INVESTIGATIONS OF THE GLOBAL NETWORK PROPERTIES OF ATTENTION AND AWARENESS

By

Franklin Douglass Godwin

Dissertation
Submitted to the Faculty of the
Graduate School of Vanderbilt University
in partial fulfillment of the requirements
for the degree of

DOCTOR OF PHILOSOPHY
in

Psychology
May, 2016
Nashville, Tennessee

Approved:

René Marois, Ph.D.
David Zald, Ph.D.
Sean Polyn, Ph.D.
Baxter Rogers, Ph.D.

ACKNOWLEDGEMENTS

This work would not have been possible without the financial support of grants provided to René Marois, John Gore, and the Vanderbilt Vision Research Center. I would not have been able to complete this research without the conversation, insight, and argument provided by the Marois lab, past and present. In particular, I would like to thank the current members of the lab: Benjamin Tamber-Rosenau, Matt Ginther, Anat Fintzi, Hana Eaton, and Lauren Hartsough. Past members of the lab who contributed to this dissertation with data or support, whether they were aware of it or not, include Suk Won Han, Chris Asplund, Justin Martin, and Paul Dux.

I am indebted to the members of my Dissertation Committee and previous committees, whose comments and guidance made this work possible. Of course, I owe the largest portion of my thanks to my advisor and mentor, René Marois. Without his efforts, little of this work would have been accomplished. I will be eternally grateful for his patience and startling attention to detail. The raw fMRI data and task timing used in Chapter IV were graciously shared by John Serences – currently at the University of California, San Diego – and Ed Awh – currently at the University of Chicago.

Throughout my entire graduate career, my family has been an unending source of pride and encouragement. Without them, I would not be where I am today. Lastly, I would like to thank my collaborator and fiancée, Julia. I am sure neither of us expected to be having conversations about graph theory over the dinner table when this journey began, but her insight and companionship have been invaluable in making it to the end.

TABLE OF CONTENTS

Page

LIST OF FIGURES

CHAPTER I

INTRODUCTION

A fundamental characteristic of human information processing is its limited capacity. Despite the myriad sources of information impinging our senses, we are only able to fully process a small portion of that information (Beck, Rees, Frith, & Lavie, 2001; Desimone & Duncan, 1995; Duncan, 1980; Duncan & Humphreys, 1989; Dux & Marois, 2009; Marois & Ivanoff, 2005). Capacity limited attentional mechanisms have been long thought to be the means by which incoming information can be filtered for the behaviorally relevant components (Broadbent, 1981; Neisser & Lazar, 1964; Triesman, 1964). Potentially important information can be identified by either its relevance to currently held goals or by its intrinsic salience (Corbetta & Shulman, 2002; Egeth, Virzi, & Garbart, 1984; Jonides, 1981; Petersen & Posner, 2012; Posner & Petersen, 1990; Yantis & Jonides, 1990). The neurobiological and phenomenological fate of selected information is the subject of a complementary question, "what mechanisms allow for conscious access to the selected information?" In other words, once selected via attentional mechanisms, how do humans become consciously aware of a particular piece of information? Despite William James's (1890) proclamation over a century ago that "everyone knows what attention is," the theoretical and experimental debate over how selective attention biases information and ultimately reaches awareness continues.

Very often, modern views of the neurobiological substrates underlying attention and awareness, as well as many other cognitive functions, take what I will describe as a "network-level" perspective. This network-level viewpoint argues that cognitive functions use a multifocal set of activations in the brain to produce flexible deployment of mental functions (Mesulam, 1990). Distinct collections of activations will underlie distinct cognitive functions, such as language, attention, memory, or executive control (Mesulam, 1990; Posner, Petersen, Fox, & Raichle, 1988). The hypothesized neurobiological bases of selective attention and conscious awareness are two primary examples of cognitive operations with widely popularized network-level (i.e. a limited set of coordinated brain regions) theories (Beck et al., 2001; Corbetta & Shulman, 2002; Corbetta, Patel, & Shulman, 2008; Desimone & Duncan, 1995; Lumer & Rees, 1999; Marois & Ivanoff, 2005; Naghavi & Nyberg, 2005). Ultimately, however, behavior is not an expression of any one cognitive process, but rather the coordination of numerous cognitive and physiological operations. Incoming perceptual information must be widely distributed to each of the various, more specialized processing systems of the brain while also maintaining a subjectively unified conscious experience of the world (Baars, 2005; Baars, 2002). Consequently, it is vital that we examine the whole brain to fully understand the processes like attention and awareness.

In this dissertation, I aim to further explore theories of the neurobiological bases of attention and awareness. In particular, I will contrast predictions made by network-level theories with those made by theories developed in the context of the whole-brain nature of behavior. To preview the conclusions presented here, these cognitive processes often generate diffuse functional changes beyond the brain regions typically identified in imaging activation studies.

Network-Level Neurobiological Theories of Awareness and Attention

How the brain begets conscious awareness has been one of the most fundamental and elusive problems in neuroscience, psychology and philosophy. Correspondingly, this problem has spawned a remarkably large number of theories that differ by the proposed extent of cortical and subcortical changes associated with awareness. Typically, researchers study awareness by contrasting the neural correlates of consciously perceived or correctly reported stimuli with the neural correlates of either masked or sub-threshold presentations. Not surprisingly, due to the heterogeneity in task demands of these studies, there is significant heterogeneity in the proposed neural correlates of conscious (Kim & Blake, 2005; Seth, Dienes, Cleeremans, Overgaard, & Pessoa, 2008). Three general classes of theories describe the range of proposed cortical recruitment necessary to produce the conscious percept of awareness and the ability to access that information: focal, network-level, and global theories. According to 'focal' theories, awareness results from the recruitment of individual cortical regions or highly circumscribed neuronal populations restricted to a single macroscopic region of the brain. Proposals vary as to the location of the minimally necessary activation, with various theories suggesting only local changes in neural activity in either the perceptual substrates (Moutoussis & Zeki, 2002; Tse, Martinez-Conde, Schlegel, & Macknik, 2005; Zeki, 2007) or in higher-level nodes of information processing pathways (Lau & Passingham, 2006) are necessary. Zeki's "micro-consciousness" proposal, for example, argues that awareness of individual visual properties, such as color and motion, requires only recruitment of the appropriate visual circuitry involved in perception. Due to the observation that color could be perceived prior to the perception of motion, he argues that there can not be a single unitary awareness (Arnold, Clifford, & Wenderoth, 2001; Moutoussis & Zeki, 2002; Zeki, 2007).

'Network-level' theories posit that awareness results from the recruitment of attention-related networks in the brain. Many studies comparing the neural correlates of seen and unseen stimuli with functional magnetic resonance imaging (fMRI) have shown recruitment of these fronto-parietal areas (Asplund, Todd, Snyder, & Marois, 2010a; Beck et al., 2001; Lumer & Rees, 1999; Lumer, Friston, & Rees, 1998; Marois & Ivanoff, 2005; Naghavi & Nyberg, 2005; Rees, Kreiman, & Koch, 2002). Numerous paradigms for examining seen and unseen stimuli with fMRI have produced attention-network or similar fronto-parietal activations, including studies of change blindness (Beck et al., 2001), binocular rivalry (Lumer et al., 1998; Lumer & Rees, 1999), the attentional blink (Asplund, Todd, Snyder, & Marois, 2010a; Feinstein, Stein, Castillo, & Paulus, 2004; Kranczioch, Debener, Schwarzbach, Goebel, & Engel, 2005; Marois, Yi, & Chun, 2004), bistable perception (Hahn et al., 2007; Kleinschmidt, Büchel, Zeki, & Frackowiak, 1998; Sterzer, Kleinschmidt, & Rees, 2009), threshold stimuli (Boly et al., 2007; Sadaghiani, Hesselmann, & Kleinschmidt, 2009), and masked stimuli (Dehaene et al., 2001; Haynes, Driver, & Rees, 2005).

Finally, a third class of models has found increased prominence within the last decades. These "global" models reflect an alternative perspective as to what comprises the necessary mechanisms for the percept of conscious awareness. These theories suggest that there are widespread changes in the activation state (Baars, 2005; Baars, 2002; Dehaene & Changeux, 2011; Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006) and functional connectivity (Edelman, 2003; Tononi, 2008; Tononi & Sporns, 2003; van Gaal & Lamme, 2012) of the brain.

A common component between these theories is that without a unified, global recruitment of cortical and subcortical mechanisms, the unified conscious experience does not form.

Whereas there is substantial disagreement as to the neurobiological substrates of awareness, theorized cortical areas controlling the allocation of attention in the brain enjoy a modicum of agreement. Two separable networks control goal-driven and stimulus-driven orienting respectively (Corbetta, 1998; Corbetta & Shulman, 2002; Yantis et al., 2002). The endogenous, or goal-driven, control of attention is attributed to the dorsal attention work, comprised principally of bilateral human frontal eye fields (FEF) and intraparietal sulcus (IPS) regions (Corbetta & Shulman, 2002; Kastner, De Weerd, Desimone, & Ungerleider, 1998; Kastner, Pinsk, De Weerd, Desimone, & Ungerleider, 1999). The dorsal network is thought to be involved in both spatial and non-spatial deployment of attention (Moore, 2003; Schall, 2002; Serences, 2004; K. G. Thompson, Biscoe, & Sato, 2005), and accordingly, electrical stimulation of the monkey FEF produces saccades (Schall, 2002). However, shifts of attention do not require overt eye movements, and stimulation below saccade threshold in the monkey FEF produces perceptual benefits without producing an eye movement (Moore, 2003). Neuroimaging in humans reveals activation in the human FEF region during covert shifts of attention in space as well as recruitment of the superior parietal lobule (Corbetta, 1998; Yantis et al., 2002). However, no spatial shift is necessary to recruit dorsal attention regions and can occur when shifting attention between overlapping objects (Serences, Schwarzbach, Courtney, Golay, & Yantis, 2004a) as well as shifts between categorical rules (Chiu & Yantis, 2009). FEF and IPS are additionally recruited following preparatory cues, reflecting their role in the preparation and execution of goal-directed attention (Corbetta, Kincade, Ollinger, McAvoy, & Shulman, 2000; Hopfinger, Buonocore, & Mangun, 2000; Kastner et al., 1999; G. L. Shulman et al., 1999), a finding confirmed through meta-analysis of both fMRI and PET studies (Wager, Jonides, & Reading, 2004).

In contrast to goal-directed attention, the stimulus-driven orienting of attention is hypothesized to be supported by a ventral network (Corbetta et al., 2008; Corbetta & Shulman, 2002) of regions including the temporo-parietal junction (TPJ), inferior frontal junction (IFJ), and anterior insula (AI) that overlaps with a related 'salience' network in AI and dorsal anterior cingulate cortex (dACC)(V. Menon, 2011; V. Menon & Uddin, 2010). Activity in ventral attention network is generally modulated by detection of unattended, salient events (Corbetta & Shulman, 2002). This attention network also shows sensitivity to capture of attention by targets occurring at unexpected locations (Arrington, Carr, Mayer, & Rao, 2000; Corbetta et al., 2000; Kincade, Abrams, Astafiev, Shulman, & Corbetta, 2005; Vossel, Thiel, & Fink, 2006), contingent capture of attention by stimuli sharing target features at an irrelevant location (Serences et al., 2005), as well as oddball stimuli (Bledowski, Prvulovic, Goebel, Zanella, & Linden, 2004; Braver, Barch, Gray, Molfese, & Snyder, 2001; Marois, Leung, & Gore, 2000; M. Stevens, Calhoun, & Kiehl, 2005). Recruitment of the ventral attention network occurs across multiple sensory modalities (Astafiev, Shulman, & Corbetta, 2006; Downar, Crawley, Mikulis, & Davis, 2000; 2002), does not require spatial shifts of attention (Marois et al., 2000), and will respond to individual features of an object (Downar et al., 2000). Completely task irrelevant, surprising images have also been shown to produce recruitment of the ventral attention network (Asplund, Todd, Snyder, & Marois, 2010a).

However, despite these well-characterized activations for goal-directed and stimulus-driven attention, both dorsal and ventral network regions will often co-activate with a larger set of regions thought to embody more generalized cognitive control networks (Duncan, 2010; M.

Fox, Corbetta, Snyder, Vincent, & Raichle, 2006). These coactivations have been referred to as a "multiple demand network" (Duncan, 2006; Duncan & Owen, 2000), "task activation ensemble" (Seeley et al., 2007),"cognitive control network" (Cole & Schneider, 2007; Cole, Pathak, & Schneider, 2010), "fronto-parietal control system" (Vincent, Kahn, Snyder, Raichle, & Buckner, 2008), and "task control network" (Dosenbach, Fair, Cohen, Schlaggar, & Petersen, 2008; Dosenbach et al., 2007; Petersen & Posner, 2012). Together, these hypothesized networks cross wide swaths of cortex, extending from posterior parietal and extrastriate cortex to anterior and lateral prefrontal regions in addition to portions of cortex within the medial walls of the brain. These "task-positive" regions (M. D. Fox et al., 2005) co-activate across a wide-ranging set of cognitively demanding tasks. The various nomenclatures refer to similar anatomical constituents but often offer different explanations of network function. For example, there is dispute as to whether these activations constitute a single network or multiple sub-networks. One theorized subdivision within the task-positive set can be placed between areas responsible for coordinating cognitive control, a broad class of functions responsible for the flexible adaptation of information processing to current goals and tasks (Cole & Schneider, 2007; Dosenbach et al., 2008; Duncan, 2006; 2010; Duncan & Owen, 2000; Uddin et al., 2010), and those responsible for the orienting of attention (Corbetta & Shulman, 2002). Each of these theories hypothesizes specialized functions for individual regions and do not necessarily preclude existence of sub-specialized connectivity within the broader network (Duncan, 2010; Hon, Epstein, Owen, & Duncan, 2006; R. Thompson & Duncan, 2009) predicted by other theories (Dosenbach et al., 2007; 2008; Seeley et al., 2007).

Utilizing Graph Theory

One thread of neuroimaging research that has reinforced network-level views of brain function comes from connectivity literature. Connectivity approaches have been one of the fastest growing segments of fMRI research over the last decade (Friston, 2011b) and have contributed significantly to our understanding of network function. For example, the discovery that many cortical areas display coordinated, spontaneous activity while not engaged in any overt task cements this notion that common activity is distributed between numerous regions (i.e. network-level), rather than being localized to circumscribed, single regions (M. D. Fox et al., 2005). Additionally, the entire global class of theories of awareness is predicated on the notion that widespread connectivity contributes to the formation of the conscious percept (Baars, 2005; Dehaene, Kerszberg, & Changeux, 1998; Tononi & Edelman, 1998; Tononi, SPORNS, & Edelman, 1994). There are numerous methodologies with which to study the connections between brain regions including analysis of the structural connections as well as functionally derived measures (Friston, 2011b; Friston, Frith, Liddle, & Frackowiak, 1993). These methodologies provide a measure of some pairwise connection, be it structural, correlational, or effective between two brain locations.

The connectivity literature presents a complex view of these networks' function. For example, both the attention and cognitive control networks described above all show resting state connectivity amongst constituent regions (Cole, Bassett, Power, Braver, & Petersen, 2014; Dosenbach et al., 2007; M. D. Fox et al., 2005; M. Fox et al., 2006; He, Chen, & Evans, 2007), but this finding is not unanimous and differential maps may be found depending on where they were seeded from (Power et al., 2011; Vincent et al., 2008). Examinations of the structural relationships between regions can provide some sense as to the anatomical and therefore potentially functional similarity between regions that may not be correlated at rest (de Schotten et al., 2011). Task-concurrent functional connectivity studies reveal that task demands can potentially cause individual regions to switch network membership (Asplund, Todd, Snyder, & Marois, 2010a; Weissman & Prado, 2012) and effective measures yield results highlighting the directional nature of the some of these connections (Friston, Moran, & Seth, 2013; Rogers, Morgan, Newton, & Gore, 2007). These are only a small fraction of the studies highlighting the complex relationships between areas that serve similar cognitive functions. The many-to-one mapping of function (e.g. the many cognitive control theories described above) to cortical region complicates the notion that individual networks may ultimately embody the necessary neural mechanisms to produce particular cognitive processes and behavior.

