and fSTS) being more involved in the processing of face features, face individuation, and gaze and expression information, respectively (Haxby et al., 2000). Neurophysiological studies also suggest differences in face coding along the posterior-anterior axis of the temporal lobe, with neurons coding for face features in more posterior regions (Freiwald et al., 2009) and for face morph vectors in more anterior regions (Leopold et al., 2006). Work in patients with prosopagnosia is also consistent with this distinction: patients with damage to posterior temporal or occipital cortex may be impaired in perceiving facial identity based on visual information, whereas patients with damage to the anterior temporal lobe are more likely to be unable to recognize an individual regardless of the source of information (Damasio et al., 1990).

One potential interpretation of our findings is that posterior areas show selectivity for individual faces, whereas selectivity in anterior face areas is related to broad classes of faces (males vs. female, old vs. young). However, previous psychophysical studies (Beale and Keil, 1995) and fMRI data from some face-selective regions (Rotshtein et al., 2005) suggest that individual face morphs along a continuum are not necessarily each represented as distinct identities, but rather as a set of two identities with a categorical boundary at some point along the morph continuum. Therefore, our findings in posterior areas could also be interpreted as these regions showing sensitivity to face *features* (rather than identities) that differed between individual morphs, whereas our findings in anterior areas may reflect sensitivity to distinct facial identities. This interpretation of our findings adds support to past literature by providing population-tuning evidence for a more feature-based code (selectivity for individual morphs) in posterior cortex and a more categorical code (selectivity for halves of the morph continuum) in anterior cortex. Although it is beyond the scope of this paper, future work could examine whether this category/exemplar distinction is represented as a continuous gradient along the anterior-posterior axis of the inferior temporal lobe or if there is a sharp boundary in this region that demarcates distinct coding mechanisms in different face areas.

2.5.3 Feature-based attention modulates voxel-level face representations

We have found that attention can selectively modify responses in voxels tuned for individual exemplars within a category. Previous studies have suggested that feature-based attention increases responses of neurons to attended features in lower-level visual regions such as V4 (McAdams and Maunsell, 2000; Motter, 1994) and MT (Martinez-Trujillo and Treue, 2004; Treue and Martinez Trujillo, 1999). Our finding that attention selectively enhances responses in voxels representing the attended face is consistent with an increase in population gain. However, it has also been suggested that attention can shift and/or

reduce the size of receptive fields of individual visual cortical neurons (Connor et al., 1996; David et al., 2008; Womelsdorf et al., 2006) and sharpen the tuning of populations of neurons (through feature-similarity gain in individual neurons; (Martinez-Trujillo and Treue, 2004)). The strongly selective nature of the attentional enhancement that we observed is also consistent with modifications in population tuning of responses. We found that attention amplified responses only in voxels selective for the attended face (or sometimes in voxels preferring the face adjacent to the attended face along the morph continuum: F5 when attending to F6). This selectivity of attention effects occurred despite broad individual tuning curves in different voxel types.

In contrast to the extensive work on attentional modulation of tuning for low-level features, considerably less is known about attention's effects on tuning for complex stimuli. Previous studies have demonstrated that attending to items within a category can enhance signals across entire category-selective brain areas (Corbetta et al., 1991; Gazzaley et al., 2005; Wojciulik et al., 1998), shift MVPA signals toward the attended category (Chen et al., 2012; Reddy et al., 2009), and cause sharper and larger release from adaptation with fMR-A in these areas (Murray and Wojciulik, 2004). By showing that attention can operate selectively on individual exemplars within a complex category, we establish a connection between the effects of attention on complex objects and on low-level visual representations (Brefczynski and DeYoe, 1999; Saproo and Serences, 2010; Serences et al., 2009; Somers et al., 1999; Tootell et al., 1998), thereby providing evidence that the effects of attention on feature representations are precisely tuned across the visual hierarchy.

2.5.4 Attentional modulation of different face selective visual areas

We observed feature-based attentional modulation across all face areas when they were grouped together and in posterior but not anterior areas when examined separately. This dissociation may be because the attention task we used emphasized subtle shifts in facial features (small morphs towards a face not on the morph continuum). We found that posterior visual areas were sensitive to small changes in stimulus properties (i.e., differences between individual faces along the morph continuum), whereas anterior areas responded to larger categorical changes. Therefore, attention may have selectively acted on those brain areas that were most relevant for performing the task. Another possibility is that the smaller number of face-selective voxels within the anterior temporal lobe (see **Table 2.1**) resulted in decreased reliability of population responses of this area, thereby limiting our ability to detect attentional modulation. This interpretation is supported by the absence of a U-shaped profile of responses in the face attention experiment (**Figure 2.5C**).

2.5.5 Specificity of tuning and attention effects to face-selective areas

To determine whether our results were unbiased and selective to face-selective areas of cortex, we performed a control set of analyses in bilateral motor and early visual cortical regions. Despite the large size of the motor cortex region, which could have favored detection of a subset of voxels with appropriate tuning properties, we found no evidence that voxels responded selectively to individual faces within this region. Furthermore, we found no evidence for attentional modulation of responses in the motor cortex control

region. These results show that our method for identifying voxels tuned for individual faces is unbiased and that the tuning and attentional modulation we report are specific to visual areas.

In addition, we used an early visual cortex control region to further investigate the specificity of our findings to higher-level visual cortex. In this region, we found strong evidence for selective tuning to individual faces in Expt. 1, an expected result given that our stimuli differ from one another in a number of low-level features. In fact, there is evidence from previous studies for very selective tuning to natural images in early visual cortex (e.g., (Kay et al., 2008)). However, the lack of a U-shaped response profile in Expt. 2 suggests that this tuning is not invariant to small changes in the stimulus configuration (e.g., rotating the faces 45 degrees, changing their contrast). This is precisely what one would expect in an area with small receptive fields tuned to low-level features. Finally, and most importantly, the attention effects we observe in higher-level face-areas were absent in the early visual control area, suggesting that the attention effects are not based on low-level features in the stimulus. Rather, attention effects were observed in areas that are specifically tuned to faces.

- SECTION 2 -
DISRUPTION

Chapter 3

DOUBLE DISSOCIATION OF TWO COGNITIVE CONTROL NETWORKS IN PATIENTS WITH FOCAL BRAIN LESIONS

3.1 Abstract

Neuroimaging studies of cognitive control have identified two distinct networks with dissociable resting state connectivity patterns. Patients with heterogeneous damage to these networks provide the first demonstration of network independence through a double dissociation of lesion location on two different measures of network integrity: functional correlations amongst network nodes and within-node graph theory network properties. The degree of network damage correlates with a decrease in functional connectivity within that network while sparing the non-lesioned network. Graph theory properties of nodes within the damaged network show evidence of dysfunction compared to the undamaged network. The effect of anatomical damage thus seems to extend beyond the lesioned area, but remains within the bounds of the existing network connections. Together this evidence supports the dual-network view of cognitive control.

3.2 Introduction

Cognitive control is required in everyday life to coordinate our thoughts and actions to achieve internal goals while still allowing for the flexibility to adjust these goals with changing task demands. While previous studies have attributed cognitive control to various prefrontal cortical regions (Badre and D'Esposito, 2007; Koechlin et al., 2003), recently it has been proposed that a dual-network architecture exists in the human brain where cognitive control depends on regions that extend beyond the frontal cortex (Dosenbach et al., 2007). In a recent cross-task analysis, Dosenbach and colleagues (Dosenbach et al., 2006) identified a number of regions active during different stages of cognitive control paradigms. Given the difficulty in isolating cognitive control networks that are simultaneously active during task performance, the investigators took advantage of the recent advent of resting state fMRI (rs-fMRI) for detecting spontaneous fluctuations between coherent brain regions. In a follow-up study, these pre-defined regions-of-interest (ROIs) obtained from the task data served as seeds in a correlation analysis of rs-fMRI data (Dosenbach et al., 2007) where graph theory and hierarchical clustering was applied to the correlation matrices. These analyses identified two distinct networks labeled as fronto-parietal (FP) and cingulo-opercular (CO) (**Figure 3.1A**). Based on their role in task data, the fronto-parietal network consists of nodes proposed to provide signals that act on a rapid time scale to initiate and adjust control whereas the cingulo-opercular network nodes act to provide signals that allow for set-maintenance over a longer time scale (Dosenbach et al., 2008; Dosenbach et al., 2007).

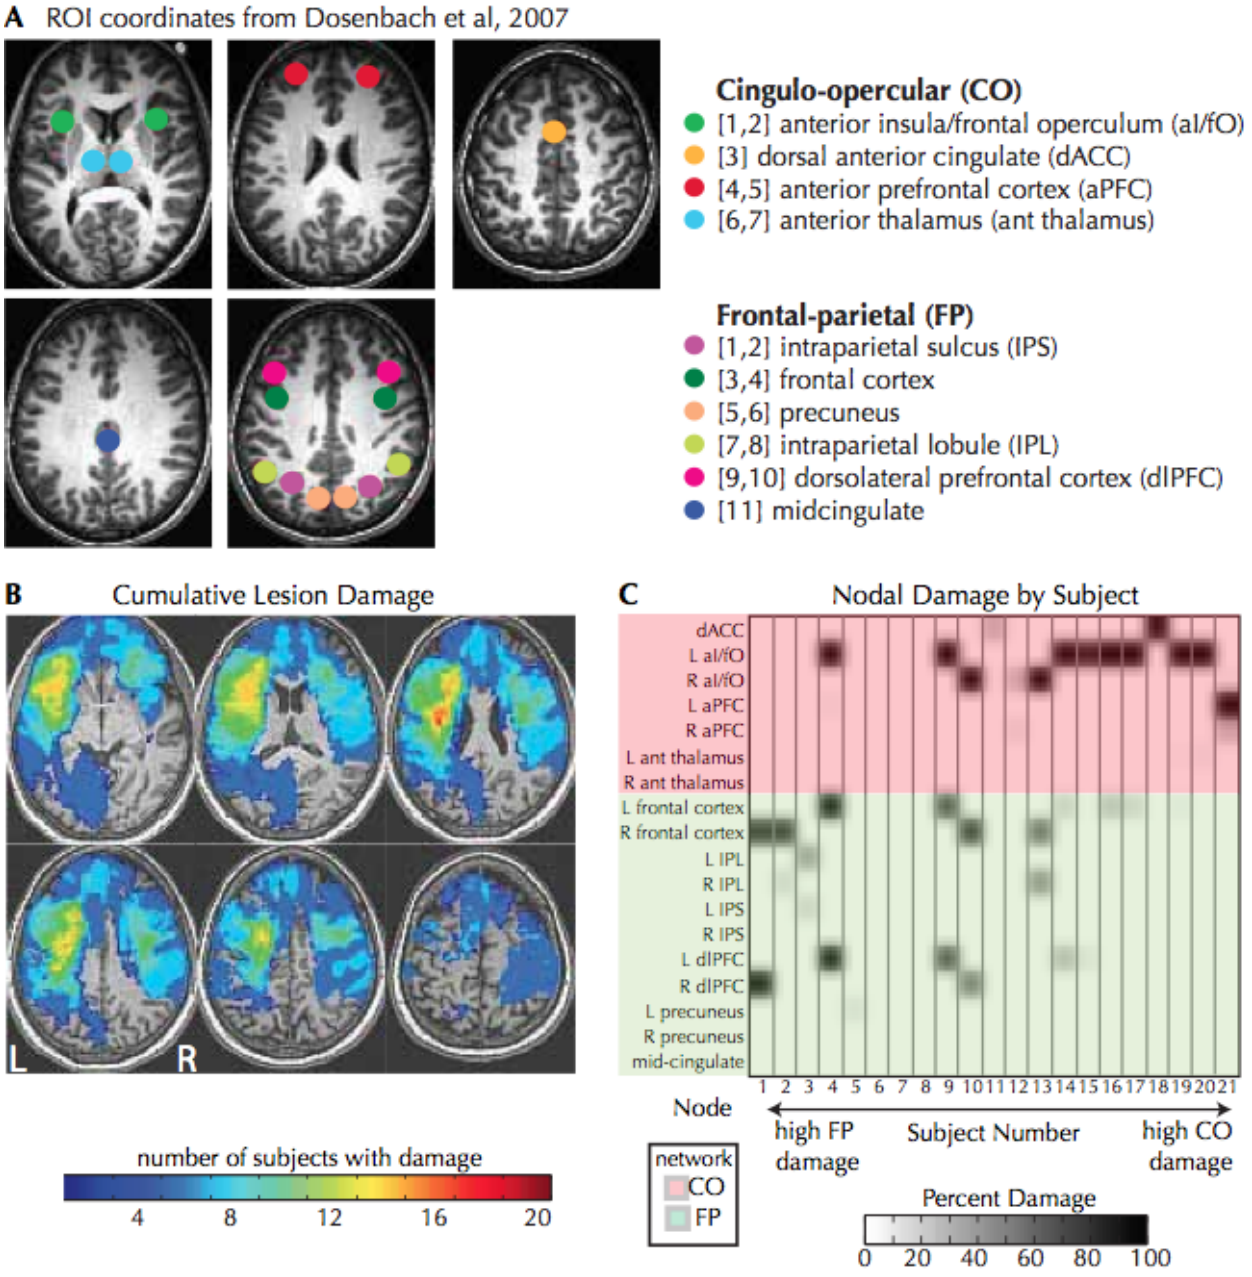

Figure 3.1: (A) Average FP and CO connectivity across subjects in the lesion and control group. (B) Average FP and CO network connectivity within each lesion patient sorted from high FP to high CO damage

Numerous studies using rs-fMRI have shown that neuronal activity is characterized by temporal correlations in blood oxygen level-dependent signal across disparate brain regions (Biswal et al., 1995; Fox et al., 2005). These fluctuations seem highly consistent over time and reflect the presence of intrinsic functional (Shehzad et al., 2009) and structural (Greicius et al., 2009) connectivity. Among these fluctuations, different networks can be distinguished, many of which show remarkable resemblance to task-related

networks (Smith et al., 2009). Using rs-fMRI, there are a number of different methods for assessing network properties, from measuring the magnitude of correlations between single nodes (He et al., 2007) to calculating mean correlation values across pre-defined network nodes (Van Dijk et al., 2010). Here we use the latter method to compare mean correlations within and between the FP and CO networks.

Although correlations between brain regions can reveal global network properties, they do not provide information about the relationship between nodes within those networks or about the small-scale regional organization within network nodes. Graph theory is a mathematical tool that has recently been applied to rs-fMRI data with the goal of quantifying the organization of network nodes both at a whole-brain and local level (Achard et al., 2006; Buckner et al., 2009; Sporns et al., 2007; Watts and Strogatz, 1998). One graph theory property that distinguishes brain networks from other non-biological networks is its characteristic small-world architecture (Watts and Strogatz, 1998) that is hypothesized to be scale-free in that it exists at both local and global levels (Eguiluz et al., 2005). In both structural and functional MRI data (Achard et al., 2006; Sporns et al., 2007), brain networks have been found to be optimized for high local and global information transfer, while maintaining low wiring costs. This organization is so consistent that disruptions in small-worldness can be used as a biomarker for distinguishing young healthy subjects from those with Alzheimer's disease (Stam et al., 2007), schizophrenia (Liu et al., 2008) and even normal aging (Achard and Bullmore, 2007).

Examining the correlation structure of resting state networks in patients with focal lesions offers the opportunity to test the hypothesized independence of the FP and CO cognitive control networks. Furthermore, the incorporation of graph theory analyses allows for the detection of small-scale changes in regional organization that may underlie more global network changes. In the current study, we scanned patients with heterogeneous focal brain damage in a rs-fMRI session. We computed time-series correlations among the ROIs in the FP and CO network as well as among voxels within each ROI. At the global network level, we found lower between- than within-network correlations across all patients and a negative relationship between within-network functional connectivity and the amount of damage sustained by the network. Simulating lesions in data from healthy control subjects resulted in less severe deficits than found in our patients. Thus, it is likely that the functional disruption captured by the decrease in connectivity extends beyond the site of anatomical damage to the remaining non-lesioned nodes of the network. Indeed, the small-worldness of these nodes was lower than those of the undamaged network supporting the notion that global network connectivity changes can be accounted for by disruptions in local network organization. The results here support the independence of these two cognitive control networks and more broadly the dual-network architecture for cognitive control.

3.3 Materials and Methods

3.3.1 Participants

Twenty-one patients with unilateral focal lesions due to stroke, mean age 58 (range = 19-83, sd +/- 14 years) and 21 healthy subjects, mean age 61 (range 19-83, sd +/- 17 years) were studied. All healthy participants were pre-screened to exclude individuals with a history of neurologic or psychiatric conditions. Informed consent was obtained from subjects in accordance with procedures approved by the Committees for Protection of Human Subjects at the University of California, Berkeley.

3.3.2 MRI Acquisition Procedures

T2*-weighted echo planar images (EPI) were collected on a whole body 3-T Siemens MAGNETOM Trio MRI scanner using a 12-channel head coil. Structural images were acquired using an axial MP-RAGE 3-D T1-weighted sequence (TR = 2,300 ms, TE = 2.98 ms, FA = 9°, 1x1x1 mm voxels) for patients and controls and an additional FLAIR image was collected for each patient to better visualize the lesion. For patients, 10 minutes of EPI data were analyzed (300 time points, TR = 2,000 ms, TE = 30 ms, 28 3.30 mm thick axial slices). For controls, 10 minutes of EPI data were analyzed in 7 subjects (435 time points, TR = 1370 ms, TE = 26 ms, 24 3.85 mm thick axial slices), and 4 minutes and 20 seconds of EPI data in 14 subjects (250 time points, TR = 1,000 ms, TE = 50 ms, 24 4.025 mm thick axial slices). All participants were instructed to stay awake with their eyes open; no other task instruction was provided.

3.3.3 MRI Preprocessing

Image preprocessing was carried out by using AFNI (Cox, 1996). The following prestatistics processing was applied: slice-time correction, removal of nonbrain structures from the EPI volumes and spatial smoothing using a 5 mm Gaussian kernel. Following Fox et al. [5], signal from movement, white matter and ventricles was regressed out. The high-resolution T1-weighted image was coregistered to the mean functional data and after that segmented using SPM5 (Wellcome Department of Cognitive Neurology, London). The template used for segmentation was derived from 152 normal subjects (MNI152; Montreal Neurological Institute, Montreal, QC, Canada). Parameters obtained from segmenting the brain were later used to normalize each individual's T1-weighted brain. This extra segmentation step was necessary for accurate registration, which is often confounded by structural brain damage.

3.3.4 Lesion mapping

Lesion masks were manually traced by EN and CG in native patient space according to visible damage on a T1-weighted anatomical scan and guided by damage and hyperintensities on a T2-weighted FLAIR image. All lesion masks were examined by MD for anatomical specificity. Individual patient masks in normalized space are shown in **Figure 3.2**.

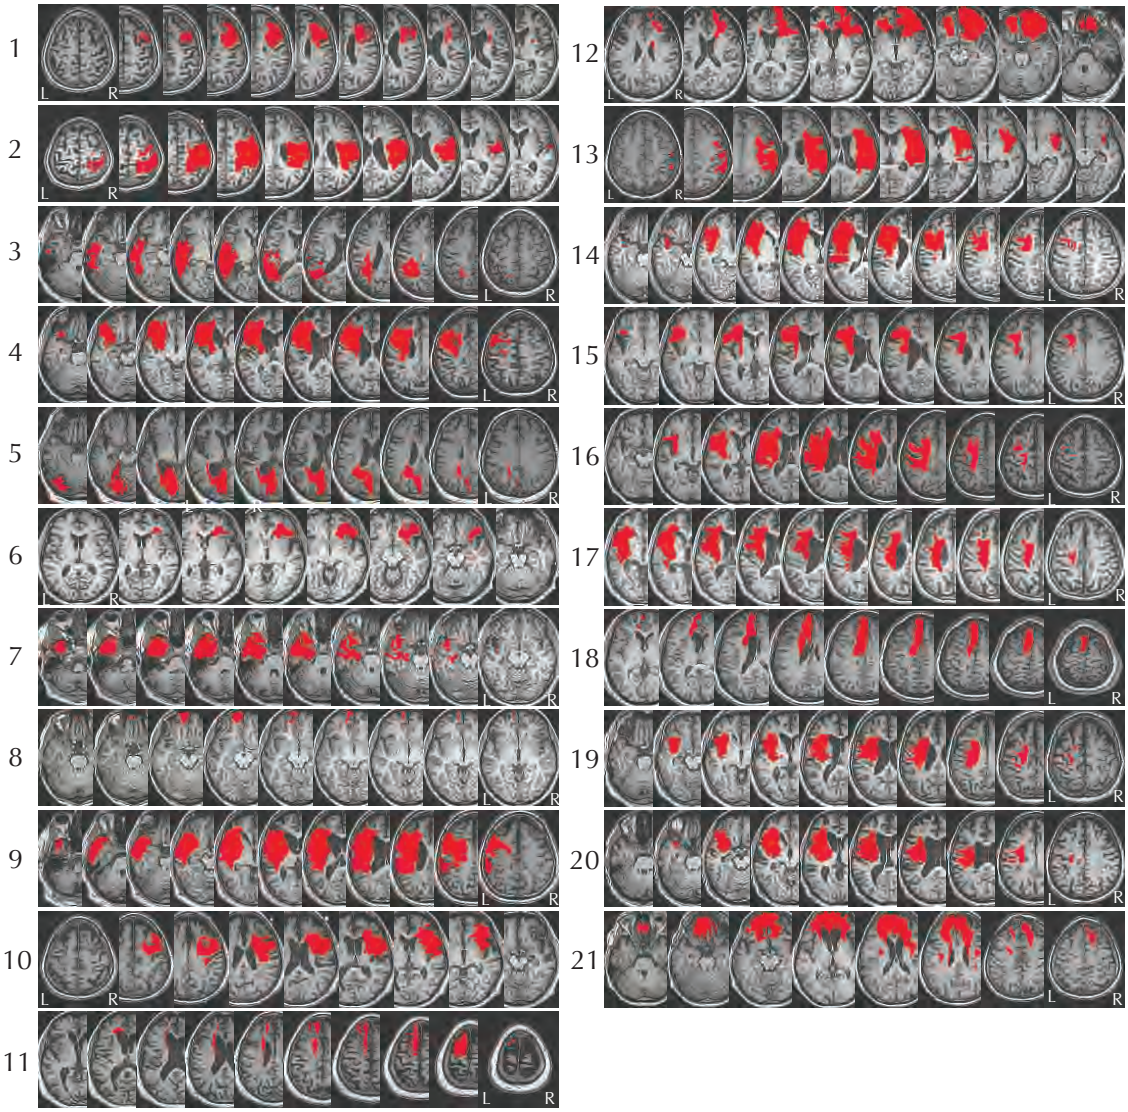

Figure 3.2: Lesion masks in normalized space for all 21 patients.

3.3.5 Percent network damage

To answer the question as to how lesions differentially affect two *a priori* defined networks (Dosenbach et al., 2007), the amount of damage for each individual was quantified by counting the percentage of voxels in each ROI that overlapped with the lesion. These percentages were then averaged over all the nodes within a network, resulting in one number for each network; a percentage of damage to that particular network. By subtracting these 'damage-scores' from one another, a number between -100 and + 100 was obtained, indicating relative damage to one network compared to the other, for each individual separately.

3.3.6 Functional connectivity

Eighteen 6-mm spheres were reverse-normalized to native space based on coordinates from Dosenbach et al. (2007). For between-ROI correlations, voxel time-series were averaged within each ROI and these averages were bandpass filtered (0.005-0.08 Hz) to remove physiological noise such as cardiac and respiratory artifact (Fox et al., 2005). Functional connectivity was assessed in each subject by correlating average time-series across seeds, resulting in an 18x18 matrix for each subject. It should be noted that pair-wise correlation values derived from this type of analyses only reflect the synchronized, or alternate, activations and deactivations in different regions and will not detect similar time-courses that are phase-shifted (i.e. these will show up as low correlations). Because of the low frequency of the filtered signal, and the nature of the resting state activity, we assumed that such phase-shifted patterns would not signal the presence of healthy interactions. For within-ROI correlations, all ROIs were matched for minimum number of voxels (74) across subjects and regions. Functional connectivity was assessed in each subject by correlating time-series across single voxels within each ROI.

Average strength of connectivity within a network was calculated by adding all fisher-transformed correlation values within a network and dividing this by the number of nodes within that network (as in He et al. (2007) and Van Dijk et al. (2010)). Apart from that average within- network strength, between-network strength was calculated by adding all the values obtained from the between-network interactions and dividing this by the total number of between-network interactions. This resulted in three values for each individual. We did not assess average strength of connectivity within each ROI as this would be highly subject to autocorrelations amongst neighboring voxels.

3.3.7 Partial correlations

A partial correlation removes the effect of a third variable from the relationship between the two variables of interest. The number of voxels in the lesion mask was calculated per patient and entered as a controlling variable in both the relative (**Figure 3.4A**) and individual network (**Figure 3.5**) correlation analyses.

3.3.8 Removing the contribution of lesion nodes from average correlations

The percent overlap between the lesion mask and individual ROIs (from **Figure 3.1C**) was used to construct a weighted-average that may more accurately reflect the contribution of each lesion node to the total network average. In this method, the correlation value from each node was multiplied by the percent damage of that node (between 0 and 1) and summed across nodes in either the FP or CO network. This sum was then divided by the sum of the percent damage values for all nodes in the network.

3.3.9 Graph Theory.

Voxel-by-voxel correlation matrices were derived for each node (ROI) as described above. These binarized 74x74 correlation matrices served as definitions for separate graphs over which graph theory metrics were assessed at a range of cost values. We did not apply graph theory techniques at the FP/CO network level because the relatively small number of nodes

(18) would produce unstable graph theory metrics. Graphs were defined by the number of nodes each connected by a number of edges to other nodes in the network. Cost values here represented the total number of edges between nodes in the graph and can be thought of as analogous to thresholding the correlation matrices to maintain an equivalent number of edges across groups. While the following computations were performed at a range of cost values (0.01 to 1), the main results are reported at cost = 0.1 which is within the range of 'ideal costs' (0.01 to 0.34) (Achard and Bullmore, 2007).

The graph metrics of interest here were the clustering coefficient (C) and the minimum path length (L) defined for each of the nodes in the network. C is a ratio that defines the proportion of possible connections that exist between the nearest neighbors of a node where high values of C imply that most of the neighbors are also neighbors of each other. The path length between any two nodes is defined as the minimum number of edges that must be traversed to form a direct connection between two nodes of interest. C and L were estimated for the whole graph by averaging over the nodal metrics within that graph. For comparison, random networks were constructed that contained the same number of nodes and edges. In a random graph, the average minimum path length is typically short and the average clustering coefficient is small. A well accepted characterization of a network that incorporates these properties is small-worldness or sigma. Sigma is the ratio of local connectedness (C, clustering coefficient) to global integration (L, path length) of a network; when this ratio is larger than 1 a network is said to demonstrate "small-worldness" (Watts and Strogatz, 1998). Lesion nodes were not included in this analysis and as such, any subjects without overlap to either the FP or CO network were not included, leaving 18 subjects in the analysis reported in **Figure 3.4B**.

3.3.10 Lesion simulation

To test the notion that the damage to a network extends beyond the damaged nodes themselves, we simulated lesions in our control group. For each patient, we randomly selected control subject to be 'matched' with that patient. In this matched simulated lesion patient, we replaced the 'lesioned' correlation values with those from the actual lesion patient. After simulating 21 patients, we calculated the group CO and FP network mean correlation values. A single simulation thus consisted of the correlation of the difference between the simulated CO and FP network means for each simulated lesion patient with the relative percent damage of the actual patient group. This correlation was recorded and the entire process repeated up to 20,000 times to acquire a stable probability of observing this correlation by chance.

3.4 Results

In this study, 21 patients with heterogeneous focal damage due to stroke or trauma (**Figure 3.1B**) underwent 10 minutes of rs-fMRI followed by high-resolution anatomical scans to enable accurate lesion mapping. All patient lesion masks can be seen superimposed on normalized T1 axial slices in **Figure 3.2**. Time-series correlations were

assessed amongst 18 pre-defined ROIs (Dosenbach et al., 2007) after standard preprocessing (Fox et al., 2005). The amount of damage sustained by each patient was quantified as the number of voxels in each ROI that overlapped with the lesion mask normalized to account for the differing network sizes. Some patients (**Figure 3.1C**) sustained significant damage to both networks, so we used a damage-difference score to represent the relative amount of damage to the networks. This relative measure allowed us to examine how the amount of damage affected the difference in connectivity between the two networks.

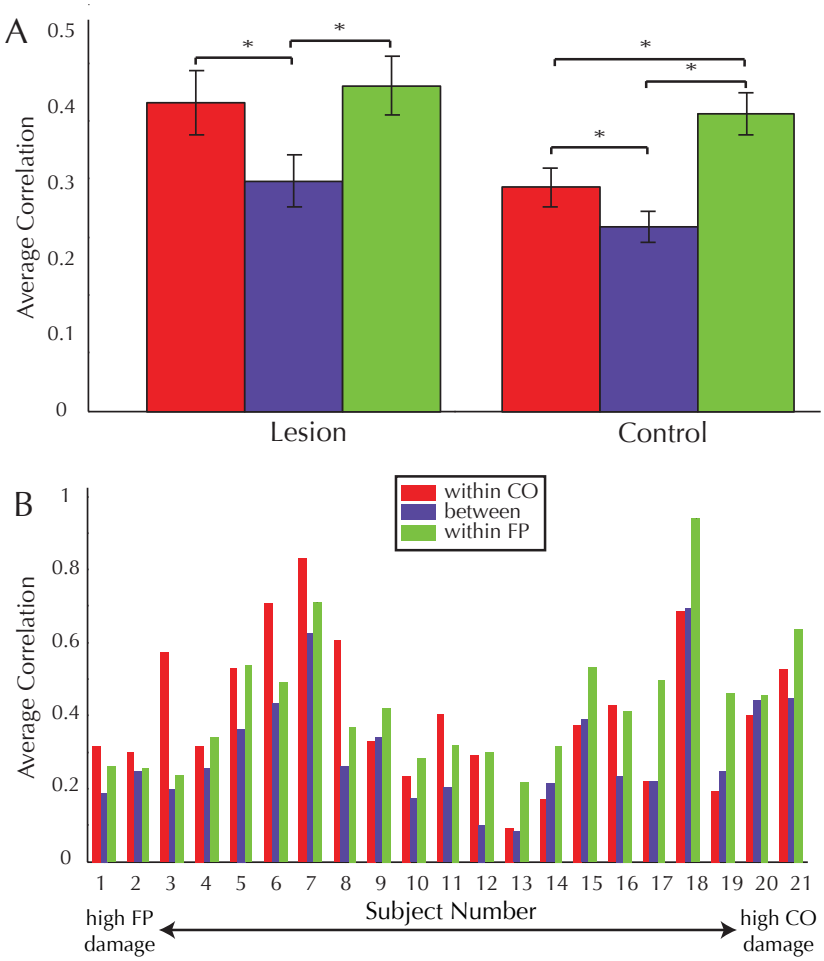

Figure 3.3: (A) Average FP and CO connectivity across subjects in the lesion and control group. (B) Average FP and CO network connectivity within each lesion patient sorted from high FP to high CO damage.

In addition to this anatomical measure of network integrity, we obtained a measure of functional connectivity within and between the CO and FP networks by averaging over the correlations amongst network nodes in each patient, a method that has been used elsewhere (Van Dijk et al., 2010). Paired t-tests (**Figure 3.3A**) revealed significantly higher average correlations within the CO network and within the FP network compared with average between-networks correlations, in the patient group [$t(20) = 3.55$; $P < 0.005$ and $t(20) = 5.94$; $P < 0.0001$, respectively] and in a group of age-matched control subjects [t

(20) = 3.4; P < 0.005 and t(20) = 9.25; P < 0.0001, respectively]. Individual patient mean correlation values are shown in **Figure 3.3B**. In previous studies of healthy populations (Dosenbach et al., 2007), the FP and CO networks exhibited higher within-network than between- network correlations. This is also true here with both lesion patients and age-matched control subjects, an observation that lends support to the dissociability of these networks in the resting state. Importantly, however, it is only by capitalizing on the extent of damage to these networks that we can address the question of whether the within-network connectivity of the FP and CO networks is independently affected by this selective damage.

To assess the extent to which the networks are independent, we correlated relative network damage with relative strength of functional connectivity within each network. Relatively more damage to a network was negatively correlated ($r = -0.64$, $P < 0.001$; **Figure 3.4A**) with relatively poorer connectivity within that network. When the size of the lesion was taken into account, this negative relationship remained significant ($r = -0.73$, $P < 0.001$). The method that we used of averaging correlation values introduces the potential confound of combining negative and positive values. However, performing the correlation analyses in **Figure 3.4A** with r^2 values did not change this relationship ($r = -0.44$, $P < 0.05$). A potential confound to our relative measure of functional connectivity is that the functional data from damaged nodes is allowed to contribute to the mean network correlation values. Because of the small number of nodes in both the CO and FP networks and the variability in the number of damaged nodes per subject, removing the damaged nodes entirely was not a viable option, as this would negatively affect the signal-to-noise ratio in the individual subject estimates. Instead, we used the percent node overlap (**Figure 3.1C**) to conduct a weighted-average, lesion-masking analysis, in which each node was weighted by the degree of lesion overlap. Correlation values from those nodes overlapping with the patients' lesions thus contributed to the network averages in proportion to our best estimate of their integrity. Repeating the correlation analysis in **Figure 3.4A** with this method yielded the same result obtained without weighted averaging ($r = -0.45$, $P < 0.05$). In addition, we performed a simulation analysis (SI Materials and Methods), which argues against the possibility that it is the damaged nodes themselves driving this negative relationship. In this simulation, the functional correlations of those nodes overlapping with regions of network damage were substituted into the correlation matrix of corresponding undamaged regions in healthy subjects. If it were the case that the low correlations from the damaged nodes were driving the negative relationship between damage and connectivity, this same relationship should have been observed in the simulated patients. However, this was not the case. Thus, we conclude that the decrease in network correlations in these patients must be due to the dysfunctional organization of the anatomically intact nodes within the damaged network.

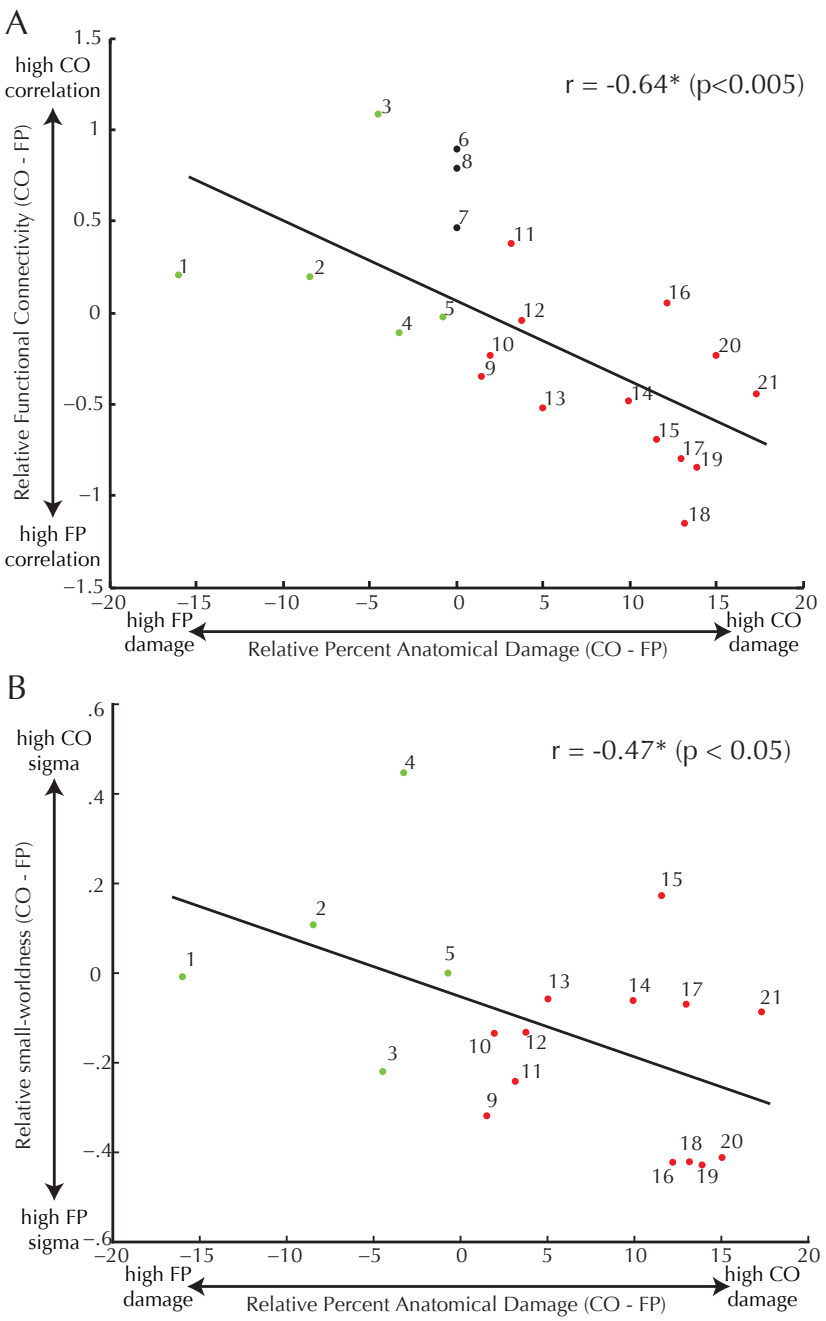

Figure 3.4: A) Relative functional connectivity of CO and FP networks versus relative CO and FP anatomical damage. (B) Relative small-worldness (sigma) of CO and FP network nodes versus relative CO and FP anatomical damage.

The correlation between lesion location and functional connectivity used relative differences and as such, could be due to increased connectivity in the undamaged network rather than decreased connectivity in the damaged network. However, examining connectivity as a function of absolute damage in each network separately rejects this claim (**Figure 3.5**). Across patients, CO damage is negatively correlated with CO (r = −0.8, P <

0.0001) and unrelated to FP connectivity (r = 0.03, P < 0.90), and FP damage is negatively correlated with FP (r = −0.44, P < 0.05) and unrelated to CO connectivity (r = −0.19, P < 0.41). Partial correlations taking into account the amount of damage to the opposite network strengthens these relationships. For CO damage, controlling for FP damage, the partial correlation with the CO network remained significant (partial r = −0.83, P < 0.0001) and remained non- significant for the FP network (partial r = 0.2, P = 0.388). For FP damage, controlling for CO damage did not change the significance of the relationship with the FP network (partial r = −0.47, P < 0.05) and the correlation with the CO network remained nonsignificant (partial r = −0.14, P = 0.559). Together, these data support our assertion that increasing damage to one network is correlated with increasingly disrupted connectivity within that network, but does not affect the other network or the connectivity between networks.

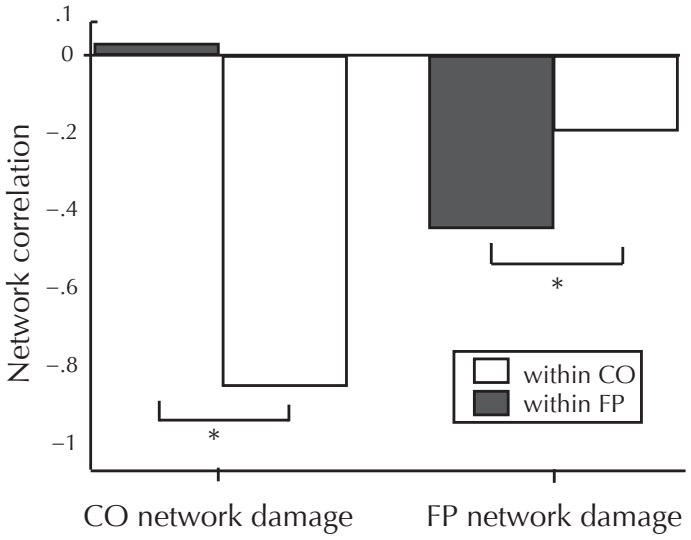

Figure 3.5: Mean correlation values within the CO or FP networks in either patients with relatively more CO network damage or patients with relatively more FP network damage. *Significant differences between within-FP and within-CO network correlations (P < 0.05).

In a parallel analysis to the ROI-to-ROI functional correlations shown in **Figure 3.4A**, we calculated within-ROI voxel-by-voxel time-series correlations to capture the local changes to FP and CO network nodes. For each ROI, we assessed graph theory metrics (clustering and path length) over these correlation matrices as well as to equivalent sized random-networks yielding a metric of small-worldness (sigma). The average sigma values across network nodes showed a similar relationship to the ROI-to-ROI functional correlations. That is, the average small-worldness of the small-scale networks contained in each individual node in the large-scale cognitive control network was decreased only within the damaged network, leaving the undamaged network relatively intact. This relationship was significant (cost = 0.1, r = -0.47, p < 0.05) within a range of cost values, but a single cost value and the best fitting regression line is shown in **Figure 3.4B**.

In summary, results from both the functional correlations across network nodes and graph theory measures within nodes indicate that anatomical damage to one cognitive control network has a specific detrimental effect on the remaining undamaged nodes in that network but no effect on the other network.

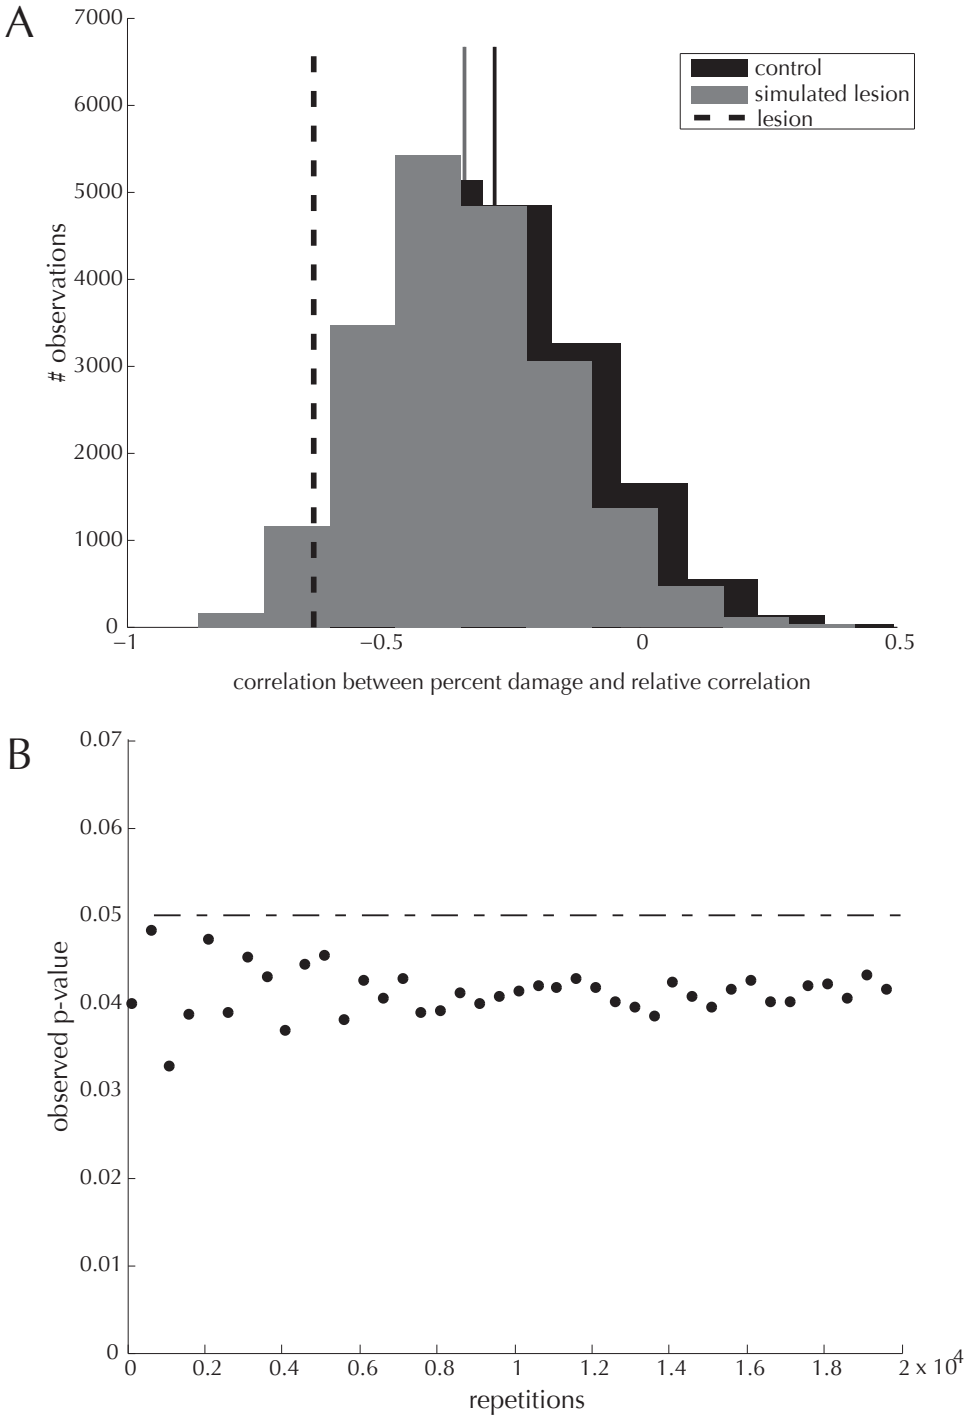

Figure 3.6: Simulated lesions. (A) Control and simulated lesion distributions of correlation values. (B) Repeated simulations show stability of this pattern of distributions where the simulated lesion group does not overlap significantly with the true lesion group. Dashed line is p = 0.05.

Simulating brain lesions in healthy control subjects supports this claim (**Figure 3.6**). We constructed a group of simulated patients by randomly selecting a control subject and inserting the 'lesioned' correlation values from each patient into the simulated patient's correlation matrix. Within this simulated patient group, we performed the same correlation analysis between relative damage and relative functional connectivity to yield an r-value. This process was repeated to create histograms of the distribution of r-values for both the simulated lesion patients and control patients (**Figure 3.6A**). While both distributions are centered over a negative r-value, the simulated patient distribution is almost entirely outside the range of the observed lesion patient correlation (black line in **Figure 3.6A**, r = -0.63). The difference between the simulated lesion and actual lesion patient distributions is quantified in **Figure 3.6B** where the observed number of r-values less than -0.64 is scaled by the number of repetitions. This p-value is quite stable for the simulated lesion patients (p ~-= 0.04) across a range of repetitions. That the probability of observing this pattern by chance is less than 5% suggests that a simulated lesion in healthy subjects is not sufficient to produce the same pattern in relative FP and CO network mean correlations as in the real patient group.

3.5 Discussion

Whether in a controlled experimental environment or in a real life situation, humans are adept at maintaining a task-relevant goal while simultaneously making moment-to-moment adjustments in behavior to achieve that goal. Such cognitive control, we and others (Badre et al., 2009; Dosenbach et al., 2008) argue, is supported by multiple independent neural systems. Here we assess network structure using multi-nodal network functional correlations and within-node graph theory metrics to provide converging evidence in support of the putative independence of these two cognitive control networks (Dosenbach et al., 2007). We have demonstrated that damage to two networks thought to support different components of cognitive control alters connectivity within the damaged network, but leaves the other network preserved. Within the damaged network, it is not only the portion of the network subjected to anatomical damage that exhibits dysfunction; rather, as we have demonstrated by within-node graph theory measurements, it is also the remote, interconnected regions that seem to be affected by this damage.

Our findings are consistent with the concept of diaschisis, which was originally formulated to describe temporary clinical deficits related to areas remote from the area of damage (Von Monakow, 1914). The concept of diaschisis has since been expanded to include neurophysiological observations of depression of activity in remote, undamaged brain sites that are functionally connected to lesion areas. The prevailing hypothesis for the underlying mechanism responsible for this phenomenon is the withdrawal of excitatory influences from anatomically connected brain regions (Finger et al., 2004), but evidence for

this mechanism from neuroimaging data in stroke patients has been mixed. Some investigators have demonstrated significant reductions in resting activity in structurally intact regions homologous to the site of damage using positron emission tomography (Fair et al., 2009) and magnetoencephalography (Guggisberg et al., 2008), as well as decreases in EEG coherence throughout the lesion hemisphere (Harmony et al., 1994) and across the entire brain (Bartolomei et al., 2006a). One study, however, found that whereas resting PET activity showed significant decreases at a distance from the lesioned site, task-evoked BOLD signal was minimally affected (Fair et al., 2009).

Although the aforementioned studies are largely consistent in their observation that functional disruptions extend beyond the region of anatomical damage, none have examined the specific effects of damage on functional networks. Perhaps the disparate results from these studies can be explained by our hypothesis that damage remains localized to the particular network affected. This has been examined in an fMRI study of patients with right hemisphere strokes who exhibited spatial neglect (He et al., 2007). In this study, functional connectivity of nodes in two attention networks was assessed at both the acute and chronic stage of recovery. Functional correlations between pairs of network nodes in both the anatomically damaged ventral attention system and intact dorsal attention system were decreased at the acute stage compared with age-matched control subjects. Over time, the dorsal system recovered its functional connectivity, but the ventral system remained impaired. The investigators posited that although only one system had sustained physical damage, because the ventral and dorsal attention systems normally interact, anatomical damage to the ventral system carried over into the dorsal system and manifested as a disruption in functional connectivity. That is, the dorsal system normally relies on the operation of the ventral system to maintain some aspects of spatial attention, so anatomical damage to the ventral system has functional consequences for the distal parts of the network even though they are remote from the damage.

Another recent study from the same group (Carter et al., 2010) examined the behavioral significance of intra- versus interhemispheric interactions following ischemic stroke. Subjects' levels of behavioral impairment were correlated with the strength of functional connectivity both within the damaged hemisphere and across to the intact hemisphere nodes. Interhemispheric interactions most strongly predicted the level of impairment in behavior, adding to the collection of evidence that anatomical damage produces network dysfunction in nonlesioned regions of the same network. However, in the current study, rather than assessing the damage to a single functional network, we used patients with damage to nodes throughout two different, predefined cognitive control networks. This disparate pattern of damage allowed comparisons of network correlations within subject rather than with a control group and, more importantly, has allowed us to assess the degree of independence of these two multinodal networks. Damage to the FP network degraded the functional connectivity of nodes only within the FP network, and damage to the CO network impaired connectivity only within CO nodes. Applying graph theory to individual nodes in these two networks provides a potential explanation for this decrement in connectivity. The small-scale organization of these brain regions only within the damaged network exhibited disruptions in small-worldness compared with the organization of nodes within the nondamaged network.

Other investigators have used graph theory to examine the effects of brain lesions on the organization of large-scale brain networks; however, to our knowledge, none of these investigators have had the advantage of using real data from patients with brain lesions. In a recent study, Alstott et al. (2009) simulated the effects of brain lesions on network properties derived using graph theory and found that the degree of disruption in network organization depends on the location of the lesion. When nodes were classified based on properties pertaining to their role in the network as a whole, removing nodes that were considered "hubs" was most detrimental. This approach treats the entire brain as a single network without considering the effect of lesions on individual subnetworks. The difficulty with using graph theory on smaller subnetworks is that many of the metrics are sensitive to the number of nodes in a graph, necessitating networks that are larger that the number of nodes in each of the networks examined here. Using graph theory at a local level as we do in the current study is one way of addressing the question of the organization of subnetworks with and without a lesion.

The observed double dissociation between the FP and CO networks using both multinodal functional correlations and within- node graph theory suggests a parallel organization whereby the function of one network may not be necessary for the other network to function adequately. Consistent with clinical observations (D'Esposito et al., 2006), the existence of independent, distributed networks supporting cognitive control may render the brain less susceptible to persistent behavioral deficits as a result of focal cortical damage. However, the finding that these networks are distinct functional units does not preclude the possibility that they flexibly interact to carry out complex cognitive control (Fan et al., 2009). That is, in a healthy individual with functionally and structurally intact cognitive control networks, either system or both could contribute toward accomplishing the goal at hand. Understanding how these parallel systems coordinate their activities to carry out complex cognitive control represents an important next step. Moreover, given the important role of these networks in a wide variety of cognitive processes, it will also be critical to relate the strength of functional connectivity to behavior in both intact and damaged networks.

Chapter 4

THE EFFECT OF THETA-BURST TMS ON COGNITIVE CONTROL NETWORKS MEASURED WITH RESTING STATE FMRI

4.1 Abstract

It has been proposed that two relatively independent cognitive control networks exist in the brain: the cingulo-opercular network (CO) and the fronto-parietal network (FP). Past work has shown that chronic brain lesions affect these networks independently. It remains unclear, however, how these two networks are affected by acute brain disruptions. To examine this, we conducted a within-subject theta-burst transcranial magnetic stimulation (TBS) experiment in healthy individuals that targeted left anterior insula/frontal operculum (L aI/fO, a region in the CO network), left dorsolateral prefrontal cortex (L dlPFC, a region in the FP network), or left primary somatosensory cortex (L S1, an experimental control region). Functional connectivity was measured in resting state fMRI scans collected before and after TBS on each day. Unlike with chronic lesions, we found that TBS was accompanied by non-selective increases in network connectivity, especially in FP network connectivity after TBS to both regions involved in cognitive control. Whole-brain connectivity results demonstrated that the L dlPFC and L aI/fO showed increased connectivity with regions in frontal, parietal, and cingulate cortex after TBS to L dlPFC and L aI/fO, but not to L S1. These results suggest that acute disruption by TBS causes more widespread and less selective changes in network connectivity than those seen after chronic brain lesions.

4.2 Introduction

Cognitive control is an essential ability that allows humans to interact with their environment and carry out complex goal-directed behaviors. In the past, it has been suggested that a distributed set of regions may contribute to cognitive control processes (Badre and D'Esposito, 2007; Dosenbach et al., 2006; Duncan and Owen, 2000; Koechlin et al., 2003), but it is not clear to what extent these systems operate independently in the brain (see review by (Dosenbach et al., 2008)). Recent work from Dosenbach and colleagues used activation data from a variety of different tasks (Dosenbach et al., 2006) combined with resting-state functional connectivity (fc) measurements (Dosenbach et al., 2007) to suggest that two separable cognitive control networks exist in the human brain (Dosenbach et al., 2008): the fronto-parietal (or FP) network and the cingulo-opercular (or CO) network (see **Figure 4.1**). The FP network is composed of dorsolateral frontal and parietal regions, the precuneus, and the mid-cingulate and is proposed to operate at fast time-scales for moment-to-moment adaptive task control. The CO network is composed of regions in the anterior insula/frontal operculum, anterior prefrontal cortex, dorsal anterior

cingulate, and thalamus, and is proposed to operate over longer time-scales for stable task maintenance. In a large population of subjects, these networks clustered separately from one another during rest, suggesting that they may operate in a somewhat segregated fashion (Dosenbach et al., 2007).

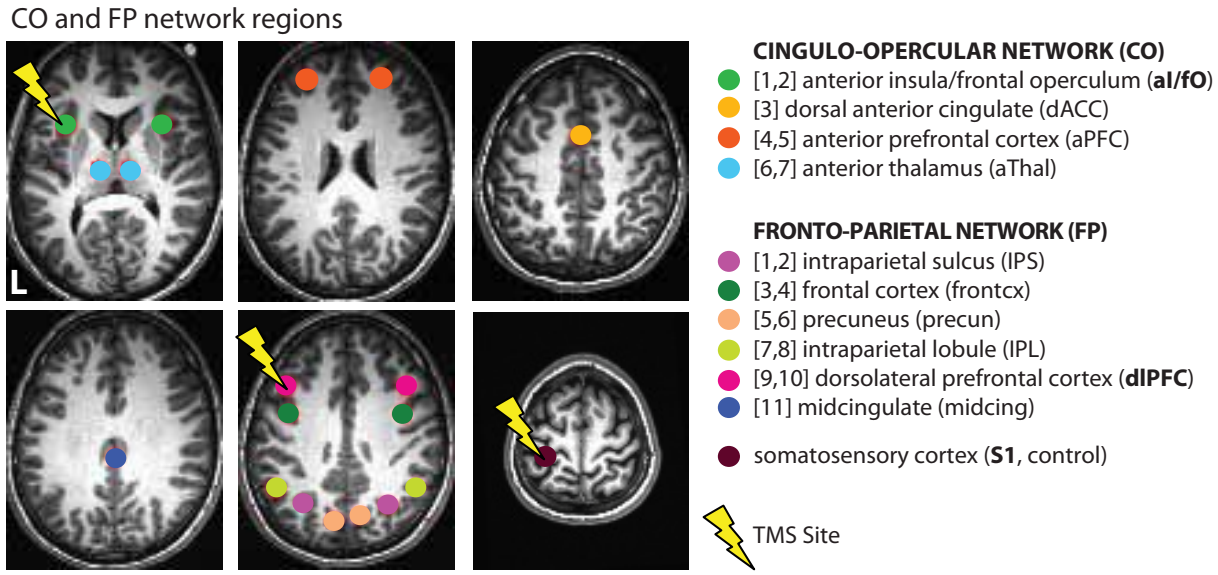

Figure 4.1: *Cognitive Control Networks.* Connectivity was examined across and within regions from the CO and FP cognitive control networks, as defined in Dosenbach et al. (2007). These measurements were taken before and after TBS to L dlPFC (a region in the FP network), L aI/fO (a region in the CO network), and L S1 (an experimental control location). Lightning bolts indicate the sites of stimulation.

More recent data from our laboratory has extended these findings by examining how these networks functionally change following chronic brain damage (Nomura et al., 2010). By examining resting state data from a group of patients who had lesions to these two networks, we demonstrated that damage to one cognitive control network caused changes in functional network interactions that extended throughout the network, beyond the site of structural damage. However, these changes were specific to the impacted network, leaving the other cognitive control network relatively intact. This finding provided strong evidence that these cognitive control networks exhibit independence in resting state data.

However, the patients in the Nomura et al (2010) study were all scanned many months (>5 m, on avg. 8 yrs) since the incident that caused their lesion, allowing ample time for long-term structural and functional reconfiguration. One previous study examining recovery after stroke to the ventral attention network found that strokes initially impacted widespread regions, causing impairments to both the ventral and the dorsal cortical attention networks. However, later during the chronic recovery period, recovery became more selectively limited to the ventral attention network (He et al., 2007). Therefore, it remains an open question how acute disruptions, such as those induced by theta-burst transcranial magnetic stimulation (TBS), will affect the organization of the CO and FP

networks. In particular, this may shed light on how these networks interact on a shorter time-scale, such as when both are activated in the service of a cognitive control task, and whether their distinct response to chronic damage represents long-term plasticity or is a signature of disruption or rapid functional reorganization that is present immediately after acute interventions.

In this study, we used TBS combined with resting state functional connectivity measurements to examine network changes following acute disruption of specific brain areas. In separate sessions, TBS was applied to the left dorsolateral prefrontal cortex (L dlPFC, a region in the FP network), to the left anterior insula/frontal operculum (L aI/fO, a region in the CO network), or to the left primary somatosensory cortex (L S1, an experimental control region that does not participate in either network). Resting state functional connectivity was measured before and after TBS for each session. If these networks react independently to acute disruption, as they do to chronic focal brain lesions, then we would expect that TBS to the L aI/fO region will only impact interactions within the CO network and TBS to the L dlPFC region will only impact interactions within the FP network, while TBS to the S1 region will not impact either network. However, it is also possible that the response to acute disruptions will show less independence and reflect more widespread changes that are not limited to the damaged region's network.

4.3 Materials and Methods

4.3.1 Participants:

Twenty-seven healthy right-handed participants completed four separate sessions with fMRI and TMS components (11 female, 18-31, mean age = 22.7 years). Participants were not on any antidepressant medications and had no history of epilepsy, anxiety, or neurological disorders. Participants provided written informed consent before the study and were compensated monetarily. The study procedures were approved by the Committee for the Protection of Human Subjects at the University of California, Berkeley.

4.3.2 Experimental Timeline:

Each participant completed four combined TMS and fMRI sessions on separate days. During the first day, participants took part in a scanning session during which a structural image, a resting state Echo Planar Imaging (EPI) functional scan, and a perfusion Arterial Spin Labeling (ASL) functional scan were acquired. In addition, resting motor thresholds (see *Transcranial Magnetic Stimulation*) were measured for each participant during this session.

On the following three days, participants began by completing a short EPI and ASL scan. After scanning, they received 40 seconds of TBS stimulation to one of three sites: L dlPFC, L aI/fO, or L S1 (session order was counterbalanced across participants). Participants then re-entered the scanner and completed two more sets of interleaved EPI and ASL scans. Approximately 10 minutes elapsed between TBS stimulation and the start of scanning. During scanning, a black fixation cross was presented on a white screen. Participants were

instructed to stay still, awake, and with their eyes open; they were told they should allow their mind to wander while trying not to focus on anything in particular. Wakefulness was monitored during all scans using an MR-compatible camera to image the eye. Each session took approximately 1 hour to complete. In a few participants, the second ASL scan was not included in the analysis due to participants either falling asleep (N=2) or the session extending beyond the scanning time limit (N=2).

4.3.3 Transcranial Magnetic Stimulation (TMS):

All transcranial magnetic stimulation occurred offline, outside of the scanner. Throughout each stimulation session, participants were seated in a comfortable chair with ear protection. Stimulation was applied using a Magstim figure-of-eight coil (outer winding diameter = 70mm; *Magstim, UK*). Pulses were delivered with a Magstim rapid stimulator with four booster modules producing biphasic pulses.

On the first day, motor thresholds were calculated using electromyography measured from electrodes on the dominant hand's first interosseous (FDI) muscle. Stimulation was applied over the hand area of the contralateral primary motor cortex, with the coil placed tangentially to the scalp and the handle pointing posteriorly. This area was identified as the area that produced the largest motor evoked potential (MEP) from the FDI when it was relaxed. Each participant's active motor threshold (AMT) was then computed by finding the minimum stimulation intensity at which 5 out of 10 trials produced an MEP in the contralateral FDI while the muscle was contracted at 20% of maximum (electromyographic signal feedback was provided on a monitor to help participant maintain this level of contraction).

On the other three days, TBS was applied to pre-defined stimulation targets (see *Stimulation Sites*, below). These targets were localized using a computerized frameless stereotaxic system (*Brainsight software, Rogue Research, Canada*), in which head position is monitored in real time with an infrared camera by using reflective markers placed on the participant's head and the TMS coil. The position of these markers was co-registered to locations on the individual participant's MRI scan that was acquired during the first day. Once the coil was placed over the stimulation target, 40 seconds of TBS was applied continuously. TBS consists of 50 Hz triplets every 200 msec. This form of TBS has been shown to produce reduced excitability of underlying cortex lasting up to 60 minutes following stimulation (Huang et al., 2005), allowing for sufficient fMRI recording time.

Stimulation Sites: Three separate sites were targeted in the three different TBS days: the left dorsolateral prefrontal cortex (L dlPFC), left anterior insula/frontal operculum (L aI/fO), and left primary somatosensory cortex (L S1; sites shown on **Figure 4.1**). These regions were defined offline in each individual. The experimental control site, L S1, was defined for each individual as a region on the post-central gyrus 10 mm away from the midline and 5 mm from the top edge of the brain. The other two sites were defined using a procedure that combines individualized anatomical and functional criteria (based on the method from (Eldaief et al., 2011)). This was done in an effort to optimize the stimulation site for each cognitive control network. Initially, a site was selected in each participant based on coordinates previously published by Dosenbach and colleagues (2007). These

coordinates, along with those from the rest of the regions in the CO and FP networks (shown in **Figure 4.1**), were reverse-normalized into native space. For each participant, a CO and FP network map was then created by conducting a whole brain functional connectivity analysis using the resting state data from the first baseline day. The time-series was averaged over all regions in the relevant network other than the target site. This averaged time-series was then correlated with every voxel in the brain (see *Functional Connectivity* for a full description of connectivity methods). For example, a CO whole brain network map would be created by averaging the time-series from all regions in the CO network other than the L aI/fO region and then computing the correlation between this average time-series and every voxel in the brain. Then, within each network map, we found the cluster of connectivity that was closest to the original reverse-normalized spatial coordinates. The peak voxel from this cluster was selected as the stimulation site and used to create a 6mm spherical region of interest (ROI). If no nearby cluster could be identified (as was the case for 4 subjects for L aI/fO and 2 subjects for L dlPFC), then the original reverse-normalized coordinates were used as the center of the ROI.

4.3.4 MRI acquisition parameters:

All MRI data were acquired from a 3T Siemens Magnetom Trio scanner using a 12-channel head coil. Structural images were acquired during the first experimental session using an MP FLASH T1-weighted scan. Echo planar imaging (EPI) scans were used to measure whole brain Blood Oxygen Level-Dependent (BOLD) signal for resting state functional connectivity analysis. EPI scans were obtained using a T2*-weighted EPI sequence (TR = 2 s, TE = 24 ms, flip angle = 60 degrees, in plane matrix = 64 x 64 pixels, each 3.5 x 3.5 mm, with 37 3.5 mm descending axial slices and a 0.7 mm slice gap). Each scan lasted approximately 6 minutes, during which 180 volumes were acquired. In between consecutive EPI scans, 5.5 min Arterial Spin Labeling scans were collected to measure cerebral perfusion and are not reported in this paper.

4.3.5 Data Preprocessing:

Preprocessing was carried out in AFNI [*http://afni.nimh.nih.gov/afni/*, (Cox, 1996)]. Images were extracted from DICOMs, slice-time corrected, realigned to the first time point of each session, co-registered to the anatomical image, and smoothed (6mm FWHM). In addition, the transformation between normal and native space was computed for the anatomical image using SPM segment (Ashburner and Friston, 2005). This was used to reverse normalize published coordinates into native space.

4.3.6 Functional connectivity:

Functional connectivity was computed using time-series correlations following the procedure described in Fox et al (2005). Voxel time-series were bandpass filtered (0.009 – 0.08 Hz). Nuisance signals from white matter, ventricles, and motion along with their temporal derivatives were regressed out (note, however, that the brain's global signal was not included as a regressor due to concerns about biases it could introduce into the correlation structure, e.g., (Saad et al., 2012)). These regressors were included in order to reduce noise in the signal from physiological, motion, and scanner artifacts. Afterwards, the time-series were averaged across all voxels within an ROI (see *Regions of interest* for definitions), and these averaged time-series were used as seeds for functional connectivity

analyses. In the ROI-to-ROI connectivity analyses, Pearson's time-series correlations were calculated among the averaged time-series from CO and FP network nodes. In the ROI-to-whole brain analyses, Pearson's time-series correlations were calculated between the average time-series of the seeded ROI (L dlPFC and L aI/fO) with each voxel in the rest of the brain. All correlation values were Fisher-transformed before statistical comparison.

4.3.7 Regions of interest:

For the three stimulation sites, regions of interest were defined based on a combination of structural and functional criteria as described under *Stimulation Sites*. In addition, the other ROIs from the CO and FP networks were defined as a 6mm sphere centered on the reverse normalized coordinates from Dosenbach et al (2007) (see **Figure 4.1** for all regions).

4.3.8 Statistical analysis and data visualization:

ROI-to-ROI analyses within pre-defined networks were compared using ANOVAs on average Fisher-transformed connectivity values within each network. A 3-factor ANOVA was used to examine changes induced by TBS across the entire CO or FP network, with factors of block (pre-TBS, post-TBS 1, and post-TBS 2), network (CO, FP, and between CO-FP), and TBS-site (aI/fO, dlPFC, and S1). When necessary, p-values accompanying the ANOVA were corrected for sphericity using the Huynh-Feldt method (indicated as [HF] in the text). Significant effects and interactions were further explored with post-hoc t-tests. In some cases, a separate ANOVA was conducted for post-TBS block 2, where effects were expected to be maximal (Hubl et al., 2008). In addition, network connectivity to the target nodes alone was separately examined using a 4-factor ANOVA with factors of block (pre-TBS, post-TBS 1, post-TBS 2), TBS-site (aI/fO, dlPFC, S1), connectivity seed (aI/fO, dlPFC), and target (CO, FP). All reported t-tests were two-tailed. ANOVAs were carried out using the SPSS package (*Version 20, Chicago, IL*). Network graphs were created using the NetworkX package in python [*http://networkx.github.io/*]. All other graphs and statistical analyses were created using python's matplotlib [*http://matplotlib.org/*] and scipy [*http://www.scipy.org/*] packages. Voxel-to-whole brain analyses were corrected for multiple comparisons using AFNI's 3dClustSim, which uses Monte Carlo simulations on $p<0.001$ uncorrected maps to estimate the minimum cluster size that would be significant (i.e., such that $p(\text{corrected})<0.05$), based on the size of the whole-brain volume and the amount of smoothness in the data.