Graph theoretical and network scientific approaches to research represent a shift towards the importance of looking at not just individual regions, not just networks, but hierarchically organized, dynamic relationships between networks (Park & Friston, 2013). Graph theory can use structural, functional, or effective connectivity to compliment studies of functional activations while expanding on the notion that coordinated activity between networks may be particularly important making connecting links between behavior and brain function (Bullmore & Sporns, 2009; Sporns, 2011). At its simplest, graph theory is a branch mathematics that deals with the pair-wise relationships between components of a network. However, as a tool for describing the interactions between connected parts, it is incredibly flexible and generalizable. Its use in the context of describing neurobiological networks has grown in popularity within the last decade (Bullmore & Sporns, 2009). In graph theoretical parlance, components of a system are referred to as “nodes” and the connections between them as “edges.” From the perspective of cognitive neuroscience, node definition depends on the methodology being used. For MRI analyses, individual voxels or regions of interest (ROIs) typically constitute the nodes of the neural system. The connections (edges) between them can be measured using any pair-wise measure of similarity (e.g. the Pearson correlation).

Remarkably, many of the same network structures observed in other self-organizing systems also appear in the brain (Watts & Strogatz, 1998). These networks, dubbed "small-world," display efficient patterns of information transfer by routing many connections through a few, highly interconnected hub regions (Sporns, Honey, & Kötter, 2007; Watts & Strogatz, 1998). Watts & Strogatz (1998) show that the United States power grid, the neural network of *C. elegans*, and a network of actors common to the same movies all display similar small-world properties. A core feature of this organization is that it appears to provide a balance for two driving features of networked systems: specialization and integration (Bullmore & Sporns, 2012; Cole et al., 2014; Power et al., 2011). Like the brain, these systems require sub-components to perform specialized computation that, when integrated with other sub-components, produced an emergent or global behavior. Ultimately, graph theoretical research provides novel utility towards the study of neurobiological systems by offering a language and analytical methods with which describe both the components and holistic structure of a network.

The individual components, nodes and edges, of these networks can be described in how they interact amongst themselves to form small, highly interconnected communities, or modules (Fig. 1). Questions about the connections made by an individual node provide a low-level entry point of analysis. The extent to which a node displays high "degree", or many connections, provides a perspective on its "centrality", or usefulness, for efficient communication (Bullmore & Sporns, 2009; 2012). A node that lies on the path between many other nodes, like a central terminal of a train route, can be described as a “hub” with vital importance for integrating the specialized functions of individual sub-networks (Bullmore & Sporns, 2009; van den Heuvel & Sporns, 2013). Analysis of the “modular” structure of the overall network examines connectivity in a all-inclusive manner, taking into account all pairwise connections between regions to generate a plausible structural or functional organization(Sporns & Betzel, 2015). The application to neuroscientific research becomes obvious due to the frequency with which the brain, at many spatial scales, is described as functioning like a network. Questions about the role of an individual cortical region and its integrative properties can be described in terms of its mathematically derived centrality and hub properties. The emergent structure of networks and how they communicate with one another can be assessed by the modularity and small-world properties they display.

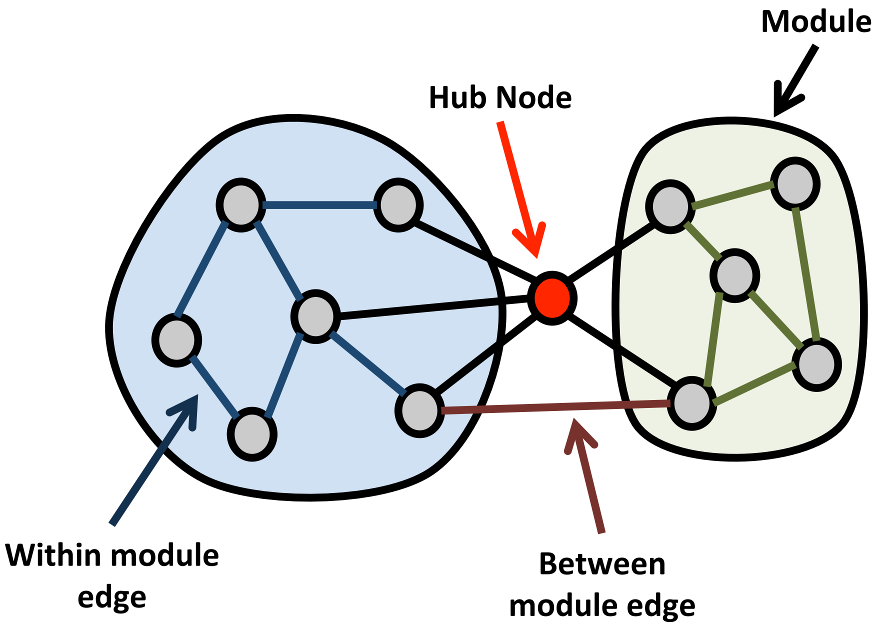

Figure 1. Schematic of basic graph theoretical structures. Nodes (circles) and edges (connecting lines) represent the basic structure of the network, potentially embodying individual regions of the brain and the functional or structural connections between them. Highly interconnected clusters of nodes form modules whereas a highly connected individual node acts as a hub.

The last decade has seen an explosion of studies focusing on the connectivity properties of the brain above and beyond just the activation patterns observed with non-invasive methods (Bullmore & Sporns, 2009; Friston, 2011a; Sporns & Betzel, 2015). Following on the heels of these advances in the study of functional connections, network science and graph theory has commensurately exploded in neuroscience. The graph theoretical perspective has begun to address connectivity differences in clinical populations (Ahmadlou, Rostami, & Sadeghi, 2012; Bachiller et al., 2014; Caeyenberghs et al., 2012; de Haan et al., 2012; King et al., 2013; Lord et al., 2011; Nomura et al., 2010; Sheffield, Repovs, Harms, Carter, Gold, MacDonald, Daniel Ragland, et al., 2015a; Sheffield, Repovs, Harms, Carter, Gold, MacDonald, Ragland, et al., 2015b) as well as task-dependent changes in connectivity in typically-developing populations (Bassett et al., 2011; Cole et al., 2014; 2013; Crossley et al., 2013; De Vico Fallani et al., 2008; Ekman, Derrfuss, Tittgemeyer, & Fiebach, 2012; Ginestet & Simmons, 2011; Hermundstad et al., 2013; Kitzbichler, Henson, Smith, Nathan, & Bullmore, 2011; Nicol et al., 2012; S. Palva, Monto, & Palva, 2010; Rosenberg et al., 2015; A. A. Stevens, Tappon, Garg, & Fair, 2012). Few graph theory studies have focused on the dynamic properties of attention or awareness. These studies have examined graph changes during tasks requiring cognitive control shifts (Cole et al., 2013; Ekman et al., 2012) and during maintenance of items in working memory (Cao et al., 2014; Ginestet & Simmons, 2011; Kitzbichler et al., 2011; S. Palva & Palva, 2012), but they have not focused on how these cognitive processes effect connectivity from a global perspective.

Specific Aims

The overarching goal of this dissertation is to extend our understanding of attention and awareness by examining the global network properties of these cognitive states using graph theoretical measures. Cognitive neuroscience research often focuses on the role of individual cortical regions and network-level activity, potentially neglecting the whole-brain changes associated with cognitive function and behavior. The experiments in this dissertation aim to expand on the role of the brain's complex functional topology by leveraging recent advancement in graph theoretical tools for studying neural connections.

Aim 1: Characterize the global network properties of visual awareness.

A theoretical debate exists as to the extent of neural activity necessary to produce a conscious percept of some piece of information. As described in the preceding sections, focal and network-level theories of awareness posit that relatively restricted activity in sensory and/or specific association cortices is sufficient to produce conscious awareness. However, several "global" theories suggest much more widespread changes are associated with awareness of a percept. Network science provides a well-suited perspective and set of tools for adjudicating between these opposing viewpoints. Chapter II of this dissertation (Godwin, Barry, & Marois, 2015) discusses the network properties of target awareness and presents evidence in favor of global theories of awareness.

Aim 2: Identify global, functional network properties of attentional capture by salient oddballs.

Highly salient visual stimuli produce a powerful orienting response, a involuntary physiological and cognitive redirection to perceptually interesting items (Pavlov, 1927; Sokolov, 1963). In the brain, detection of task-irrelevant, oddball stimuli has been shown to consistently recruit a ventral fronto-parietal network consisting of the anterior insula (AI), temporo-parietal junction (TPJ), and the inferior frontal junction (IFJ). However, the orienting response to novel or salient stimuli is comprised of a wide-ranging set of physiological and cognitive processes recruited in order to orient to and evaluate novel stimuli (Beatty, 1982; Han & Marois, 2014; Kahneman & Beatty, 1966; Tracy et al., 2000). Chapter III examines the functional topology associated with the detection of a task-irrelevant, oddball stimulus to characterize the global functional changes associated with the orienting response.

Aim 3: Examine the differential global network properties of task relevant and irrelevant information.

Chapters II and III respectively examine the global network function related to relevant targets and irrelevant oddballs. However, these studies are not designed to test the relationship of these influences of information processing concurrently. Chapter IV of this dissertation tests hypotheses the patterns of global topological changes associated with the detection of task-relevant and irrelevant stimuli by using results from Chapters II and III as functional markers of task relevant and task irrelevant processing. These analyses aim to examine whether topology changes contribute to the process of biasing information or suppressing distractors during an attention task. The simultaneous examination of both relevant and irrelevant competition for cortical representation better reflects the experience of the natural world where innumerable relevant and irrelevant sources of information impinge our senses.

The recruitment of specific, circumscribed networks plays a prominent role in theories of attention and cognitive control. There are, however, numerous disagreements as to what individual regions are recruited, how these areas communicate to form networks, and how those network coordinate activity to produce coherent behavior. The burgeoning literature and use of network science to study cognitive function provides a vital linking component between the function of individual regions or networks and the rest of the brain during behavior.

CHAPTER II

BREAKDOWN OF THE BRAIN'S FUNCTIONAL NETWORK MODULARITY WITH AWARENESS

This chapter is based on Godwin et al. (2015).

The debate over the neurobiological basis of conscious awareness best encapsulates the tensions between network-level and connectionist-based, global theories. Within this debate are theories that run the entire gamut of proposed neural correlates of consciousness. Three broad classes of models have been proposed to explain the neural basis of awareness, and these classes primarily differ on the predicted extent of neural information changes associated with conscious perception. According to 'focal' theories, awareness results from local changes in neural activity in either the perceptual substrates (Moutoussis & Zeki, 2002; Tse et al., 2005; Zeki, 2007) or in higher-level nodes of information processing pathways (Lau & Passingham, 2006). By contrast, 'network-level' theories posit that awareness is tightly associated with activation of parieto-frontal attention networks of the brain (Asplund, Todd, Snyder, & Marois, 2010a; Beck et al., 2001; Lumer et al., 1998; Lumer & Rees, 1999; Marois & Ivanoff, 2005; Naghavi & Nyberg, 2005; Rees et al., 2002). Finally, 'global' models propose that awareness results from widespread changes in the activation state (Baars, 2005; Baars, 2002; Dehaene et al., 2006; Dehaene & Changeux, 2011) and functional connectivity (Edelman, 2003; Tononi, 2008; Tononi & Sporns, 2003; van Gaal & Lamme, 2012) of the brain. While there is strong experimental support for network-level theories there is scant experimental evidence in favor of truly sweeping, widespread changes in brain activity with conscious perception despite the fact that global scale models have recently come to prominence in the theoretical landscape of this field.

Using a graph theoretical approach applied to ultra-high field functional magnetic resonance imaging (fMRI) data, here we experimentally tested a key tenet of global theories: the widespread emergence of large-scale functional connectivity with awareness. Graph theory analyses are ideal tools to test global models of awareness because they can provide concise measures of the integration and segregation of interconnected nodes of a system (Bullmore & Sporns, 2012). Applied to functional imaging data, we treat individual brain ROIs as nodes, functional connectivity between as edges, and functional brain networks as interconnected modules of nodes. When examining a large set of ROIs that encompass the different networks of the human cerebral cortex (Pettersson-Yeo, Allen, Benetti, McGuire, & Mechelli, 2011; Power et al., 2011), we can apply graph theory analyses to estimate the extent to which key measures of global information processing are altered by the state of awareness. This approach has been previously applied to study differences in cognitive states (Bassett et al., 2011; Cao et al., 2014; Cole et al., 2013; Ekman et al., 2012; Ferri, Rundo, Bruni, Terzano, & Stam, 2008; Ginestet & Simmons, 2011; Hermundstad et al., 2013; Kitzbichler et al., 2011; A. A. Stevens et al., 2012). Although recent studies have taken advantage of graph theory analysis to examine the connectivity patterns that precedes a conscious event (Weisz et al., 2014) or following pharmacologically induced loss of consciousness (Schröter et al., 2012), this approach has yet to be used for characterizing the topology associated with conscious target perception *per se*, a necessary test for global theories of awareness.

If the changes with awareness are truly global, one should see such changes even if the task does not require complex discrimination, identification, and semantic processes that may recruit vast extents of cortical tissue that have not necessarily been associated with conscious perception; in other words, these global changes should appear even for the simple conscious detection of a flashed disk. For this reason, we had participants perform an elementary masked target detection task (Fig. 2) while being scanned at ultra-high field (7 Tesla [T]). The task included three trial types: Forward-Masked, Backward-Masked, and No-Target conditions. In the Forward-Masked (paracontrast) condition, a 133ms-duration annular mask offset 33ms prior to the target (a disk whose exterior border coincided with the interior border of the annulus), presented for 33ms. In the Backward-Masked (metacontrast) condition the order of mask-target presentation was reversed while keeping all timing parameters the same. Under such conditions, forward masking of targets has been shown to impair target detection more than backward masking (Lefton & Newman, 1976; Schiller & Smith, 1966). Consequently, the mask/target orderings provided a manipulation of target awareness while maintaining the same mask and target presentation times across both Forward- and Backward-Masked conditions. Because on each trial participants made a detection response about the presence or absence of the target followed by a confidence rating on their response, subjects' performance could be assessed on both an objective (discriminability index, d') and subjective (confidence rating) measure of awareness (Seth et al., 2008). In turn, only trials in which the target was either seen (Aware) or unseen (Unaware) at high confidence levels were used for analysis of brain imaging data. Finally, because the report of the percept was 12 seconds removed from the stimulus presentations (Fig. 2), the task design precluded initiation of the motor response itself from influencing estimates of awareness. Although response selection and motor preparation processes likely occur during this period, similar preparation would occur across all conditions.

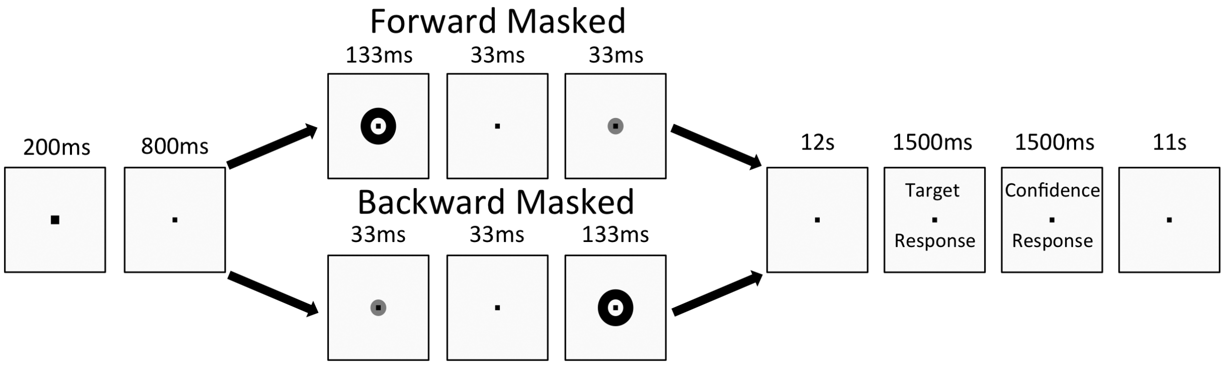

Figure 2. Schematic of behavioral paradigm with Forward-Masked and Backward-Masked trial types (No-Target trials not shown). On each trial, participants responded whether or not they detected the target stimulus and indicated a confidence rating for their answer (see Methods).