4.4 Results

4.4.1 Identification of CO and FP networks within individuals before TBS

Using the resting-state fMRI data collected before TBS, the CO and FP networks were dissociable in each individual subject (see **Figure 4.2**). ROI-to-ROI connectivity analyses demonstrated that within network connectivity was significantly higher than between network connectivity for both the CO network [$t(26) = 15.16$, $p<10^{-13}$] and the FP network [$t(26) = 13.53$, $p<10^{-12}$] and that this pattern was evident across individual subjects

(**Figure 4.2A and B**). Furthermore, seeded whole brain analyses also produced a similar result: seeding either L aI/fO or L dlPFC (the regions that were subsequently stimulated) produced a map of regions that recapitulated, respectively, the CO and FP networks (see **Figure 4.2C**, displayed at ([t(26) > 10.29, p[uncorr]<10^{-10}, p[FDR corrected]<10^{-5}]). These results suggest that, prior to stimulation, the targeted regions were integrated within functionally correlated and dissociable networks.

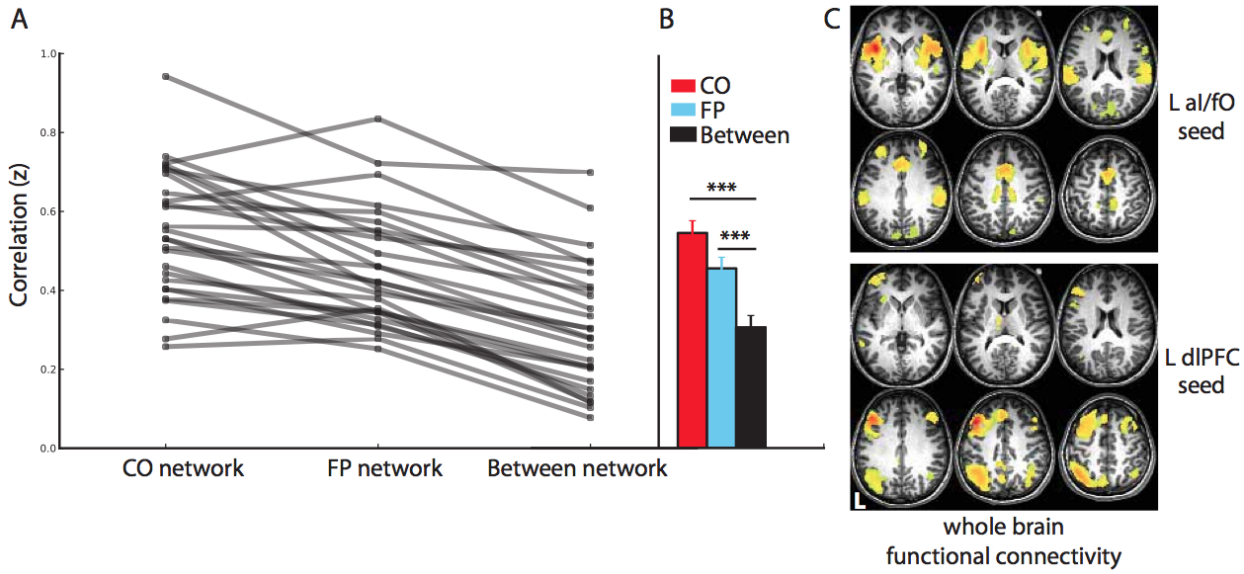

Figure 4.2: *Connectivity in the CO and FP networks before TBS*. (A) Connectivity within the CO and FP networks and across the two networks are shown for each individual (gray lines). Across all individuals, within network connectivity was higher than between network connectivity. (B) Statistics across the group of subjects confirm these findings, with connectivity within the CO network (red) and within the FP network (blue) significantly higher than connectivity across the two networks (black; ****, p<0.001*). (C) Whole brain functional connectivity analyses seeded in the locations targeted for TBS recapitulated the CO and FP networks: CO network regions were functionally connected with the L aI/fO region (*top*), and FP network regions were functionally connected with the L dlPFC region (*bottom; in this and subsequent figures L marks the left hemisphere of the brain*).

4.4.2 The effect of TBS on CO and FP networks

The influence of TBS on the average connectivity of each network was measured with a 3-way ANOVA, with factors of block (pre-TBS, post-TBS 1, and post-TBS 2), network (CO, FP, and CO-FP), and TBS site (aI/fO, dlPFC, and S1), after averaging over all pairwise connections within the CO network, within the FP network, or across the two networks (see **Figure 4.3**). This ANOVA revealed a trend toward an interaction between block, network, and TBS site [F(8,208) = 2.115, p[HF]=0.064] suggesting that TBS may have had weak selective effects on different ROIs. Focusing on block 2 (which occurred approximately 20 min after TBS, when past findings would suggest stronger effects on fMRI signals (Hubl et al., 2008)), we also found a trend for a significant three-way interaction [F(4,109) = 2.47, p[HF]=0.074]. These results seem to be mainly driven by increases in FP network functional connectivity after TBS to either cognitive control network site [aI/fO

TBS: t(26) = 3.45, p<0.01; dlPFC TBS: t(26) = 3.57, p<0.01]. However, TBS to either of these sites did not a significantly greater increase in FP network connectivity than TBS to S1 [aI/fO TBS vs S1 TBS: t(26) = 1.65, p=0.11; dlPFC TBS vs S1 TBS: t(26) = 1.24, p=0.23]. In addition, TBS to each of the three sites caused increases in connectivity between the CO and FP networks [aI/fO TBS: t(26) = 2.11, p<0.05; dlPFC TBS: t(26) = 2.97, p<0.01; S1 TBS: t(26) = 1.78, p=0.084].

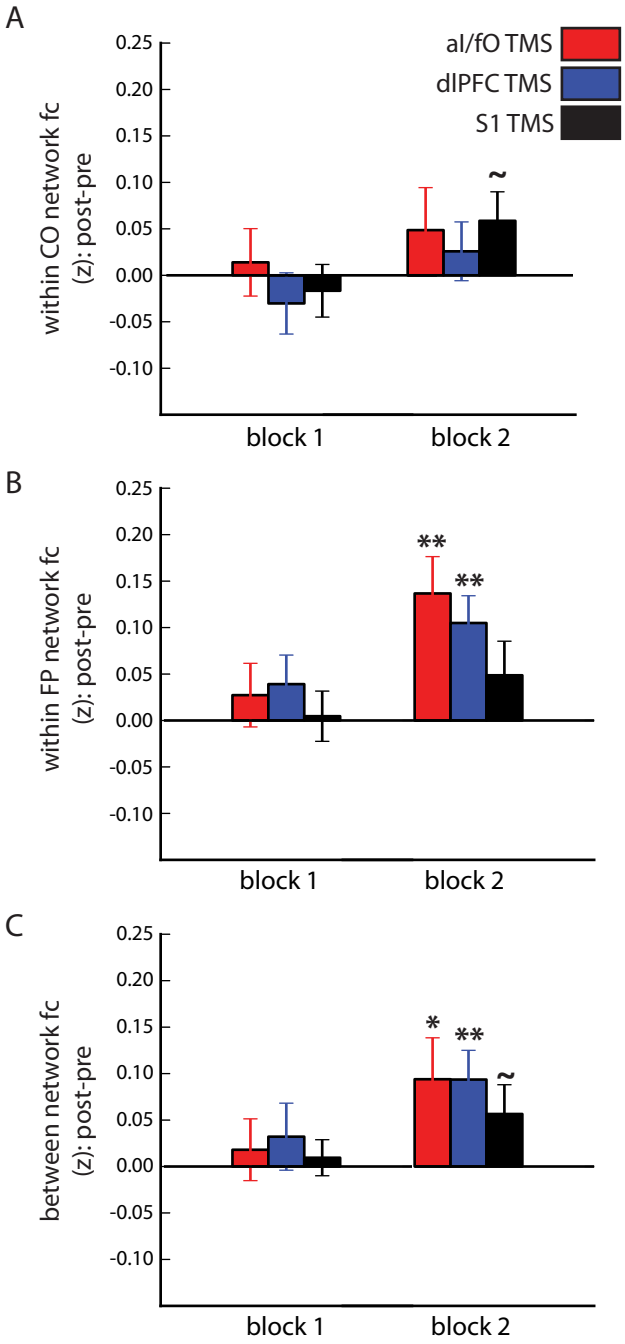

Figure 4.3: *Connectivity in the CO and FP networks after TBS*. Connectivity after TBS is shown within the CO network (A), FP network (B) and between the two networks (C). Different sites of TBS (red = aI/fO, blue =

dlPFC, black = S1) did not have systematically different effects on within- and between-network connectivity, although connectivity generally tended to increase after TBS (Post-TBS vs Pre-TBS: *** p<0.01, * p<0.05, ~ p<0.10*).

Network connectivity measures increased significantly over time, regardless of the site of TBS [main effect of block: $F(2,52) = 8.60$, $p<0.001$]. There was also an interaction between block and network [$F(4,104) = 3.251$, $p[HF]<0.03$], which was driven by larger increases in network connectivity after TBS within the FP network and between the FP and CO network, compared with the connectivity within the CO network [average change in functional connectivity after TBS, FP vs. CO: $t(26) = 2.45$, $p<0.03$; FP-CO vs CO: $t(26) = 2.65$, $p<0.02$]. Finally, a main effect of network connectivity magnitude [$F(2,52) = 128.2$, $p<10^{-20}$] was found, recapitulating the persistence of stronger within- than between-network connectivity across the entire fMRI session [within CO vs. between network connectivity: $t(26)=17.62$, $p<10^{-15}$; within FP vs. between network connectivity: $t(26)=13.01$, $p<10^{-12}$].

Figure 4.4 plots the changes for every ROI to ROI connection within or across the CO and FP networks in post-TBS block 2. The changes summarized in **Figure 4.3** are not specific to any single connection in these networks, but rather seem to be generally present across all network connections. In particular, for analyses restricted to examining connectivity of only the stimulated ROI to other ROIs in the CO and FP networks, rather than summarizing connectivity across the entire network, the pattern is in the same direction as the network-wide effects but did not reach significance [ANOVA with factors of block (pre, post1, or post2), TBS site (aI/fO, dlPFC, or S1), seed (aI/fO or dlPFC), and network (CO, FP), interaction between block x TBS site x seed x network: $F(4,104) = 1.45$, $p[HF] = 0.23$]. It therefore appears as though focal acute disruptions from TBS are causing widespread, non-specific effects, with generalized increases in connectivity, particularly in connectivity within the FP network after TBS to either cognitive control network region.

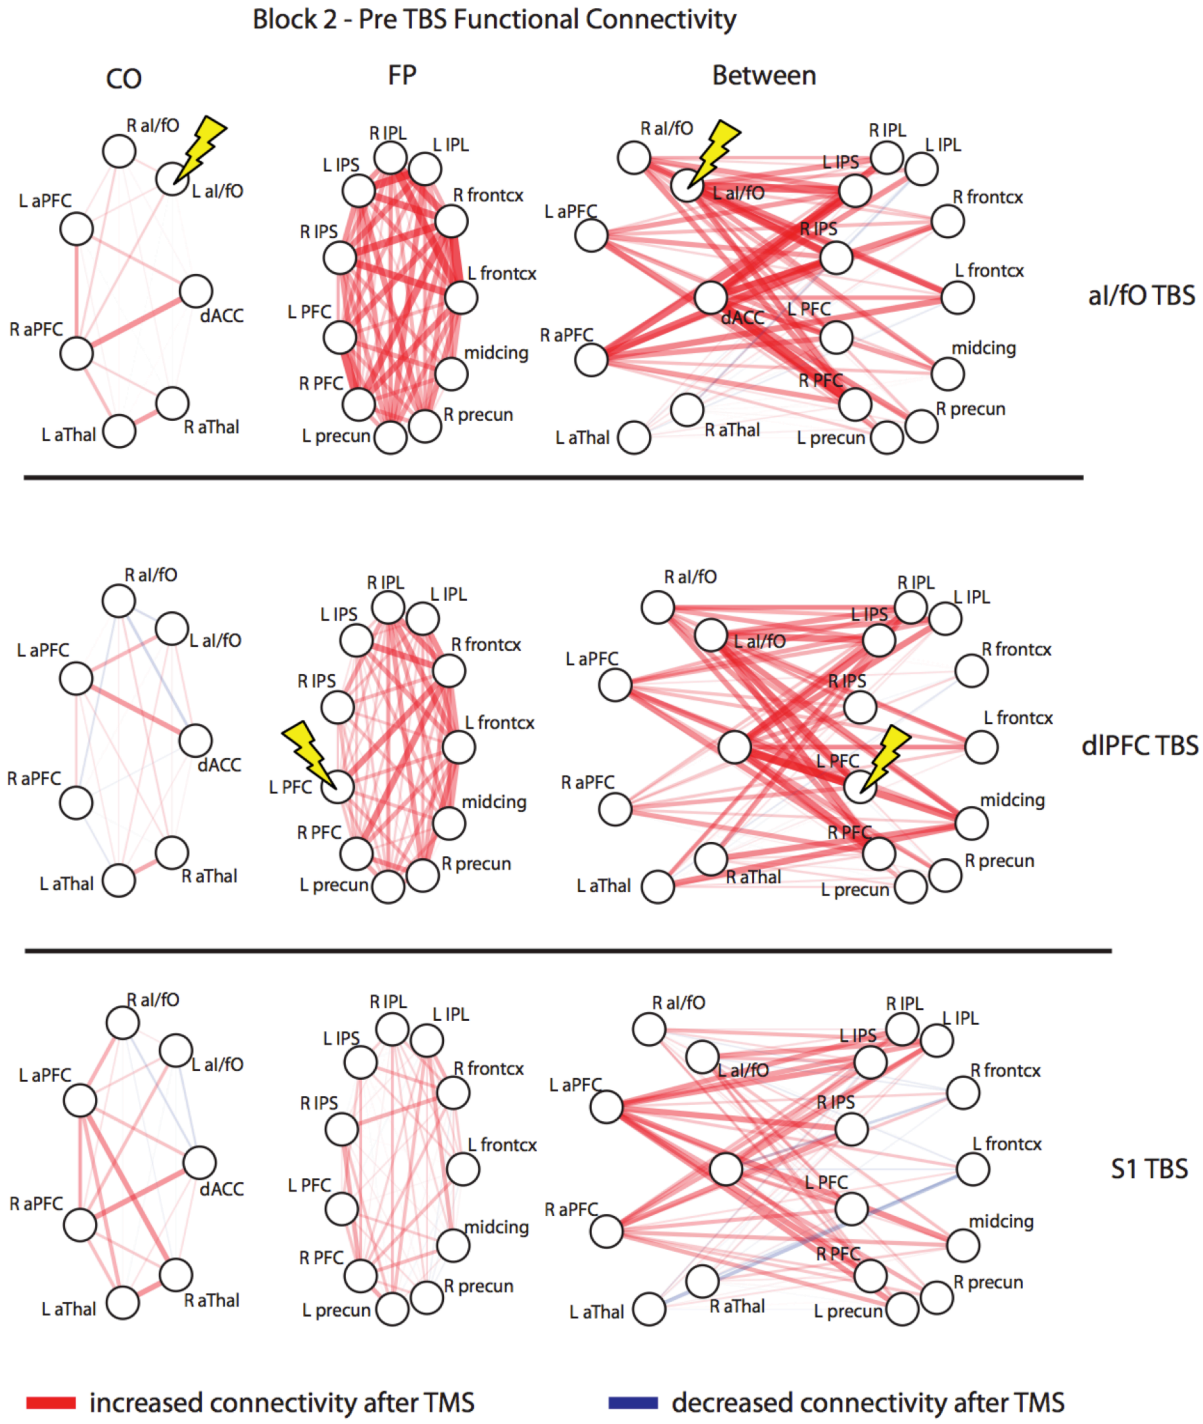

Figure 4.4: *Network interactions after TBS*. Connectivity changes in the second block after TBS are shown for each individual connection with the CO network (*left column*), the FP network (*middle column*), and between the two networks (*right column*) after TBS to the L aI/fO (*top row*), L dlPFC (*middle row*), and L S1 (*bottom row*). As can be seen, TBS generally increased connectivity across all nodes of the network, with no unique effects seen at the stimulation site. Edge color represents the sign of change after TBS (red = increases, blue =

decreases) and the line thickness and opacity indicates the magnitude of change. Lightning bolts highlight the region targeted in each row.

4.4.3 The effects of TBS across the whole brain

In a final exploratory analysis we examined the changes in connectivity of the L aI/fO and L dlPFC seed to the whole brain, after TBS to any of the three targeted locations (see **Figure 4.5**). Here we focus again on the second block after TBS (no significant cortical clusters were observed during the first block at p[corr]<0.05, minimum cluster size threshold range = 56-61 voxels). During this time window, TBS to either of the cognitive control network sites caused widespread increases in the connectivity of the aI/fO and dlPFC seeds to lateral frontal, parietal, and cingulate cortices, not limited to regions within the CO and FP network ROIs [p[corr]<0.05, minimum cluster size threshold range = 59-72 voxels].

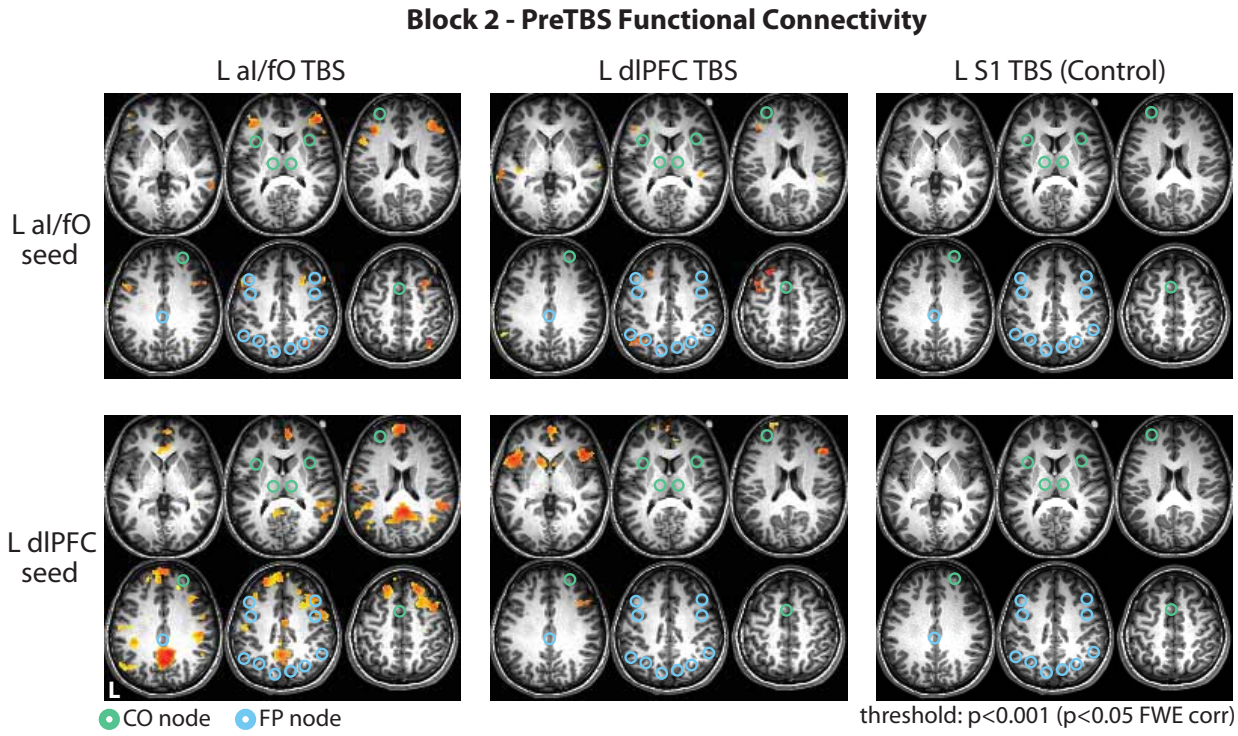

Figure 4.5: *Whole brain changes in connectivity after TBS*. Changes in connectivity of the L aI/fO region (*top row*) and L dlPFC region (*bottom row*) to the whole brain in the second block after TBS are shown following TBS to L aI/fO (*left column*), L dlPFC (*middle column*), and L S1 (*right column*). TBS to L aI/fO and L dlPFC led to widespread increases in connectivity across lateral frontal, parietal, and cingulate regions (see ***Tables 4.1 and 4.2***). No significant changes were seen after TBS to L S1. Connectivity values are thresholded at p<0.001 (p<0.05 FWE cluster correction; red values indicate increases in connectivity after TBS). Nodes from the CO and FP networks are overlaid on the brain maps in green and cyan, respectively.

	Clusters	Peak coordinate (MNI, LPI)	Size (voxels, 2 mm^3)
aI/fO connectivity	R inferior frontal sulcus/middle frontal gyrus	+40.0 +38.0 +14.0	638
	L inferior frontal sulcus/middle frontal gyrus	-40.0 +32.0 +14.0	368
	L posterior inferior frontal sulcus/middle frontal gyrus	-46.0 +8.0 +36.0	266
	R intraparietal sulcus	+38.0 -62.0 +52.0	187
	R superior frontal gyrus	+30.0 +16.0 +46.0	149
	R lateral orbitofrontal cortex	+34.0 +38.0 -10.0	142
	L rostral superior temporal gyrus	-54.0 -6.0 -10.0	84
	R middle temporal gyrus	+64.0 -46.0 +0.0	72
dlPFC connectivity	Posterior cingulate cortex	+8.0 -50.0 +30.0	2008
	R superior frontal gyrus	+42.0 +14.0 +52.0	1737
	L rostral superior frontal gyrus	-22.0 +46.0 +44.0	1416
	Anterior cingulate cortex	-4.0 +52.0 -8.0	1055
	R supramarginal gyrus/angular gyrus	+62.0 -46.0 +22.0	637
	L supramarginal gyrus/angular gyrus	-44.0 -42.0 +22.0	260
	Rostral precuneus	-0.0 -48.0 +64.0	239
	R anterior supramarginal gyrus	+50.0 -24.0 +30.0	179
	L precentral gyrus	-60.0 -20.0 +20.0	163
	L posterior angular gyrus	-54.0 -68.0 +24.0	150
	R anterior middle temporal gyrus	+66.0 -12.0 -8.0	117
	R inferior anterior insula	+32.0 +22.0 -16.0	102

Table 4.1: *Clusters showing increased connectivity after L aI/fO TBS*. This table lists all clusters that showed significant changes in connectivity in the second block after L aI/fO TBS ($p<0.001$ uncorrected, $p<0.05$ FWE cluster-size corrected). The top half of the table lists regions that exhibited changes in connectivity with L aI/fO, and the bottom half of the table lists regions that exhibited changes in connectivity with L dlPFC. No significant clusters showed decreased connectivity after L aI/fO TBS.

aI/FO stimulation. TBS to L aI/fO led to increases in connectivity of L aI/fO, with the largest clusters in bilateral regions of lateral inferior and middle frontal gyrus (including clusters near the dlPFC site of the FP network) and a cluster in the right intraparietal sulcus (proximal to the IPS node of the FP network). In addition, there were smaller clusters in the right superior frontal gyrus, orbitofrontal cortex, posterior temporal lobe, and left anterior temporal lobe. TBS to aI/fO also increased connectivity of the L dlPFC, primarily with regions in the anterior and posterior cingulate cortex, bilateral supramarginal/angular gyrus (together these sites are often characterized as part of the default mode network (Buckner et al., 2008; Fox et al., 2005; Raichle et al., 2001)), and the bilateral superior frontal gyrus. Small clusters of increased connectivity were also seen in left precentral gyrus, the right inferior anterior insula, and the midcingulate (see **Table 4.1** for all clusters).

	Clusters	Peak coordinate (MNI, LPI)	Size (voxels, 2 mm^3)
aI/fO connectivity	L angular gyrus/intraparietal sulcus	-48.0 -70.0 +36.0	160
	L dorsal premotor cortex	-42.0 +12.0 +44.0	144
	L middle frontal gyrus	-42.0 +26.0 +24.0	135
	L superior frontal gyrus	-24.0 +30.0 +48.0	135
	L middle temporal gyrus	-64.0 -34.0 +4.0	101
	R supramarginal gyrus	+44.0 -32.0 +18.0	69
	L superior temporal gyrus	-42.0 -24.0 +0.0	64
dlPFC connectivity	L anterior insula/inferior frontal gyrus	-46.0 +18.0 -4.0	1024
	R anterior insula/inferior frontal gyrus	+44.0 +24.0 +0.0	535
	Medial superior frontal gyrus	-0.0 +6.0 +64.0	246
	Anterior cingulate	+2.0 +58.0 +4.0	237
	R middle frontal gyrus	+46.0 +28.0 +22.0	208
	L anterior superior frontal gyrus	-20.0 +60.0 +10.0	91

Table 4.2: *Clusters showing increased connectivity after L dlPFC TBS*. This table lists all clusters that exhibited significant changes in connectivity in the second block after L dlPFC TBS ($p<0.001$ uncorrected, $p<0.05$ FWE

cluster-size corrected). The same conventions are used in this Table as in ***Table 4.1***. No significant clusters showed decreased connectivity after L dlPFC TBS.

dlPFC stimulation. TBS to L dlPFC, on the other hand, caused increases in connectivity of the L aI/fO with a few clusters in the left dorsal premotor cortex, middle and superior frontal gyrus, and angular gyrus. Smaller clusters of connectivity increases were also found in the left middle and superior temporal gyri and the right supramarginal gyrus. L dlPFC increased in connectivity with bilateral regions of the anterior insula/inferior frontal gyrus (near the aI/fO node of the CO network), anterior cingulate, and medial superior frontal gyrus. In addition, small increases were seen in the right middle frontal gyrus and anterior superior frontal gyrus (see **Table 4.2** for all clusters).

S1 stimulation. TBS to S1 did not cause any significant changes in connectivity across the whole brain, suggesting that these widespread effects may be somewhat specific to stimulation of cognitive control network regions (**Figure 4.5**).

Comparison between stimulation conditions. Direct comparisons between aI/fO TBS and S1 TBS at the whole brain level ($p(uncorr) < 0.001$, $p(corr) < 0.05$) showed that aI/fO TBS increased connectivity between dlPFC and default mode regions significantly more than S1 TBS (with significant clusters in the posterior cingulate, anterior cingulate, angular gyrus, and anterior temporal cortex, along with a small region in the frontal cortex that showed relatively decreased connectivity after aI/fO TBS; see **Table 4.3**). No other significant effects were seen at the whole-brain level for direct comparisons between either aI/fO TBS or dlPFC TBS and S1 TBS.

However, if a more relaxed uncorrected threshold was used ($p(uncorrected) < 0.01$, p(FWE corrected)<0.05, minimum cluster sizes 186-192 voxels), then aI/fO TBS showed significantly greater increases in connectivity than S1 TBS between aI/fO and several clusters of the bilateral middle frontal gyrus (along with a small subcortical cluster near the caudate nucleus). Also, at this threshold, dlPFC TBS caused significantly greater increases in connectivity than S1 TBS between dlPFC and a large cluster in the left anterior insula, along with clusters in the right temporal lobe. Similar clusters to these were also found at slightly different uncorrected thresholds ($p(uncorr) < 0.005$, p(corr)<0.05) and at the standard uncorrected threshold when using the average of both post-TBS blocks, suggesting that they may represent stable regions of increases in connectivity that are slightly weaker in magnitude and therefore not detected at the standard uncorrected threshold in block 2. See **Table 4.3** for a full list of clusters surviving direct comparisons. No clusters reached significance in the direct comparisons between TBS to aI/fO and TBS to dlPFC (even at the alternative thresholds used above).

	Clusters	Peak coordinate (MNI, LPI)	Size (voxels, 2 mm^3)
aI/fO vs S1 TBS			
aI/fO connectivity	* R middle frontal gyrus	+40.0 +46.0 +14.0	266
	* L caudal middle frontal gyrus	-50.0 +6.0 +36.0	244
	* R lateral orbitofrontal cortex	+32.0 +46.0 -8.0	236
	* L middle frontal gyrus	-38.0 +32.0 +14.0	210
	* Bilateral caudate nucleus	-12.0 -2.0 -10.0	208
	* R caudal middle frontal gyrus	+54.0 +18.0 +36.0	208
dlPFC connectivity	Posterior cingulate cortex	-2.0 -50.0 +34.0	415
	R anterior middle temporal gyrus	+66.0 -14.0 -8.0	246
	Rostral anterior cingulate cortex	-4.0 +52.0 -10.0	195
	Anterior cingulate cortex	-8.0 +26.0 -4.0	90
	L angular gyrus	-54.0 -66.0 +30.0	75
	Anterior cingulate cortex	+4.0 +30.0 -18.0	67
	R anterior middle frontal gyrus, white matter (*negative*)	+26.0 +42.0, +2.0	62
dlPFC vs S1 TBS			
aI/fO connectivity	None		
dlPFC connectivity	* L anterior insula/inferior frontal gyrus	-52.0 +22.0 -2.0	398
	* R superior temporal sulcus	+54.0 -10.0 -8.0	300
	* R temporal pole	+48.0 +24.0 -26.0	291

* Alternative threshold: p(uncorrected) < 0.01, p(FWE corrected) < 0.05

Table 4.3 *Clusters showing significantly different changes in connectivity after TBS to different sites*. This table lists clusters that exhibited significantly different changes in connectivity between the different TBS conditions in the second block after TBS. The first section of the table lists significant clusters in the comparison between aI/fO TBS and S1 TBS, and the second section lists significant clusters in the comparison between dlPFC and S1 TBS. No significant differences were seen in the direct comparisons between aI/fO and dlPFC TBS. In each section, changes in connectivity with the aI/fO ROI are listed first, followed by changes in connectivity with the dlPFC ROI.

Summary: Intriguingly, both the L aI/fO and, especially, the L dlPFC showed enhanced connectivity with regions very proximally located to sites in the opposing network (FP for L aI/fO and CO for L dlPFC) when TBS was applied to those regions. This may suggest that TBS is increasing connectivity between the two networks in regions that are not completely captured by the ROI-to-ROI analysis previously conducted. In addition, TBS to aI/fO caused a strong increase of the coupling between dlPFC and regions associated with the default mode network. No significant changes were seen after S1 TBS, and the effects of aI/fO TBS and dlPFC TBS were seen in direct contrasts with S1 TBS, although not always at standard uncorrected thresholds. No clusters survived direct comparisons between aI/fO and dlPFC TBS, suggesting that, although TBS caused widespread changes in the connectivity of these regions, the profiles of connectivity changes did not differ dramatically between the two regions.

4.5 Discussion

Here we have shown that TBS to a node in either the CO or FP network causes non-specific changes in network connectivity both between and within the CO and FP networks. ROI-to-ROI analyses did not reveal any significant selective changes based on TBS site. However, ROI-to-voxel analyses revealed that many different regions in frontal, parietal, and cingulate cortex showed increased connectivity after TBS to aI/fO (in the CO network) or dlPFC (in the FP network) which was not observed after TBS to a control site in somatosensory cortex. This suggests that acute disruption to cognitive control networks causes widespread changes in network connectivity that are not confined to the networks themselves.

4.5.1 Acute and chronic changes in network interactions after nodal disruption

Many previous studies have suggested that focal damage can modify functional connectivity across regions remote from the lesion site and throughout distributed networks (Carter et al., 2010; Grefkes et al., 2008; He et al., 2007; Mintzopoulos et al., 2009; Nomura et al., 2010; Ovadia-Caro et al., 2013; Sharma et al., 2009; van Meer et al., 2010; Wang et al., 2010; Warren et al., 2009). Moreover, after damage, changes in behavioral recovery over time track with changes in functional connectivity in particular subsystems (attention: (Carter et al., 2010; He et al., 2007), somatomotor: (Carter et al., 2010; Sharma et al., 2009; Wang et al., 2010), somatomotor in rats: (van Meer et al., 2010), multiple cortical networks: (Ovadia-Caro et al., 2013)), suggesting that modified connectivity may serve as a sensitive marker (and perhaps one of the driving mechanisms) of reorganization after disruption.

However, the network changes observed here following TBS were distinct from those we reported after focal brain lesions to either the CO or FP networks (Nomura et al., 2010). First, in our past study, chronic damage to the CO network selectively decreased the magnitude of connections among intact CO network nodes but did not impact the FP

network. Similarly, chronic damage to the FP network decreased the magnitude of connectivity among intact FP network nodes without impacting the CO network. In the present study, however, TBS to either cognitive control region increased the connectivity within the FP network (along with connectivity between the two networks) and across a widespread set of regions not confined to either network. Differences in the results from these two studies may in part be due to the amount of time between disruption of nodal function and scanning (short in the case of TBS and long-term in the case of the focal brain lesions, e.g. greater than 6 months). Previous work examining recovery of the dorsal and ventral attention networks after strokes impacting the ventral attention network (He et al., 2007) reported disruption of the connectivity of both networks during acute periods soon after an infarct. However, later in recovery, disruptions became more selectively limited to the ventral attention network. These results mirror the selective nature of disruptions seen in our chronic lesion patients (Nomura et al., 2010) and the diffuse, non-selective effects of disruptions seen here acutely after TBS.

Furthermore, it is unclear why the general pattern of effects in this study would be an increase in connectivity after TBS rather than a decrease, given the large amount of behavioral and physiological evidence suggesting that continuous TBS has inhibitory effects on local neural processing (reviewed in (Cardenas-Morales et al., 2011; Hoogendam et al., 2010; Thut and Pascual-Leone, 2010)). This seemingly inverted direction of effects has been seen in other studies that used excitatory TMS but observed predominantly decreased connectivity (e.g., (Bilek et al., 2013; Eldaief et al., 2011), and in studies that used inhibitory 1 Hz TMS and found remote increases in connectivity (Eldaief et al., 2011; Vercammen et al., 2010)see (Fox et al., 2012a) for a review of connectivity changes in fMRI-TMS studies). One possibility is that our findings and the others described here reflect compensatory mechanisms in network activity following TMS, in this case up-regulating distributed activity after decreases in local activity caused by TBS.

In addition, although most past studies of changes in functional connectivity after focal lesions have shown decreases relative to matched controls (He et al., 2007; Nomura et al., 2010), some studies have demonstrated that the recovery period after damage is characterized by both increases and decreases in connectivity, depending on the connection examined (Mintzopoulos et al., 2009; Park et al., 2011; Wang et al., 2010). In addition, experimental strokes induced in rats, although generally causing decreases in interhemispheric connectivity, also led to some increases in functional connectivity within the contralesional hemisphere (van Meer et al., 2010). These results suggest that disruption and recovery can lead to a varied pattern of changes that involve a mixture of up- and down-regulation of connectivity over time.

4.5.2 Increases in near-cross-network coupling seen in whole brain analysis

The changes after TBS in our study were better captured by our whole-brain analysis than the ROI-to-ROI analysis. In fact, the whole-brain analysis suggested that there were increases in “near-cross-network" coupling in regions adjacent to the examined ROIs that were not evident in ROI-based analysis. For example, TBS to dlPFC increased the connectivity between dlPFC and a region in the anterior insula/frontal operculum that was

ventral to the aI/fO ROI defined a priori in the CO network. Similarly, TBS to aI/fO led to increases in connectivity to sites in the middle frontal gyrus and posterior parietal cortex that were near the dlPFC and IPS regions, respectively, of the FP network.