Methods

Participants

Twenty-eight individuals (aged 18-33; fifteen females), recruited from the Vanderbilt University community, participated in the study. All participants had normal or corrected to normal vision. The Vanderbilt University Institutional Review Board approved the experimental protocol and informed consent was obtained from all subjects. Data from four participants were excluded due to technical difficulties (two participants), excessive head motion (> 6 mm; one participant), and failure to follow task instructions (one participant).

Behavioral Paradigm

Stimuli were presented using the Psychophysics Toolbox (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007; Pelli, 1997) in MATLAB (The MathWorks, Natick, MA, USA). All stimuli were shown overlaid on a white background with a persistent, centered, black fixation square (.25° visual angle). Participants were instructed to monitor each trial for a target stimulus, ignoring a mask stimulus. The target stimulus, a filled gray disc (1° visual angle), was presented at the center of the screen for 33 ms. The mask was a centered, black annulus (from 1° inner edge to 2° outer edge) surrounding the target disk and was shown for 133 ms.

All participants saw three conditions, psuedorandomly mixed within runs: Forward-Masked, Backward-Masked, and No Target. Forward-Masked and Backward-Masked trial types are shown in Figure 2. In addition, fourteen of the twenty-four participants saw rare oddball images in 5% of all trials. These surprise trials were not analyzed for the current study.

Each trial began with enlargement of the fixation square for 200ms, cuing the participants for the upcoming target and/or mask presentation 800 ms after the fixation returned to its standard size. An interstimulus interval (ISI) of 33 ms separated the mask and target on all target-present trials. When the mask precedes the target (forward or paracontrast masking) at these timing parameters, target detection is severely impaired. By contrast, when the mask follows the target (backward or metacontrast masking), target detection improves (Lefton & Newman, 1976; Schiller & Smith, 1966). A 12 s fixation interval followed the stimulus presentations to allow the BOLD signal for target detection to be dissociated from the target response BOLD signal (see fMRI methods below). Following the 12 s interval, participants responded to an on-screen prompt (1.5 s) whether they had detected the target stimulus, using one of two right-handed button presses for “yes” or “no”. Participants were then prompted (1.5 s) to provide a rating, on a scale of 1 to 5, of how confident they were in their previous detection response (1 = No confidence; 5 = Total confidence) with a left-handed button press. The rating scale remained on screen for the duration of each prompt. The next trial began following another 11 s fixation period.

This stimulus presentation paradigm afforded several advantages for assessing the changes in global functional connectivity with target awareness. First, the reversed mask/target orderings provided a manipulation of target awareness while maintaining identical mask and target presentation durations across both Forward and Backward Masked conditions. This consistency across conditions allowed examination of robust effects of target awareness without differences in overall physical stimulation. Moreover, our paradigm yielded robust numbers of trials in which the subjects were highly confident they either did or did not see the target (Kunimoto, Miller, & Pashler, 2001). In addition, because all stimuli were presented at fixation, the task required no spatial shifts of attention or eye movements. Finally, by using a very simple stimulus target-mask presentation paradigm that only required rudimentary target stimulus detection (brief percept of a disk), our manipulation provides a strong test of the global theories of awareness because it is unlikely to evoke widespread activation associated with identification, discrimination, or semantic processing (as may occur, for instance, with the attentional blink paradigm).

Twenty-one of the participants completed between four and five fMRI runs (three completed four runs whereas the remaining 18 completed five runs), each consisting of twenty trials split between the Backward-Masked, Forward-Masked and No-Target trials (35%, 35%, 30% respectively). Three additional participants completed between five and eight fMRI (one completing five runs, one completing six runs, and one completing eight runs) runs consisting of fifteen trials each (5 of each of the three trial types). Trial types were presented in a pseudo-randomized order in each run. Trials during which participants failed to provide either a detection response or a confidence rating were excluded from analysis (<1% of all trials).

fMRI methods

Stimuli were presented using an Avotec SV-6011 projector (Avotec, Inc. Stuart, FL.) back-projected onto a screen inside the scanner. Participants lying in the scanner viewed the screen through a mirror mounted to the head coil. The experiment was performed at the Vanderbilt University Institute of Imaging Sciences on a 7T Philips Achieva MRI system to benefit from high sensitivity to blood oxygenation level dependent (BOLD) signals (Gati, Menon, Ugurbil, & Rutt, 1997; Ogawa et al., 1993) while providing full-brain coverage at conventional imaging resolutions. Whole brain, anatomical T1 weighted images were acquired with a 1x1x1 mm voxel resolution. Functional T2* images were acquired using a 3D PRESTO (Principles of Echo Shifting using a Train of Observations) sequence with 3x3x3 mm voxels and a 1 s volume acquisition time (Barry et al., 2011). Scan parameters consisted of a 10 ms repetition time (TR), 14 ms echo time (TE), 10° flip angle, 216 mm x 216 mm in-plane field of view (FOV), 72 × 72 matrix, and 40 slices (covering 120 mm superior-inferior). Each functional scan included either 410 brain volumes (Subjects 1-12) or 545 brain volumes (Subjects 13-24).

Data preprocessing was performed using Brain Voyager QX 2.3 (Brain Innovation, Maastricht, The Netherlands) and included 3D head motion correction and linear trend removal. Functional and anatomical runs were co-registered and transformed into standard Talairach space (Talairach & Tournoux, 1988).

Functional Connectivity Analysis

Cortical regions of interest (ROIs) were first defined from the set of coordinates reported in Power et al. (2011), in which authors identified nodes based on resting-state connectivity data with strong overlap with known functional-network systems. These nodes were parceled into sub-graphs based on community detection algorithms resulting in strong concordance with previously identified functional networks. Coordinates, originally reported in MNI space, were converted to Talairach space using the mni2tal.m function in MATLAB (Matthew Brett). Four-mm radius spheres were drawn around each coordinate to create a list of 264 ROIs to serve as nodes for graph theoretical analyses.

In order to examine connectivity differences associated with target awareness with maximum power, Hit and Miss trials were each collapsed across both Forward and Backward Masked trials, yielding a total of 486 high-confidence Hit trials and 276 high-confidence Miss trials. High-confidence trials were defined as those with "4" or "5" ratings on the 5-point confidence scale. Given that each condition (Hit and Miss) was comprised of trials primarily originating from one masking condition (Backwards and Forward, respectively), we also performed the Hit vs. Miss analyses on high-confidence trials within each masking condition to confirm that the results were not due to masking differences. Although the results of this analysis by masking conditions are consistent with the main findings, they suffer from low power because several subjects lacked a sufficient number of high-confidence trials of one trial type to make these comparisons statistically meaningful. Consequently, to further rule out the possibility that changes in connectivity differences with awareness may be due to masking differences, we also compared False Alarm and Correct Rejection trials of the no-target condition, as these two trial types are associated with very difference percepts but identical physical presentations.

Task-dependent functional connectivity between nodes was estimated using the generalized psychophysiological interaction (gPPI) method (McLaren, Ries, Xu, & Johnson, 2012). This method aims to account for task-activation effects and non-task-specific correlations separately from task-induced connectivity differences (Gitelman, Penny, Ashburner, & Friston, 2003). GPPI is considered to be more powerful than a “standard” PPI analysis (Cisler, Bush, & Steele, 2014) and has been shown to be a better estimate of task-based functional connectivity than a cross-correlation coefficient (Kim & Horwitz, 2008). It utilizes the general linear model (GLM) with three regressor types: condition-specific task regressors (analogous to those used in standard fMRI GLMs), a “seed” timecourse regressor, and condition-specific interaction regressors. The seed timecourse regressor is included to capture signal variance in regions that show correlations with the seed region outside of task periods of interest. Interaction regressors were created using the deconvolution method described in McLaren et al. (2012). In essence, interaction regressors predict correlation with seed region signal, but only during task-relevant timepoints (capturing an interaction of seed and condition factors). Our GLMs included condition-specific regressors for high and low confidence Hit and Miss trials. Each GLM consisted of condition-specific regressors (modeled as events locked to target/mask presentation), a seed node BOLD time course, and condition-specific interaction regressors, all in addition to task regressors of no interest for fixation and response periods. Parameter estimates for high confidence interaction regressors were recorded and organized for graph theoretical analyses.

A separate GLM was run was for each of the 264 possible seeds for each subject. For each seed region, averaged parameter estimates for gPPI interaction regressors were extracted from and averaged across voxels in the 263 remaining ROIs providing a task-modulated measure of weighted (as opposed to binary) connectivity between each region pair. Larger parameter estimates indicate greater signal coherence between seed and target ROI and thus greater functional connectivity. No directionality was assumed in the data and reciprocal pairwise connections were averaged to generate the final, symmetric graph. PPI analyses were performed using Brain Voyager QX 2.3 and custom MATLAB software in addition to the BVQXtools and NeuroElf toolboxes.

Graph Theoretical Analysis

Graph theoretical analyses were performed using the Brain Connectivity Toolbox (Rubinov & Sporns, 2010). Graphs (i.e. ROIs and their functional connections) were constructed and analyzed on an individual subject basis for each condition of interest (e.g. high-confidence hit, high-confidence miss trials). As noted in Rubinov & Sporns (2010), network properties often differ based on the number of nodes, connections, and degree distribution of a network. To construct network measures that allow comparisons across conditions/graphs that might differ along any of these properties, graph theory metrics are "normalized" by comparing – on a per subject basis – each of these metrics against a null hypothesis network; i.e. a network with a randomized topology that otherwise conserves the size, density, and degree distribution of the original network. All graph theoretic measures were normalized via division by their respective metric averaged across 100 random graphs generated via the Brain Connectivity Toolbox function "randmio_und.m" Statistical significance of normalized graph theoretic measures was performed via Wilcoxon sign-rank tests given the lack of evidence concerning the normal distribution of graph theoretic metrics (Bullmore & Sporns, 2009). Because graph theoretic metrics can be threshold-dependent (van Wijk, Stam, & Daffertshofer, 2010), it is important to examine graph measurements over a range of possible connection densities. In accordance with prior studies (Cao et al., 2014; Dosenbach et al., 2008; Rubinov et al., 2009), results were obtained for graphs density thresholded at the top 10% to 30% of individual subject connections, with 5% steps, for a total of five threshold levels. Results are presented for the top 10% of connections. Thresholded matrices were rescaled to the range [0,1] (Rubinov & Sporns, 2010). This normalization procedure was calculated by dividing each connectivity value by the maximum value in the graph in order to rescale values to a similar range as correlations.

Following the approach of previous work in the field, no correction for multiple comparisons across the multiple measures of network properties were applied because of the non-independence of these measures (Cao et al., 2014; Rubinov & Sporns, 2010; 2011) and because the current study aims only to test the global patterns of topology, rather than the independent properties of individual nodes or ROIs. Correction for multiple comparisons is normally applied when testing several ROIs (Bassett, Nelson, Mueller, Camchong, & Lim, 2012; Lynall et al., 2010; Schröter et al., 2012) or if statistical significance (on a per voxel basis) acts as a thresholding step (Bassett et al., 2009; E. T. Bullmore & Bassett, 2011; Cole et al., 2014; Dosenbach et al., 2007; Rubinov & Sporns, 2011). The statistical reliability of our findings was provided by a test of replication; the same results were obtained in two independent data analyses: comparisons of Hits vs. Misses as well as the comparison of False Alarms vs. Correct Rejections.

In order to measure changes in the graph properties of connections across the brain's neural network as a function of our awareness manipulation, we describe connectivity based on measures of functional segregation, functional integration, and centrality (Rubinov & Sporns, 2010). To estimate the functional segregation of the brain's connectivity, we assessed "modularity", or the amount to which a graph (i.e. our entire ROI set) can be divided into non-overlapping modules (i.e. sub-networks of ROIs), via Newman's spectral algorithm for weighted matrices (M. Newman, 2004). Weighted modularity (Q^w) was calculated by:

$$Q^w = \frac{1}{l^w}\sum\nolimits_{i,j\in N}\left[w_{i,j} - \frac{k_i^w k_j^w}{l^w}\right]\partial_{m_i,m_j}$$

with weighted connections between nodes i and j ($w_{i,j}$), sum of all weights in a graph (l^w), weighted degree of a node (k_i), module containing node i (m_i), and $\partial_{m_i,m_j} = 1$ if $m_i = m_j$. Based on the modules identified by Newman's algorithm we examined the "participation coefficient" (y^w), or the degree to which nodes (i.e. ROIs) connect with nodes in other functional modules (sub-networks). Participation provides a measure of centrality per node; that is, a measure of a node's importance in inter-modular communication. Nodes with high participation increase global integration by facilitating between-module communications. The *participation coefficient* was calculated by:

$$y_i^w = 1 - \sum\nolimits_{m\in M}\left(\frac{k_i^w(m)}{k_i^w}\right)^2$$

where $k_i^w(m)$ is the weighted degree of connections between node i and nodes in module m.

We additionally measured each node's "clustering coefficient", a measure of the degree of segregation present in a network that estimates the extent to which connectivity is clustered around each node irrespective of its module membership. The clustering coefficient (C^w), where t_i is the number of triangles around node I, is calculated with the following formula:

$$C^w = \frac{1}{n}\sum\nolimits_{i\in N}\frac{2t_i^w}{k_i(k_i - 1)}$$

Finally, “average path length” (L^w), a measure of network integration, provided a statistic describing the functional distance between nodes, where $d_{i,j}$ is the shortest path length between nodes i and j, computed with a inverse mapping of weight to length.

$$L^w = \frac{1}{n}\sum_{i\in N} \frac{\sum_{j\in N, j\neq i} d_{ij}^w}{n - 1}$$

Further descriptions of how these metrics are calculated and their implications can be found in Rubinov & Sporns (2010).

As a measure of non-parametric effect size, we report the dependent measures probability of superiority, or PS_{dep} (Grissom & Kim, 2012). PS_{dep} is defined as

$$PS_{dep} = \frac{n_+}{N}$$

where n_+ is the number of positive difference scores, discarding ties and N is the sample size. PS_{dep} ranges from 0 to 1 and can be interpreted simply as the proportion of paired samples showing greater values in one condition compared to the other.

BOLD Activation Maps for Awareness

In order to isolate regions activated by awareness of the target stimulus, a separate model was defined (a non-PPI, GLM analysis) for trials in which the participant successfully detected the target (hits) and trials when participants reported not having seen the target (misses), collapsing across Forward- and Backward-Masked trials to increase power. In addition, high-confidence trials and low-confidence trials were modeled separately, resulting in four separate regressors: high-confidence hits, high-confidence misses, low-confidence hits, and low confidence misses. To isolate awareness-sensitive regions, high-confidence hit trials were directly compared to high-confidence miss trials. Group random-effect statistical parametric maps (SPMs) of the activation differences between Aware and Unaware trials were then created and thresholded at $p < .05$, cluster corrected for multiple comparisons (Forman et al., 1995) utilizing the Brain Voyager cluster correction plugin.

Results by Masking Condition

Because most of the Target Aware trials came from the Backward-Masked condition (~83%) whereas the Target Unaware trials primarily originated from the Forward-Masked condition (~84%), we assessed whether the results obtained for the data pooled across masking conditions also held when the analytical tests between Target Aware and Unaware trials were compared within each masking condition. This analysis is statistically limited, however, by the low trial number in some of the conditions: In the Backward-Masked condition, there were a total of 403 high-confidence Aware trials compared to only 45 Unaware trials, whereas the Forward-Masked condition included few (83) high-confidence Aware trials compared to high-confidence Unaware trials (231). As a result, 7 of the subjects for the Backward-Masked condition and 12 of the subjects for the forward-masked condition were excluded from these analyses because they had no or too few (1 or 0 instance per run) high-confidence aware or high-confidence unaware trials to perform the graph theory analyses with the gPPI method.

Results

Target discriminability was greater under the Backward-Masked condition than the Forward-Masked condition (main effect of masking condition, F(1,92) = 35.51, p < .001). Moreover, this difference in target discriminability between the Backward and Forward conditions increased at the highest confidence ratings (main effect of confidence rating, Fig. 3, F(4,92) = 9.82, p <.001). These results not only demonstrate that the masking manipulation was successful in affecting target detection but also indicate that high confidence ratings provide a robust means of distinguishing between seen and unseen targets.