These findings may imply that the ROIs used in our ROI-to-ROI analysis were non-optimal and did not effectively capture the CO and FP networks. However, these same ROIs during pre-TBS periods revealed the CO and FP networks (indicated by higher within- than between-network connectivity and seed-based analyses) in every individual. This suggests that our findings are not due to selection of poor ROIs. One possibility is that this extended connectivity following TBS reflects additional/expanded regions in the CO and FP network. These may include regions that normally participate in these networks (but are outside of the arbitrarily sized spherical ROI) or regions that normally exhibit sub-threshold/poorer network allegiance. A comparison of the patterns of connectivity before TBS shown in **Figure 4.2C** to the regions showing enhanced connectivity after TBS in **Figure 4.5** supports this notion: baseline network connectivity extends into the sites of post-TBS near-cross-network coupling. One past study has suggested that there may be gradients in the connectivity patterns within the extent of ROIs classically defined as components of the frontoparietal control and default mode networks and that these gradients vary depending on participants' age (Anderson et al., 2011). It is possible, therefore, that TBS may act as another factor that operates on the steepness of these gradients, thus effectively expanding the region of the ROI that passes threshold connectivity.

An increase in cross-network coupling may also indicate that additional regions in the opposing network are recruited in a compensatory fashion after TBS. Cooperation between these two normally independent networks may be necessary to maintain function in the acute phase of nodal disruption. Alternatively, increased inter-connectivity among these networks may be a sign of network dysfunction resulting from TBS, whereby these networks are no longer able to sufficiently suppress interactions among one another. Further studies, with behavioral measures associated with each of these networks, would help to tease apart these possibilities.

4.5.3 aI/fO TBS increased the coupling between dlPFC and regions in the default mode network

In addition to an increase in near-cross-network coupling, TBS to aI/fO caused dlPFC to become more tightly coupled with nodes in the default mode network (DMN), such as the posterior and anterior cingulate and angular gyrus (Buckner et al., 2008). Interestingly, TBS to dlPFC only produced relatively small changes in the connectivity of aI/fO that were not significantly different from changes induced by S1 TBS. This may suggest a hierarchy to the relationship between these two networks, with TBS to CO network regions able to more strongly influence FP (and default mode) network regions than the reverse (at least during the acute period after TBS).

In the past, regions in extended frontoparietal cognitive control networks have been found to couple dynamically with the DMN under particular task demands (Spreng et al., 2010) or in a manner that was correlated with individual differences in dopamine levels (Dang et al.,

2012) and age (Spreng and Schacter, 2012). Recent studies have suggested that aI/fO may act as a critical node that mediates the relationship between remote regions in the dorsal frontoparietal cognitive control network and the default mode network. In particular, Sridiharan and colleagues (2008) found that the anterior insula causally influenced connectivity of the DMN and a "central executive network" (a dorsal frontoparietal cognitive control network similar to the FP network) in the context of three separate tasks. These finding prompted the creation of a triple-network model (Menon, 2011; Menon and Uddin, 2010) that proposes a hierarchical relationship among these networks, with the salience network (similar to the CO network, with a prominent node in the anterior insula) causally influencing the DMN and the central executive network. In support of this model, recent findings in patients with schizophrenia (Manoliu et al., 2013), frontotemporal dementia (Chiong et al., 2013), and traumatic brain injury (Bonnelle et al., 2012) all suggest that the anterior insula (or broader salience network) influences DMN interactions. Within this context, then, our findings support the idea that TBS to aI/fO is another type of perturbation of the system that can cause dynamic changes in the coupling between the FP and default mode networks.

Another, not mutually exclusive, possibility is that changes in connectivity found following aI/fO and dlPFC TBS may be driven by the different network-level roles these regions play within the CO and FP networks, respectively. Damage to regions with many between-network connections (connectors) has been shown in the past to cause more widespread impacts on a variety of different networks, whereas this is not true of damage to regions with many within-network connections (within-module hubs; e.g., (Gratton et al., 2012)). If the aI/fO region acts as more of a connector region, linking the CO to other networks, and the dlPFC region acts as more of a within-network hub node with links predominantly to other areas in its own network, this might explain why damage to aI/fO would have relatively greater impact on FP network connectivity than damage to dlPFC has on CO network connectivity.

4.5.4 Generalized increases in connectivity after TBS

Regardless of the site of TBS (dlPFC, aI/fO, or S1), the ROI-to-ROI analysis revealed increases in connectivity both within and between the two cognitive control networks. These increases were significantly higher in FP and between network connectivity than CO network connectivity. These non-specific effects point to the clear need for other TBS locations to act as experimental control stimulation sites in studies seeking to make strong statements regarding the specificity of TBS effects. Past studies of dynamic connectivity have also reported increased network activity throughout scan sessions (Allen et al., 2012), and studies of sleep/drowsiness have also suggested that network connectivity may generally increase across the recording period (Fukunaga et al., 2006).

In particular, only a few prior studies have examined changes in functional connectivity during resting state fMRI after TMS, and none of these studies have used different stimulation sites to examine the specificity of their findings to the stimulation location (Eldaief et al., 2011; van Meer et al., 2010; Vercammen et al., 2010). One study applied 1 Hz TMS over dlPFC and observed decreases in resting state connectivity between the DMN and

the lateral temporal cortex compared with sham stimulation, but did not examine other sites or networks to determine whether this effect was selective to the particular stimulation condition (van der Werf et al., 2010). Another recent study stimulated a node in the default mode network with either 1 Hz or 20 Hz TMS and observed either increases (for 1 Hz TMS) or decreases (for 20 Hz TMS) in connectivity across default mode locations. This study controlled for both the specificity of effects to stimulation frequency (note again the inverted direction of effects in this study) and network examined, but did not test for the specificity in stimulation site (Eldaief et al., 2011). Therefore, our study design is unique in that it allows us to assess specificity in stimulation site effects on resting state connectivity. Critically, this allowed us to observe the different effects (or, for certain contrasts, the lack thereof) that result from stimulating distinct, putatively independent, networks and experimental control locations.

4.5.5 Methodological limitations

Several difficulties arise in conducting a combined resting state fMRI – TBS study. First, a large amount of variability exists in TMS effects across individuals, and these may have a wide variety of causes (e.g., previous activity, differences in genetic markers, time-of-day, gender, age, etc.; see review by (Ridding and Ziemann, 2010)). Here, we used a large sample size (relative to other TBS studies) and a within-subjects design to mitigate these effects.

Second, differences may exist in the anatomical locations of stimulation sites across individuals that may contribute to variability in TBS effects. Here, we used a method adapted from that of Eldaief et al (2011) to target subject-specific network regions (see also (Fox et al., 2012b)). In support of the success of this method, we found that before TBS, the proposed networks were identifiable in all individuals. This suggests that variability in these effects is not driven by imprecise targeting of stimulation sites across individuals.

Third, although TMS is often characterized as a method for creating temporary "virtual lesions" (e.g., (Pascual-Leone et al., 1999)), TMS and chronic focal lesions differ along a number of dimensions. Whereas TBS is fairly selective in the location that is targeted, the size and location of lesions after brain injury differ across individuals and is not under experimental control. However, in our previous lesion study (Nomura et al., 2010), the most common locations for lesions overlapped with aI/fO and dlPFC regions, thus increasing the similarity between the acute and chronic approaches in this regard. In addition, past studies have suggested that TBS impacts underlying cortical activity through inhibitory mechanisms that increase long-term depression of synapses ((Huang et al., 2005; Huang et al., 2011), reviewed by (Cardenas-Morales et al., 2011)). Lesions, instead, destroy large populations of neurons as well as adjacent white matter tracts. These different consequences of TBS and lesions may also have fundamentally different effects on network activity. Consistent with this, lesions lead to decreases in network connectivity (Nomura et al., 2010), while TBS generally leads to increases in functional connectivity in the present study (see discussion on this point in previous section). However, one past study has directly compared the effects of inhibitory 1 Hz TMS to that of focal lesions to the

same locations in prefrontal cortex (Miller et al., 2011) and reported similar results in functionally coupled regions.

Despite the inherent limitations associated with each of these techniques, the use of combined TBS-fMRI methodology, a large within-subjects design with multiple stimulation sites, and the ability to compare these results directly to previous work in focal lesion patients allows us to greatly extend our understanding of how focal disruptions affect network communication. The results obtained in this study suggest that TBS causes fundamentally different changes in network connectivity than chronic focal lesions, demonstrating more widespread and less selective effects on network organization. These differences suggest that the study of acute versus chronic brain perturbations may highlight unique aspects of the dynamic nature of network communication, allowing us to make more causal and specific claims about interactions across these regions.

Chapter 5

INDEXING THE VARIABILITY IN THE FUNCTIONAL BRAIN EFFECTS OF THETA BURST TMS: A COMBINED BOLD AND ASL FUNCTIONAL MRI STUDY

5.1 Abstract

Transcranial Magnetic Stimulation (TMS) is an important tool for testing causal relationships in cognitive neuroscience research. However, the efficacy of TMS can be quite variable across individuals and difficult to measure, especially in studies without well-characterized behavioral effects. Here, we examined whether physiological measures from fMRI index variability in the effects of TMS. Each of 27 participants completed four combined TMS-fMRI sessions during which both standard EPI and perfusion scans were recorded during a resting state. In each session after the first baseline day, continuous theta-burst TMS was applied to one of three locations: left dorsolateral prefrontal cortex (L dlPFC), left anterior insula/frontal operculum (L aI/fO), or left primary somatosensory cortex (L S1). Blood Oxygen Level-Dependent (BOLD) and Cerebral Blood Flow (CBF) signal changes were measured both before and after TMS on each day from a series of interleaved resting state and perfusion scans. We found that TMS caused a trend towards decreased mean BOLD signal and significantly increased BOLD signal standard deviation under the coil, but these changes were not specific to a particular location. In contrast, mean CBF increased under the coil in a site-specific manner. No consistent changes were observed in CBF standard deviation. In addition, TMS-induced changes in mean CBF were selectively related to TMS-induced changes in functional connectivity. Subject groups with different directions of CBF change after TMS showed strikingly different patterns of change in the functional interactions of stimulated networks. These results suggest that CBF may be a useful independent marker of individual differences in the effects of TMS.

5.2 Introduction

Transcranial magnetic stimulation (TMS) is a powerful technique that allows temporary disruption or enhancement of function of a targeted brain region in healthy individuals (Dayan et al., 2013; Hoogendam et al., 2010). The combination of TMS and functional Magnetic Resonance Imaging (fMRI) allows for strong inferences to be made regarding the causal interactions among brain areas (Bestmann et al., 2008; Driver et al., 2009). This multi-modal approach can allow neuroscientists to step beyond the mere correlational arguments that can be made with measures of functional connectivity from fMRI data, potentially leading to significant advancements in our understanding of the dynamic nature of brain networks.

However, past studies have found that the effect of TMS on brain function can be extremely variable across individuals (Ridding and Ziemann, 2010), including in the physiological changes occurring locally at the stimulation site (e.g., (Gangitano et al., 2002; Hamada et al., 2013; Hamidi et al., 2009; Kozel et al., 2000; Maeda et al., 2000; Martin et al., 2006; Muller-Dahlhaus et al., 2008). Even after employing individualized TMS parameters (e.g., motor thresholds, stimulation sites), the influence of TMS on underlying tissue can differ due to a number of other reasons (Ridding and Ziemann, 2010). These include factors related to biophysical interactions between TMS and the target region (e.g, the distance between the skull surface and cortex, the precise folding pattern of the underlying cortex, the orientation at which the TMS coil is held (reviewed by (Wagner et al., 2009), the previous state of the participant (e.g., interactions with ongoing fluctuations (Hamidi et al., 2009), previous activity (Gentner et al., 2008), time of day (Sale et al., 2008), hormonal state (Inghilleri et al., 2004)), and more stable, intrinsic differences among participants (the individual's age (Muller-Dahlhaus et al., 2008), genetic makeup (Cheeran et al., 2008), and the ability to activate distinct neuronal populations (Hamada et al., 2013)).

In both primary motor and sensory cortical regions, the effect of TMS on underlying cortex can be estimated by measuring either muscle activity or calculating sensitivity thresholds after TMS (Dayan et al., 2013). In multi-modal association regions, in contrast, there is typically no direct behavioral measure of the functional properties of a single region, which complicates attempts to index TMS variability. Resting-state fMRI is becoming an increasingly popular method used to characterize the functional roles of multi-modal regions in large-scale "intrinsic" networks that emerge in the absence of any task instruction (Fox and Raichle, 2007). Resting-state data are easy to collect, reproducible within and across individuals (Damoiseaux et al., 2006; Shehzad et al., 2009), and can be used to measure a large number of clinically and behaviorally relevant networks in a short period of time (Greicius, 2008; Sadaghiani et al., 2010; Smith et al., 2009; Zhang and Raichle, 2010). Therefore, the application of TMS to multi-modal regions in intrinsic networks is an extremely useful way to increase our understanding of both large-scale network interactions and the functional roles of individual multi-modal regions. However, given the lack of behavioral measures during the resting state, the inherent variability in the effects of TMS on cortical function has posed a significant challenge.

Three past studies have successfully found TMS-induced changes in brain networks during resting-state fMRI ((Eldaief et al., 2011; van der Werf et al., 2010; Vercammen et al., 2010); reviewed by (Fox et al., 2012a). However, despite the lack of a control target region to assess regional specificity of the effects, the changes caused by TMS were only found in a small subset of regions remote from stimulation (e.g., medial temporal cortex after dlPFC TMS (van der Werf et al., 2010), insula after temporo-parietal junction TMS (Vercammen et al., 2010)), or in predicted locations with quite variable responses across individuals (Eldaief et al., 2011), *see Fig 5*; *note also that these results were in an unpredicted direction*). Therefore, it is imperative to find a physiological metric that can be recorded with fMRI that accounts for variability in the effect of TMS. Functional MRI can be used to measure both the Blood Oxygen Level-Dependent (BOLD) signal and regional Cerebral Blood Flow (rCBF). The BOLD signal is the most widely used measure in fMRI, due to its relatively high signal-to-noise ratio and temporal resolution. However, rCBF measures of brain perfusion

may have several advantages over the BOLD signal, including a unit of measurement that refers to a physical quantity (ml/g/min, unlike the arbitrary units of the BOLD), a stronger ability to detect slow changes in neuronal activity, and better comparability across scan sessions (Aguirre et al., 2005). These advantages make perfusion fMRI a strong candidate for measuring TMS-induced physiological changes.

Here we examined whether TMS would induce consistent, functionally-relevant changes in BOLD or rCBF measures, collected in interleaved scans. We used continuous theta burst TMS (cTBS), a TMS protocol that produces long-lasting inhibition of underlying cortical activity (Huang et al., 2005), in order to target three different locations: left dorsolateral prefrontal cortex (*L-dlPFC*), left anterior-insula/frontal operculum (*L-aI/f0*), or left primary somatosensory cortex (*L-S1*). The two frontal regions represent areas thought to be critically involved in separable large-scale cognitive control networks (Dosenbach et al., 2008; Dosenbach et al., 2007).

We first explored whether there were general changes in fMRI signals under the TMS site. Crucially, we then examined how variability in these measures was related to variability in the functional properties (i.e., functional connectivity) of the targeted regions. Using three different stimulation sites allowed us to determine both whether TMS-induced changes were specific to the targeted region and how generalizable our metrics were to a variety of different brain sites.

5.3 Materials & Methods

5.3.1 Participants

Twenty-seven healthy right-handed participants (11 female, 18-31, mean age = 22.7 years) completed four separate combined fMRI and TMS sessions. All participants provided written informed consent before the study in accordance with procedures approved by the Committee for the Protection of Human Subjects at the University of California, Berkeley and were compensated monetarily for participation.

5.3.2 Experimental Timeline

See **Figure 5.1** for a depiction of the experimental timeline. The first session was used to acquire an anatomical image for each participant and to map subject-specific TMS sites and motor thresholds. The subsequent three sessions included a pre-TMS period, during which one echo-planar imaging (EPI) and one arterial spin labeling (ASL) scan were collected (*see Section 5.3.5*). This was followed by a 40-second period of theta-burst TMS outside the scanner, after which participants re-entered the scanner and completed two blocks each of alternating EPI and ASL scans. A different TMS site was targeted in each session (with order counterbalanced across participants): left dlPFC, left aI/fO, and left S1. During scanning, a fixation cross was presented to the screen and participants were instructed to remain awake with their eyes open and to allow their mind to wander without focusing on any particular topic. During all scans, the participant's wakefulness was monitored using an MR-compatible eye-tracking camera. Each session took approximately one hour to

complete, and separate sessions were always completed at least a day apart. In a few participants (N=1 for aI/fO TMS, N=1 for dlPFC TMS, and N=2 for S1 TMS), the second perfusion scan was missing due to participants either falling asleep (N=2) or sessions exceeding their time limit (N=2).

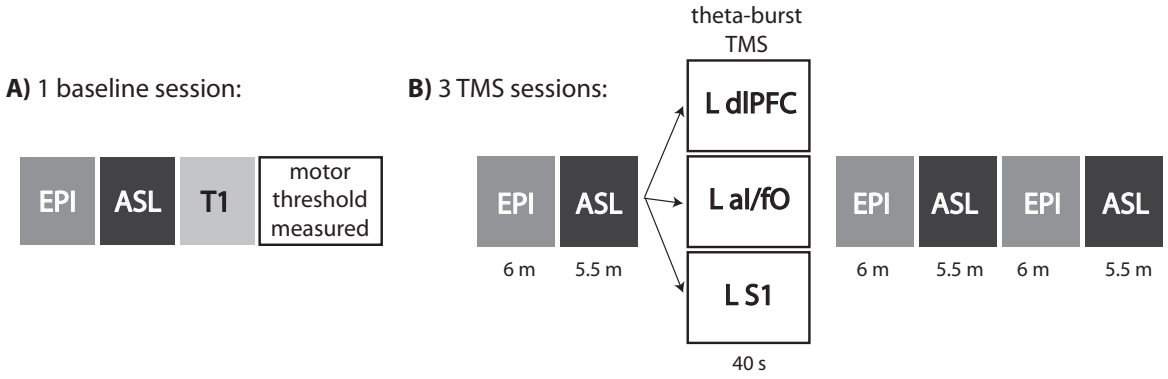

Figure 5.1 *Experimental Timeline*. Participants (N=27) completed 4 separate sessions. (A) The first session was used to acquire baseline EPI and ASL fMRI measures, a high-resolution anatomical MRI scan, and resting motor thresholds using EMG. (B) The remaining 3 fMRI scans began with the collection of pre-TMS EPI and ASL scans. These were followed by 40 s of theta-burst TMS to one of three target locations (L dlPFC, L aI/fO, or L S1). Following TMS, two EPI and two ASL scans were collected in an interleaved fashion.

5.3.3 Transcranial magnetic stimulation

All TMS stimulation occurred outside of the scanner. Subjects were seated in a comfortable chair throughout all TMS sessions. TMS was applied using a handheld figure-eight coil (outer winding diameter = 70 mm), and all pulses were delivered to the coil using a Magstim rapid stimulator with four booster modules producing biphasic pulses (*Magstim, UK*).

Motor thresholds were calculated during the first session. Electromyography was used to record potentials from electrodes placed on the first dorsal interosseous (FDI) muscle from the dominant hand. Stimulation was applied over the hand representation within primary motor cortex, with the coil placed tangentially to the scalp and the handle of the coil pointing posteriorly. This area was defined as the scalp location where stimulation led to the largest motor-evoked potential (MEP) from the contralateral FDI while the participant was relaxed. To define the stimulation intensity, the active motor threshold (AMT) was calculated in each participant. This was defined as the minimum intensity of single-pulse stimulation required to produce a detectable MEP from the contralateral FDI on 5 out of 10 trials while the participant maintained contraction at 20% of their maximum in the FDI. EMG signal feedback was provided through visual feedback on a computer monitor to help the participant maintain contraction consistently at this level.

Stimulation sites (defined as described below) were localized in each participant using a computerized frameless stereotaxic system (*Brainsight software, Rogue Research, Canada*). In this system, head position is defined in real time by using reflective markers that are

placed on the participant's head and imaged with an infrared camera. These positions are co-registered with anatomical locations placed on the individual participant's structural MRI scan (acquired during the first session). Markers are placed both on the participant's head and on the coil so that the relative position of the coil to the participant's head (and MRI) can be tracked to precisely place the coil with respect to the targeted anatomical locations.

Once the coil was placed at the targeted location in each participant, 40 seconds of continuous theta-burst stimulation was applied at 80% of each individual's AMT. During theta-burst TMS, 50 Hz trains of three TMS pulses are provided every 200 ms continuously over a period of 40 s (total of 600 pulses). This form of repetitive TMS was selected to produce long-lasting reductions in cortical excitability that persist for up to 60 minutes following stimulation (as described by (Huang et al., 2005) to allow for sufficient fMRI recording time, with the stimulation parameters used in this study matching those used by Huang and colleagues (2005).

5.3.4 Stimulation sites

Three sites of interest were targeted in each individual (see **Figure 5.2A**): the left dorsolateral prefrontal cortex (L dlPFC), left anterior insula/frontal operculum (L aI/fO), and left primary somatosensory cortex (L S1). Stimulation sites were defined for each individual based on a combination of anatomical and functional criteria (modeled after (Eldaief et al., 2011). The control site (L S1) was established anatomically, by centering a spherical region-of-interest (ROI) over the left postcentral gyrus in each individual. Each participant's left S1 was identified on the anatomical scan, and ROIs were drawn as spheres with a radius of 6 mm centered 10 mm away from the midline and 6 mm from the top edge of the brain. The other two sites were based on coordinates from the fronto-parietal (L dlPFC) and cingulo-opercular (L aI/fO) cognitive control networks first described by Dosenbach and colleagues (Dosenbach et al., 2008; Dosenbach et al., 2007) and modified for each individual participant through a functional connectivity analysis (see **Figure 5.2B**). We reverse-normalized all 18 regions in the Dosenbach cognitive control networks (11 fronto-parietal and 7 cingulo-opercular) into individual subject space.

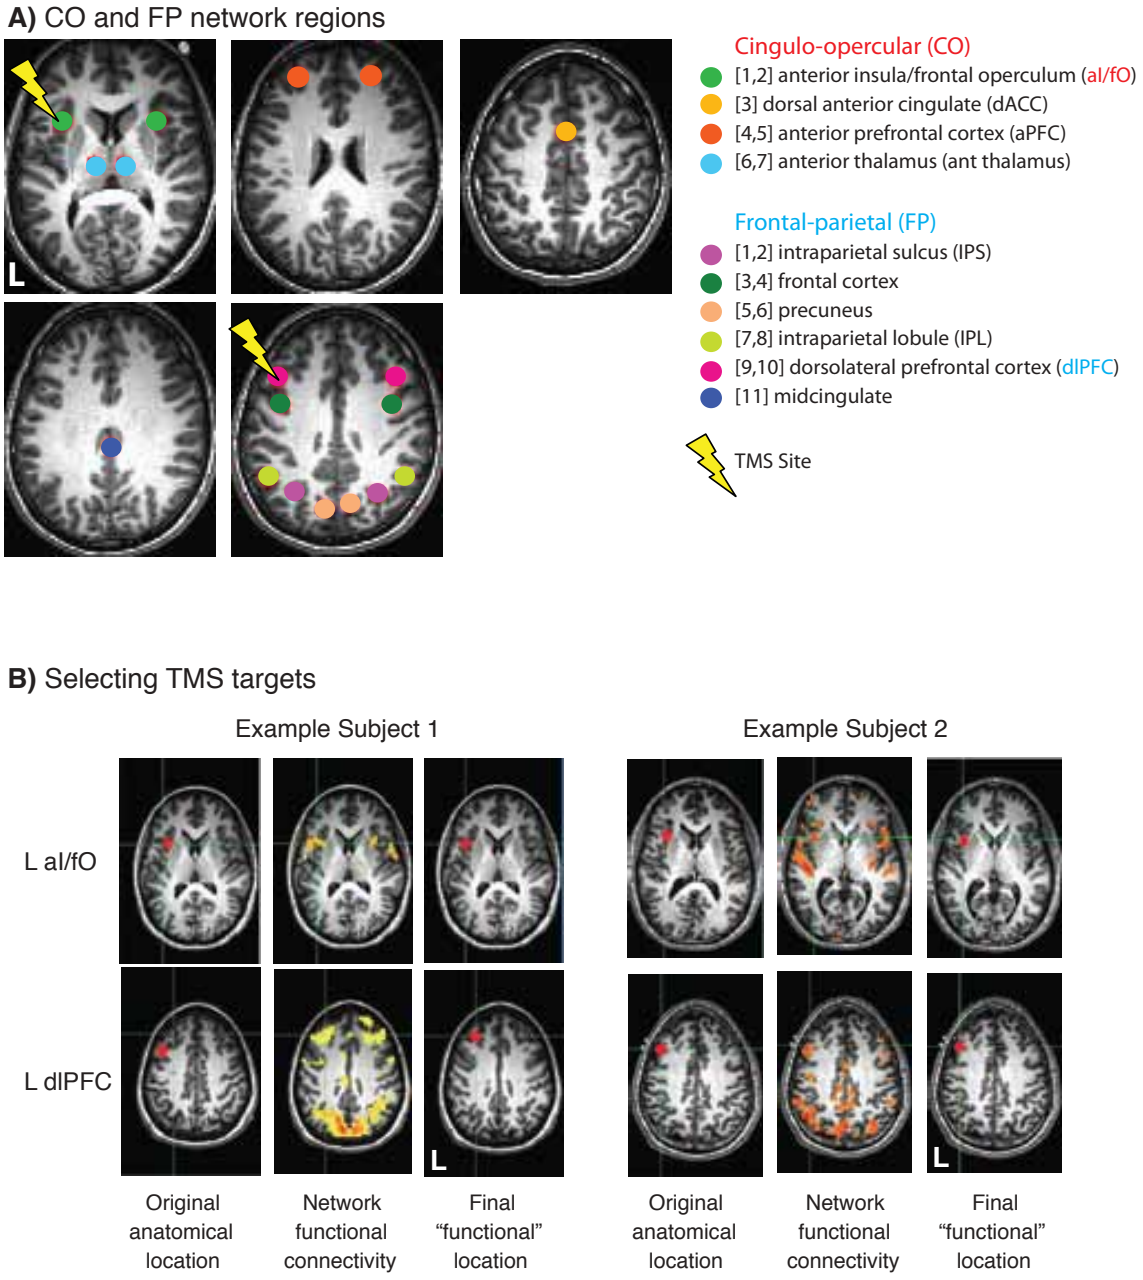

Figure 5.2 *Regions of Interest*. (A) Two of the TMS target regions were part of intrinsic cognitive control networks: the L dlPFC (in the fronto-parietal or FP network) and the L aI/fO (in the cingulo-opercular or CO network). The other nodes from these networks are also displayed. In addition, S1 was used as a control target region (*not shown*). (B) The L dlPFC and L aI/fO TMS targets were functionally localized in each subject by finding the cluster of within-network functional connectivity closest to previously published coordinates (Dosenbach et al., 2007), *see Section 5.3.7*). In some subjects (e.g., *Example Subject 2*), these clusters were already near the predefined coordinates, but in others this process shifted target coordinates (e.g., *Example Subject 1*).

In each participant, a whole-brain functional connectivity analysis was then conducted using the resting-state EPI fMRI data from the baseline session. The whole-brain functional connectivity analysis was conducted for each network separately, using the average time-series from all regions in each network (excluding the target stimulation site). For example, a cingulo-opercular whole-brain functional connectivity map was created by averaging the time-series from all cingulo-opercular regions except the L aI/fO ROI and then computing Pearson's correlation coefficients between this average time-series and that of every voxel of the brain. Similarly, a fronto-parietal whole-brain connectivity map was created using the average time-series from all fronto-parietal regions except the L dlPFC ROI (*see Section 5.3.6 below, for more details on data preprocessing and functional connectivity measures*).

Then, within each network map (cingulo-opercular for the L aI/fO node and fronto-parietal for the L dlPFC node) we found the cluster of connectivity closest to the originally specified coordinates for the target region (*see* **Figure 5.2B** for examples from two participants). The peak of this cluster was selected and used to create a spherical ROI (6 mm) that served as the target for stimulation by TMS. If no nearby cluster could be found (as occurred for 4 subjects for L aI/fO and 2 subjects for L dlPFC), the original reverse-normalized coordinates were used as the center of the ROI. These subject-specific ROIs were defined in an effort to optimize the stimulation site for functionally impacting each cognitive control network in each individual.

5.3.5 MRI acquisition parameters

All Magnetic Resonance Imaging (MRI) data were acquired from a Siemens MAGNETON Trio 3-Tesla scanner using a 12-channel head coil. Structural images were acquired during the first session using a whole-brain MP FLASH T1-weighted scan. Two forms of whole-brain functional MRI scans were acquired in each session: EPI scans for measuring resting-state functional connectivity and ASL scans for measuring resting-state perfusion.

EPI scans were obtained using a T2*-weighted pulse sequence sensitive to blood oxygenation level-dependent (BOLD) contrast (TR = 2 s, TE = 24 ms, flip angle = 60 degrees, in-plane matrix = 64 x 64 pixels each 3.5 x 3.5 mm, with 37 3.5-mm descending axial slices and a 0.7 mm slice gap), and each scan consisted of 180 time points (6 minutes). Perfusion scans were collected using a pseudo-continuous ASL sequence with a standard EPI readout and PACE prospective motion correction (Dai et al., 2008; Wu et al., 2007) (TR = 4 s, TE = 11 ms, flip angle = 90 degrees, in-plane matrix = 64 x 64 pixels each 3.4 x 3.4 mm with 20 5-mm descending axial slices and a 1 mm slice gap). Each scan consisted of 80 images (40 control/label pairs) and lasted about 5.3 minutes. Presaturation of fat was applied for both EPI and ASL scans.

5.3.6 Data processing

<u>Perfusion ASL fMRI:</u>

Preprocessing: Preprocessing of the ASL scans was performed in SPM8 [*http://www.fil.ion.ucl.ac.uk/spm/software/spm8/*]. Data were extracted from DICOMs and

then realigned to the first time-point of the entire session (day) using a rigid body (6 parameters) method.

CBF calculation: Cerebral blood flow (CBF) was calculated with software implementing a simple subtraction procedure and equation (1) from Wang et al (2005). This generated CBF values for each time point. For mean CBF calculations, these values were averaged across all time-points within a scan.

Post CBF processing: After CBF calculation, the individual and mean CBF maps were coregistered to the anatomical image using the mean BOLD image produced from the ASL calculations to improve alignment (all alignments were manually checked; if coregistration with the anatomical image was poor, an individual motion-corrected ASL image was used as the target image). This step included resampling of data into 1 mm^3 isotropic voxels. Data were then smoothed (6 mm FWHM) and gray-matter masked to exclude signals from high-intensity blood vessels outside of the brain. Finally, images were normalized by the global gray-matter CBF in each scan to produce regional CBF (rCBF) maps for the scan.

Resting state EPI fMRI

Preprocessing: Preprocessing for the resting-state EPI scans was carried out in AFNI [*http://afni.nimh.nih.gov/afni/* (Cox, 1996)]. Images were extracted from DICOMs, and data were slice-time corrected (quantic Lagrange polynomial interpolation). Functional images were realigned to the first time-point in each session (using a rigid body 6-parameter method) and were coregistered to the anatomical image (using AFNI's align_epi_anat.py). For BOLD time-series analyses, the time-series were highpass filtered (0.01 Hz) and smoothed (6 mm FWHM).

Functional connectivity: Functional connectivity was computed using time-series correlations as in Fox et al (2005). Voxel time-series were bandpass filtered (0.009 – 0.08 Hz), in order to minimize physiological noise such as respiratory or cardiac artifacts, and smoothed (6mm FWHM). In addition, we regressed out nuisance signals from subject-specific white matter and ventricle masks, each subject's motion parameters, and their temporal derivatives. Functional connectivity was assessed by computing Pearson's correlation coefficients between the average time-series from each ROI and every other voxel in the brain (for whole-brain analyses to determine TMS stimulation sites; see *Section 5.3.4*) or between the ROI and all other ROIs in the same cognitive control network (in order to assess the functional properties of the ROI). These connectivity values were then Fisher-transformed before statistical analysis.

5.3.7 Regions of Interest

The transformation from normal to native space was computed using SPM's segment function (Ashburner and Friston, 2005); this was used to reverse-normalize coordinates into native space. For the three sites targeted by TMS, regions of interest were determined as described under *Section 5.3.4*, based on a combination of anatomical and functional criteria for two cognitive control network ROIs (Dosenbach et al., 2007) and a somatosensory cortex ROI. Additional ROIs used for functional connectivity analyses were defined using the reverse-normalized coordinates from Dosenbach (2007). Homologous ROIs were defined in the same manner (for R dlPFC and R aI/fO) or by reflecting the coordinates across the midline (for R S1).

5.3.8 Statistical Analysis

Repeated-measures ANOVAs were carried out using SPSS (*Version 20, Chicago, IL*). First, an ANOVA with factors of time-point (pre-TMS or average of post-TMS blocks) and TMS condition (dlPFC, aI/fO, and S1) was performed in order to determine whether TMS consistently impacted BOLD or CBF measures under the coil. If a consistent change was found, this was further investigated with a repeated-measures ANOVA, with factors of time-point (pre-TMS vs post-TMS average), TMS condition (dlPFC, aI/fO, or S1), and ROI (dlPFC, aI/fO, or S1), in order to determine whether the TMS effects showed selectivity to the targeted TMS sites. Analyses initially compared the TMS effect averaged over the two post-TMS blocks and then investigated the duration of the effects by treating them separately for each block. In participants missing their second perfusion scan (see *Section 5.3.2*), the averages represented the values from their first post-TMS perfusion scan, and these participants were omitted from the block-specific analyses.

In addition, subsequent correlation analyses were conducted using Scipy [*http://www.scipy.org/*] to determine whether there was a relationship between the changes in CBF measures and changes in functional connectivity of targeted nodes. Correlation values in these analyses represent Spearman's (rho) instead of Pearson's correlation coefficients, resulting in less influence from potential outliers. Differences between correlation values were evaluated using the formula described in Steiger (1980) after using the conversion from Spearman's to Pearson's coefficients described in Myers & Sirois (2004). All reported t-tests and correlations are accompanied by two-tailed p-values. Plots and network graphs were created using the NetworkX [*http://networkx.github.io/*] and Matplotlib [http://matplotlib.org/] packages in ipython [*http://ipython.org/*].