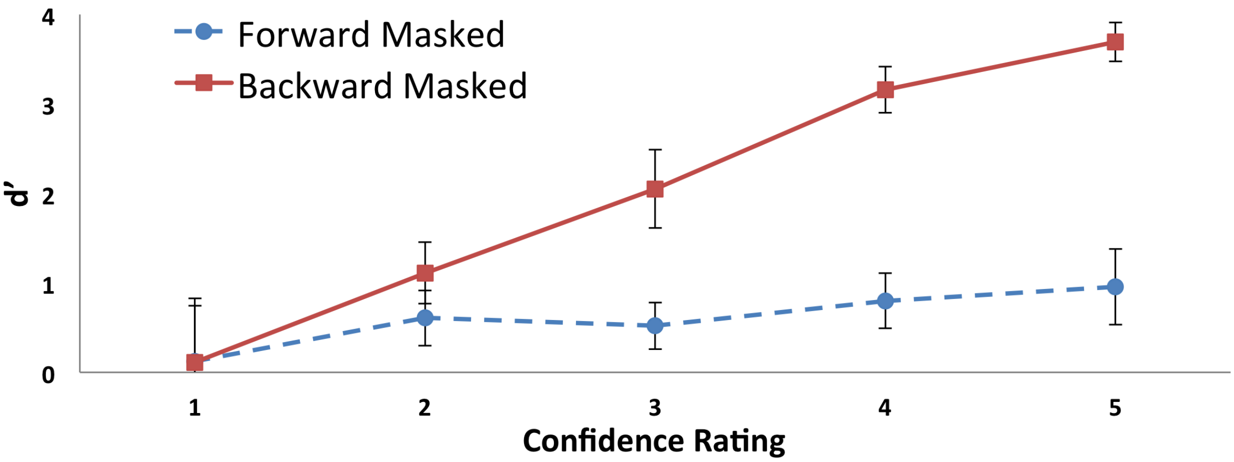

Figure 3. Target stimuli were seen reliably more often in Backward-Masked trials compared to Forward- Masked trials, especially at high confidence ratings. Global connectivity analysis was confined to the high-confidence trials ("4" and "5" ratings) to ensure potent differences between aware and unaware states. Error bars show within-group standard error.

To assess whether consciously Aware and Unaware target states were associated with distinct global patterns of functional connectivity, we compared the differences in graph theoretical metrics between seen and unseen trials for high confidence ratings only. In order to increase the number of trials entering into the analysis, each trial type was pooled across both the forward and backward-masked conditions and classified as Target Aware (seen) or Unaware (unseen) trial types. We examined whether the results obtained for the data pooled across masking conditions also held for comparisons within each masking conditions.

Figure 4 illustrates weighted connection matrices for all pair-wise PPI parameter estimates, averaged across subjects, between each of the 264 cortical nodes for Target Aware and Unaware conditions organized using the Power et al. (2011) parcellation. Projections of the nodes and edges onto two-dimensional cortical representations in Figure 4 highlight the widespread differences in functional PPI strengths between Aware and Unaware states. It is difficult to draw firm conclusions about connectivity changes with awareness from visual inspection of these matrices or projections alone. Hence, we quantitatively assessed network topology changes with awareness by estimating key graph theoretical metrics – network segregation, integration and centrality – based on the top 10% of connection strengths (Fig. 4), though similar results were obtained using a range of strength thresholds. If awareness is associated with widespread increases in cortical functional connectivity, it would likely be accompanied by decreased network segregation, increased network integration, and increased node centrality (i.e. highly interactive nodes that facilitate functional integration).

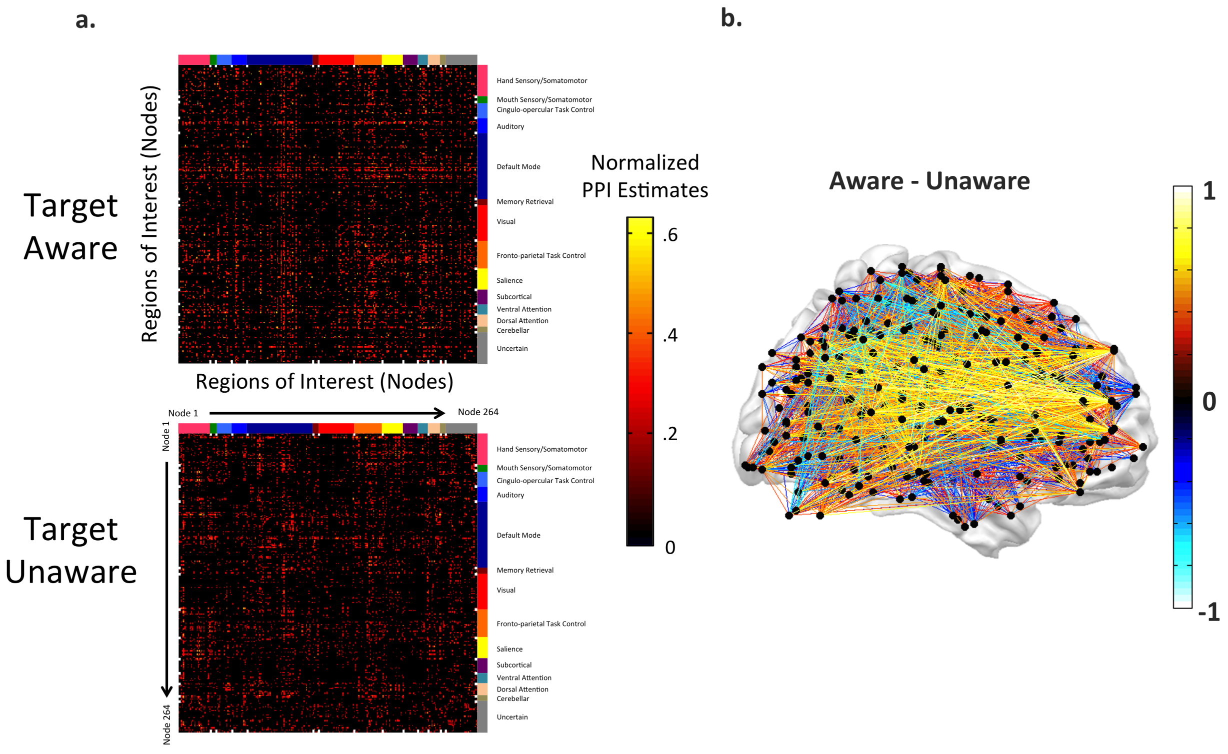

Figure 4. Group-averaged, symmetric connectivity graphs for Target Aware and Unaware conditions thresholded at the top 10%. Node coordinates were derived from Power, et al. (2011). (a) Aware (top) and Unaware (bottom) matrices are organized by the 14 network assignments described by Power and colleagues. Heat scale indicates magnitude of interaction regressor (seed x condition; PPI) parameter estimates. (b) Differences between subject-averaged connectivity matrices plotted for Aware minus Unaware. Values were derived by subtracted subject-averaged graphs shown in Figure 3a. Hot colors depict stronger connections for the Aware than Unaware condition, whereas cool colors depict Unaware connections greater than Aware. Plots are shown overlaid on a surface projection for reference to anatomical direction and general location. Difference scores with the greatest absolute value are plotted above weaker differences.

Functional modularity, a measure of the ability to segregate the connectivity patterns into clearly distinct networks, decreased with target awareness (Fig. 5, Target Aware vs. Unaware; Wilcoxon's signed-rank test, $p = .043$, $PS_{dep} = .75$). Moreover, the average participation coefficient, a value assessing between-network connectivity strengths, was greater in the Target Aware condition than in the Unaware condition (Wilcoxon's signed-rank test, $p = .009$, $PS_{dep} = .67$). Typically, functional network topologies are complex, with more long distance connections than lattice (i.e. serially connected) networks but fewer than randomly-connected networks

(Bullmore & Sporns, 2012). This complex topology is thought to be a functionally and metabolically efficient middle ground between random and lattice organizations. Changes in functional modularity and participation toward more "random-like" (i.e. a normalized value of 1) organizations suggest a shift along this efficiency spectrum to favor longer distance connections at the expense of segregation of functional networks (Kitzbichler et al., 2011).

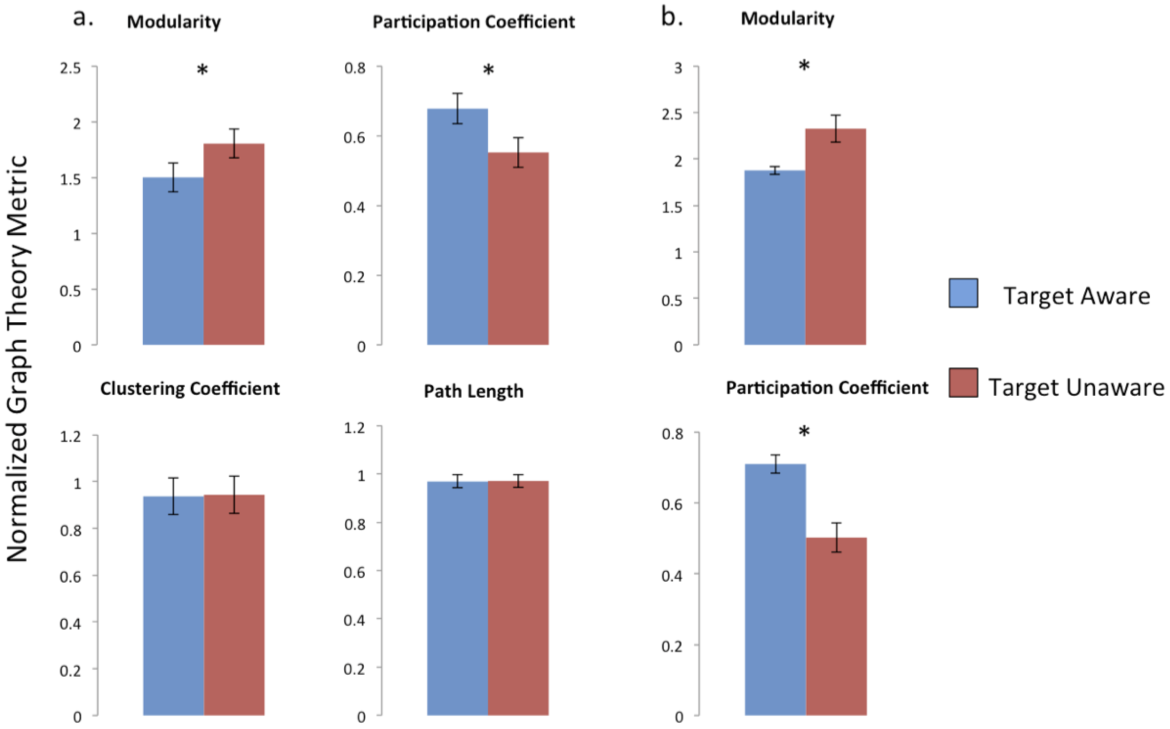

Figure 5. Effect of Awareness manipulation on graph theoretic measurements for Target Aware and Unaware trials. (a) Comparison of Target aware and Unaware trials (collapsed across target-present conditions). Significant differences between Aware vs. Unaware conditions were only observed for Modularity and Participation metrics. (b) Comparison of false alarm trials (Aware) to correct rejection trials (Unaware) from the No-Target condition. All y-axis values represent the ratio of the observed graph theory metric to the corresponding random graph metrics (see Methods). Error bars represent within-group standard error. Asterisks indicate significant differences between groups at $p < .05$.

Modularity and average participation coefficient were the only metrics to exhibit changes between Aware and Unaware conditions. The average clustering coefficient, a different measure of network segregation that estimates the degree to which neighboring nodes tend to interconnect with one another, was not affected by the awareness manipulation (Wilcoxon's signed-rank test, $p = .689$, $PS_{dep} = .54$). The average path length, a metric of integration that measures the average functional distance between two nodes, did not show differences between Aware and Unaware conditions (Wilcoxon's signed-rank test, $p = .753$, $PS_{dep} = .54$). Unlike modularity and participation coefficient, neither of these metrics takes into account module membership. The finding that modularity and participation coefficient – measures which are sensitive to changes in inter-modular communication as opposed to changes in individual node connectivity – are those that are altered with awareness strongly suggests that awareness is associated with a breakdown of the brain's network modularity. Similar results were observed across a range of connectivity thresholds (Fig. 6).

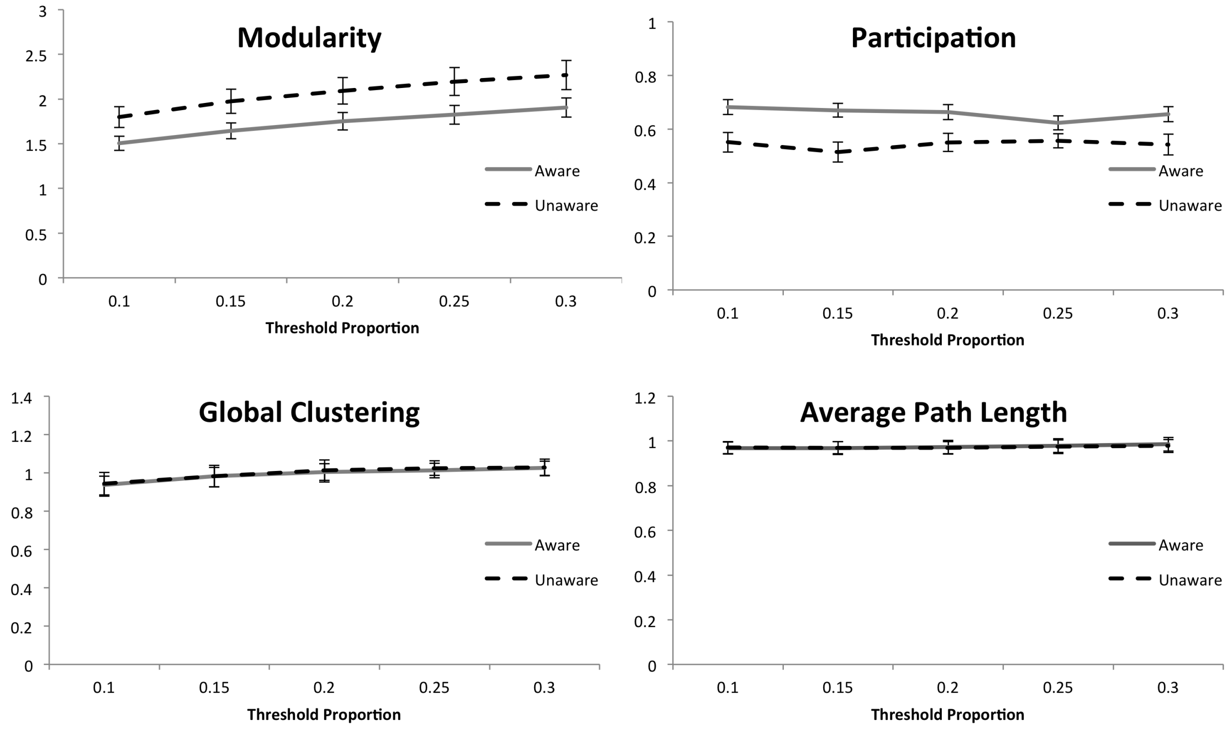

Figure 6. Graph theoretical measures across multiple threshold levels of the main fMRI experiment data. Relationships between Aware and Unaware graph metrics do not change across the top 10% to 30% of connectivity values as assessed by PPI strength. Error bars represent within-group standard error.

The hypothesis that conscious target perception is associated with whole-brain functional connectivity changes would be strongly validated if it were replicated in an independent trial data set: high-confidence false alarms and correct rejections in no-target trials. These two trial types have physically identical stimulus presentations (no targets) and identical levels of subjective confidence, but very distinct percepts (Aware vs. Unaware). False alarm trials showed lower modularity compared to correct rejection trials (Wilcoxon's signed-rank test, $p = .007$, $PS_{dep} = .67$), and average participation coefficients were higher during false alarm trials compared to correct rejections as well (Wilcoxon's signed-rank test, $p = .0002$, $PS_{dep} = .79$). Absent a target, the highly confident awareness of a percept appears sufficient to alter the brain's network topology. This might take place because the false alarm trials correspond to those with structured internal noise sufficient to trigger the simple percept of a flash (Ress & Heeger, 2003). These results provide converging evidence that reported perception of the target stimulus, regardless of masking manipulations or target presence, produces decreased functional modularity while increasing average participation.

It is conceivable that the functional connectivity changes observed with awareness arise from large connectivity changes between a small number of networks, rather than truly global topological changes. If one module or a small set of modules showed greater between-network connectivity changes compared to the others, it should be detectable by testing for an interaction between networks and target awareness state in participation coefficient (a measure of inter-module communication). To test this possibility, we utilized the fourteen consensus networks identified by Power et al. (2011). Using the participation coefficients (calculated either based on module membership detected by our modularity calculations or based on the Power et al. (2011) networks as the source of module membership), we estimated the average participation coefficient for nodes in each of the fourteen Power et al. networks for each subject in the Target Aware and Unaware conditions. A repeated-measures ANOVA showed no significant interaction between condition (Aware vs. Unaware) and network for either of the analytical methods (Fs (13,299) < 1.03, ps > .43). It is unlikely that the global modularity and participation differences observed in the present data between the Aware and Unaware conditions primarily stem from changes in a restricted set of networks.

If not driven by a small number of networks, perhaps our global metrics are instead skewed by massive connectivity changes occurring with awareness in a small number of nodes.

To examine this possibility, we tested the integrity of the global network to removal of the most highly interconnected nodes on a per-subject basis (Bullmore & Sporns, 2009). If these nodes were singularly responsible for what appeared to be global effects, we should no longer observe difference in modularity and participation coefficient between Aware and Unaware conditions following removal of these highly connected nodes. Excluded nodes were defined as the regions with the greatest summed weighted connection strengths (i.e. the sum of all pairwise connections made by that particular node, thresholded at the top 1% (3 nodes) or 10% (27 nodes) of the most interconnected nodes).

Even after removal of these nodes, modularity still differed between Aware and Unaware, with greater modularity in the Unaware compared to Aware conditions (1%: $p = .046$; $PS_{dep} = 0.71$, 10% $p = .046$; $PS_{dep} = 0.71$). Correspondingly, functional participation still increased with Awareness after this targeted attack (1%: $p = .01$, $PS_{dep} = 0.67$; marginally at 10% $p = .063$, $PS_{dep} = 0.58$). No significant differences were found for the clustering coefficient (1%, $p = .732$, $PS_{dep} = .54$; 10%, $p = .775$, $PS_{dep} = .5$) or average path length (1%, $p = .753$, $PS_{dep} = .54$; 10%, $p = .65$, $PS_{dep} = .58$). The effects of Awareness on graph theory metrics do not appear to be driven by a small subset of nodes showing the highest connectivity changes.

A final possibility that we considered was that the global connectivity changes were driven by the brain regions that showed significant BOLD amplitude changes with awareness (as identified in SPMs of the contrast of high-confidence Aware vs. high-confidence Unaware conditions; see Figure 7). To address this issue, we performed a similar targeted attack analysis as above, excluding nodes that overlapped with activated voxels in the SPMs. After removal of the eight nodes that were identified as overlapping with foci activated with the Awareness manipulation, modularity was still greater in the Unaware condition ($p = .043$, $PS_{dep} = 0.67$), whereas participation was greater in the Aware condition ($p = .007$, $PS_{dep} = 0.67$). Again, we found no significant differences between average path length ($p = .753$, $PS_{dep} = .54$) and the clustering coefficient ($p = .797$, $PS_{dep} = .5$) between conditions. The brain regions that showed significant BOLD amplitude changes do not solely drive the global connectivity changes observed with the Awareness manipulation.

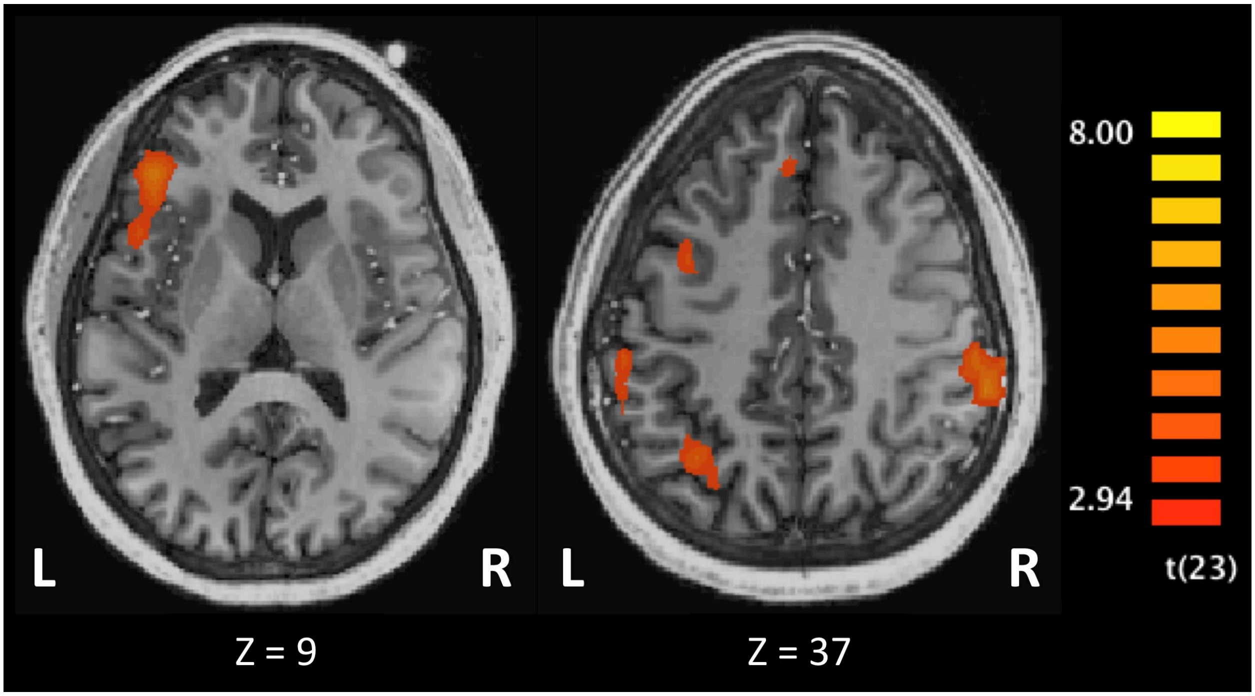

Figure 7. BOLD Amplitude SPM for Awareness. Foci showing activation in the contrast of Aware > Unaware (high confidence trials collapsed across forward-and backward- masked conditions) are presented. These activations include VLPFC (-42,32,12), PreSMA (- 1,13,51), left MFG (-36,1,46), left and right IPL (-56,-44,40; 59,-38,35), and left IPS (-34,- 59,38). All coordinates (peak activations) are reported in Talairach space. Maps thresholded at $p < .05$, cluster corrected for multiple comparisons.

Taken together, these results suggest that target awareness is associated with degradation of modularity in the brain's functional networks via an increase in the participation coefficient without changes in clustering coefficient. Awareness may be associated with a widespread increase in functional connectivity across modules rather than within modules. These results are also in line with reports that manipulations of working memory load can increase intermodular communication and decrease modularity in the absence of global efficiency changes (Ginestet & Simmons, 2011), lending credence to our conclusion that decreased functional modularity with awareness results from widespread increased inter-modular connectivity.

BOLD Amplitude Patterns of Awareness

In order to contrast the global connectivity changes with those obtained when measuring amplitude differences in activation with awareness, we contrasted the SPM for Target Aware vs. Unaware (high-confidence) trials (collapsed across masking conditions). This map revealed activations in a few scattered fronto-parietal locations – left ventrolateral prefrontal cortex (VLPFC), pre-supplementary motor area (PreSMA), left middle frontal gyrus (MFG), bilateral inferior parietal lobule (IPL), and left intraparietal sulcus (IPS) (Figure 7). There were no regions more activated in the Unaware compared to the Aware condition. This pattern of activations, which is much more circumscribed than the global changes in connectivity that we observed, is generally consistent with previous evidence of fronto-parietal activations with target awareness (Lau & Passingham, 2006; Marois et al., 2000; 2004; Tse et al., 2005).

Results by Masking Condition

The pattern of results observed, when splitting the data across masking directions, is similar to that obtained in the main analysis for the top 10% of connections (Figure 8). Specifically both backward-masked and forward-masked tended to show a decrease in modularity and an increase in participation in Aware relative to Unaware trials, with no observable trends for clustering and distance metrics. However, as evidenced in the graph, only a subset of the results reached statistical significance when broken down by masking condition. For backward-masked trials (metacontrast), modularity was significantly greater in the Unaware condition compared to the Aware condition ($p = .002$, $PS_{dep} = .88$). No other comparisons using this data broken out by masking condition reached statistical significance (p's > .1). This is not surprising given that this breakdown by masking condition suffers from low power due to the small number of trials and subjects that could be included with this analysis. Nevertheless, the overall pattern of results, coupled with the False Alarm and Correct Rejection comparisons, suggest that the graph theory measures reported in this chapter reflect genuine changes in global connectivity with Awareness.

Analyses excluding PPI deconvolution step

We have analyzed the global data using PPI excluding the deconvolution step, as described in MacLaren et al. (2012). For this analysis, the interaction predictors were calculated as the timepoint-by-timepoint product of the task regressor and the extracted seed timecourse. All other aspects of the graph theoretical analysis were identical to those presented above. With this reanalysis at the top 10% of connection strengths, the pattern of results for functional Modularity remained the same as in the earlier analysis, with the average modularity in the Aware condition marginally lower than in the Unaware condition (Wilcoxon's signed-rank test, $p = .05$, $PS_{dep} = .71$). Participation was significantly greater in the Aware compared to the Unaware condition (Wilcoxon's signed-rank test, $p = .005$, $PS_{dep} = .75$), as in previous analyses. The clustering coefficient showed a trend for being greater in the Aware condition than in the Unaware condition (Wilcoxon's signed-rank test, $p = .07$, $PS_{dep} = .63$). The average functional path length did not significantly differ between conditions (Wilcoxon's signed-rank test, $p = .15$,

$PS_{dep} = .54$). Thus, without the deconvolution step, we see similar effects on modularity, participation coefficient, and functional path length to what was presented above. We also see a marginal difference in the average clustering coefficient, a trend that was not observed in any other iteration of our analysis.

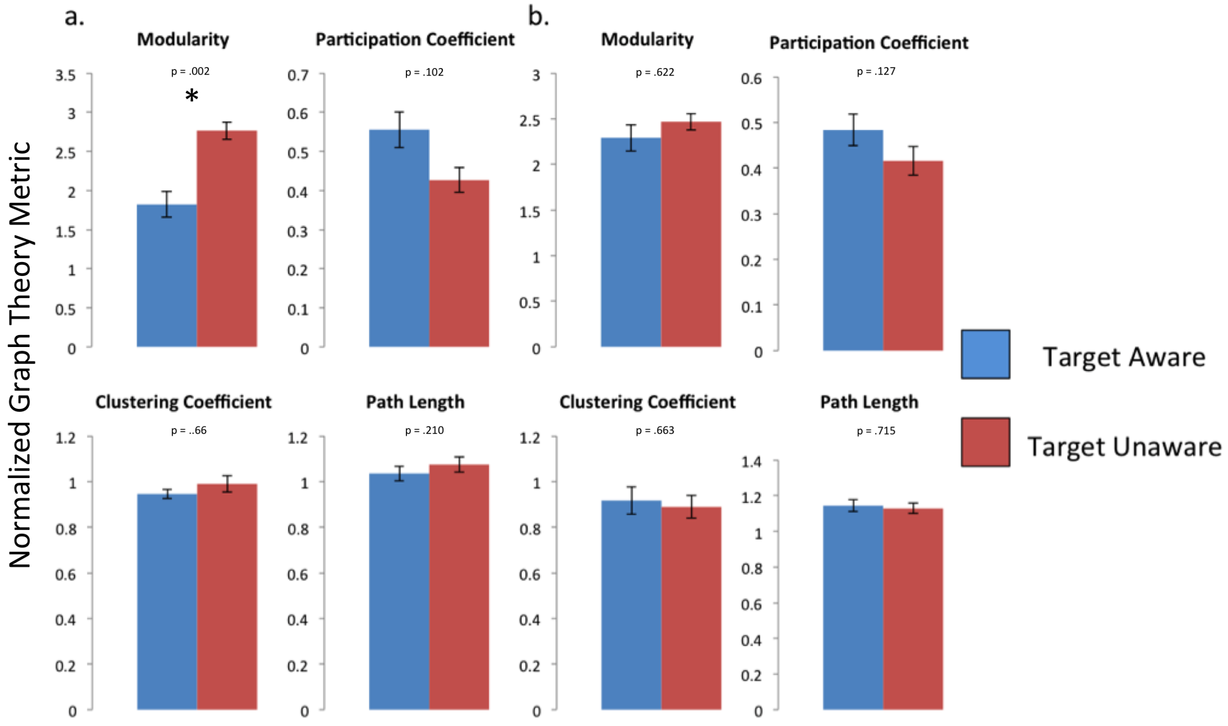

Figure 8. Differences in graph metrics for Aware and Unaware trials in forward- and backward-masked conditions. (a) Results for backward-masked trials only. (b) Results for forward-masked trials only. Error bars represent within-group standard error.

Pearson correlation analysis

To confirm our results were not simply due to an artifact of the PPI procedure, we reanalyzed the data with a Pearson Correlation method using Jochen Weber's NeuroElf toolbox for MATLAB and custom MATLAB scripts. Using the "vtc.RegressOut" function included in the toolbox, we removed task-related variance as well as several additional, z-transformed, nuisance variables (timecourses from white matter and ventricle ROIs) on a run-by-run basis. Following this regression step, we generated condition-specific timecourses by concatenating 12-volume timecourses (with a 1 s volume acquisition time) corresponding to the post-target delay in our design. In order to cancel out the contribution of intrinsic functional connectivity in these timecourses, we analyzed the differences between Aware and Unaware connectivity for individual subject connectivity matrices. We Fisher z-transformed the Pearson correlated matrices, calculated difference matrices (Aware - Unaware) per subject and performed graph theoretical analyses on these difference matrices, testing vs. a baseline as in a paired t-test. Because our analyses utilize normalized (by null graph) values, a baseline value of 1 would indicate the matrix is comprised of random data points. Using a paired t-test vs. 1, we found significant modularity (mean = 1.4391; $p = .04 \times 10^{-9}$), participation (mean = .8110; p = .000001), and clustering effects (mean = .6951; $p = .02 \times 10^{-6}$). There was no significant difference for the average path length (mean = 1.0212; p = .2). While this analysis confirms a difference in modularity and participation between conditions that we observed with the PPI method, interpretation of graph theory metrics based on a matrix of Pearson correlation difference scores is difficult in the context of the present design. Specifically, whereas the gPPI analysis examines the topology of connectivity matrices under the conditions of target awareness and unawareness, this Pearson correlation method requires the examination of a difference matrix to account for the influences of intrinsic connectivity. While some individual connectivity strengths may not significantly differ between these conditions, they still contribute to the overall topology of the network in conjunction with stronger connectivity differences. When calculating difference matrices, we likely lose a substantial contributing factor to graph topology due to the subtraction of connectivity strengths. This complicates interpretation of the graph theoretical properties of difference matrices and makes extrapolation of our findings to the properties of awareness more difficult. Consequently, the main claims of this paper are based on the gPPI connectivity analysis described above.

Head Motion

Head motion has been shown to have substantial impacts on estimates of functional connectivity (Power, Barnes, Snyder, Schlaggar, & Petersen, 2012; Van Dijk, Sabuncu, & Buckner, 2012). In order to ensure that our analyses were not confounded by a condition-specific pattern of head movement, we examined head motion during the 12-second period following target presentation. We compared the root-mean-square of movement deviations in six estimated motion parameters (Van Dijk et al., 2012). No significant interaction was observed between conditions in our main analysis (high-confidence hits and misses) and these estimated movement parameters (F(1,5) = .76, p = .58).