5.4 Results

5.4.1 BOLD signal and rCBF measurements under the TMS site

We first examined whether there were characteristic changes under the TMS site in the mean or standard deviation of the BOLD signal and CBF values at the three targeted locations (aI/fO, dlPFC, S1; **Figure 5.3**).

BOLD

Relative to the pre-TMS baseline, there was a weak trend across TMS sites for mean BOLD to decrease under the coil (in the average of post-TMS blocks; *main effect of block: $F(1,26) = 2.595$, $p<0.12$*), but there was no significant change in mean BOLD signal in any individual TMS site [*all post-hoc t-tests $t(26) < 1.32$, $p>0.19$*] (see **Figure 5.3A**). In contrast, there was a significant increase in the standard deviation of the BOLD signal following TMS under the coil across the three stimulation sites [*main effect of block: $F(1,26) = 11.496$, $p<0.002$*] and in two of the three TMS sites individually [*aI/fO TMS: $t(26) = 3.39$, $p<0.005$; dlPFC TMS: $t(26) = 2.82$, $p<0.01$; S1 TMS: $t(26) = 1.65$, $p<0.12$*] (**Figure 5.3A**). However, these changes in standard deviation of the BOLD signal were also observed in the non-stimulated ROIs (**Figure 5.3B**), suggesting that they were not reflecting the local, site-selective, effects of TMS but might instead be non-specific temporal trends. This was supported by the results

from a separate ANOVA testing whether TMS-induced changes were seen selectively only under the coil or across all tested ROI's [*non-significant interaction of block, ROI, and TMS site: F(4,104) = 1.32, p<0.27*]).

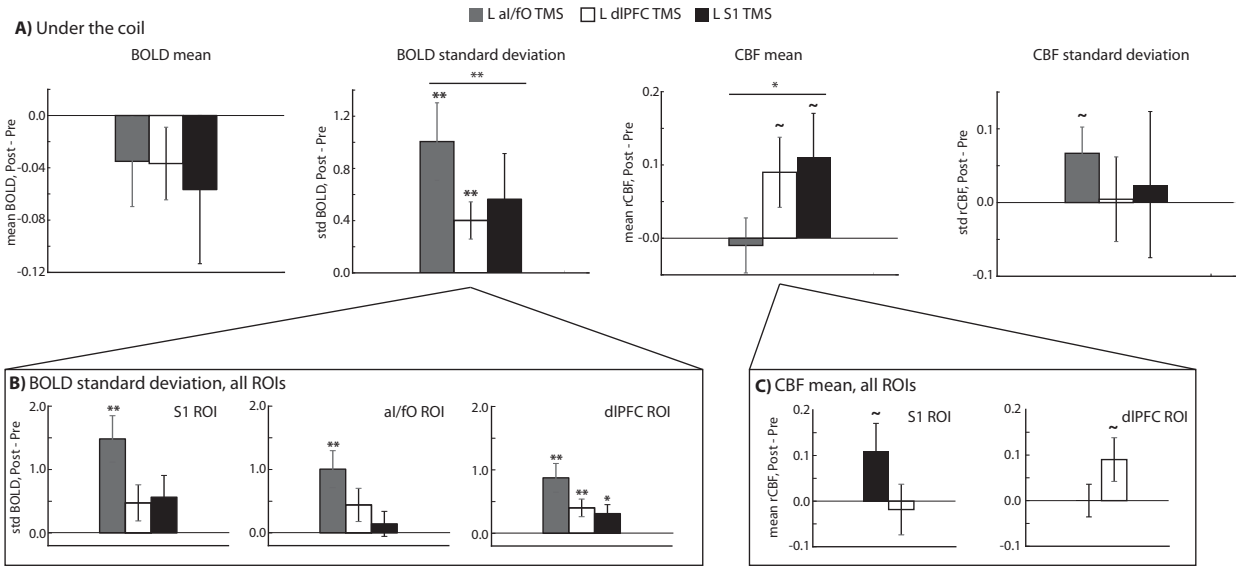

Figure 5.3 *BOLD and CBF measures after TMS across all sites.* (A) Changes in BOLD signal (mean and standard deviation) and CBF (mean and standard deviation) under the TMS site are shown for the three targeted locations. For the standard deviation of the BOLD signal (B) and the mean CBF (C), where consistent effects of TMS were seen under the coil, site selectivity was examined by comparing changes across the different targeted locations and regions of interest. (*** p<0.01, * p<0.05, ~ p<0.10; horizontal lines in A indicate a significant main effect of block*).

CBF

Relative to the pre-TMS baseline, mean rCBF increased under the coil after TMS across TMS sites (**Figure 5.3A**; [*Main effect of block, F(1,26)=4.61, p<0.05*]). Specifically, increases in CBF under the coil were found after TMS to left dlPFC [*first block after TMS: t(26)=3.10, p<0.005; average of both post-TMS blocks: t(26)=1.88, p=0.071*] and left S1 [*second block: t(24)=1.94, p=0.064; average of both blocks: t(26)=1.82, p=0.08*] but not the left aI/fO region [*average and all blocks individually: p>0.69*]. Interestingly, these changes were partially mirrored in the nodes homologous to the TMS site, in the contralateral hemisphere (*[main effect of block: F(1,26)=4.07, p=0.054]*). At the level of individual ROIs, rCBF increased in the right dlPFC after left dlPFC TMS [*first block after TMS: t(26)=3.52, p<0.002; average of both blocks: t(26)=2.82, p<0.01*]. The standard deviation of the rCBF measurement, in contrast, did not show consistent changes under the coil after TMS (**Figure 5.3A** [*main effect of block F(1,26)=0.63, p>0.43]).*

Furthermore, for the dlPFC and S1 TMS conditions, the change in mean rCBF under the coil was selective to the TMS site **Figure 5.3C** [*ANOVA with dlPFC and S1 conditions showed an interaction between block, ROI, and TMS site F(1,26)=4.40, p<0.05]*). In the dlPFC region, dlPFC TMS increased activity significantly more than S1 TMS during the first block after

TMS [*t(26)=2.45 p<0.03*], and in the S1 region, S1 TMS showed a weak trend toward increasing activity more than dlPFC TMS after both blocks *[average of both blocks: t(26)=1.57, p<0.13*]. This suggests that for these two sites, the increases in rCBF were specific to the target of TMS.

5.4.2 Individual variability of fMRI measurements under the TMS site

TMS-induced changes in both BOLD and CBF were extremely variable across individuals (see **Figure 5.4**). Although both the standard deviation of the BOLD signal and mean CBF showed some consistent TMS effects on the group average (see *Section 5.4.2*), individual participants showed both increases and decreases in any given session. A question that arises from this observation is whether the variability in the local effects of TMS is functionally relevant or simply a consequence of imprecise measurements. One way to examine this is by determining whether the TMS-induced variability in these local variables is related to changes in the functional connectivity of our stimulated regions. Since the two frontal sites in our experiment were selected based on their roles within intrinsic cognitive control networks, we used these network interactions to probe the functional consequences of variability in TMS effects.

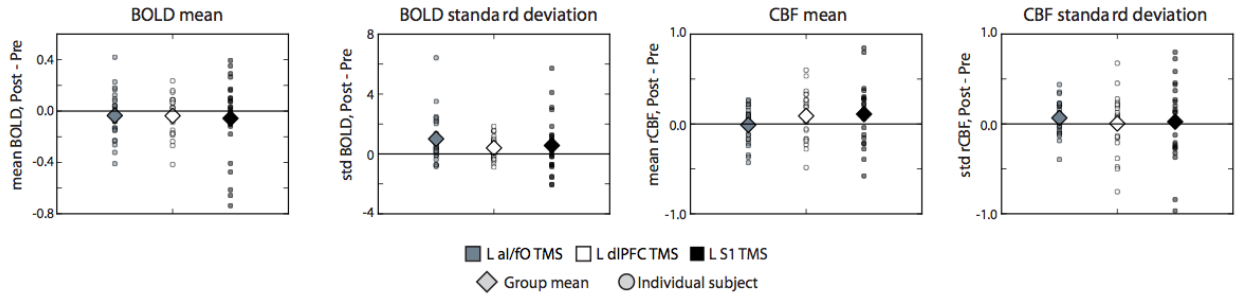

Figure 5.4: *Individual variability in BOLD and CBF effects.* Changes in BOLD signal (mean and standard deviation) and CBF (mean and standard deviation) under the coil are shown for each individual for each TMS conditions. As can be seen, individual variability was high, with any given measure containing subsets of participants showing increases and decreases.

5.4.3 Individual differences

<u>Correlations with functional connectivity.</u>
For both the aI/fO and dlPFC ROIs, changes in rCBF after TMS were negatively related to changes in functional connectivity between the stimulation site and its own network [*average of both post-TMS blocks, aI/fO to CO: r_S = -0.59, p<0.005;* dlPFC to FP: r_S = *-0.39, p<0.05*] (see **Figure 5.5A, left**). That is, increases in perfusion after TMS were related to decreases in the functional connectivity between the targeted region and its own network. These relationships were weaker when each post-TMS block was analyzed separately [*aI/fO to CO network, Bl1: r_s = -0.30, p=0.12, Bl2: r_s = -0.57, p<0.005; dlPFC to FP network, Bl1: r_s = -0.37, p=0.061, Bl2: r_s = -0.23, p=0.27*].

The relationship between changes in CBF and functional connectivity was present in the aI/fO ROI only after TMS to aI/fO, not after TMS to S1 [r_S = *-0.01, p > 0.97*], and these two relationships were significantly different from one another [*Z = -2.44, p<0.02*] (**Figure 5.5A, right**). Similarly, CBF and functional connectivity for the dlPFC ROI were only related after TMS to dlPFC, not after TMS to S1 [r_s = *-0.03, p > 0.87*]. Although the direct comparison between these correlations did not reach significance [*Z = -1.37, p=0.171*], the relationships trended toward being significantly different from one another during the first block after TMS [S1 TMS: r_S = *0.14, p>0.49*; S1 TMS vs dlPFC TMS: *Z = -1.92, p=0.055*]. This suggests that variability in the local CBF effect of TMS is selectively related to variability in the functional properties of the stimulated region.

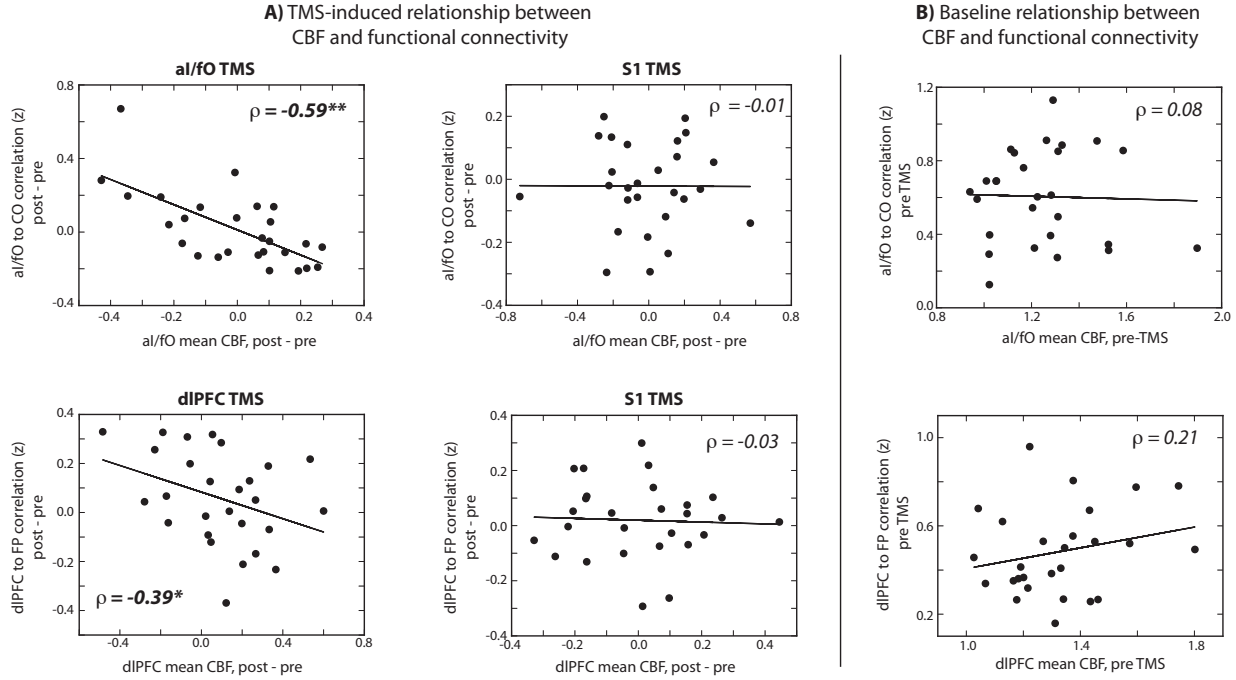

Figure 5.5: *Relationship between CBF and connectivity*. (A) Changes in aI/fO CBF were related to changes in connectivity of aI/fO to the CO network after TMS to aI/fO (*top-left*) but not after TMS to S1 (*top-right*). Similarly, changes in dlPFC CBF were related to changes in connectivity of dlPFC to the FP network after TMS to dlPFC (*bottom-left*) but not after TMS to S1 (*bottom-right*). (B) There was no relationship at baseline between CBF and functional connectivity for either the aI/fO (*top*) or dlPFC (*bottom*) node.

Furthermore, rCBF was not significantly related to the functional connectivity of each node when TMS was not applied. During the pre-TMS period, the rCBF of aI/fO was not related to the functional connectivity of the node to its own network [*aI/fO ROI: r_S = 0.08, p>0.69*]. Similarly, the rCBF of dlPFC was not related to the functional connectivity of the node to its own network before TMS [*dlPFC ROI: r_S = 0.21, p>0.29*] (**Figure 5.5B**).

There was no relationship between the TMS-induced changes in either mean BOLD or CBF standard deviation and functional connectivity in any condition [*mean BOLD: all |r| < 0.22, p>0.27; CBF standard deviation: besides a weak trend in the aI/fO condition (r=-0.36,*

p=0.067), all |r| < 0.14, p>0.49]. Changes in the standard deviation of the BOLD signal, however, were always positively related to changes in functional connectivity, regardless of TMS site [*all r > 0.31, p < 0.12; except for the aI/fO ROI after S1 TMS, all conditions showed a relationship at p<0.05, including the cross conditions, e.g., aI/fO ROI after dlPFC TMS*], suggesting that these correlations do not reflect a TMS-induced effect. This consistent relationship may not be surprising, since these measures are calculated based on the same time series, and in fact may simply reflect an inherent relationship between temporal standard deviation and functional connectivity measures.

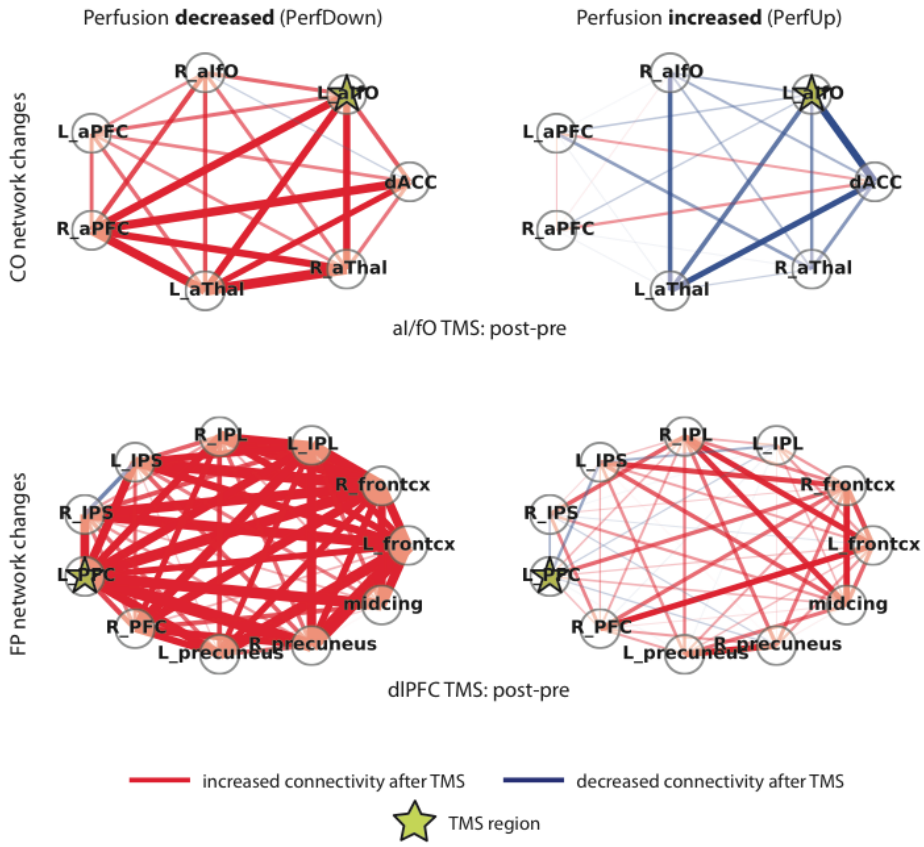

Figure 5.6: *Connectivity in groups selected based on changes in CBF following TMS.* One group contained participants in which CBF decreased under the site of TMS (*PerfDown*, *left column*), and the other contained participants in which CBF increased under the site of TMS (*PerfUp*, *right column*). Changes in functional connectivity within the CO network after aI/fO TMS (*top row*) and within the FP network after dlPFC TMS (*bottom row*) are shown. Colors indicate the direction of change after TMS (*red = increases, blue = decreases*), and the width and transparency of the lines indicate the magnitude of changes (*wider/more opaque = higher magnitude changes*). Stars mark the TMS stimulation locations.

<u>Participant subgroups.</u>

When participants were selected based on whether perfusion increased (*PerfUp* group) or decreased (*PerfDown* group) under the site of TMS, strikingly different effects were seen in the connectivity of the TMS node with all of the other nodes in its functional network (**Figure 5.6**). After aI/fO TMS, the PerfDown group trended toward increased connectivity

between the aI/fO node and the CO network [*t(12) = 1.91, p=0.08*] compared to the pre-TMS period. Instead, the PerfUp group showed significantly decreased connectivity with that node [*t(13) = -2.38, p<0.04*], and these groups showed significantly different changes in within-network functional connectivity from one another [*t(25) = -2.84, p<0.01*]. In addition, these groups showed significantly different changes across the entire CO network [*t(25) = -2.25, p<0.04*]. Similarly, after dlPFC TMS, the PerfDown showed significantly increased connectivity between L dlPFC and the FP network [*t(7) = 3.64, p<0.01*], as well as in general between all nodes across the entire network [*t(7) = 3.93, p<0.01*]. The changes in connectivity between dlPFC and its network (FP) were also significantly different between the PerfUp and PerfDown groups [*t(25) = -2.48, p<0.03*].

Relationship across sessions:

Many different factors may contribute to the variability in TMS effects. This variability may be primarily driven by stable differences in individual subject's susceptibility to TMS or by time-varying factors such as fluctuations in participants' states, previous experiences, and the location of the targeted region (which differed across the three TMS session in our study). If certain participants are simply more susceptible to TMS than others, regardless of TMS site, then we might expect to see a relationship between the changes in perfusion under the coil across different sites. However, no such relationship was observed for any combination of stimulated locations (**Figure 5.7** [*all* $|r_S| < 0.27, p>0.18$]). This suggests that global individual-based variables are not strongly driving the variability seen in the response to TMS.

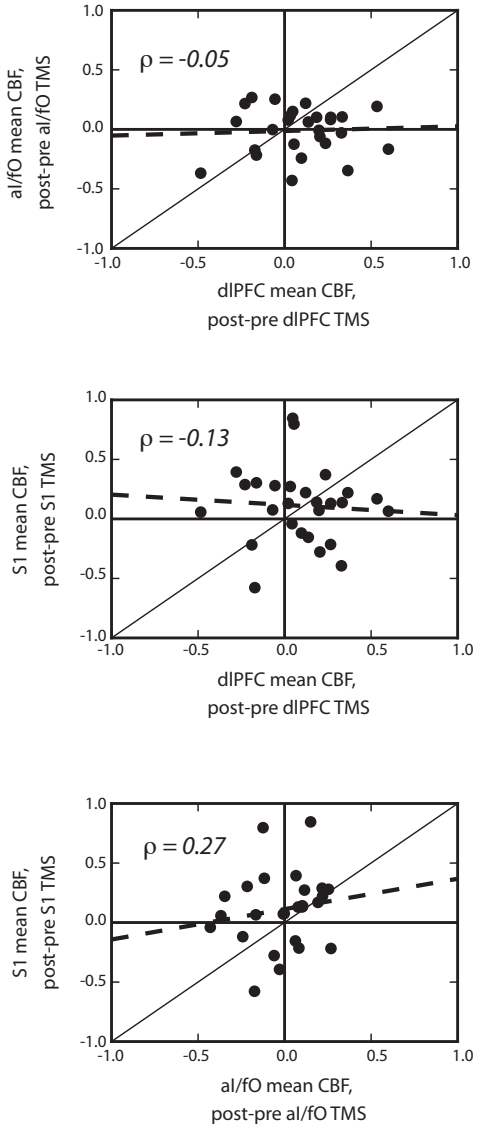

Figure 5.7 Relationship between TMS effects across sites. This figure shows the relationship between changes in CBF under the coil for different TMS sites targeted on separate days, with each individual point representing a single participant. Changes in CBF under the coil were not strongly correlated across different TMS locations (*dlPFC to aI/fO, top: r = -0.06; S1 to dlPFC, middle: r = -0.14; S1 to aI/fO, bottom: r = 0.26*). Dashed lines indicate the linear regression fit, and black diagonal lines indicate equality between the TMS effects on perfusion.

5.5 Discussion

The effects of repetitive TMS can be quite variable across individuals and TMS sites (Ridding and Ziemann, 2010), and it is often not clear how to directly measure this variability, especially when stimulation is applied over non-primary cortical regions during a resting state without corresponding behavioral measures. Here we examined whether fMRI BOLD or rCBF measures reflect consistent, selective, changes under the site of

stimulation and whether these changes were linked to the functional properties of that region. Importantly, we tested three stimulation sites, allowing us to examine the generality of these findings across TMS locations. We found that TMS tends to increase mean rCBF under the site of stimulation in a site-selective manner. Furthermore, these increases were related to changes in the interactions of the targeted region with other nodes in its functional network. BOLD measures also showed changes under the site of stimulation (increased standard deviation, trend for decreased mean), but these were not selective for a particular stimulation site. Moreover, rCBF standard deviation did not change consistently after TMS. Our findings suggest that mean rCBF measures may be particularly useful independent variables for selecting participants from groups that have differing functional effects of repetitive TMS.

5.5.1 Increased CBF under the TMS site

In our study, CBF averaged across individuals increased under the TMS site after both L dlPFC and L S1 stimulation. Changes in mean CBF after TMS to the L aI/fO were much more variable in direction than TMS to the other two regions, and this may be because this region is further from the coil and therefore more difficult to access with TMS (*see Section 5.5.2*). The increased mean CBF after TMS to L dlPFC and L S1 seems potentially counterintuitive, given the substantial physiological and behavioral evidence suggesting that the primary impact of continuous theta burst TMS is inhibitory (Cardenas-Morales et al., 2010; Hoogendam et al., 2010; Thut and Pascual-Leone, 2010). Theta-burst TMS decreases the magnitude of evoked potentials in motor (Huang et al., 2005) and somatosensory (Ishikawa et al., 2007) regions, decreases motor excitability (Di Lazzaro et al., 2005), increases saccade latencies (Nyffeler et al., 2006), and increases phosphene thresholds (Franca et al., 2006). Furthermore, studies of the physiological underpinnings of theta-burst TMS have suggested that it alters inhibitory systems, as assessed with electrophysiological recordings from the spinal cord (Di Lazzaro et al., 2005) and MR spectroscopy measurements (Stagg et al., 2009) in humans, and with electrophysiological recordings and measurement of protein expression in rats (Benali et al., 2011). In addition, some studies indicate that the inhibitory actions of theta-burst TMS may occur via long-term depression of synapses (Huang et al., 2005; Huang et al., 2011), and there is pharmacological evidence for the involvement of an NMDA receptor-dependent mechanism (Huang et al., 2007; Teo et al., 2007).

However, past studies using inhibitory TMS protocols have also shown evidence for increased CBF under the coil [Paired-pulse TMS: (Strafella and Paus, 2001); 1 Hz: (Fox et al., 1997; Knoch et al., 2006; Lee et al., 2003; Rounis et al., 2005); cTBS: (Orosz et al., 2012)]. A number of potential explanations have been put forth for these findings. One possibility is that more neurons are recruited to maintain performance at a constant level after inhibitory TMS to a region (Orosz et al., 2012), but this explanation may be difficult to interpret during a resting-state condition such as the one used in this study, given that there is no explicit behavioral task. Others (Lee et al., 2003) have instead proposed that, while theta-burst TMS may drive an inhibitory decrease in post-synaptic activity, excitatory pre-synaptic activity may be up-regulated in a compensatory fashion, leading to the observed increase in CBF. Finally, it is also possible (as suggested by (Lee et al., 2003) that

the inhibition caused by theta-burst TMS may selectively increase the activity of metabolically demanding inhibitory neurons, which then increases CBF while concomitantly decreasing activity in excitatory neurons (Lauritzen, 2001).

In addition, in our study, the increased CBF in the targeted region was accompanied by increases in CBF of the homologous node in the contralateral hemisphere, especially for the L dlPFC condition. Many past studies have reported changes in CBF in locations far from the coil after TMS at different frequencies (reviewed by (Paus, 2005; Reithler et al., 2011)) including some that have shown increased CBF in sites homologous to stimulation (after 1 Hz TMS: (Knoch et al., 2006); after paired associative stimulation: (Strafella and Paus, 2001)), although the directionality and location of these effects is often variable. Similarly, several papers have examined whether other effects (e.g., motor evoked potentials) are altered in the hemisphere contralateral to TBS stimulation and found a mixture of effects that were similar in the two hemispheres, no effects in the contralateral hemisphere, or opposite effects in the two hemispheres, depending on the exact sites of stimulation and measurement techniques (Di Lazzaro et al., 2011; Ishikawa et al., 2007; Stefan et al., 2008; Suppa et al., 2008). Our results add to this literature, indicating that, for TBS over these target sites during rest, the rCBF of the homologous region changes in the same direction as the stimulated region, suggesting that a similar physiological mechanism may also be at play in the homologous region.

5.5.2 Local changes in CBF after TMS are related to changes in the functional properties of the stimulated region

Following TMS, both BOLD and CBF measures were highly variable across participants in our study. Variability in the effects of TMS have been reported across a wide variety of different stimulation protocols and locations (Gangitano et al., 2002; Hamada et al., 2013; Hamidi et al., 2009; Kozel et al., 2000; Maeda et al., 2000; Martin et al., 2006; Muller-Dahlhaus et al., 2008). Many factors have been proposed to contribute to this variability (Ridding and Ziemann, 2010), and these can broadly be split into 1) static subject-specific variables (e.g., individuals' age (Muller-Dahlhaus et al., 2008), genetic makeup (Cheeran et al., 2008), and variability in the activation of inhibitory interneurons of individuals (Hamada et al., 2013)), and 2) fluctuating, site- or state-specific variables (e.g., previous activity (Gentner et al., 2008), relationship between TMS and underlying electrophysiological fluctuations (Hamidi et al., 2009; Johnson et al., 2010; Johnson et al., 2012), geometry of underlying neural tissue (Wagner et al., 2009)). All of these factors could have contributed to the variability observed in our study. However, the lack of a relationship between the changes in CBF under the coil across different sites suggests that subject-specific variables are unlikely to be a primary determinant of the variability in our study (**Figure 5.7**). This stands in contrast to a recent report that the ability to activate distinct neural populations accounts for a large amount of the variability in TBS effects in primary motor cortex (Hamada et al., 2013); however it is possible that these interactions between TMS and neuronal subpopulations may differ across target sites, in addition to individuals, and therefore may not be expected to correlate across stimulation days.

To determine whether the local variability we measured in CBF was indexing functionally

meaningful changes in the effects of TMS on the brain (rather than measurement variability), we examined whether there was a relationship between CBF variability and changes in functional connectivity. We found a specific association between these measurements that was evident only after TMS. Our findings raise the possibility that TMS may affect individual subjects (and their underlying neural interactions) in qualitatively different fashions, perhaps along a functional continuum. On an individual subject level, rCBF may either increase or decrease after theta-burst TMS. This variability in the effect of TMS on rCBF is coupled with the functional integration of the stimulated region in its own intrinsic network. Interestingly, this relationship was selective to changes induced by TMS: it was not present after TMS to remote locations, nor was there a relationship between rCBF and functional connectivity before TMS. Groups of participants selected based on the direction of local perfusion changes after TMS showed significantly different changes in network connectivity. This suggests that multiple mechanisms (including those proposed in *Section 5.5.1*) may cause the relative balance between inhibitory and excitatory functional changes to vary across individuals after stimulation and that this may further vary depending on the exact stimulated location and state of the participant.

This coupling between changes in rCBF and in functional connectivity was particularly strong after TMS to the L aI/fO, a region that is deeper and more difficult to stimulate than the other target areas. Although this general region has been successfully targeted with TMS in the past (Higo et al., 2011), we found particularly strong variability in this area in the direction of mean CBF effects after TMS (dlPFC and S1 showed a great deal of variability as well, but the average change in mean CBF under the coil was generally positive). This may partly be due to the deep position and variability in the cortical folding (and therefore in the orientations of cortical pyramidal neurons) in the aI/fO region. It is likely that the more superficial extents of this region were the primary areas stimulated by TMS and that the extent and magnitude of activation in these superficial areas differed across participants (e.g., **Figure 5.2b**). This may have caused this node to be successfully targeted in only a portion of individuals or the node to have been targeted with different mechanisms across individuals, depending on the cortical folding and distance between the coil and stimulated region. The relationship between changes in CBF and changes in functional connectivity of this region support this notion.

Further complicating the mechanistic interpretation of the CBF effects after TMS, we found that *increased* CBF was related to relatively *decreased* functional connectivity of the stimulated region. Although functional connectivity tended to non-specifically increase across these networks after TMS ((Gratton et al., submitted-a); **Figure 5.6**), perfusion increases were related to relatively smaller increases in connectivity (or even slightly decreased connectivity, in the case of aI/fO). Alternatively, this could be stated as decreases in perfusion being related to relatively larger increases in functional connectivity. One possibility is that greater decreases in local activity (as indexed by decreases in CBF) promote greater increases in compensatory network activity (as indexed by an increase in resting-state functional connectivity measurements). However, if CBF increases are actually indexing local decreases in output activity (accompanied by increased input activity from inhibitory neurons or compensatory pre-synaptic activity; e.g., (Lee et al., 2003), see *Section 5.5.1*), then relatively decreased functional connectivity may be seen as a

natural consequence of these changes, with both measures indicative of reduced neural processing. Future research utilizing invasive recordings may help to tease apart these different possibilities.

Three past studies have examined the effect of repetitive TMS on resting-state connectivity ((Eldaief et al., 2011; van der Werf et al., 2010; Vercammen et al., 2010), reviewed by (Fox et al., 2012a). Vercammen et al. (2010) and Van der Werf et al. (2010) used 1 Hz stimulation over the temporo-parietal junction or dlPFC, respectively, and compared the effects to sham stimulation conditions. Both studies only found small changes in the functional connectivity of a small set of remote regions (increased functional connectivity between TPJ and the insula (Vercammen et al., 2010) or decreased connectivity between the default mode and the lateral temporal cortex (van der Werf et al., 2010)). Eldaief and colleagues stimulated the inferior parietal lobule (IPL, a component of the default mode network) with either 1 Hz or 20 Hz TMS. They found no consistent changes after 1 Hz TMS but decreased connectivity of IPL with the rest of its network after 20 Hz TMS, even though this type of TMS is thought to have a local excitatory action (Eldaief et al., 2011). Therefore, as a group, these studies do not yield a clear consensus view of how TMS modifies network interactions. This may reflect the large amount of variability in the efficacy of TMS which may have exaggerated influences in resting-state studies, where the lack of behavioral measures means that TMS-induced variability can not be tracked with performance.

In our study, TMS-induced changes in functional connectivity and CBF measurements were associated with each other, but this relationship was not evident prior to TMS. One past study (Liang et al., 2013) has suggested that, at baseline, group-averaged functional connectivity and CBF values are positively related across regions, especially for long-range functional connections. Our study may not have had enough power to detect this baseline relationship (in addition, we did not examine the relationship for group-averaged values across individual voxels). However, in all cases examined in Liang et al. (2013), the relationship between CBF magnitude and functional connectivity values was positive. This is not consistent with the relationship seen between TMS-induced changes in functional connectivity and perfusion measures in our study, which were negatively related. This may suggest that, although local metabolic activity and long-range functional connections are generally in agreement with one another, this relationship may become fundamentally altered by external disruption. Past studies have found that local disruptions from focal brain lesions can also cause systematic changes in the functional interactions of resting state networks (Carter et al., 2010; Gratton et al., 2012; He et al., 2007; Nomura et al., 2010). It is possible that other factors, such as task-based variables or internally driven goals, may differentially alter local metabolism and global functional coupling, selectively modulating this interaction at particular moments in time (Liang et al., 2013).

5.5.3 Changes in other physiological metrics after TMS

In addition to the changes in CBF discussed above, TMS also altered other physiological metrics that we examined in this study, most notably an increased standard deviation of the BOLD signal and a trend toward decreased mean of the BOLD signal. However, unlike for mean CBF, in these cases the physiological changes that were seen were not specific to

the TMS location but instead were equally present across all three measurement sites. This suggests that these effects may be driven by non-specific changes such as generalized increases in alertness or increased discomfort from nerve stimulation after TMS to a particular location. This result demonstrates that including control locations, both for stimulation and for measurement of TMS effects, may be particularly important for dissociating non-specific physiological changes from those driven by the disruption of a specific brain region. Although the inferences from many TMS studies are of this sort (e.g., TMS to location A causes impairment B, implicating location A in that activity), only the minority of TMS studies have compared stimulation across different sites.

5.5.4 Conclusions

Together, our results suggest that rCBF is useful for measuring the local effects of TMS. In particular, we found that individual variability in CBF is selectively related to changes in functional connectivity of the stimulated region. BOLD signal measures, however, either did not show group based changes (mean BOLD) or exhibited effects that were not specific to the site of stimulation (BOLD standard deviation). Finally, we found that variability in TMS effects was not correlated across experimental sessions on different days, suggesting that is variability is not due to stable, individual-specific factors.

Chapter 6

FOCAL BRAIN LESIONS TO CRITICAL LOCATIONS CAUSE WIDESPREAD DISRUPTION OF THE MODULAR ORGANIZATION OF THE BRAIN

6.1 Abstract

Although it is generally assumed that brain damage predominantly affects only the function of the damaged region, here we show that focal damage to critical locations causes disruption of network organization throughout the brain. Using resting state functional magnetic resonance imaging, we assessed whole-brain network structure in patients with focal brain lesions. Only damage to those brain regions important for communication between sub-networks (e.g. "connectors") – but not to those brain regions important for communication within sub-networks (e.g. "hubs") – led to decreases in modularity, a measure of the integrity of network organization. Critically, this network dysfunction extended into the structurally intact hemisphere. Thus, focal brain damage can have a widespread, non-local, impact on brain network organization when there is damage to regions important for the communication between networks. These findings fundamentally revise our understanding of the remote effects of focal brain damage, and may explain numerous puzzling cases of functional deficits that are observed following brain injury.

6.2 Introduction

The brain can be thought of as a complex network composed of functionally separable sets of regions, referred to here as sub-networks or modules, which supports both local processing within and distributed processing across modules. The study of patients with focal brain lesions provides a unique approach for understanding this organization, allowing one to link cognitive functions to specific brain regions (e.g. the inability to recognize faces after damage to the fusiform gyrus implicates this area in face processing). Yet, despite some rough correspondence between the location of a lesion and the resultant cognitive deficits, an explanation for why focal lesions can also lead to non-specific deficits continues to elude neuroscientists and clinicians. Localization of function relies on the assumption that undamaged brain regions continue to function normally, suggesting that deficits are solely attributable to the damaged tissue (Farah, 1994). Arguments against this assumption have pointed to (i) specific cognitive deficits that occur following focal white matter damage (i.e., disconnection syndromes such as alexia without agraphia; (Geschwind, 1965a, b)), (ii) neurophysiological changes remote from a lesion (referred to as diaschisis; (Feeney and Baron, 1986; Price et al., 2001)) and (iii) parallel distributed processing models of complex cognitive functions (Farah, 1994). These observations provide suggestive evidence that cognitive functions are driven not only by local processing but also by interactions between a distributed set of brain regions (Mesulam, 1990).

Despite 150 years of neuropsychological evidence for local and distributed effects of focal brain damage, a mechanistic framework to reconcile these observations is lacking. While recent studies have examined changes in brain activity distant from the area of damage, these have been limited to studying changes within specific brain subsystems (e.g., regions subserving cognitive control (Nomura et al., 2010), attention (Carter et al., 2010; He et al., 2007), somatomotor function (Carter et al., 2010; Gerloff et al., 2006; Grefkes et al., 2008; Mintzopoulos et al., 2009; Sharma et al., 2009) and language (Price et al., 2001; Warren et al., 2009)), because it has proven difficult to precisely characterize the global changes in brain organization that may occur following focal damage. Recently, mathematical tools based on graph theory have emerged as a method to quantify large-scale network properties of the brain (Bullmore and Sporns, 2009; Guye et al., 2010; Sporns, 2010b), to parcellate these networks into modules, and to identify the roles of individual brain regions within this structure (He et al., 2009b; Meunier et al., 2010; Meunier et al., 2009). Initial studies simulating brain damage point to a potential mechanism describing why only some focal lesions may show robust long-distance effects (Alstott et al., 2009; He et al., 2009b; Honey and Sporns, 2008; Sporns et al., 2007; Young et al., 2000). In these analyses, simulated damage to brain regions that connect different network modules (connectors) predicted widespread effects on whole brain communication, whereas damage to areas that connect regions within a module (hubs) predicted more local, non-extensive effects (He et al., 2009b; Honey and Sporns, 2008); see **Figure 6.1** for a schematic depiction of these nodal roles). The extent to which damage affects large-scale brain organization can be measured with Newman's Modularity (Q; (Newman and Girvan, 2004), a comparison between the number of connections within a module to the number of connections between modules. Modularity quantifies the ability of the brain to differentiate into separable sub-networks, and is an essential property found in many complex systems that allows the system to easily evolve, develop, and engage in flexible, dynamic behaviors (Meunier et al., 2010).

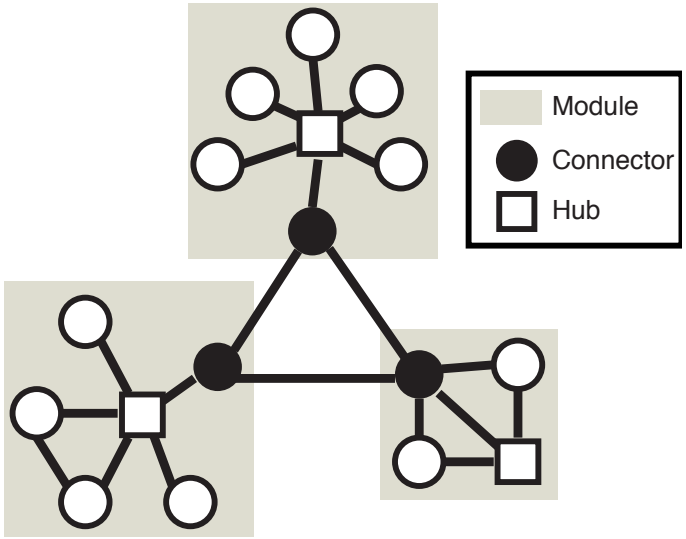

Figure 6.1: *Schematic of nodal roles.* Schematic illustrating the role of connectors (black circles) and hubs (white squares) in modular organization. Hubs have many within-module connections and connectors have

many between-module connections. Note that any brain region may have both connector- and hub-like properties and can have continuous values for hubness and connectorness (see Figure 6.5).

We empirically tested the hypothesis generated by studies using simulated lesion data by collecting resting state functional magnetic resonance imaging (rs-fMRI) data from 35 patients with focal brain lesions and a set of 24 healthy control subjects. Studies of spontaneous coherent fluctuations at rest consistently identify stable (Shehzad et al., 2009) intrinsic functional networks that, in a short fMRI recording session, recapitulate a number of sub-networks normally engaged by a variety of different tasks (Smith et al., 2009). Here, we sought to investigate the differential vulnerability of brain regions with a range of nodal properties within these sub-networks by studying patients with focal damage to a variety of areas distributed throughout the brain (**Figure 6.2**). This enabled us to test how damage to regions with different nodal properties (e.g. hubness or connectorness) influenced large-scale network structure, without regard for potential differences in the specific anatomical site of damage and neuropsychological characteristics of each patient.

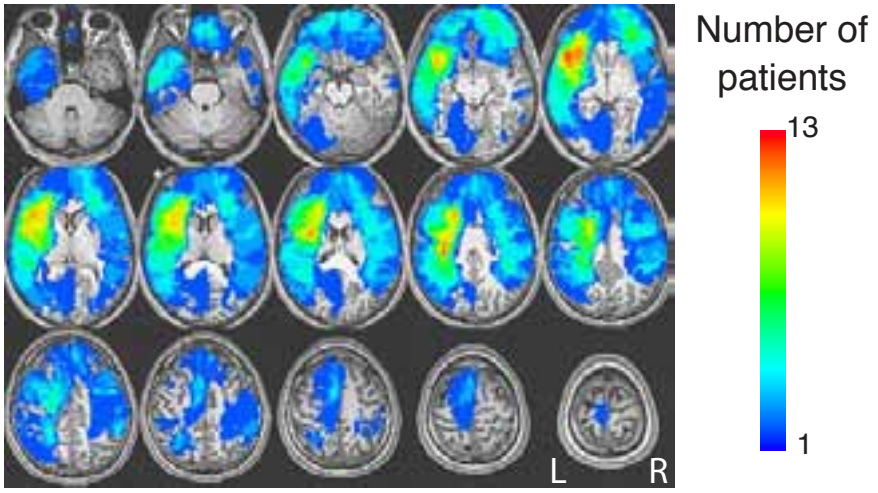

Figure 6.2: *Overlap of individual patient lesions.* An overlap plot of the location of brain damage in 35 patients with focal lesions. The color designates the number of patients with damage to a particular brain region. On this and subsequent figures: R = right hemisphere, L = left hemisphere.

6.3 Materials and Methods

6.3.1 Subjects

Thirty-five patients (age range 17-84, mean age=60) with focal lesions due to stroke (n=25), traumatic brain injury (n=6), and tumors (n=4), and 24 healthy subjects (age range 18-37, mean age = 24) were studied. Four of the 35 lesion patients had bilateral lesions and were excluded from the single hemisphere analyses, leaving 31 unilateral patients (21 left and 10 right). All healthy subjects were pre-screened to exclude individuals with a history of neurologic or psychiatric conditions. Informed consent was obtained from subjects in

accordance with procedures approved by the Committees for Protection of Human Subjects at the University of California, Berkeley.

6.3.2 MRI Acquisition Procedure

T2*-weighted echo planar images (EPI) were collected on a whole body 3-T Siemens MAGNETOM Trio MRI scanner using a 12-channel head coil. Structural images were acquired using an axial MP-RAGE 3-D T1-weighted sequence (TR = 2,300 ms, TE = 2.98 ms, FA = 9°, 1x1x1 mm voxels) for patients and controls. An additional FLAIR image was collected for each patient to better localize the lesion. For patients, 10 minutes of EPI data were analyzed (300 time points, TR = 2,000 ms, TE = 30 ms, 28 3.30 mm thick axial slices). For controls, ten blocks of 4 minutes and 20 seconds of EPI data were analyzed (217 time points, TR = 1,370 ms, TE = 50 ms, 24 3.85 mm thick axial slices). All subjects were instructed to stay awake with their eyes open. No other task instruction was provided.

6.3.3 MRI Pre-processing

Image preprocessing was carried out in AFNI (Cox, 1996). The following pre-statistics processing was applied: removal of non-brain structures from the EPI volumes, and spatial smoothing using a 5 mm Gaussian kernel. Following Fox et al. (Fox et al., 2005), signal from movement, white matter and ventricles was regressed out. The high-resolution T1-weighted image was coregistered to the mean functional data and then segmented using SPM5 (Friston et al., 2007). The template used for segmentation was derived from 152 normal subjects (MNI152; Montreal Neurological Institute, Montreal, QC, Canada) and relied upon probabilistic tissue maps. Spatial transformation parameters to translate between native and standard space were obtained from the brain segmentation and used to normalize each individual's T1-weighted brain. This extra segmentation step was necessary for accurate registration, which is often confounded by structural brain damage.

6.3.4 Lesion Mapping

Lesion masks were manually traced by authors E.N. and C.G. in native patient space according to visible damage on a T1-weighted anatomical scan and guided by damage and hyperintensities on a T2-weighted FLAIR image. All lesion masks were examined by M.D. for anatomical specificity.

6.3.5 Regions of Interest

Ninety cortical and subcortical regions from the Automated Anatomical Labeling (AAL) Atlas (Tzourio-Mazoyer et al., 2002) were reverse normalized into individual subject space to create subject-specific regions of interest (ROIs). ROIs in which EPIs covered less than 25% of voxels within an ROI in any subject within a group were excluded from further analysis, resulting in 90 regions in control subjects and 88 regions in patients.

For single hemisphere analyses, the standard 90 AAL atlas ROIs were divided into two equally-sized halves (superior and inferior) to ensure that a sufficient number of ROIs would be present within each hemisphere for graph theory analyses. Although ROI size was not matched between whole brain and single hemisphere analyses, results were only compared between each single hemisphere or between whole-brain graphs, not across these two classes. As describe, regions with less than 25% EPI coverage within an ROI in

either hemisphere in any single subject in a group were excluded, resulting in 87 regions in controls and 86 regions in lesion patients in each hemisphere.

6.3.6 Functional Connectivity.

Voxel time-series were averaged within each ROI. These averages were bandpass filtered (0.005-0.08 Hz) to remove physiological noise such as cardiac and respiratory artifacts (Fox et al., 2005). Functional connectivity was assessed in each subject by computing time-series correlations between average time-series across all pairs of ROIs, resulting in a correlation matrix for each control subject and lesion patient. Similarly, for single hemisphere analyses, all of the pairs of averaged time-series from ROIs within that hemisphere were used to construct a correlation matrix. Separate correlation matrices were constructed for each hemisphere in each participant.

It should be noted that pair-wise correlation values derived from this type of analysis will only reflect the synchronized, or alternate, activations and deactivations in different regions and will not detect similar time-courses that are phase-shifted (i.e. these will show up as low correlations). Because of the low frequency of the filtered signal, and the nature of the resting state activity, we assumed that such phase-shifted patterns would not signal the presence of healthy interactions.

6.3.7 Graph theory, simulated annealing, and modularity

ROI-by-ROI correlation matrices were derived for each participant as described above and Fisher-transformed to produce normally distributed values. We thresholded each correlation matrix to create an adjacency matrix (a matrix of 1s and 0s reflecting above-threshold correlations and below-threshold correlations respectively), which then served as the basis for defining an unweighted, undirected graph. Graphs are defined as a set of nodes (ROIs from the correlation matrix) connected by a number of edges (correlation values above a threshold). Graphs were created and analyzed using the NetworkX Python package (Hagberg et al., 2008).

The number of edges was equated between individual subject graphs by thresholding at a set cost, rather than correlation, value. A cost value represents the fraction of total possible edges that is present in the graph. By thresholding in this manner, we are able to directly compare the pattern of connections that exists across different graphs, removing differences in the graph structure that would only result from differences in correlation magnitudes across individuals (Achard and Bullmore, 2007; Bartolomei et al., 2006b). All analyses were done over a range of possible cost values (0.05 - 0.25), in order to determine the stability of the results. These values are within the ideal cost range (approximately 0.01 - 0.30) where the sparsity of the graph is considered optimal, since within this range many graph theory metrics, including small-worldness, are maximal (Achard and Bullmore, 2007; Bassett et al., 2008; Bullmore and Bassett, 2011). The main results in the manuscript are presented at a cost of 0.15 or 0.20, points in the center of the ideal cost range (although results were largely consistent across costs).

This procedure of creating graphs from correlation matrices was done for each lesion and control participant separately, by using their individual correlation matrices (control

subject blocks were averaged to create a single matrix), and for the control subjects as a group, by averaging the Fisher-transformed correlation matrices across all control subjects to create a Control Template. At the end of this step, we produced a separate graph for each lesion patient, control participant, and the Control Template at each cost value for the whole brain and for each individual hemisphere.

Each graph created above was partitioned into separate sub-networks, or modules, derived by optimizing Newman's Modularity (Newman and Girvan, 2004) using simulated annealing (Guimera and Nunes Amaral, 2005). Simulated annealing is a data driven method that makes no assumptions about the number of modules to be found in a graph. This method is computationally slow but very robust at finding the optimal modularity structure (Danon et al., 2005; Guimera and Nunes Amaral, 2005); for this size of graphs, the computational cost was not a limiting factor. Modularity is the ratio of the number of within- compared to between- module connections and reflects the strength of a graph's modular organization. Modularity, or *Q*, is defined as:

$$Q = \sum_{i=1}^{m} \left(e_{ii} - a_i^2 \right)$$

where e_{ii} is the fraction of edges that connect two nodes within a module, *i*, a_i is the fraction of edges connecting a node in module *i* to any other node, and *m* is the total number of modules. Modularity will be 1 if all edges fall within a module and will be 0 if there are no more connections within a module than would be expected by chance (Newman and Girvan, 2004).

At each cost, three modularity-optimized partitions were derived for each patient, for each control subject, and for the Control Template (the whole brain network and each hemisphere network independently). These optimized partitions were then evaluated using modularity and compared to one another using a normalized mutual information (MI) measure (Danon et al., 2005). *MI* quantifies the similarity of two partitions and is defined as:

$$MI(A,B) = \frac{-2\sum_{i=1}^{m_A}\sum_{j=1}^{m_B} N_{ij} \log\left(\frac{N_{ij}N}{N_{i.}N_{.j}} \right)}{\sum_{i=1}^{m_A} N_i \log\left(\frac{N_{i.}}{N} \right) + \sum_{j=1}^{m_B} N_{.j} \log\left(\frac{N_{.j}}{N} \right)}$$

where m_A is the number of communities in one partition, *A*, and m_B is the number of communities in a second partition, *B*. This measure is based on defining a confusion $m_A \times m_B$ matrix, *N*, with rows corresponding to modules in partition *A* and columns corresponding to modules in partition *B*. The elements of *N*, N_{ij}, then quantify the number of nodes that were in module *i* of partition *A* and module *j* of partition *B*. *MI* will be 1 if the

two partitions are identical; low values can result either from modules containing different nodes or the presence of different numbers of modules (Danon et al., 2005).

We compared modularity and MI values between control subjects and lesion patients in the whole brain using a two-tailed two-sample t-test and between the two hemispheres in each subject group using a paired-sample t-test. Due to the a-priori hypothesis that lesioned hemispheres should have more disruption than non-lesioned hemispheres, these tests were conducted with a one-tailed test.

Hemisphere effects between patient and control groups were tested with a two-sample t-test after matching the groups for the hemispheres tested (i.e. the same number of left and right hemispheres went into the patient and control subject analyses).

6.3.8 Nodal Roles, damage scores, and correlations with modularity.

From the Control Template partition we calculated metrics reflecting the roles of each node. The participation coefficient (PC) is a measure of the number of *inter-module connections* for each node normalized by their expected value. The PC value for each node *i*, PC_i, is defined as:

$$PC_i = 1 - \sum_{s=1}^{N_M} \left(\frac{k_{is}}{k_i} \right)^2$$

where k_i is the total number of connections to node *i* and k_{is} is the number of connections between node *i* and nodes in module *s*. If a node has connections uniformly distributed to all modules, then its PC value will be 1; if, on the other hand, if its links are concentrated within its own module, its PC value will be 0 (Guimera and Nunes Amaral, 2005).

The within-module degree (WD) is a z-scored measure of the number of *intra-module connections* to each node. The WD value for each node *i*, WD_i, is defined as:

$$WD_i = \frac{k_i - \overline{k}_{s_i}}{\sigma_{k_{s_i}}}$$

where k_i is the number of connections between node *i* and other nodes in its module s_i, $\overline{k_{s_i}}$ is the average degree of all nodes in s_i, and $\sigma_{k_{s_i}}$ is the standard deviation of the degree of all nodes in s_i. This gives a relative (z-scored) measure of how well connected any node is to other nodes within its own module (Guimera and Nunes Amaral, 2005).

Nodes with higher PC values are considered *connectors* and nodes with higher WD values are considered modular *hubs* (sometimes called "provincial" or "connector" hubs depending on their corresponding PC value; (Guimera and Nunes Amaral, 2005). It should be noted, however, that the term "hubs" in the literature has also been applied to nodes scoring high on a variety of different properties (e.g. high on a variety of measures of

degree or centrality, see for example (Bassett et al., 2008; Buckner et al., 2009; Sporns, 2010a)).

In each patient, we used the Control Template nodal roles to calculate separate PC- and WD- damage scores. PC-damage scores were calculated by averaging the PC nodal values from the Control Template for each AAL ROI that was lesioned in that patient and then multiplying that value by the percent damage sustained by the node. Similarly, WD-damage scores were the average of the WD nodal values from the Control Template for each AAL ROI that was lesioned in that patient, weighted by the amount of damage to each node.

To study the relationship between the healthy role of nodes that are damaged and the resultant modular strength in patients we examined the linear relationship between damage scores and modularity using a Pearson's correlation. The significance of correlations at each threshold was tested with a non-parametric permutation test (two-tailed). The significance of the differences between WD and PC correlations was tested using the z-statistic described by Steiger (Steiger, 1980) to compare dependent correlations (two-tailed test).

For individual hemisphere analyses, the mean and standard deviation of individual control subjects was used to normalize the modularity values from each patient based on the hemisphere damaged. Hemispheric values of modularity (m) were normalized in patients to account for potential hemispheric differences using the formula:

$$m_z = \frac{m - \overline{m}_{controls}}{\sigma_{controls}}$$

Here m is the original patient modularity, $\overline{m}_{controls}$ is the average modularity of control subjects in the hemisphere under examination, $\sigma_{controls}$ is the standard deviation in modularity of control subjects in the hemisphere, and m_z is the resultant normalized modularity score. These values were then correlated with the PC- and WD-damage scores as described previously.

6.4 Results

We estimated functional connectivity in resting state fMRI using time-series correlations (Fox et al., 2005) between different brain regions defined from the AAL digital brain atlas in both patients with focal lesions and healthy controls. Functional connectivity graphs are defined as a set of nodes (regions within the AAL atlas) connected by edges (pairs of nodes for which connectivity values passed a threshold, *see Methods*). Graphs from individual patients were compared to a Control Template graph formed from the averaged functional connectivity data of healthy controls.

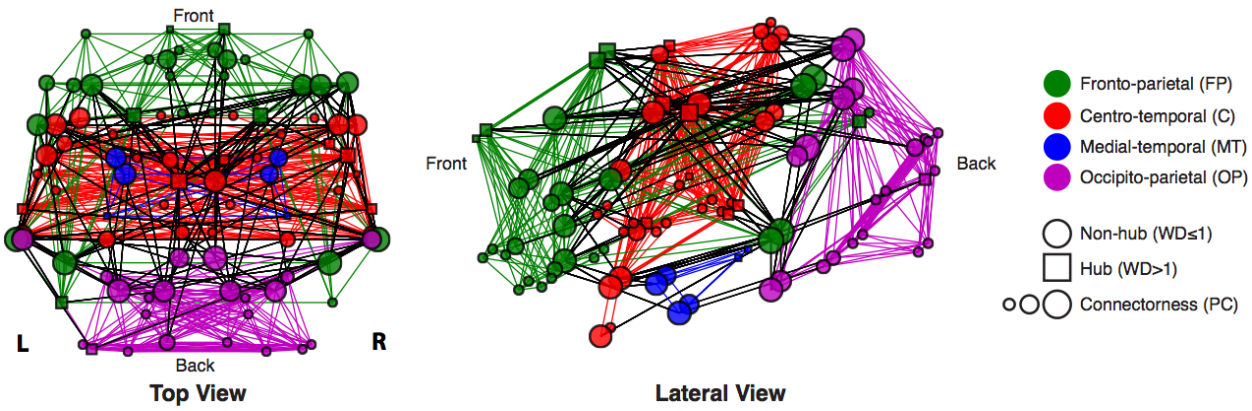

Figure 6.3: *Modular organization of the healthy Control Template.* Top (Left) and Lateral (Right) views of the Control Template graph created from average rs-fMRI data from 24 healthy subjects. Brain regions (nodes) are represented by squares (hubs, WD>1) or circles (non-hubs, WD≤1) and node size represents connectorness (PC). The modularity-optimized partition of the graph resulted in four modules: fronto-parietal (FP), centro-temporal (C), medial-temporal (MT), and occipito-parietal (OP). Within-module edges match module color and between-module edges are black. This partitioning scheme was consistent across thresholds and present in both hemispheres when analyzed separately (see Figure 6.7). R = right hemisphere, L = left hemisphere.

We found four modules in the healthy control network: a fronto-parietal, centro-temporal, occipital-parietal, and medial-temporal module (**Figure 6.3**). This modular organization, which is similar to that observed in other rs-fMRI studies in healthy individuals (He et al., 2009b; Meunier et al., 2009), was consistent across individual subjects. The Mutual Information (a measure of the similarity between two graph partitions that ranges from 0 for unrelated partitions to 1 for exactly the same partition) was 0.59 on average between the Control Template graph and the individual healthy subject graphs, and was stronger in the right hemisphere than the left [$t(23)=-2.15$, $p = 0.04$]. Across patients, damage was distributed across all modules and many patients had damage to more than one module (**Figure 6.4**). No clear relationship between modularity and the pattern of damage across modules was observed. On average, patients had lower modularity values than controls [$t(57)=4.14$, $p = 10^{-4}$]. As expected, the lesioned hemisphere was significantly less modular [$t(30)=1.93$, $p = 0.03$] and was less similar to the Control Template [$t(30)=2.48$, $p = 0.01$] than the non-lesioned hemisphere. However, the non-lesioned hemisphere was also significantly less modular than the corresponding hemisphere in controls [$t(53)=3.75$, $p = 10^{-6}$], suggesting that dysfunction extended into the intact hemisphere in patients.

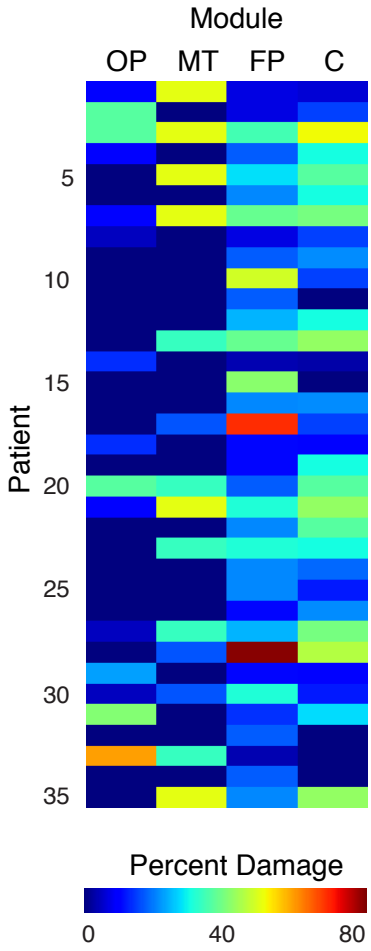

Figure 6.4: *Damage to each module.* Patients had focal brain damage to a heterogeneous set of cortical regions that overlapped with multiple modules. Color indicates percentage of damage to the module and patients are sorted based on their modularity scores, ranging from low (top) to high (bottom) levels of modularity. FP - fronto-parietal, C - Centro-temporal, MT - medial-temporal, OP - occipito-parietal.

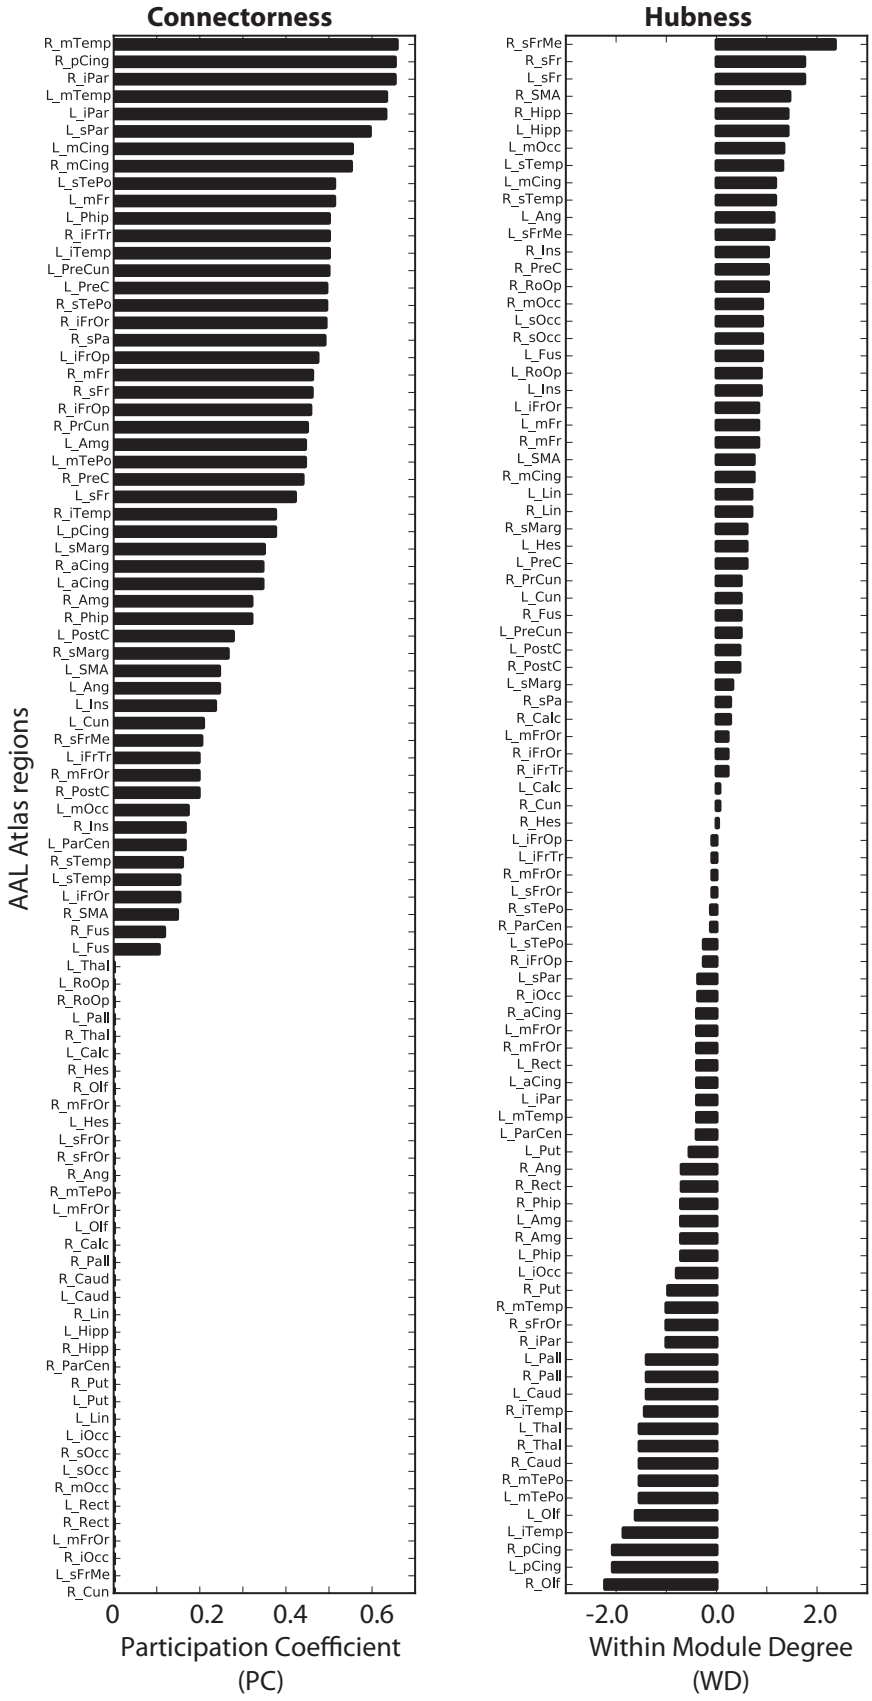
Connectorness
Hubness
AAL Atlas regions
R_mTemp
R_pCing
R_iPar
L_mTemp
L_iPar
L_sPar
L_mCing
R_mCing
L_sTePo
L_mFr
L_Phip
R_iFrTr
L_iTemp
L_PreCun
L_PreC
R_sTePo
R_iFrOr
R_sPa
L_iFrOp
R_mFr
R_sFr
R_iFrOp
R_PrCun
L_Amg
L_mTePo
R_PreC
L_sFr
R_iTemp
L_pCing
L_sMarg
R_aCing
L_aCing
R_Amg
R_Phip
L_PostC
R_sMarg
L_SMA
L_Ang
L_Ins
L_Cun
R_sFrMe
L_iFrTr
R_mFrOr
R_PostC
L_mOcc
R_Ins
L_ParCen
R_sTemp
L_sTemp
L_iFrOr
R_SMA
R_Fus
L_Fus
L_Thal
L_RoOp
R_RoOp
L_Pall
R_Thal
L_Calc
R_Hes
R_Olf
R_mFrOr
L_Hes
L_sFrOr
R_sFrOr
R_Ang
R_mTePo
L_mFrOr
L_Olf
R_Calc
R_Pall
R_Caud
L_Caud
R_Lin
L_Hipp
R_Hipp
R_ParCen
R_Put
L_Put
L_Lin
L_iOcc
R_sOcc
L_sOcc
R_mOcc
L_Rect
R_Rect
L_mFrOr
R_iOcc
L_sFrMe
R_Cun
0
0.2
0.4
0.6
Participation Coefficient
(PC)
R_sFrMe
R_sFr
L_sFr
R_SMA
R_Hipp
L_Hipp
L_mOcc
L_sTemp
L_mCing
R_sTemp
L_Ang
L_sFrMe
R_Ins
R_PreC
R_RoOp
R_mOcc
L_sOcc
R_sOcc
L_Fus
L_RoOp
L_Ins
L_iFrOr
L_mFr
R_mFr
L_SMA
R_mCing
L_Lin
R_Lin
R_sMarg
L_Hes
L_PreC
R_PrCun
L_Cun
R_Fus
L_PreCun
L_PostC
R_PostC
L_sMarg
R_sPa
R_Calc
L_mFrOr
R_iFrOr
R_iFrTr
L_Calc
R_Cun
R_Hes
L_iFrOp
L_iFrTr
R_mFrOr
L_sFrOr
R_sTePo
R_ParCen
L_sTePo
R_iFrOp
L_sPar
R_iOcc
R_aCing
L_mFrOr
R_mFrOr
L_Rect
L_aCing
L_iPar
L_mTemp
L_ParCen
L_Put
R_Ang
R_Rect
R_Phip
L_Amg
R_Amg
L_Phip
L_iOcc
R_Put
R_mTemp
R_sFrOr
R_iPar
L_Pall
R_Pall
L_Caud
R_iTemp
L_Thal
R_Thal
R_Caud
R_mTePo
L_mTePo
L_Olf
L_iTemp
R_pCing
L_pCing
R_Olf
-2.0
0.0
2.0
Within Module Degree
(WD)

Figure 6.5: *Nodal Roles.* Nodes in the Control Template were characterized based on their Participation Coefficient (PC; left), a measure of the number of inter-module connections, and Within-module Degree (WD; right), a measure of the number of intra-module connections. Individual nodes from different modules had a wide range of values, providing a continuum over which patient damage could be assessed. In the figure, AAL region names correspond to abbreviations reported in the AAL atlas (Tzourio-Mazoyer et al., 2002); R = right hemisphere, L = left hemisphere).

For the set of damaged nodes in each patient, we assessed the normal role of those nodes within a healthy brain (defined by the Control Template) based on two characteristics: *within-module degree* (WD) and *participation coefficient* (PC; **Figure 6.5**). WD is a measure of the number of intra-module connections of a node and PC is a measure of the number of inter-module connections of a node (Guimera and Nunes Amaral, 2005). Nodes with high PC values are *connectors* and nodes with high WD values are *hubs* (see **Figure 6.1** for illustration). We created a PC and WD damage score by averaging these healthy PC or WD values from the Control Template for all nodes that were damaged in each patient, weighted by the amount of damage to each node. We then examined the relationship between patients' weighted damage scores and modularity after damage. Patients' PC damage score correlated negatively with modularity (**Figure 6.6**), suggesting that modular organization decreased with increasing amounts of damage to connectors [R=-0.41, p = 0.02, cost = 0.15]. No such relationship, however, was seen between WD damage scores and modularity [R=0.14]. Thus, increasing damage to connectors predicted decreases in modularity significantly more than damage to hubs [Z = 2.37, p = 0.02].