Discussion

A key finding of the present study is the selective effect of target awareness on a specific subset of graph theory metrics associated with inter-modular connectivity: While modularity and inter-modular participation indices were affected by the awareness manipulation, there were no changes in clustering coefficient and characteristic path length. Interestingly, the latter two parameters are used to estimate a network's small-worldness (Rubinov & Sporns, 2010), i.e., a network's tendency to exhibit both high functional segregation and efficiency. The average path length (and the related global efficiency metric) is largely insensitive to multiple, long paths in larger networks (Estrada & Hatano, 2008), precisely the pattern of connections that appear to be most affected by awareness based on participation differences observed in our data. Additionally, existence of small-worldness is defined by strong local clustering (Humphries & Gurney, 2008) that may not be affected by the cross-modular changes seen here with awareness. Awareness of a simple percept may not significantly affect small-worldness because both of the measures used to compute this network feature are more sensitive to intra-modular local connections, whereas it evidently is the inter-modular connections that are preferentially affected by conscious target perception. As previous studies have identified broad increases in long-distance oscillatory synchrony as a marker of consciousness (Gross et al., 2004; King et al., 2013; E. F. Lau, Phillips, & Poeppel, 2008; Melloni et al., 2007; Schröter et al., 2012), it is conceivable that the increase in participation and breakdown of modularity that we report are a result of changes in long-distance functional synchrony following target awareness.

As a whole, the present findings provide strong evidence in support of global theories of awareness, in which the conscious perception of a stimulus is associated with whole-brain dynamic alterations in functional connectivity. Such results may explain why awareness is a unitary phenomenal experience (Seth, Izhikevich, Reeke, & Edelman, 2006), and suggest the means by which information at the focus of attention is broadcasted across the cerebral cortex. It is noteworthy that the changes associated with awareness observed here are more extensive than those reported on the basis of BOLD amplitude changes, which tend to show scattered cortical activation foci with awareness (Dehaene et al., 2001; Haynes et al., 2005; Lau & Passingham, 2006; Marois et al., 2000; 2004; Tse et al., 2005). Such a BOLD amplitude pattern of fronto-parietal activity was in fact observed in our own data when comparing Target Aware and Unaware conditions (see Figure 7). As such, these results imply – aside from differences in sensitivity between BOLD amplitude and functional connectivity approaches – that connectivity differences may reflect latent changes in the functional state of distant brain structures, priming them for target-related activation should the task or environmental situation call upon it.

Importantly, even though our results reveal global summary changes in modularity and participation that are not driven by massive connectivity changes in a single network or small number of nodes, they do not discount focal or network-level contributions to awareness. Instead, we view these global connectivity changes as a complementary aspect of awareness to the recruitment of fronto-parietal regions observed in numerous univariate studies of BOLD amplitude changes (Asplund, Todd, Snyder, & Marois, 2010a; Beck et al., 2001; Dehaene et al., 2001; Haynes et al., 2005; Lau & Passingham, 2006; Lumer et al., 1998; Lumer & Rees, 1999; Marois et al., 2000; 2004; Marois & Ivanoff, 2005; Naghavi & Nyberg, 2005; Rees et al., 2002; Tse et al., 2005). It is only once we are able to integrate the findings obtained at different levels of analysis that we will have a comprehensive understanding of the neural basis of awareness.

Finally, while our findings suggest that changes in a single network do not drive the global functional connectivity changes, not all network inter-connectivity is similarly engaged. It may very well be, for example, that visual and auditory awareness would reveal a different balance in functional weights across the brain's networks. Nevertheless, while more studies are evidently required to account for the differentiable aspects of consciousness (Tononi & Edelman, 1998), our results reveal a possible mechanism supporting the integrative nature of this mental state.

CHAPTER III

INCREASE IN THE BRAIN'S FUNCTIONAL SMALL-WORLDNESS WITH ATTENTION CAPTURE

The previous chapter considered the case in which consciously perceived information was goal relevant. However, salient but task-irrelevant stimuli often enjoy a privileged, distracting access to our mental faculties. Generally, novel and salient stimuli in the environment powerfully grab attention. The existence of such a stimulus-driven orienting response to unexpected events has clearly adaptive behavioral purposes and produces wide-ranging responses, both cognitively and physiologically (Pavlov, 1927; Sokolov, 1963). In particular, the onset of novel stimuli in the environment produces numerous physiological changes, including sympathetic nervous system activation, increased arousal, pupillary dilation, and galvanic skin response (Beatty, 1982; Kahneman & Beatty, 1966; Tracy et al., 2000). These stimuli also mobilize a host of related cognitive processes, including (but not limited to) the orienting of attention, evaluative processes, and planning of appropriate behavioral responses (Han & Marois, 2014; Kahneman, 1973; Sara & Bouret, 2012). Altogether, these processes prepare an individual to interact with the attention-grabbing event in accordance with its behavioral meaning and value. As such, the cumulative response to unexpected or sudden events in the environment can be considered a global one, recruiting aspects of both body and mind.

In contrast to the global effects resulting from the capture of attention by a surprise stimulus, neurobiological studies of attention have mostly focused on its selective neural properties (Chawla, Rees, & Friston, 1999; Corbetta & Shulman, 2002; Desimone & Duncan, 1995; Kastner et al., 1999; Reynolds & Chelazzi, 2004; Reynolds, Chelazzi, & Desimone, 1999). However, even such focal effects can trigger a cascade of neurophysiological changes that propagate across brain regions, for example via oscillatory mechanisms (Fries, Reynolds, Rorie, & Desimone, 2001; Womelsdorf, Fries, Mitra, & Desimone, 2005). Indeed, widespread changes in neurophysiological activity are an integral part of an influential model of attention – the biased competition model – in which it is postulated that attention biases the competition for representation of stimuli or events in the brain (Desimone & Duncan, 1995). Yet, evidence that attention produces truly global neural changes is scant. Attention networks of the brain have been well delineated, including a ventrolateral network composed of the temporo-parietal junction (TPJ), inferior frontal junction (IFJ) and anterior insula (AI)(Corbetta & Shulman, 2002; Downar et al., 2000; Marois et al., 2000; Serences et al., 2005) that overlaps with a related 'salience' network in AI and dorsal anterior cingulate cortex (dACC) (V. Menon, 2011; V. Menon & Uddin, 2010), as well as a dorsal attention network consisting of core components in the frontal eye field and dorsal parietal cortex (Beck et al., 2001; Chica, Valero-Cabre, Paz-Alonso, & Bartolomeo, 2014; Corbetta & Shulman, 2002; Kastner et al., 1999; G. L. Shulman et al., 1999). These studies, as well as many others, have ultimately produced a network-level view of attention, in which the control of attention results from the coordinated function of a circumscribed set of brain areas. However, as noted above, numerous physiological and cognitive processes take place following the capture of attention by an unexpected event in order to bring about the swift and systemic changes that are necessary to insure a prompt analysis, physiological reaction, and behavioral response to the event (Pavlov, 1927; Sokolov, 1963). Given the wide-ranging extent of these changes, they may likely originate from well beyond the confines of the attention networks. Moreover, the rapidity and flexibility at which an individual may have to interact with the surprising event could be optimized if cortical resources are primed to address this event. Thus, we hypothesize that attention capture by a surprising event involves widespread alterations in the functional topology of the human brain to facilitate the rapid exchange of information across the cerebral cortex.

Using functional magnetic resonance imaging (fMRI) combined with graph theoretical analyses, we sought to examine whether detection of salient, novel, task-irrelevant “oddball” events produced functional connectivity changes throughout the brain or whether they were mostly confined to the attention/salience networks. Graph theory is ideal for testing global functional connectivity predictions as it provides concise measures of the functional topology of an entire system (Bullmore & Sporns, 2009). Participants performed a demanding goal-oriented visual search task in which they monitored a rapid stream of distractors for the presentation of target letters, with the task disrupted by the rare presentations of unexpected, salient oddball images. Such oddballs powerfully capture attention and produce a strong orienting response (Asplund, Todd, Snyder, & Marois, 2010a; Asplund, Todd, Snyder, Gilbert, & Marois, 2010b; Braver et al., 2001; Marois et al., 2000; M. Stevens et al., 2005) that quickly habituates after the first few oddball instances (Asplund, Todd, Snyder, & Marois, 2010a; Asplund, Todd, Snyder, Gilbert, & Marois, 2010b; Sokolov, 1963). The rapid habituation of the orienting response makes a specific prediction about the life span of the changes in functional connectivity that may accompany attention capture. Specifically, repeated ‘oddball’ presentations should attenuate the changes in the brain’s functional topology brought about by the initial presentations, despite the continued presence of these task-irrelevant events. Consequently, we aim to examine two factors of the functional topology following oddball presentation: (1) whether the initial presentations of oddballs produce a global shift in the pattern of functional connectivity rather than network-level (e.g. attention network) changes and (2) whether such global changes mirror the behavioral habituation pattern that is observed with repeated presentations of the surprising events.

Methods

This study presents novel connectivity and graph theory analyses of data acquired in Asplund et al. (2010b). The graph analytical methods have been previously described in detail in Godwin et al. (2015) and in Chapter II of this dissertation.

Participants

Thirty-one individuals (12 females) recruited from the Vanderbilt University community participated in the study. All participants had normal or corrected-to-normal vision. The Vanderbilt University Institutional Review Board approved the experimental protocol, and informed consent was obtained from all subjects. Data from one participant was excluded due to technical problems during the scan session.

Behavioral Paradigm

The main task (Fig. 9) consisted of searching for a target letter (‘X’) presented in a rapid serial visual presentation (RSVP) of distractor letters (all letters were shown in white Helvetica font, 1.8° x 1.8°, overlaid on a dark gray background). Trials were 5.4s seconds long and consisted of a 3.4 s letter RSVP followed by a 2 s response period and a 2.6 s inter-trial-interval (ITI). RSVP periods consisted of 31 letters chosen randomly, without possibility of two identical letters occurring in a row, from a possible set of 20 with vowels being excluded from the possible set. All stimuli were shown for 100 ms with a 10 ms inter-stimulus-interval (ISI). An onscreen prompt directed participants to initiate a right index finger button press response to indicate "target present" or a right middle finger response to indicate 'target absent.' ITIs displayed a white fixation cross on the screen for the duration. The target was present on 77.5% of the 40 total trials presented over the course of one fMRI run per participant.

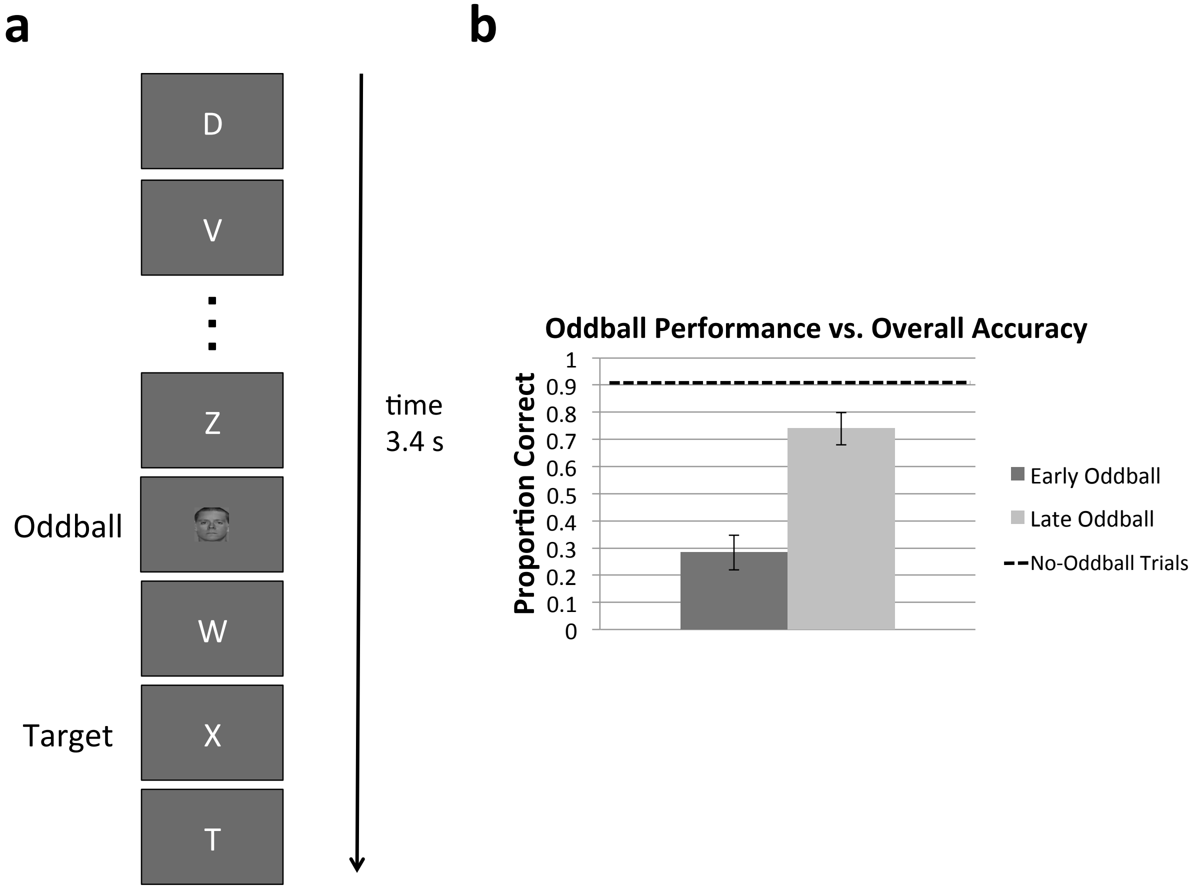

Figure 9. (a) Schematic of behavioral design. Each trial (8 s) began with a 3.4 s RSVP of 31 letters, with each letter presented for 100 ms with a 10 ms inter-stimulus-interval. Participants had 2 s to respond following each trial as to the presence of a target, indicated by a prompt on screen. A 2.6 s inter-trial-interval followed each response period while a white fixation cross remained on the screen. Size of stimuli shown in figure are not to scale. (b) Rates of target detection for Early (first two) oddball trials, late (last four) oddball trials and no-oddball trials.

On a subset of trials (6 out of 40) participants were shown a salient, task-irrelevant 'oddball' during the RSVP at fixation in place of a distractor letter stimulus. Oddballs consisted of grayscale faces, each distinct from one another, and were shown during both target-present trials (5) and no-target trials (1). On target-present trials, oddball images were displayed 330 ms prior to the target. Although only faces were used in the current study, we have previously shown that similar behavioral effects are obtained for other oddball types (Asplund, Todd, Snyder, & Marois, 2010a). The order of oddball trials, including the target-less oddball trial, was pseudo-randomized and was balanced across subjects. Oddball trials were separated by a minimum of two trials containing a target with no oddball stimulus. Only target-present trials were practiced prior to the fMRI session and participants needed to be proficient at 80% accuracy prior to scanning.

All stimuli were presented using Psychophysics Toolbox extensions (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997) in MATLAB (MathWorks).

fMRI Methods

The study was carried out in a 3T GE MRI system at the Vanderbilt University Institute of Imaging Sciences (VUIIS). Stimuli were back-projected from an LCD projector (Avotec, Inc.) onto a screen viewed by the participant lying in the scanner via a mirror. Nineteen 7-mm-thick axial slices were taken parallel to the anterior commissure–posterior commissure (AC-PC) line (0 mm skip; 3.75 × 3.75 mm in-plane). T2*-weighted image acquisition consisted of a 25 ms echo time, 70° flip angle, 240 mm FOV, 64 × 64 matrix, and a 2,000 ms repetition time. 166 brain volumes were collected during the functional scan, although the first 6 volumes were discarded to allow for signal stabilization.

Data preprocessing was performed using Brain Voyager QX 2.6 (Brain Innovation) and included 3D head-motion correction, linear trend removal, and a correction for slice acquisition timing. Functional and anatomical runs were co-registered and transformed into standard Talairach space (Talairach & Tournoux, 1988). Prior to the functional connectivity stage of analysis, six estimates of motion and the first derivatives of those estimates were regressed out of the data to account for the potential effects of head motion on connectivity estimates (Power et al., 2012).

Functional Connectivity Analysis

Regions of interest were defined according to the results of a previously reported analysis of resting state functional connectivity (Power et al., 2011), in which authors identified 264 nodes with strong overlap with known functional-network systems. Conversion of these MNIcoordinates into Talairach coordinate space was performed in the current study via the mni2tal.m function. Six-millimeter radius spheres were drawn around each coordinate to create a list of 264 ROIs to serve as nodes for graph theoretical analyses.

To assess how connectivity differed as a function of task, connections between nodes were measured using the generalized psychophysiological interaction (gPPI) method (McLaren et al., 2012). In using gPPI, connections are estimated based on correlations in blood oxygen level dependent (BOLD) signal in such a way as to separately estimate task activation effects and non-task specific correlations in order to isolate task-dependent changes in connectivity (Cisler et al., 2014; Gitelman et al., 2003; Kim & Horwitz, 2008). Three general classes of predictors are included in the general linear model (GLM) to simultaneously account for respective sources of variance: predictors of task activation, a seed node timecourse, and condition-specific interaction timecourses. Predictors of task activation are analogous to those used in standard fMRI GLMs. Interaction regressors predict changes in connectivity isolated to periods following the presentation of the respective stimuli while zeroing out predictions at other timepoints. The parameter estimates corresponding to the interaction predictors serve as our task-modulated measurement of functional connectivity.

In order to produce a full connectivity matrix, a separate GLM was performed for each of the 264 seeds (nodal ROIs). For each GLM, a timecourse was extracted from a seed region, included as a predictor in the GLM, and utilized to produce a seed-specific interaction regressor with task condition. Increases in the interaction (PPI) parameter estimates for a particular voxel in a given GLM indicates an increase in the synchrony of the seed timecourse and the target voxel specific to the particular task condition. We assumed no directionality in these predictors and, consequently, averaged across the diagonal in our per-subject connectivity matrices, thus producing a symmetric graph (Rubinov & Sporns, 2010).

GLMs included condition-specific regressors for the first two oddball trials ("Early" oddballs), the next four oddball trials ("Late" oddballs), target-only trials, and no-target trials. The decision to distinguish the initial two oddballs from the later four was based on the pattern of behavioral results indicating an exponential decrement in oddballs' capacity to capture attention with repeated presentations, with the first two presentations capturing more attention than the next four, as judged by target detection performance (Asplund, Todd, Snyder, & Marois, 2010a; Asplund, Todd, Snyder, Gilbert, & Marois, 2010b). In addition to the two oddball conditions (Early and Late oddballs), parameter estimates for interaction regressors to be input into the graph theoretical analyses were also obtained for the Search trials condition, as no-target trials provide a baseline for assessing connectivity and topology in the absence of distracting oddball images. Based on the comparison of target detection in the Late oddballs vs. Search trials (see Fig. 9), the Late oddballs still capture attention (see Results for a comparison of Early and Late oddball topologies against all oddball-absent trials). gPPI analyses were conducted using custom scripts in the BrainVoyager QX 2.6 and MATLAB environments.

Graph Theoretical Analysis

Graphs were constructed for each of our conditions of interest (Early oddballs, Late oddballs, and Search trials) and examined via the Brain Connectivity Toolbox (Rubinov & Sporns, 2010) for graph theoretical analyses. However, several initial steps were performed prior to the calculation of graph theoretic statistics. In particular, one must account for the effects of graph metrics due to the number of nodes, connections, and distribution of connections in a graph (Rubinov & Sporns, 2010). First, graphs were proportionally thresholded (to exclude the potential influence of noise connections on graph metrics) in a range from the top 5% strongest connections to 25% strongest connections at 5% intervals, similar to previously reported work (Cao et al., 2014; Dosenbach et al., 2007; Godwin et al., 2015; Rubinov et al., 2009). We present the results for top 10% of connections here, as well as validation with alternate thresholds. Following thresholding, matrices were rescaled to the range [0,1] by dividing connectivity values by the maximum value on a per-subject graph basis. Finally, following calculation of graph metrics, all values were "normalized" by comparing said values to those derived from the average of 100 null reference graphs. Reference graphs were constructed by utilizing the Brain Connectivity Toolbox function "randmio_und." This normalization step allows for the comparison of graphs across thresholds, which contain variable connection densities and can affect values of graph metrics (van Wijk et al., 2010).

Corrections for multiple comparisons of graph metrics was performed by gating comparison of pairwise conditions by the significant result of an omnibus one-way ANOVA per metric examined (Rubinov & Sporns, 2011). Given the non-independence of the network properties examined, no correction for multiple comparisons was performed based on the number of metrics tested, similar to previous work in the field (Cao et al., 2014; Godwin et al., 2015; Rubinov & Sporns, 2010; 2011). Additionally, the current study aims to assess globally averaged measures of topology, as opposed to testing the significance of individual nodes, negating the need for corrections for multiple comparisons based on the number of nodes (Bassett et al., 2011; Lynall et al., 2010; Schröter et al., 2012) or significance per voxel as a thresholding measure (Bassett et al., 2009; E. T. Bullmore & Bassett, 2011; Cole et al., 2014; Dosenbach et al., 2007; Rubinov & Sporns, 2011).

To get a more complete understanding of the connectivity changes associated with attention capture, we tested metrics of functional segregation, functional integration, and centrality of the graph system to test the "integration" and "small-worldness" hypotheses (Godwin et al., 2015; Rubinov & Sporns, 2010).

The integration hypothesis predicts changes in modularity and participation coefficient as a function of target awareness (Rubinov & Sporns, 2010). We utilize Newman's spectral algorithm for weighted matrices to determine the modularity of graphs, a measure assessing the degree to which connections can be segregated into non-overlapping "modules" with a maximal amount of within-module connections while minimizing between module connections. Weighted modularity (Q^w) was calculated by:

$$Q^w = \frac{1}{l^w} \sum_{i,j \in N} \left[w_{i,j} - \frac{k_i^w k_j^w}{l^w} \right] \partial_{m_i, m_j}$$

with weighted connections between nodes i and j ($w_{i,j}$), sum of all weights in a graph (l^w), weighted degree of a node (k_i), module containing node i (m_i), and $\partial_{m_i,m_j} = 1$ if $m_i = m_j$. The degree to which modularity increases reflects the extent of isolated processing within modules, independent of communication with nodes in other modules, and hence the segregation of these islands of information processing.

Utilizing the modules defined by the modularity algorithm, we can calculate the degree to which nodes communicate across modules; i.e. the participation coefficient (y_i^w) of those nodes, a measure of centrality of the graph's system. Nodes with high participation increase global integration by facilitating between-module communications. Specifically,

$$y_i^w = 1 - \sum_{m \in M} \left(\frac{k_i^w(m)}{k_i^w} \right)^2$$

where $k_i^w(m)$ is the weighted degree of connections between node i and nodes in module m.

The "small-worldness" hypothesis predicts changes in the balance between the demands for functional integration necessary for efficient global communication and functional segregation necessary for specialized processing (Bullmore & Sporns, 2009; Watts & Strogatz, 1998), potentially uniquely occurring in the clustering coefficient and characteristic path length.

The clustering coefficient highlights the degree to which information is shared across triads of nodes, reflecting local sites of dense interconnection. The clustering coefficient (C^w), where t_i is the number of triangles around node I, is calculated with the following formula:

$$C^w = \frac{1}{n}\sum_{i\in N} \frac{2t_i^w}{k_i(k_i - 1)}$$

The characteristic path length, a metric of functional integration reflecting the degree to which information is spread widely throughout the graph, measures the functional distance between any two nodes. Specifically, the characteristic path length (L^w) is defined the functional distance between nodes, where $d_{i,j}$ is the shortest path length between nodes *i* and *j*, computed with a inverse mapping of weight to length.

$$L^w = \frac{1}{n}\sum_{i\in N} \frac{\sum_{j\in N, j\neq i} d_{ij}^w}{n - 1}$$

Finally, Small-Worldness (S^w) can be calculated by the ratio of the normalized (by null network) clustering coefficient to the average path length:

$$S^w = \frac{C^w}{C^{rand}} \bigg/ \frac{L^w}{L^{rand}}$$

where C^{rand} and L^{rand} are clustering and average path length calculated on comparable random networks (Humphries & Gurney, 2008).

Results

Target detection in the fMRI run was worse in the first two oddball trials than in the last four ($t_{(29)}$ = 6.036, p < .001; Fig. 9b). Importantly, target detection in the last four oddball trials was also worse than in target-only trials ($t_{(29)}$ = 56.064, p < .001). Thus, although attenuated, attention capture still took place in the last four oddballs. Taken together, these behavioral results indicate that oddballs powerfully capture attention, and that the strength of this attention capture is contingent on the number of times oddball are presented.

The rapid attenuation of attention capture with serial oddball presentations provides a powerful means to assess whether potential global connectivity changes reflect attention capture *per se* rather than the physical presentation of the oddballs, as the latter is equated in all oddball trials. Indeed, based on the behavioral results, one would expect a temporal order effect of oddball presentations onto global brain connectivity, with largest effects with the first two presentations, followed by the next four oddball presentations, which is in turn followed by trials consisting of only goal-driven processes (Search trials).

To examine the effects of attention capture on global functional connectivity patterns, we assessed the pair-wise functional connectivity across 264 nodal regions of interest of the cerebral cortex (Power et al., 2011) via the generalized psychophysiological interaction (PPI) method (McLaren et al., 2012). From the pair-wise functional connections across these 264 nodes, we assessed summary properties of this ensemble functional cortical network (i.e. the graph) using graph theoretical metrics based on the top 10% strongest functional connections. Given the extensive behavioral effects produced by the initial presentations of salient, task-irrelevant stimuli, we hypothesized that these would be accompanied with pervasive changes in functional connectivity, reflecting the powerful capture of attention by the unexpected oddballs. We further predicted that such putative global connectivity changes would be attenuated with additional oddball presentations, though these changes would still be present relative to trials with no oddball presentations. Consistent with this hypothesis, a node-based search for such pattern of functional connectivity changes revealed a widespread distribution of progressively decreasing connectivity strengths with the three conditions of interest, namely the first two (Early) oddballs, last four (Late) oddballs, and No oddballs (Search trials; Fig. 10) with no discernable structure in the topological arrangement of these changes.

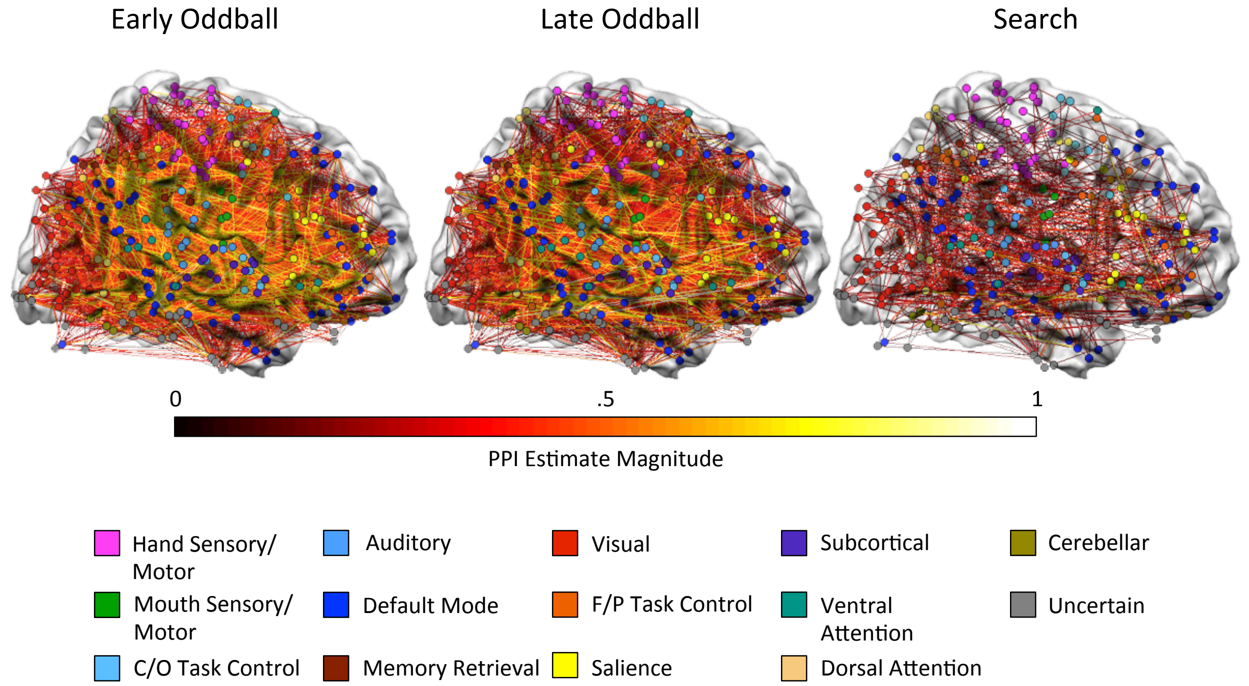

Figure 10. Subject-averaged connectivity for edges showing connectivity strength changes consistent with the behavioral pattern of Early oddball > Late oddball > Search trials, with hotter colored edges denoting increased magnitude of interaction regressor (seed x condition; PPI) parameter estimates. Heat scale indicates magnitude of interaction regressor (seed x condition; PPI) parameter estimates after rescaling to the range [0,1], as was performed for graph theoretical analyses. Node coordinates and network memberships were derived from Power et al. (2011), with nodes color-coded according to purported network membership. Node alignment with inflated surface is approximate and displayed for illustration purposes only. Abbreviations: C/O – Cingulo-opercular; F/P – Fronto-parietal.

It is difficult, however, to draw any further conclusions from the visual inspection of the changes in connectivity patterns rendered into two dimensions. To provide a quantitative and global assessment of the changes in functional connectivity patterns that occur with attention capture by oddballs, we submitted the PPI matrices data to graph theory analysis. Specifically, using graph theory metrics of functional segregation, integration, and centrality, we tested two competing models by which attention capture may affect the brain's global functional topology, namely via changes in "integration" versus "small-worldness." According to the integration model, suggested by a recent study about the global connectivity changes that take place with conscious target detection (Godwin et al., 2015), attention capture would lead to a breakdown of the brain's functional modularity (i.e. the ability to separate the graph into non-overlapping modules or networks), with such breakdown resulting from an increase in the functional participation of inter-modular connections. Alternatively, the small-worldness model instead argues that attention capture by oddball presentations would trigger an increase in the brain's functional small-worldness, a composite measure of the balance between functional specialization and efficient information propagation in the brain (Bullmore & Sporns, 2012; Watts & Strogatz, 1998). According to this model, the capture of attention by the presentation of oddballs would lead to oddball information being more efficiently processed locally and distributed globally. For both of these hypotheses, the predicted graph theoretical effects of oddball presentations would be greatest with Early oddballs, with diminished effects for Late oddballs and even lesser effects during Search trials.

The graph theoretical properties of Early oddball presentations, Late oddball presentations, as well as Search trials for the top 10% of functional connections are depicted in Figure 11. To address the first proposed account of global functional connectivity change – integration, or decreased modularity caused by an increase in participation coefficient – we looked for main effects of condition in modularity and participation coefficient. Functional modularity did show a main effect of condition ($F_{(2,29)} = 5.324$, $p = .008$), but it was counter to the pattern predicted in the integration hypothesis; modularity was higher in Early oddball trials than in Late oddball trials ($t_{(29)} = 2.249$, $p = .032$) and than in Search trials ($t_{(29)} = 2.98$, $p = .006$). No significant difference was observed between late oddball trials and search trials ($t_{(29)} = .825$, p

= .416). The greater modularity for the Early oddball condition suggests that functional modules were more easily segregated in the first few oddball presentations than in other trial types. The participation coefficient, which measures cross-modular connections defined by the modularity algorithm, showed no significant main effect of condition ($F_{(2,29)} = 2.102$, $p = .131$). Taken together, these results suggest that the capture of attention by the presentations of novel, salient oddballs does not break down the brain's network modularity; if anything, these oddballs increase it. Rather than causing global increases in functional integration, the initial oddball presentations – when the behavioral capture effect is most pronounced – displayed significantly more segregation compared to Search trials.