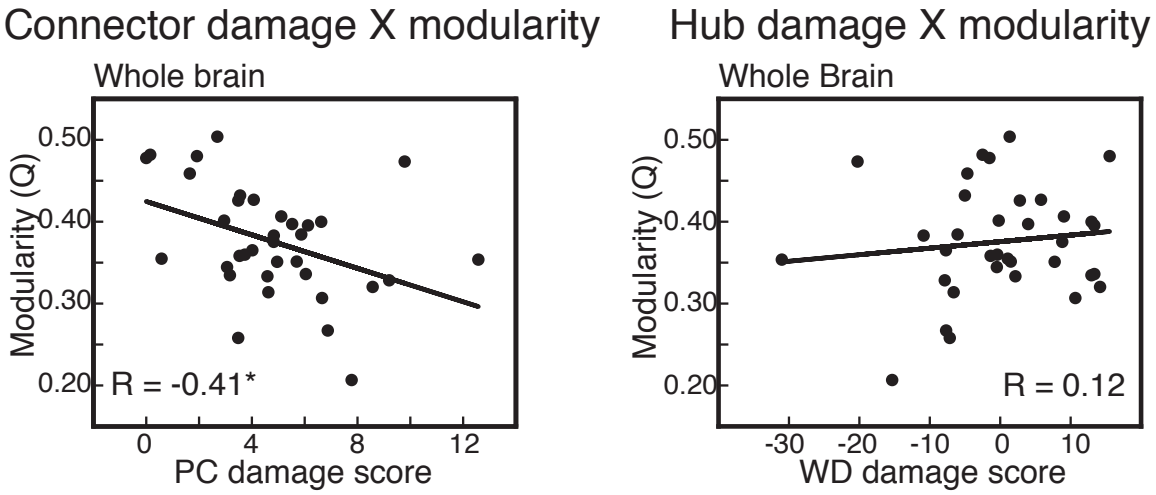

Figure 6.6: *PC and WD damage score correlations with modularity across the whole brain*. (Left) PC damage score was negatively correlated with the modularity of individual subjects, across the whole brain (cost = 0.15, r = -0.41, p = 0.02). (Right) WD damage score was not related to modularity and the correlation between PC and modularity was more negative than the correlation between WD and modularity. These relationships were consistent across a range of cost thresholds.

Mathematically, modularity is inversely proportional to the number of connections between modules (*see Methods*), suggesting that having fewer nodes with high PC values

should lead to increased modularity. We observed this inverse relationship in healthy control subjects, where those subjects who had fewer connector nodes had higher modularity values (R=0.49, p<0.01**,** cost = 0.15). However, in patients with focal lesions we observed that a loss of connectors was related to decreases, rather than increases, in modularity. This indicates that in patients the loss of connectors does not simply cause the loss of a select number of edges, but rather leads to a reorganization of the graph that is detrimental to its modular structure.

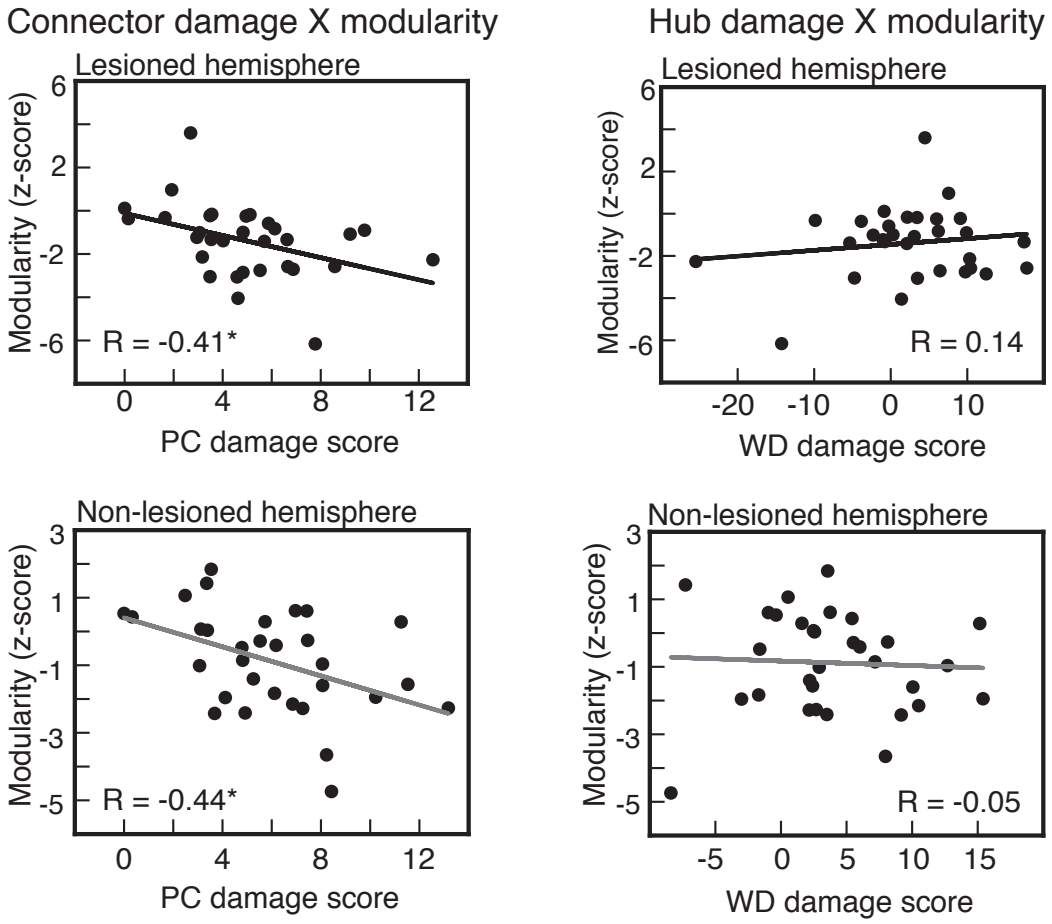

Figure 6.7: *PC and WD damage score correlations with modularity for each hemisphere*. (Left-column) PC damage score was negatively correlated with the modularity of individual subjects in the (Top) lesioned (cost = 0.15, r = -0.41, p = 0.02) and (Bottom) non-lesioned hemisphere (cost = 0.20, r = -0.44, p = 0.01). (Right-column) WD damage score was not related to modularity in either hemisphere. The correlation between PC and modularity was more negative than the correlation between WD and modularity in both hemispheres and these relationships were consistent across a range of cost thresholds.

To determine whether this network reorganization extended to regions remote from the lesion, we examined the relationship between damage and modularity in graphs constructed independently for the lesioned and non-lesioned hemisphere. Given the hemispheric differences in modularity in control subjects, we first normalized the patients' modularity values according to the side on which the lesion occurred to correct for

hemispheric asymmetries (*see Methods*). Similar to the whole brain results, we observed a negative relationship between PC damage score and normalized modularity in the lesioned hemisphere [R=-0.41, p = 0.02]. Critically, this relationship was also present in the non-lesioned hemisphere [R=-0.44, p = 0.01], demonstrating that damage to connectors disrupts network structure even in the intact hemisphere. In both hemispheres, PC damage score was more negatively correlated with modularity than WD damage score (**Figure 6.7**; Lesioned: Z=2.16, p = 0.03; Non-lesioned: Z=1.95, p = 0.05) demonstrating a specific relationship between connector damage and modular organization. Examples of individual patients with low and high connector damage are shown in **Figure 6.8**, clearly emphasizing the global and hemispheric disruptions that result from damage to connectors.

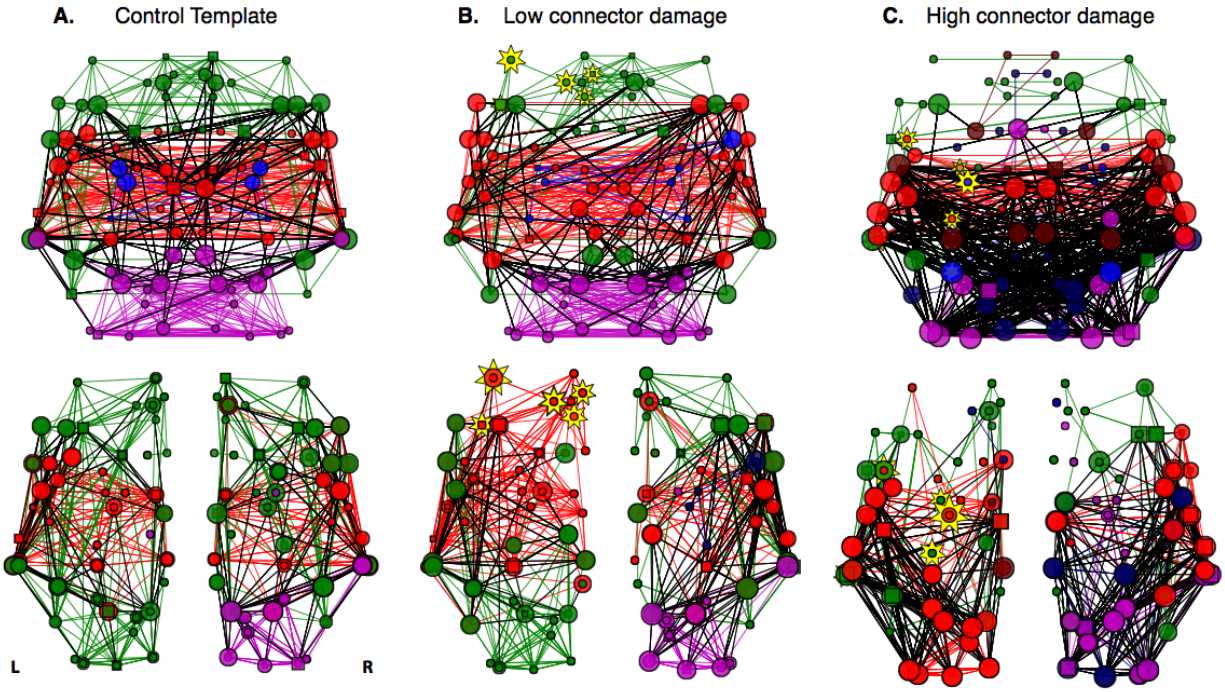

Figure 6.8*: Examples of individual patients with low and high PC damage scores.* (A) Control Template graphs from the average of the healthy control subjects across (Top) the whole brain and (Bottom) each hemisphere separately. (B) Patients with low connector damage tended to have preserved modular structure across the whole brain (top; Q=0.48) and both hemispheres (bottom, lesioned: Q=0.36, non-lesioned: Q=0.41). (C) Patients with high connector damage, however, had highly disrupted modular organization across the whole brain (top; Q=0.21) and both hemispheres (bottom; Q=0.20; non-lesioned Q=0.19). Plotting conventions follow Figure 6.3. Module colors are assigned to match Control Template modules with the highest number of overlapping nodes. Yellow stars represent lesioned nodes, with size of the star proportional to the percent damage to that node (these two patients were approximately matched for the sizes of their lesions).

To ensure that the relationship between brain damage and modularity was specific to the nodal role of the damaged region and the network structure of the lesioned brain, we performed several additional analyses to rule out alternate explanations. First, we determined that the correlations between damage scores and modularity were consistent across graph thresholds (costs), demonstrating that the results were not an artifact of the

particular threshold applied for analyses. The correlation between modularity and PC damage score for whole-brain analyses was r = −.36, −.28, −.41, −.36, −.30, respectively, for costs = 0.05, 0.10, 0.15, 0.20, 0.25 (for the lesioned hemisphere, r = −.14, −.20, −.41, −.36, −.27, and for the nonlesioned hemisphere, r = −.09, −.24, −.33, −.44, −.38). Second, it is possible that changes in modularity are related to the extent of tissue damage, rather than the nodal role of the damaged region. However, we found that correlations between modularity and lesion size were non-significant (R = -0.095, p=0.58 at cost=0.15; not significant at any threshold), suggesting that nodal role determines network disruption. Thirdly, lesions to connectors may simply decrease the average magnitude of connectivity of the brain, rather than modifying the modular structure of those connections. However, the correlation between PC damage score and the average functional connectivity magnitude in each graph was not significant (R = 0.091, p=0.59 at cost=0.15, not significant at any threshold), suggesting that it was the pattern of connections across the brain, rather than their overall strength, that was modified. Finally, we tested whether the individual thresholds applied to achieve the same edge density across patients (*See Methods*) could explain the relationship with connector damage. Again, these correlations were not significant across any cost (R = 0.071, p=0.68 at cost=0.15). Taken together, these analyses strengthen our finding that damage to critical sites represented by connectors results in a disruption in modular organization of the brain.

6.5 Discussion

Here, we analyzed rs-fMRI data from a population of patients with focal brain damage to examine the effects of lesions on the structure of large-scale functional brain networks. This study provides direct evidence that global changes in brain organization occur specifically after focal damage to brain regions that are normally necessary for communication between subnetworks. Importantly, this disruption in hemispheric network organization extends to an intact hemisphere without any structural damage. This demonstrates that connectors (e.g. brain regions that connect different modules), but not hubs (e.g. brain regions that connect regions within a module), are crucial for the maintenance of large-scale modular structure.

Our findings confirm predictions derived from studies simulating the effects of local damage on brain networks (He et al., 2009b; Honey and Sporns, 2008) that specifically implicated connector damage with disruption in measures of network organization. Furthermore, these results provide the first empirical quantification of the large-scale, long-distance, functional effects that can be caused by focal lesions to critical network locations. Although previous studies of patient populations (Carter et al., 2010; Gerloff et al., 2006; Grefkes et al., 2008; He et al., 2007; Nomura et al., 2010; Price et al., 2001; Sharma et al., 2009) have suggested that focal lesions may impact regions remote from the site of structural damage, the focus in these studies was on changes in connectivity magnitude among a small number of brain regions. Studies examining changes in whole-brain functional connectivity networks that occur following traumatic brain injury (Cao and Slobounov, 2010; Castellanos et al., 2010; Nakamura et al., 2009) and tumors (Bartolomei

et al., 2006b; Bosma et al., 2009; Guggisberg et al., 2008) have shown evidence for disruption in these functional networks. In addition, a recent study (Crofts et al., 2011) using diffusion MRI tractography after focal lesions to the basal ganglia and internal capsule found widespread changes in the organization of white matter connections extending into the intact hemisphere, suggesting that structural changes may, in part, underlie the reorganization of large-scale functional networks observed in our study. We were able to expand on this finding both by quantifying the effects of lesions on network structure across the whole brain and by linking the extent of these remote effects to the network role of the damaged region.

A variety of studies have attempted to link diseases with changes in the nodal roles in specific regions (see reviews by (Guye et al., 2010; Pievani et al., 2011; Sporns, 2010b)). For example, global changes in network structure are accompanied by local changes in nodal roles of specific regions in multiple sclerosis (where regional efficiency and correlation strengths of several regions including the insula were correlated with white matter deposition (He et al., 2009a), in traumatic brain injury (where the degree of nodes was decreased in frontal and increased in parietal and occipital regions; (Cao and Slobounov, 2010)), in schizophrenia (where hubs defined by a variety of properties shifted away from frontal and cingulate regions; (Bassett et al., 2008)), and in Alzheimer's disease (where decreases in clustering of the hippocampus was seen in patients (Supekar et al., 2008)).

In a similar approach to ours, studies investigating network structure in healthy controls and those with Alzheimer's disease have found that global changes in large scale network structure of the brain (measured by clustering, path length, and phase-lag index) were predicted by the selective damage to nodes with many connections (high degree nodes) in healthy controls on the basis of modeling results (Stam et al., 2009). In confirmation of the modeling results, another study found that high degree regions in control subjects correspond to the locations of amyloid plaque deposition (Buckner et al., 2009). Thus, although there has been quite a bit of variability in the specific nodal roles and network metrics examined, studies of the network structure in disease have suggested that changes in the properties of specific nodes are linked to widespread network disruption. Our study builds upon these findings, by linking structural damage to specific node types with the magnitude of widespread network disruption.

More broadly, this study adds to a growing literature proposing that function, or dysfunction, of individual brain regions cannot be considered in isolation but rather must be placed within the context of the brain's large-scale network organization (McIntosh, 2000; Mesulam, 1990; Sporns, 2010b). Here we provide evidence for a general principle of brain function where it is the network role, rather than the anatomical location, of a region that determines its importance for organizing modular brain structure and, consequently, the impact that damage to the region will have.

In addition, our findings offer a reinterpretation of perplexing neuropsychological observations present throughout the literature which show that focal lesions can often be accompanied by unexpected or widespread functional deficits that are not predicted from

the local functional properties of the area (Devinsky and D'Esposito, 2001). Furthermore, our findings suggest that an understanding of the network role of the damaged region could be critical for determining a patient's prognosis after brain injury. We predict that damage to hub regions that are central to sub-networks subserving specific cognitive functions will produce specific cognitive deficits, whereas a wider range of functional deficits will result from damage to connector regions that are central to the coordination between multiple sub-networks. Future work will need to examine the extent to which recovery processes reshape the brain's functional networks, potentially through a reassignment of nodal roles to compensate for damage.

Chapter 7

CONCLUSIONS AND FUTURE DIRECTIONS

In this dissertation, I have provided concrete examples of how local regional specialization and network integration interact with one another. This occurred both as a result of performing a complex cognitive task (Section 1, sustained visual attention) and after focal disruptions (Section 2, lesions and TMS). These examples show that network activity can precisely modulate local representations but that changing local properties themselves can also modulate network structure. In particular, the final chapter of this dissertation suggests that interactions between these two levels may be critically dictated by the roles of local regions within this large-scale network structure.

Follow-up projects for Sections 1 and 2 are in progress and will help to further delineate these interactions. For ***Section 1***, I am currently completing analysis on a follow-up project that examines how local changes in the representations of individual faces interact with changes in remote large-scale attention networks that were measured simultaneously. In addition, we have collected data on an analogous attention task with lower-level visual stimuli (Gabor patches at different orientations) to allow the interactions between local and network activity to be directly compared for higher- and lower-level visual processing. A final visual attention project has used a double-blind psychopharmacology experiment to explore how tuning for spatial representations is affected by different neurotransmitter systems in the brain. Preliminary data indicate that cholinergic, but not dopaminergic or noradrenergic, interventions selectively tune spatial representations during a sustained attention task. The data from these follow-up experiments will help better define the network and pharmacological mechanisms that underlie sustained visual attention.

For ***Section 2***, we are currently conducting two follow-up projects. First, we are examining whether acute disruptions from TMS to connector regions also cause rapid modular reorganization or if, like with the network-level project, the changes that result from acute disruption are different from those that result from chronic brain lesions. Second, we are examining whether disruptions lead to systematic compensatory changes in the local roles of regions within large-scale networks, above and beyond the global reorganization seen in Chapter 5.

In addition to these direct follow-ups, this research raises a number of questions for future exploration. First, it would be interesting to characterize the relationship between different measures of local properties (e.g., nodal roles within a network and voxel-based tuning) and how these relationships are affected by the region a network is in or the task a participant is completing. Second, it will be essential to better characterize the dynamic nature of large-scale networks, since most studies that have examined them in detail (including the work here) have assumed that they remain static over a 5-minute recording period. We know that the brain functions on many different time scales, and it is therefore

important to better understand which networks and regions show the most flexibility in their activations, local properties, and network participation over time. Finally, the measures used in this dissertation include the assumption that large-scale networks are spatially separable, but there may be regions that overlap (or directly participate) in more than one network. In the future, novel methods can be employed to define these regions and determine (a) whether they show distinct local properties, (b) whether these vary depending on the cognitive context, and (c) how variability in the amount or regions of overlap predicts inter-individual differences in performance.

References

Achard, S., and Bullmore, E. (2007). Efficiency and cost of economical brain functional networks. PLoS Comput Biol *3*, e17.

Achard, S., Salvador, R., Whitcher, B., Suckling, J., and Bullmore, E. (2006). A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs. J Neurosci *26*, 63-72.

Aguirre, G.K., Detre, J.A., and Wang, J. (2005). Perfusion fMRI for functional neuroimaging. Int Rev Neurobiol *66*, 213-236.

Allen, E.A., Damaraju, E., Plis, S.M., Erhardt, E.B., Eichele, T., and Calhoun, V.D. (2012). Tracking Whole-Brain Connectivity Dynamics in the Resting State. Cerebral cortex.

Alstott, J., Breakspear, M., Hagmann, P., Cammoun, L., and Sporns, O. (2009). Modeling the impact of lesions in the human brain. PLoS Comput Biol *5*, e1000408.

Anderson, J.S., Ferguson, M.A., Lopez-Larson, M., and Yurgelun-Todd, D. (2011). Connectivity gradients between the default mode and attention control networks. Brain Connect *1*, 147-157.

Andrews, T.J., Davies-Thompson, J., Kingstone, A., and Young, A.W. (2010). Internal and external features of the face are represented holistically in face-selective regions of visual cortex. J Neurosci *30*, 3544-3552.

Ashburner, J., and Friston, K.J. (2005). Unified segmentation. NeuroImage *26*, 839-851.

Badre, D., and D'Esposito, M. (2007). Functional magnetic resonance imaging evidence for a hierarchical organization of the prefrontal cortex. Journal of cognitive neuroscience *19*, 2082-2099.

Badre, D., Hoffman, J., Cooney, J.W., and D'Esposito, M. (2009). Hierarchical cognitive control deficits following damage to the human frontal lobe. Nature neuroscience *12*, 515-522.

Bartels, A., Logothetis, N.K., and Moutoussis, K. (2008). fMRI and its interpretations: an illustration on directional selectivity in area V5/MT. Trends Neurosci *31*, 444-453.

Bartolomei, F., Bosma, I., Klein, M., Baayen, J.C., Reijneveld, J.C., Postma, T.J., Heimans, J.J., van Dijk, B.W., de Munck, J.C., de Jongh, A.*, et al.* (2006a). Disturbed functional connectivity in brain tumour patients: evaluation by graph analysis of synchronization matrices. Clin Neurophysiol *117*, 2039-2049.

Bartolomei, F., Bosma, I., Klein, M., Baayen, J.C., Reijneveld, J.C., Postma, T.J., Heimans, J.J., van Dijk, B.W., de Munck, J.C., de Jongh, A.*, et al.* (2006b). How do brain tumors alter functional connectivity? A magnetoencephalography study. Ann Neurol *59*, 128-138.

Bassett, D.S., Bullmore, E., Verchinski, B.A., Mattay, V.S., Weinberger, D.R., and Meyer-Lindenberg, A. (2008). Hierarchical organization of human cortical networks in health and schizophrenia. J Neurosci *28*, 9239-9248.

Beale, J.M., and Keil, F.C. (1995). Categorical effects in the perception of faces. Cognition *57*, 217-239.

Benali, A., Trippe, J., Weiler, E., Mix, A., Petrasch-Parwez, E., Girzalsky, W., Eysel, U.T., Erdmann, R., and Funke, K. (2011). Theta-burst transcranial magnetic stimulation alters cortical inhibition. J Neurosci *31*, 1193-1203.

Bestmann, S., Ruff, C.C., Blankenburg, F., Weiskopf, N., Driver, J., and Rothwell, J.C. (2008). Mapping causal interregional influences with concurrent TMS-fMRI. Experimental brain research Experimentelle Hirnforschung Experimentation cerebrale *191*, 383-402.

Bilek, E., Schafer, A., Ochs, E., Esslinger, C., Zangl, M., Plichta, M.M., Braun, U., Kirsch, P., Schulze, T.G., Rietschel, M.*, et al.* (2013). Application of high-frequency repetitive transcranial magnetic stimulation to the DLPFC alters human prefrontal-hippocampal functional interaction. J Neurosci *33*, 7050-7056.

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

Bland, J.M., and Kerry, S.M. (1998). Statistics notes. Weighted comparison of means. Bmj *316*, 129.

Bonnelle, V., Ham, T.E., Leech, R., Kinnunen, K.M., Mehta, M.A., Greenwood, R.J., and Sharp, D.J. (2012). Salience network integrity predicts default mode network function after traumatic brain injury. Proc Natl Acad Sci U S A *109*, 4690-4695.

Bosma, I., Reijneveld, J.C., Klein, M., Douw, L., van Dijk, B.W., Heimans, J.J., and Stam, C.J. (2009). Disturbed functional brain networks and neurocognitive function in low-grade glioma patients: a graph theoretical analysis of resting-state MEG. Nonlinear Biomed Phys *3*, 9.

Brefczynski, J.A., and DeYoe, E.A. (1999). A physiological correlate of the 'spotlight' of visual attention. Nat Neurosci *2*, 370-374.

Buckner, R.L., Andrews-Hanna, J.R., and Schacter, D.L. (2008). The brain's default network: anatomy, function, and relevance to disease. Ann N Y Acad Sci *1124*, 1-38.

Buckner, R.L., Sepulcre, J., Talukdar, T., Krienen, F.M., Liu, H., Hedden, T., Andrews-Hanna, J.R., Sperling, R.A., and Johnson, K.A. (2009). Cortical hubs revealed by intrinsic functional connectivity: mapping, assessment of stability, and relation to Alzheimer's disease. J Neurosci *29*, 1860-1873.

Bullmore, E., and Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. Nat Rev Neurosci *10*, 186-198.

Bullmore, E.T., and Bassett, D.S. (2011). Brain graphs: graphical models of the human brain connectome. Annu Rev Clin Psychol *7*, 113-140.

Cao, C., and Slobounov, S. (2010). Alteration of cortical functional connectivity as a result of traumatic brain injury revealed by graph theory, ICA, and sLORETA analyses of EEG signals. IEEE Trans Neural Syst Rehabil Eng *18*, 11-19.

Cardenas-Morales, L., Gron, G., and Kammer, T. (2011). Exploring the after-effects of theta burst magnetic stimulation on the human motor cortex: a functional imaging study. Hum Brain Mapp *32*, 1948-1960.

Cardenas-Morales, L., Nowak, D.A., Kammer, T., Wolf, R.C., and Schonfeldt-Lecuona, C. (2010). Mechanisms and applications of theta-burst rTMS on the human motor cortex. Brain Topogr *22*, 294-306.

Carter, A.R., Astafiev, S.V., Lang, C.E., Connor, L.T., Rengachary, J., Strube, M.J., Pope, D.L., Shulman, G.L., and Corbetta, M. (2010). Resting interhemispheric functional magnetic resonance imaging connectivity predicts performance after stroke. Ann Neurol *67*, 365-375.

Castellanos, N.P., Paul, N., Ordonez, V.E., Demuynck, O., Bajo, R., Campo, P., Bilbao, A., Ortiz, T., del-Pozo, F., and Maestu, F. (2010). Reorganization of functional connectivity as a correlate of cognitive recovery in acquired brain injury. Brain *133*, 2365-2381.

Catani, M., Dell'acqua, F., Bizzi, A., Forkel, S.J., Williams, S.C., Simmons, A., Murphy, D.G., and Thiebaut de Schotten, M. (2012). Beyond cortical localization in clinico-anatomical correlation. Cortex *48*, 1262-1287.

Cheeran, B., Talelli, P., Mori, F., Koch, G., Suppa, A., Edwards, M., Houlden, H., Bhatia, K., Greenwood, R., and Rothwell, J.C. (2008). A common polymorphism in the brain-derived neurotrophic factor gene (BDNF) modulates human cortical plasticity and the response to rTMS. J Physiol *586*, 5717-5725.

Chen, A.J., Britton, M., Turner, G.R., Vytlacil, J., Thompson, T.W., and D'Esposito, M. (2012). Goal-directed attention alters the tuning of object-based representations in extrastriate cortex. Front Hum Neurosci *6*, 187.

Chiong, W., Wilson, S.M., D'Esposito, M., Kayser, A.S., Grossman, S.N., Poorzand, P., Seeley, W.W., Miller, B.L., and Rankin, K.P. (2013). The salience network causally influences default mode network activity during moral reasoning. Brain *136*, 1929-1941.

Connor, C.E., Gallant, J.L., Preddie, D.C., and Van Essen, D.C. (1996). Responses in area V4 depend on the spatial relationship between stimulus and attention. J Neurophysiol *75*, 1306-1308.

Corbetta, M., Miezin, F.M., Dobmeyer, S., Shulman, G.L., and Petersen, S.E. (1991). Selective and divided attention during visual discriminations of shape, color, and speed: functional anatomy by positron emission tomography. J Neurosci *11*, 2383-2402.

Corbetta, M., and Shulman, G.L. (2002). Control of goal-directed and stimulus-driven attention in the brain. Nat Rev Neurosci *3*, 201-215.

Cox, R.W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. Computers and biomedical research, an international journal *29*, 162-173.

Crofts, J.J., Higham, D.J., Bosnell, R., Jbabdi, S., Matthews, P.M., Behrens, T.E., and Johansen-Berg, H. (2011). Network analysis detects changes in the contralesional hemisphere following stroke. NeuroImage *54*, 161-169.

D'Esposito, M., Cooney, J.W., Gazzaley, A., Gibbs, S.E., and Postle, B.R. (2006). Is the prefrontal cortex necessary for delay task performance? Evidence from lesion and FMRI data. J Int Neuropsychol Soc *12*, 248-260.

Dai, W., Garcia, D., de Bazelaire, C., and Alsop, D.C. (2008). Continuous flow-driven inversion for arterial spin labeling using pulsed radio frequency and gradient fields. Magn Reson Med *60*, 1488-1497.

Damasio, A.R., Tranel, D., and Damasio, H. (1990). Face agnosia and the neural substrates of memory. Annu Rev Neurosci *13*, 89-109.

Damoiseaux, J.S., Rombouts, S.A., Barkhof, F., Scheltens, P., Stam, C.J., Smith, S.M., and Beckmann, C.F. (2006). Consistent resting-state networks across healthy subjects. Proc Natl Acad Sci U S A *103*, 13848-13853.

Dang, L.C., Donde, A., Madison, C., O'Neil, J.P., and Jagust, W.J. (2012). Striatal dopamine influences the default mode network to affect shifting between object features. Journal of cognitive neuroscience *24*, 1960-1970.

Danon, L., Díaz-Guilera, A., Duch, J., and Arenas, A. (2005). Comparing community structure identification. Journal of Statistical Mechanics: Theory and Experiment *2005*, P09008-P09008.

David, S.V., Hayden, B.Y., Mazer, J.A., and Gallant, J.L. (2008). Attention to stimulus features shifts spectral tuning of V4 neurons during natural vision. Neuron *59*, 509-521.

Dayan, E., Censor, N., Buch, E.R., Sandrini, M., and Cohen, L.G. (2013). Noninvasive brain stimulation: from physiology to network dynamics and back. Nature neuroscience *16*, 838-844.

Desikan, R.S., Segonne, F., Fischl, B., Quinn, B.T., Dickerson, B.C., Blacker, D., Buckner, R.L., Dale, A.M., Maguire, R.P., Hyman, B.T.*, et al.* (2006). An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. NeuroImage *31*, 968-980.

Devinsky, O., and D'Esposito, M. (2001). Neurology of Cognitive and Behavioral Disorders (New York, NY, Oxford University Press).

Di Lazzaro, V., Dileone, M., Pilato, F., Capone, F., Musumeci, G., Ranieri, F., Ricci, V., Bria, P., Di Iorio, R., de Waure, C.*, et al.* (2011). Modulation of motor cortex neuronal networks by rTMS: comparison of local and remote effects of six different protocols of stimulation. J Neurophysiol *105*, 2150-2156.

Di Lazzaro, V., Pilato, F., Saturno, E., Oliviero, A., Dileone, M., Mazzone, P., Insola, A., Tonali, P.A., Ranieri, F., Huang, Y.Z.*, et al.* (2005). Theta-burst repetitive transcranial magnetic stimulation suppresses specific excitatory circuits in the human motor cortex. J Physiol *565*, 945-950.

Dosenbach, N.U., Fair, D.A., Cohen, A.L., Schlaggar, B.L., and Petersen, S.E. (2008). A dual-networks architecture of top-down control. Trends in cognitive sciences *12*, 99-105.

Dosenbach, N.U., Fair, D.A., Miezin, F.M., Cohen, A.L., Wenger, K.K., Dosenbach, R.A., Fox, M.D., Snyder, A.Z., Vincent, J.L., Raichle, M.E.*, et al.* (2007). Distinct brain networks for adaptive and stable task control in humans. Proc Natl Acad Sci U S A *104*, 11073-11078.

Dosenbach, N.U., Visscher, K.M., Palmer, E.D., Miezin, F.M., Wenger, K.K., Kang, H.C., Burgund, E.D., Grimes, A.L., Schlaggar, B.L., and Petersen, S.E. (2006). A core system for the implementation of task sets. Neuron *50*, 799-812.

Driver, J., Blankenburg, F., Bestmann, S., Vanduffel, W., and Ruff, C.C. (2009). Concurrent brain-stimulation and neuroimaging for studies of cognition. Trends in cognitive sciences *13*, 319-327.

Duncan, J., and Owen, A.M. (2000). Common regions of the human frontal lobe recruited by diverse cognitive demands. Trends in neurosciences *23*, 475-483.

Eguiluz, V.M., Chialvo, D.R., Cecchi, G.A., Baliki, M., and Apkarian, A.V. (2005). Scale-free brain functional networks. Phys Rev Lett *94*, 018102.

Eldaief, M.C., Halko, M.A., Buckner, R.L., and Pascual-Leone, A. (2011). Transcranial magnetic stimulation modulates the brain's intrinsic activity in a frequency-dependent manner. Proc Natl Acad Sci U S A *108*, 21229-21234.

Engel, S.A., Rumelhart, D.E., Wandell, B.A., Lee, A.T., Glover, G.H., Chichilnisky, E.J., and Shadlen, M.N. (1994). fMRI of human visual cortex. Nature *369*, 525.

Fair, D.A., Snyder, A.Z., Connor, L.T., Nardos, B., and Corbetta, M. (2009). Task-evoked BOLD responses are normal in areas of diaschisis after stroke. Neurorehabil Neural Repair *23*, 52-57.

Fairhall, S.L., and Ishai, A. (2007). Effective connectivity within the distributed cortical network for face perception. Cereb Cortex *17*, 2400-2406.

Fan, J., Gu, X., Guise, K.G., Liu, X., Fossella, J., Wang, H., and Posner, M.I. (2009). Testing the behavioral interaction and integration of attentional networks. Brain Cogn *70*, 209-220.

Farah, M.J. (1994). Neuropsychological Inference with an Interactive Brain - a Critique of the Locality Assumption. Behav Brain Sci *17*, 43-61.

Feeney, D.M., and Baron, J.C. (1986). Diaschisis. Stroke *17*, 817-830.

Finger, S., Koehler, P.J., and Jagella, C. (2004). The Monakow concept of diaschisis: origins and perspectives. Arch Neurol *61*, 283-288.

Fox, M.D., Halko, M.A., Eldaief, M.C., and Pascual-Leone, A. (2012a). Measuring and manipulating brain connectivity with resting state functional connectivity magnetic resonance imaging (fcMRI) and transcranial magnetic stimulation (TMS). NeuroImage *62*, 2232-2243.

Fox, M.D., Liu, H., and Pascual-Leone, A. (2012b). Identification of reproducible individualized targets for treatment of depression with TMS based on intrinsic connectivity. NeuroImage *66C*, 151-160.

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

Fox, M.D., Snyder, A.Z., Vincent, J.L., Corbetta, M., Van Essen, D.C., and Raichle, M.E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. Proc Natl Acad Sci U S A *102*, 9673-9678.

Fox, P., Ingham, R., George, M.S., Mayberg, H., Ingham, J., Roby, J., Martin, C., and Jerabek, P. (1997). Imaging human intra-cerebral connectivity by PET during TMS. Neuroreport *8*, 2787-2791.

Franca, M., Koch, G., Mochizuki, H., Huang, Y.Z., and Rothwell, J.C. (2006). Effects of theta burst stimulation protocols on phosphene threshold. Clin Neurophysiol *117*, 1808-1813.

Freiwald, W.A., Tsao, D.Y., and Livingstone, M.S. (2009). A face feature space in the macaque temporal lobe. Nat Neurosci *12*, 1187-1196.

Friston, K.J., Ashburner, J., Kiebel, S.J., Nichols, T.E., and Penny, W.D. (2007). Statistical Parametric Mapping: The Analysis of Functional Brain Images (Academic Press).

Friston, K.J., and Price, C.J. (2011). Modules and brain mapping. Cogn Neuropsychol *28*, 241-250.

Fukunaga, M., Horovitz, S.G., van Gelderen, P., de Zwart, J.A., Jansma, J.M., Ikonomidou, V.N., Chu, R., Deckers, R.H., Leopold, D.A., and Duyn, J.H. (2006). Large-amplitude, spatially correlated fluctuations in BOLD fMRI signals during extended rest and early sleep stages. Magn Reson Imaging *24*, 979-992.

Gandhi, S.P., Heeger, D.J., and Boynton, G.M. (1999). Spatial attention affects brain activity in human primary visual cortex. Proc Natl Acad Sci U S A *96*, 3314-3319.

Gangitano, M., Valero-Cabre, A., Tormos, J.M., Mottaghy, F.M., Romero, J.R., and Pascual-Leone, A. (2002). Modulation of input-output curves by low and high frequency repetitive transcranial magnetic stimulation of the motor cortex. Clin Neurophysiol *113*, 1249-1257.

Gazzaley, A., Cooney, J.W., McEvoy, K., Knight, R.T., and D'Esposito, M. (2005). Top-down enhancement and suppression of the magnitude and speed of neural activity. J Cogn Neurosci *17*, 507-517.

Gentner, R., Wankerl, K., Reinsberger, C., Zeller, D., and Classen, J. (2008). Depression of human corticospinal excitability induced by magnetic theta-burst stimulation: evidence of rapid polarity-reversing metaplasticity. Cerebral cortex *18*, 2046-2053.

Gerloff, C., Bushara, K., Sailer, A., Wassermann, E.M., Chen, R., Matsuoka, T., Waldvogel, D., Wittenberg, G.F., Ishii, K., Cohen, L.G.*, et al.* (2006). Multimodal imaging of brain reorganization in motor areas of the contralesional hemisphere of well recovered patients after capsular stroke. Brain *129*, 791-808.

Geschwind, N. (1965a). Disconnexion syndromes in animals and man. I. Brain *88*, 237-294.

Geschwind, N. (1965b). Disconnexion syndromes in animals and man. II. Brain *88*, 585-644.

Gilaie-Dotan, S., and Malach, R. (2007). Sub-exemplar shape tuning in human face-related areas. Cereb Cortex *17*, 325-338.

Gratton, C., Lee, T., Nomura, E.M., and D'Esposito, M. (submitted-a). The effect of theta-burst TMS on cognitive control networks measured with resting state fMRI.

Gratton, C., Lee, T.G., Nomura, E.M., and D'Esposito, M. (submitted-b). Indexing the variability in the functional brain effects of theta burst TMS: a combined BOLD and ASL functional MRI study.

Gratton, C., Nomura, E.M., Perez, F., and D'Esposito, M. (2012). Focal brain lesions to critical locations cause widespread disruption of the modular organization of the brain. J Cognitive Neurosci *24*, 1275-1285.

Gratton, C., Sreenivasan, K.K., Silver, M.A., and D'Esposito, M. (2013). Attention selectively modifies the representation of individual faces in the human brain. J Neurosci *33*, 6979-6989.

Grefkes, C., Nowak, D.A., Eickhoff, S.B., Dafotakis, M., Kust, J., Karbe, H., and Fink, G.R. (2008). Cortical connectivity after subcortical stroke assessed with functional magnetic resonance imaging. Ann Neurol *63*, 236-246.

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

Greicius, M.D., Supekar, K., Menon, V., and Dougherty, R.F. (2009). Resting-state functional connectivity reflects structural connectivity in the default mode network. Cerebral cortex *19*, 72-78.

Gross, C.G. (2011). Face Neurons. In The Oxford Handbook of Face Perception, A.J. Calder, G. Rhodes, M.H. Johnson, and J.V. Haxby, eds. (Oxford, Oxford University Press), pp. 51-76.

Gross, C.G., Rocha-Miranda, C.E., and Bender, D.B. (1972). Visual properties of neurons in inferotemporal cortex of the Macaque. Journal of neurophysiology *35*, 96-111.

Guggisberg, A.G., Honma, S.M., Findlay, A.M., Dalal, S.S., Kirsch, H.E., Berger, M.S., and Nagarajan, S.S. (2008). Mapping functional connectivity in patients with brain lesions. Ann Neurol *63*, 193-203.

Guimera, R., and Nunes Amaral, L.A. (2005). Functional cartography of complex metabolic networks. Nature *433*, 895-900.

Guye, M., Bettus, G., Bartolomei, F., and Cozzone, P.J. (2010). Graph theoretical analysis of structural and functional connectivity MRI in normal and pathological brain networks. Magma *23*, 409-421.

Hagberg, A.A., Schult, D.A., and Swart, P.J. (2008). Exploring network structure, dynamics, and function using NetworkX. In Proceedings of the 7th Python in Science Conference (SciPy2008), G. Varoquaux, T. Vaught, and J. Millman, eds. (Pasadena, CA USA), pp. pp. 11–15.

Hamada, M., Murase, N., Hasan, A., Balaratnam, M., and Rothwell, J.C. (2013). The role of interneuron networks in driving human motor cortical plasticity. Cerebral cortex *23*, 1593-1605.

Hamidi, M., Slagter, H.A., Tononi, G., and Postle, B.R. (2009). Repetitive Transcranial Magnetic Stimulation Affects behavior by Biasing Endogenous Cortical Oscillations. Front Integr Neurosci *3*, 14.

Harmony, T., Marosi, E., Fernandez, T., Bernal, J., Silva, J., Rodriguez, M., Reyes, A., Fernandez-Bouzas, A., Alonso, M., Mendizabal, R., *et al.* (1994). EEG coherences in patients with brain lesions. Int J Neurosci *74*, 203-226.

Harris, A., and Aguirre, G.K. (2008). The representation of parts and wholes in face-selective cortex. J Cogn Neurosci *20*, 863-878.

Haxby, J.V., Gobbini, M.I., Furey, M.L., Ishai, A., Schouten, J.L., and Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. Science *293*, 2425-2430.

Haxby, J.V., Hoffman, E.A., and Gobbini, M.I. (2000). The distributed human neural system for face perception. Trends in cognitive sciences *4*, 223-233.

He, B.J., Snyder, A.Z., Vincent, J.L., Epstein, A., Shulman, G.L., and Corbetta, M. (2007). Breakdown of functional connectivity in frontoparietal networks underlies behavioral deficits in spatial neglect. Neuron *53*, 905-918.

He, Y., Dagher, A., Chen, Z., Charil, A., Zijdenbos, A., Worsley, K., and Evans, A. (2009a). Impaired small-world efficiency in structural cortical networks in multiple sclerosis associated with white matter lesion load. Brain *132*, 3366-3379.

He, Y., Wang, J., Wang, L., Chen, Z.J., Yan, C., Yang, H., Tang, H., Zhu, C., Gong, Q., Zang, Y., *et al.* (2009b). Uncovering intrinsic modular organization of spontaneous brain activity in humans. PloS one *4*, e5226.

Higo, T., Mars, R.B., Boorman, E.D., Buch, E.R., and Rushworth, M.F. (2011). Distributed and causal influence of frontal operculum in task control. Proc Natl Acad Sci U S A *108*, 4230-4235.

Honey, C.J., and Sporns, O. (2008). Dynamical consequences of lesions in cortical networks. Hum Brain Mapp *29*, 802-809.

Honey, C.J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.P., Meuli, R., and Hagmann, P. (2009). Predicting human resting-state functional connectivity from structural connectivity. Proc Natl Acad Sci U S A *106*, 2035-2040.

Hoogendam, J.M., Ramakers, G.M., and Di Lazzaro, V. (2010). Physiology of repetitive transcranial magnetic stimulation of the human brain. Brain Stimul *3*, 95-118.

Huang, Y.Z., Chen, R.S., Rothwell, J.C., and Wen, H.Y. (2007). The after-effect of human theta burst stimulation is NMDA receptor dependent. Clin Neurophysiol *118*, 1028-1032.

Huang, Y.Z., Edwards, M.J., Rounis, E., Bhatia, K.P., and Rothwell, J.C. (2005). Theta burst stimulation of the human motor cortex. Neuron *45*, 201-206.

Huang, Y.Z., Rothwell, J.C., Chen, R.S., Lu, C.S., and Chuang, W.L. (2011). The theoretical model of theta burst form of repetitive transcranial magnetic stimulation. Clin Neurophysiol *122*, 1011-1018.

Hubl, D., Nyffeler, T., Wurtz, P., Chaves, S., Pflugshaupt, T., Luthi, M., von Wartburg, R., Wiest, R., Dierks, T., Strik, W.K., *et al.* (2008). Time course of blood oxygenation level-dependent signal response after theta burst transcranial magnetic stimulation of the frontal eye field. Neuroscience *151*, 921-928.

Inghilleri, M., Conte, A., Curra, A., Frasca, V., Lorenzano, C., and Berardelli, A. (2004). Ovarian hormones and cortical excitability. An rTMS study in humans. Clin Neurophysiol *115*, 1063-1068.

Ishikawa, S., Matsunaga, K., Nakanishi, R., Kawahira, K., Murayama, N., Tsuji, S., Huang, Y.Z., and Rothwell, J.C. (2007). Effect of theta burst stimulation over the human sensorimotor cortex on motor and somatosensory evoked potentials. Clin Neurophysiol *118*, 1033-1043.

Johnson, J.S., Hamidi, M., and Postle, B.R. (2010). Using EEG to explore how rTMS produces its effects on behavior. Brain Topogr *22*, 281-293.

Johnson, J.S., Kundu, B., Casali, A.G., and Postle, B.R. (2012). Task-dependent changes in cortical excitability and effective connectivity: a combined TMS-EEG study. J Neurophysiol *107*, 2383-2392.

Kadosh, K.C., Henson, R.N.A., Kadosh, R.C., Johnson, M.H., and Dick, F. (2010). Task-dependent Activation of Face-sensitive Cortex: An fMRI Adaptation Study. J Cogn Neurosci *22*, 903-917.

Kanwisher, N., McDermott, J., and Chun, M.M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. The Journal of neuroscience : the official journal of the Society for Neuroscience *17*, 4302-4311.

Kay, K.N., Naselaris, T., Prenger, R.J., and Gallant, J.L. (2008). Identifying natural images from human brain activity. Nature *452*, 352-355.

Knoch, D., Treyer, V., Regard, M., Muri, R.M., Buck, A., and Weber, B. (2006). Lateralized and frequency-dependent effects of prefrontal rTMS on regional cerebral blood flow. NeuroImage *31*, 641-648.

Koechlin, E., Ody, C., and Kouneiher, F. (2003). The architecture of cognitive control in the human prefrontal cortex. Science *302*, 1181-1185.

Kozel, F.A., Nahas, Z., deBrux, C., Molloy, M., Lorberbaum, J.P., Bohning, D., Risch, S.C., and George, M.S. (2000). How coil-cortex distance relates to age, motor threshold, and antidepressant response to repetitive transcranial magnetic stimulation. J Neuropsychiatry Clin Neurosci *12*, 376-384.

Kriegeskorte, N., Formisano, E., Sorger, B., and Goebel, R. (2007). Individual faces elicit distinct response patterns in human anterior temporal cortex. Proc Natl Acad Sci U S A *104*, 20600-20605.

Lauritzen, M. (2001). Relationship of spikes, synaptic activity, and local changes of cerebral blood flow. J Cereb Blood Flow Metab *21*, 1367-1383.

Lee, L., Siebner, H.R., Rowe, J.B., Rizzo, V., Rothwell, J.C., Frackowiak, R.S., and Friston, K.J. (2003). Acute remapping within the motor system induced by low-frequency repetitive transcranial magnetic stimulation. J Neurosci *23*, 5308-5318.

Lee, T.G., and D'Esposito, M. (2012). The dynamic nature of top-down signals originating from prefrontal cortex: a combined fMRI-TMS study. J Neurosci *32*, 15458-15466.

Leopold, D.A., Bondar, I.V., and Giese, M.A. (2006). Norm-based face encoding by single neurons in the monkey inferotemporal cortex. Nature *442*, 572-575.

Liang, Z., Li, T., King, J., and Zhang, N. (2013). Mapping thalamocortical networks in rat brain using resting-state functional connectivity. NeuroImage.

Liu, Y., Liang, M., Zhou, Y., He, Y., Hao, Y., Song, M., Yu, C., Liu, H., Liu, Z., and Jiang, T. (2008). Disrupted small-world networks in schizophrenia. Brain *131*, 945-961.

Loffler, G., Yourganov, G., Wilkinson, F., and Wilson, H.R. (2005). fMRI evidence for the neural representation of faces. Nat Neurosci *8*, 1386-1390.

Maeda, F., Keenan, J.P., Tormos, J.M., Topka, H., and Pascual-Leone, A. (2000). Interindividual variability of the modulatory effects of repetitive transcranial magnetic stimulation on cortical excitability. Experimental brain research Experimentelle Hirnforschung Experimentation cerebrale *133*, 425-430.

Manoliu, A., Riedl, V., Zherdin, A., Muhlau, M., Schwerthoffer, D., Scherr, M., Peters, H., Zimmer, C., Forstl, H., Bauml, J.*, et al.* (2013). Aberrant Dependence of Default Mode/Central Executive Network Interactions on Anterior Insular Salience Network Activity in Schizophrenia. Schizophr Bull.

Martin, P.G., Gandevia, S.C., and Taylor, J.L. (2006). Theta burst stimulation does not reliably depress all regions of the human motor cortex. Clin Neurophysiol *117*, 2684-2690.

Martinez-Trujillo, J.C., and Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. Curr Biol *14*, 744-751.

McAdams, C.J., and Maunsell, J.H. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. J Neurosci *19*, 431-441.

McAdams, C.J., and Maunsell, J.H. (2000). Attention to both space and feature modulates neuronal responses in macaque area V4. J Neurophysiol *83*, 1751-1755.

McCarthy, G., Puce, A., Gore, J.C., and Allison, T. (1997). Face-specific processing in the human fusiform gyrus. J Cognitive Neurosci *9*, 605-610.

McIntosh, A.R. (2000). Towards a network theory of cognition. Neural Netw *13*, 861-870.

Menon, V. (2011). Large-scale brain networks and psychopathology: a unifying triple network model. Trends in cognitive sciences *15*, 483-506.

Menon, V., and Uddin, L.Q. (2010). Saliency, switching, attention and control: a network model of insula function. Brain Struct Funct *214*, 655-667.

Mesulam, M.M. (1990). Large-scale neurocognitive networks and distributed processing for attention, language, and memory. Ann Neurol *28*, 597-613.

Meunier, D., Lambiotte, R., and Bullmore, E.T. (2010). Modular and hierarchically modular organization of brain networks. Front Neurosci *4*, 200.

Meunier, D., Lambiotte, R., Fornito, A., Ersche, K.D., and Bullmore, E.T. (2009). Hierarchical modularity in human brain functional networks. Front Neuroinform *3*, 37.

Miller, B.T., Vytlacil, J., Fegen, D., Pradhan, S., and D'Esposito, M. (2011). The prefrontal cortex modulates category selectivity in human extrastriate cortex. Journal of cognitive neuroscience *23*, 1-10.

Minear, M., and Park, D.C. (2004). A lifespan database of adult facial stimuli. Behavior research methods, instruments, & computers : a journal of the Psychonomic Society, Inc *36*, 630-633.

Mintzopoulos, D., Astrakas, L.G., Khanicheh, A., Konstas, A.A., Singhal, A., Moskowitz, M.A., Rosen, B.R., and Tzika, A.A. (2009). Connectivity alterations assessed by combining fMRI and MR-compatible hand robots in chronic stroke. NeuroImage *47 Suppl 2*, T90-97.

Moeller, S., Freiwald, W.A., and Tsao, D.Y. (2008). Patches with links: a unified system for processing faces in the macaque temporal lobe. Science *320*, 1355-1359.

Motter, B.C. (1994). Neural correlates of attentive selection for color or luminance in extrastriate area V4. J Neurosci *14*, 2178-2189.

Muller-Dahlhaus, J.F., Orekhov, Y., Liu, Y., and Ziemann, U. (2008). Interindividual variability and age-dependency of motor cortical plasticity induced by paired associative stimulation. Experimental brain research Experimentelle Hirnforschung Experimentation cerebrale *187*, 467-475.

Mur, M., Ruff, D.A., Bodurka, J., Bandettini, P.A., and Kriegeskorte, N. (2010). Face-identity change activation outside the face system: "release from adaptation" may not always indicate neuronal selectivity. Cereb Cortex *20*, 2027-2042.

Murray, S.O., and Wojciulik, E. (2004). Attention increases neural selectivity in the human lateral occipital complex. Nat Neurosci *7*, 70-74.

Myers, L., and Sirois, M.J. (2004). Spearman Correlation Coefficients, Differences between. Encyclopedia of Statistical Sciences, 1-2.

Nakamura, T., Hillary, F.G., and Biswal, B.B. (2009). Resting network plasticity following brain injury. PloS one *4*, e8220.

Naselaris, T., Prenger, R.J., Kay, K.N., Oliver, M., and Gallant, J.L. (2009). Bayesian reconstruction of natural images from human brain activity. Neuron *63*, 902-915.

Natu, V.S., Jiang, F., Narvekar, A., Keshvari, S., Blanz, V., and O'Toole, A.J. (2010). Dissociable neural patterns of facial identity across changes in viewpoint. J Cogn Neurosci *22*, 1570-1582.

Nestor, A., Plaut, D.C., and Behrmann, M. (2011). Unraveling the distributed neural code of facial identity through spatiotemporal pattern analysis. Proc Natl Acad Sci U S A *108*, 9998-10003.

Newman, M., and Girvan, M. (2004). Finding and evaluating community structure in networks. Physical Review E *69*.

Nomura, E.M., Gratton, C., Visser, R.M., Kayser, A., Perez, F., and D'Esposito, M. (2010). Double dissociation of two cognitive control networks in patients with focal brain lesions. Proc Natl Acad Sci U S A *107*, 12017-12022.

Nyffeler, T., Wurtz, P., Luscher, H.R., Hess, C.W., Senn, W., Pflugshaupt, T., von Wartburg, R., Luthi, M., and Muri, R.M. (2006). Repetitive TMS over the human oculomotor cortex: comparison of 1-Hz and theta burst stimulation. Neuroscience letters *409*, 57-60.

Orosz, A., Jann, K., Wirth, M., Wiest, R., Dierks, T., and Federspiel, A. (2012). Theta burst TMS increases cerebral blood flow in the primary motor cortex during motor performance as assessed by arterial spin labeling (ASL). NeuroImage *61*, 599-605.

Ovadia-Caro, S., Villringer, K., Fiebach, J., Jungehulsing, G.J., van der Meer, E., Margulies, D.S., and Villringer, A. (2013). Longitudinal effects of lesions on functional networks after stroke. J Cereb Blood Flow Metab.

Park, C.H., Chang, W.H., Ohn, S.H., Kim, S.T., Bang, O.Y., Pascual-Leone, A., and Kim, Y.H. (2011). Longitudinal changes of resting-state functional connectivity during motor recovery after stroke. Stroke *42*, 1357-1362.

Pascual-Leone, A., Tarazona, F., and Catala, M.D. (1999). Applications of transcranial magnetic stimulation in studies on motor learning. Electroencephalogr Clin Neurophysiol Suppl *51*, 157-161.

Paus, T. (2005). Inferring causality in brain images: a perturbation approach. Philosophical transactions of the Royal Society of London Series B, Biological sciences *360*, 1109-1114.

Petersen, S.E., Fox, P.T., Posner, M.I., Mintun, M., and Raichle, M.E. (1988). Positron emission tomographic studies of the cortical anatomy of single-word processing. Nature *331*, 585-589.

Pievani, M., de Haan, W., Wu, T., Seeley, W.W., and Frisoni, G.B. (2011). Functional network disruption in the degenerative dementias. Lancet Neurol *10*, 829-843.

Pourtois, G., Schwartz, S., Seghier, M.L., Lazeyras, F., and Vuilleumier, P. (2005). Portraits or people? Distinct representations of face identity in the human visual cortex. J Cogn Neurosci *17*, 1043-1057.

Price, C.J., Warburton, E.A., Moore, C.J., Frackowiak, R.S., and Friston, K.J. (2001). Dynamic diaschisis: anatomically remote and context-sensitive human brain lesions. Journal of cognitive neuroscience *13*, 419-429.

Quiroga, R.Q., Reddy, L., Kreiman, G., Koch, C., and Fried, I. (2005). Invariant visual representation by single neurons in the human brain. Nature *435*, 1102-1107.

R Development Core Team (2011). R: A Language and Environment for Statistical Computing (Vienna, Austria, R Foundation for Statistical Computing).

Raichle, M.E., MacLeod, A.M., Snyder, A.Z., Powers, W.J., Gusnard, D.A., and Shulman, G.L. (2001). A default mode of brain function. Proc Natl Acad Sci U S A *98*, 676-682.

Reddy, L., Kanwisher, N.G., and VanRullen, R. (2009). Attention and biased competition in multi-voxel object representations. Proc Natl Acad Sci U S A *106*, 21447-21452.

Reithler, J., Peters, J.C., and Sack, A.T. (2011). Multimodal transcranial magnetic stimulation: using concurrent neuroimaging to reveal the neural network dynamics of noninvasive brain stimulation. Prog Neurobiol *94*, 149-165.

Ridding, M.C., and Ziemann, U. (2010). Determinants of the induction of cortical plasticity by non-invasive brain stimulation in healthy subjects. J Physiol *588*, 2291-2304.

Rotshtein, P., Henson, R.N., Treves, A., Driver, J., and Dolan, R.J. (2005). Morphing Marilyn into Maggie dissociates physical and identity face representations in the brain. Nat Neurosci *8*, 107-113.

Rounis, E., Lee, L., Siebner, H.R., Rowe, J.B., Friston, K.J., Rothwell, J.C., and Frackowiak, R.S. (2005). Frequency specific changes in regional cerebral blood flow and motor system connectivity following rTMS to the primary motor cortex. NeuroImage *26*, 164-176.

Ruff, C.C., Bestmann, S., Blankenburg, F., Bjoertomt, O., Josephs, O., Weiskopf, N., Deichmann, R., and Driver, J. (2008). Distinct causal influences of parietal versus frontal areas on human visual cortex: evidence from concurrent TMS-fMRI. Cerebral cortex *18*, 817-827.

Ruff, C.C., Blankenburg, F., Bjoertomt, O., Bestmann, S., Freeman, E., Haynes, J.D., Rees, G., Josephs, O., Deichmann, R., and Driver, J. (2006). Concurrent TMS-fMRI and psychophysics reveal frontal influences on human retinotopic visual cortex. Current biology : CB *16*, 1479-1488.

Saad, Z.S., Gotts, S.J., Murphy, K., Chen, G., Jo, H.J., Martin, A., and Cox, R.W. (2012). Trouble at rest: how correlation patterns and group differences become distorted after global signal regression. Brain Connect *2*, 25-32.

Sadaghiani, S., Hesselmann, G., Friston, K.J., and Kleinschmidt, A. (2010). The relation of ongoing brain activity, evoked neural responses, and cognition. Front Syst Neurosci *4*, 20.

Sale, M.V., Ridding, M.C., and Nordstrom, M.A. (2008). Cortisol inhibits neuroplasticity induction in human motor cortex. J Neurosci *28*, 8285-8293.

Saproo, S., and Serences, J.T. (2010). Spatial attention improves the quality of population codes in human visual cortex. J Neurophysiol *104*, 885-895.

Sawamura, H., Orban, G.A., and Vogels, R. (2006). Selectivity of neuronal adaptation does not match response selectivity: a single-cell study of the FMRI adaptation paradigm. Neuron *49*, 307-318.

Serences, J.T., Saproo, S., Scolari, M., Ho, T., and Muftuler, L.T. (2009). Estimating the influence of attention on population codes in human visual cortex using voxel-based tuning functions. NeuroImage *44*, 223-231.

Sharma, N., Baron, J.C., and Rowe, J.B. (2009). Motor imagery after stroke: relating outcome to motor network connectivity. Ann Neurol *66*, 604-616.

Shehzad, Z., Kelly, A.M., Reiss, P.T., Gee, D.G., Gotimer, K., Uddin, L.Q., Lee, S.H., Margulies, D.S., Roy, A.K., Biswal, B.B.*, et al.* (2009). The resting brain: unconstrained yet reliable. Cerebral cortex *19*, 2209-2229.

Silver, M.A., and Kastner, S. (2009). Topographic maps in human frontal and parietal cortex. Trends in cognitive sciences *13*, 488-495.

Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.E., Toro, R., Laird, A.R.*, et al.* (2009). Correspondence of the brain's functional architecture during activation and rest. Proc Natl Acad Sci U S A *106*, 13040-13045.

Somers, D.C., Dale, A.M., Seiffert, A.E., and Tootell, R.B. (1999). Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. Proc Natl Acad Sci U S A *96*, 1663-1668.

Sporns, O. (2010a). Brain Network Disease. In Networks of the brain (Cambridge, MIT Press), pp. 207-231.

Sporns, O. (2010b). Networks of the brain (Cambridge, MIT Press).

Sporns, O., Honey, C.J., and Kotter, R. (2007). Identification and classification of hubs in brain networks. PloS one *2*, e1049.

Spreng, R.N., and Schacter, D.L. (2012). Default network modulation and large-scale network interactivity in healthy young and old adults. Cerebral cortex *22*, 2610-2621.

Spreng, R.N., Stevens, W.D., Chamberlain, J.P., Gilmore, A.W., and Schacter, D.L. (2010). Default network activity, coupled with the frontoparietal control network, supports goal-directed cognition. NeuroImage *53*, 303-317.

Sridharan, D., Levitin, D.J., and Menon, V. (2008). A critical role for the right fronto-insular cortex in switching between central-executive and default-mode networks. Proc Natl Acad Sci U S A *105*, 12569-12574.

Stagg, C.J., Wylezinska, M., Matthews, P.M., Johansen-Berg, H., Jezzard, P., Rothwell, J.C., and Bestmann, S. (2009). Neurochemical effects of theta burst stimulation as assessed by magnetic resonance spectroscopy. J Neurophysiol *101*, 2872-2877.

Stam, C.J., de Haan, W., Daffertshofer, A., Jones, B.F., Manshanden, I., van Cappellen van Walsum, A.M., Montez, T., Verbunt, J.P., de Munck, J.C., van Dijk, B.W*., et al.* (2009). Graph theoretical analysis of magnetoencephalographic functional connectivity in Alzheimer's disease. Brain *132*, 213-224.

Stam, C.J., Jones, B.F., Nolte, G., Breakspear, M., and Scheltens, P. (2007). Small-world networks and functional connectivity in Alzheimer's disease. Cerebral cortex *17*, 92-99.

Stefan, K., Gentner, R., Zeller, D., Dang, S., and Classen, J. (2008). Theta-burst stimulation: remote physiological and local behavioral after-effects. NeuroImage *40*, 265-274.

Steiger, J.H. (1980). Tests for Comparing Elements of a Correlation Matrix. Psychol Bull *87*, 245-251.

Strafella, A.P., and Paus, T. (2001). Cerebral blood-flow changes induced by paired-pulse transcranial magnetic stimulation of the primary motor cortex. J Neurophysiol *85*, 2624-2629.

Supekar, K., Menon, V., Rubin, D., Musen, M., and Greicius, M.D. (2008). Network analysis of intrinsic functional brain connectivity in Alzheimer's disease. PLoS Comput Biol *4*, e1000100.

Suppa, A., Ortu, E., Zafar, N., Deriu, F., Paulus, W., Berardelli, A., and Rothwell, J.C. (2008). Theta burst stimulation induces after-effects on contralateral primary motor cortex excitability in humans. J Physiol *586*, 4489-4500.

Teo, J.T., Swayne, O.B., and Rothwell, J.C. (2007). Further evidence for NMDA-dependence of the after-effects of human theta burst stimulation. Clin Neurophysiol *118*, 1649-1651.

Thut, G., and Pascual-Leone, A. (2010). A review of combined TMS-EEG studies to characterize lasting effects of repetitive TMS and assess their usefulness in cognitive and clinical neuroscience. Brain Topogr *22*, 219-232.

Tootell, R.B., Hadjikhani, N., Hall, E.K., Marrett, S., Vanduffel, W., Vaughan, J.T., and Dale, A.M. (1998). The retinotopy of visual spatial attention. Neuron *21*, 1409-1422.

Treue, S., and Martinez Trujillo, J.C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. Nature *399*, 575-579.

Treue, S., and Maunsell, J.H. (1999). Effects of attention on the processing of motion in macaque middle temporal and medial superior temporal visual cortical areas. J Neurosci *19*, 7591-7602.

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., and Joliot, M. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. NeuroImage *15*, 273-289.

Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. The Quarterly journal of experimental psychology A, Human experimental psychology *43*, 161-204.

van der Werf, Y.D., Sanz-Arigita, E.J., Menning, S., and van den Heuvel, O.A. (2010). Modulating spontaneous brain activity using repetitive transcranial magnetic stimulation. BMC Neurosci *11*, 145.

Van Dijk, K.R., Hedden, T., Venkataraman, A., Evans, K.C., Lazar, S.W., and Buckner, R.L. (2010). Intrinsic functional connectivity as a tool for human connectomics: theory, properties, and optimization. J Neurophysiol *103*, 297-321.

van Meer, M.P., van der Marel, K., Wang, K., Otte, W.M., El Bouazati, S., Roeling, T.A., Viergever, M.A., Berkelbach van der Sprenkel, J.W., and Dijkhuizen, R.M. (2010). Recovery of sensorimotor function after experimental stroke correlates with restoration of resting-state interhemispheric functional connectivity. J Neurosci *30*, 3964-3972.

Vercammen, A., Knegtering, H., Liemburg, E.J., den Boer, J.A., and Aleman, A. (2010). Functional connectivity of the temporo-parietal region in schizophrenia: effects of rTMS treatment of auditory hallucinations. J Psychiatr Res *44*, 725-731.

Von Monakow, C. (1914). Localization in the cerebrum and the degeneration of functions through conical sources (Wiesbaden, Germany, JF Bergmann).

Wagner, T., Rushmore, J., Eden, U., and Valero-Cabre, A. (2009). Biophysical foundations underlying TMS: setting the stage for an effective use of neurostimulation in the cognitive neurosciences. Cortex *45*, 1025-1034.

Wang, J., Zhang, Y., Wolf, R.L., Roc, A.C., Alsop, D.C., and Detre, J.A. (2005). Amplitude-modulated continuous arterial spin-labeling 3.0-T perfusion MR imaging with a single coil: feasibility study. Radiology *235*, 218-228.

Wang, L., Yu, C., Chen, H., Qin, W., He, Y., Fan, F., Zhang, Y., Wang, M., Li, K., Zang, Y., *et al.* (2010). Dynamic functional reorganization of the motor execution network after stroke. Brain *133*, 1224-1238.

Warren, J.E., Crinion, J.T., Lambon Ralph, M.A., and Wise, R.J. (2009). Anterior temporal lobe connectivity correlates with functional outcome after aphasic stroke. Brain *132*, 3428-3442.

Watts, D.J., and Strogatz, S.H. (1998). Collective dynamics of 'small-world' networks. Nature *393*, 440-442.

Winston, J.S., Henson, R.N., Fine-Goulden, M.R., and Dolan, R.J. (2004). fMRI-adaptation reveals dissociable neural representations of identity and expression in face perception. Journal of neurophysiology *92*, 1830-1839.

Wojciulik, E., Kanwisher, N., and Driver, J. (1998). Covert visual attention modulates face-specific activity in the human fusiform gyrus: fMRI study. J Neurophysiol *79*, 1574-1578.

Womelsdorf, T., Anton-Erxleben, K., Pieper, F., and Treue, S. (2006). Dynamic shifts of visual receptive fields in cortical area MT by spatial attention. Nat Neurosci *9*, 1156-1160.

Wu, W.C., Fernandez-Seara, M., Detre, J.A., Wehrli, F.W., and Wang, J. (2007). A theoretical and experimental investigation of the tagging efficiency of pseudocontinuous arterial spin labeling. Magn Reson Med *58*, 1020-1027.

Young, M.P., Hilgetag, C.C., and Scannell, J.W. (2000). On imputing function to structure from the behavioural effects of brain lesions. Philosophical transactions of the Royal Society of London Series B, Biological sciences *355*, 147-161.

Zanto, T.P., Rubens, M.T., Thangavel, A., and Gazzaley, A. (2011). Causal role of the prefrontal cortex in top-down modulation of visual processing and working memory. Nature neuroscience *14*, 656-661.

Zeki, S., Watson, J.D., Lueck, C.J., Friston, K.J., Kennard, C., and Frackowiak, R.S. (1991). A direct demonstration of functional specialization in human visual cortex. J Neurosci *11*, 641-649.

Zhang, D., and Raichle, M.E. (2010). Disease and the brain's dark energy. Nat Rev Neurol *6*, 15-28.

Zhang, H., Tian, J., Liu, J., Li, J., and Lee, K. (2009). Intrinsically organized network for face perception during the resting state. Neuroscience letters *454*, 1-5.