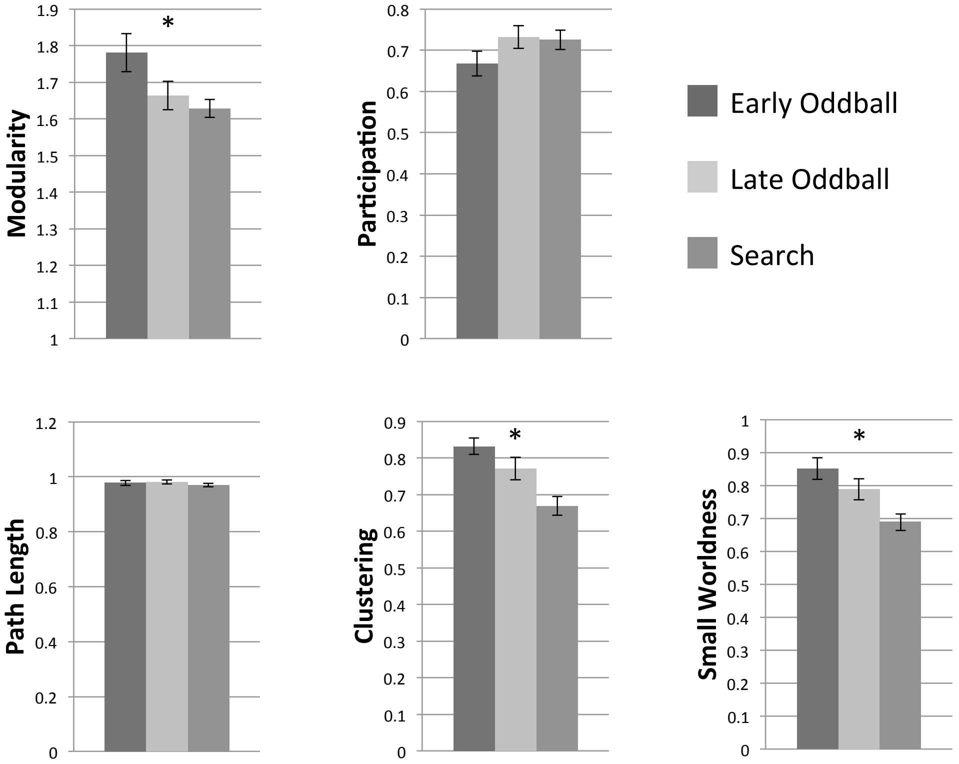

Figure 11. Effect of oddball manipulation on selected graph theoretic measures for Early oddball, Late oddball, and search events. All y-axis values represent the ratio of the observed graph theory metric to the corresponding random graph metrics (see Methods). Asterisks indicate significant main effects of condition at p < .05. Error bars represent within-group standard error.

To address the second proposed account of global changes in functional connectivity with attention capture, we first assessed the metrics that are factored in the computation of small-worldness: the clustering coefficient and characteristic path length. The clustering coefficient showed a main effect of condition ($F_{(2,29)} = 7.91$, $p < .001$), with higher clustering coefficient in Early and Late oddballs than in Search trials ($t_{(29)} = 3.599$, $p = .001$ and $t_{(29)} = 2.742$, $p = .01$, respectively), and a non-significant trend between the Early and Late oddball trials ($t_{(29)} = 1.471$, $p = .152$). For characteristic path length, there were no main effect of condition ($F_{(2,29)} = .547$, $p = .566$). Small-worldness, measured as the ratio of the normalized clustering coefficient over the normalized characteristic path length (Humphries & Gurney, 2008; Watts & Strogatz, 1998), showed a main effect of condition ($F_{(2,29)} = 7.539$, $p = .001$). Specifically, small-worldness was higher during Early and Late oddball presentations than during Search trials ($t_{(29)} = 3.521$, $p = .001$; $t_{(29)} = 2.531$, $p = .017$), with a non-significant trend between Early and Late oddball conditions ($t_{(29)} = 1.537$, $p = .135$).

The estimates of small-worldness for Early oddball trials, Late oddball trials and Search trials tended to match the behavioral data pattern (i,e. $\text{Oddball}_{early} > \text{Oddball}_{late} > \text{Search}$). In order to quantify the amount of evidence for this pattern in the graph theory data, we adopted a Bayesian inequality constrained analysis of variance (Klugkist, Laudy, & Hoijtink, 2005; Mulder et al., 2009). Specifically, we assessed how much more evidence there is for the constrained inequality hypothesis ($\text{Oddball}_{early} > \text{Oddball}_{late} > \text{Search}$; hereafter $H_{oddball}$) than for the alternative hypothesis that there is no effect of condition (i.e. $\text{Oddball}_{early} = \text{Oddball}_{late} = \text{Search}$; hereafter H_{null}). This method of comparison using constrained hypotheses generates allows for comparison across multiple constrained hypotheses by computing a Bayes factor relative to a common, unconstrained null. We found that small-worldness showed substantially more evidence (Bayes factor = 5.45) for the hypothesis $H_{oddball}$, compared to the hypothesis H_{null} (Bayes factor = 0.88). There was also more evidence for $H_{oddball}$ than for the alternative hypothesis that $\text{Oddball}_{early} = \text{Oddball}_{late} > \text{Search}$ (Bayes factor = 2.97).

It is conceivable that the small-worldness results were caused by large changes in a restricted set of nodes in the overall graph – such as those corresponding to attention networks associated with the orienting of attention to unexpected stimuli – instead of emanating from changes distributed throughout the cerebral cortex. To address this issue, we examined the graph metrics of small-worldness and clustering coefficient per network rather than for the entire graph. If one or a few networks dominate changes in clustering or small-worldness with condition, there should be an interaction between network and condition to the graph theory metrics. As shown in Figure 12, however, there were no interactions between network and condition for either clustering ($F_{(26,754)} = .918$, $p = .583$) or small-worldness ($F_{(26,754)} = .978$, $p = .496$); all networks appear to show the same trend observed in the global results, with mean clustering coefficients and small-worldness greatest for Early oddballs and smallest for the Search trials. Thus, it does not appear that network-specific modulations in graph metrics drive the overall global results in clustering coefficient or small-worldness.

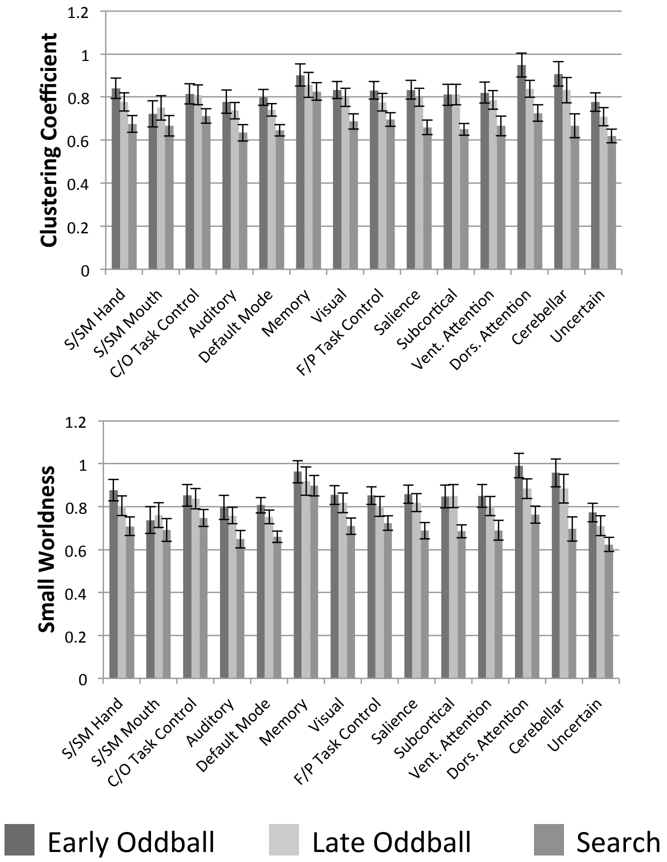

Figure 12. Effect of oddball manipulation on graph theoretic measures of clustering coefficient and small-worldness averaged across Power et al. (2011) modules. No interaction between network and condition was observed for any of the examined metrics. Error bars represent within-group standard error.

Another, and more direct, means to assess whether the global results of small-worldness may be driven by the attention networks is to employ a targeted attack procedure (Bullmore & Sporns, 2009), which consists in assessing the impact of excluding the ventral and dorsal attention networks, respectively, from the graph theory analysis. If excluding these nodes from the graph lead to significant changes in small-worldness (and clustering coefficient), then one can conclude that the attention networks are critical to the global changes in small-worldness with oddball presentations. Aside from excluding these nodes and the edges they formed with other portions of the overall graph, the analytical approach was identical to that of the globally averaged analysis described above. When excluding the ventral attention network, main effects of condition for the top 10% of connections were still observed for clustering coefficient ($F_{(2,29)} = 7.51$, $p = .001$) and small-worldness ($F_{(2,29)} = 7.178$, $p = .002$), with similar pair-wise differences as observed with the entire graph. Similarly, targeted attack of the dorsal attention network yielded significant effects of condition for clustering ($F_{(2,29)} = 6.109$, $p = .004$) and small worldness ($F_{(2,29)} = 5.809$, $p = .005$), with similar pair-wise differences as observed with the entire graph. Thus, the results of the targeted attacks of the ventral and dorsal attention networks mirror those obtained with the entire graph. These findings suggest that the global changes in functional connectivity with the presentations of novel, salient oddballs do not uniquely occur in attention networks. This is not to say, however, that the attention networks do not play critical roles in the responses to attention-capturing events (Asplund, Todd, Snyder, & Marois, 2010a). Rather, our findings suggest that distinct from BOLD amplitude-based changes in these attentional networks, there is also a more diffuse, global set of changes that accompany detection of task irrelevant, oddballs.

Head motion

Head motion has been shown to have significant effects on functional connectivity and graph theoretical measures (Drakesmith et al., 2015; Power et al., 2011; Van Dijk et al., 2012). We compared the root mean square of movement deviations in six estimated motion parameters (Van Dijk et al., 2012) to ensure that our analyses were not confounded by a condition-specific pattern of head movement. No significant interaction was observed between conditions in our main analysis (early oddballs, late oddballs, and search trials) and these estimated movement parameters ($F_{(2, 5)} = 1.48$, $p = 0.147$), suggesting that our principal graph theoretical results are not differentially affected by head motion.

No-oddball trials

Early and late oddballs were compared against trials with no oddball or target presented in order to compare the stimulus-driven capture effects of the oddball vs. an uninterrupted goal-driven state. However, in order to determine whether the inclusion of target-present trials would change the pattern of results observed, we recomputed PPI values and re-ran the globally averaged analyses to compare early oddballs, late oddballs, and all non-oddball trials (both target-present and target-absent).

Using this comparison, modularity no longer displayed a main effect of condition ($F_{(2,29)} = 1.574$, $p = .216$). However, the clustering coefficient ($F_{(2,29)} = 6.029$, $p = .004$) and small-worldness ($F_{(2,29)} = 5.422$, $p = .007$) showed similar main effects of condition as observed in the main analysis. The participation coefficient ($F_{(2,29)} = 1.393$, $p = .257$) and characteristic path length ($F_{(2,29)} = .303$, $p = .740$) all showed no effect of condition, similar to the main analysis.

As for pairwise comparisons, the clustering coefficient showed significant differences between early oddball trials and search trials ($t_{(29)} = 3.055$, $p = .005$) as well as between early and no-oddball trials ($t_{(29)} = 2.393$, $p = .023$), with no significant differences between early and late oddball clustering ($t_{(29)} = 1.208$, $p = .237$). Small-worldness showed significant differences between early oddball and no-oddball trials ($t_{(29)} = 2.921$, $p = .007$) as well as late oddball and no-oddball trials ($t_{(29)} = 2.345$, $p = .026$), with no significant difference between early and late oddballs ($t_{(29)} = 1.025$, $p = .314$). These results largely mirror the pairwise differences observed in the main analysis, supporting our conclusions that measures of segregation differ amongst conditions as a result of the detection of oddball stimuli.

Thresholding

Thresholding of graphs ultimately serves to exclude noise connections from consideration when calculating summary graph theoretical measures; even a small number of noise connections can have drastic effects on descriptions of graphs (Drakesmith et al., 2015). However, various graph measures are sensitive to the number of nodes and edges included in this calculation; consequently, analyzing only a single threshold places a constraint on the interpretability of the metrics (van Wijk et al., 2010). To consider the possibility that our results might be unique to a particular graph threshold, we repeated our globally averaged analyses at several additional thresholds ranging from the top 5% of connections to 25%, at 5% intervals. The results of these analyses are depicted in Figure 13. To test whether our results were unique to the 10% threshold described in the main analysis, threshold was included as an additional factor in a two-way analysis of variance (ANOVA). We examined several aspects of the results at these additional thresholds. First, do we replicate the main effects of condition? Next, is there an interaction between condition and threshold for our metrics? If an interaction exists, does this interaction conflict with our originally presented pattern of results?

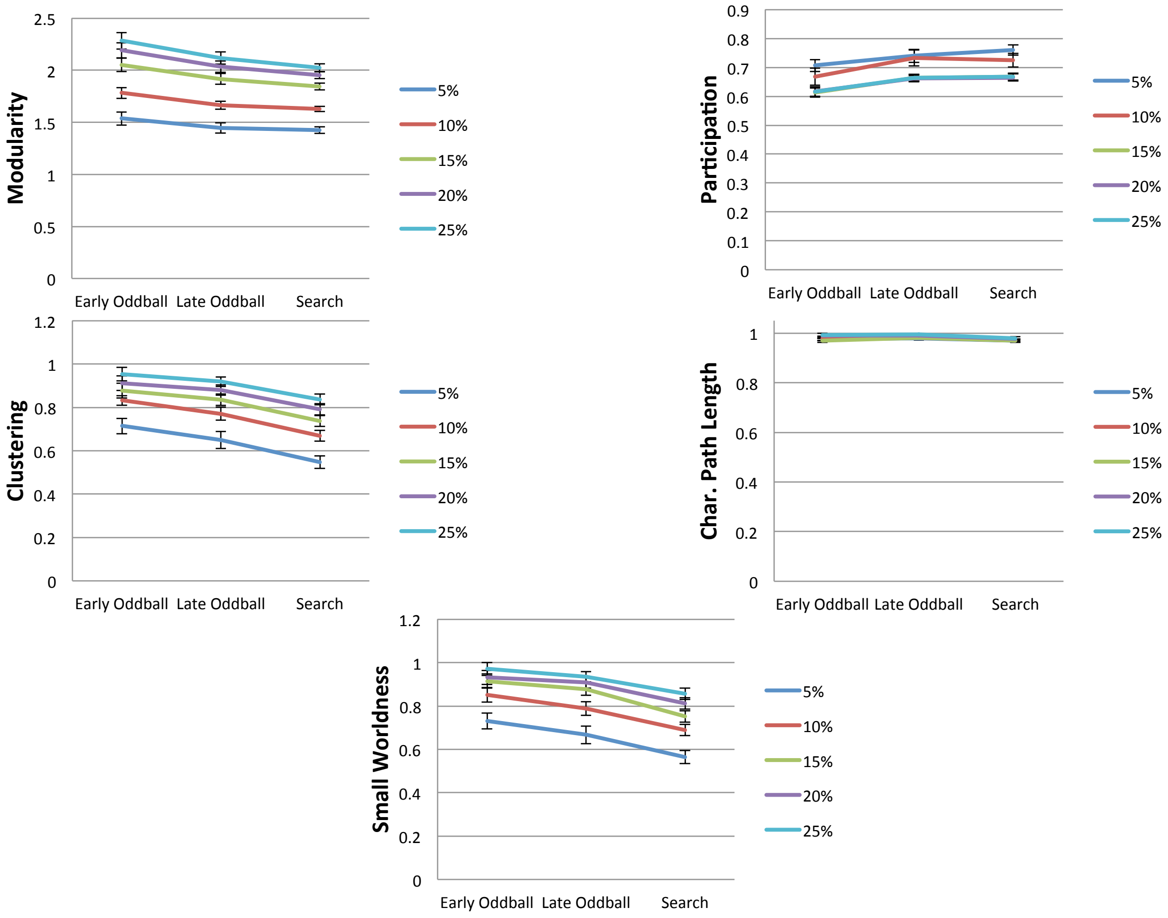

Figure 13. Graph metrics calculated across the range of the top 5% of connections to 25% of connections, at 5% increments. The effects of condition on modularity, participation, clustering, characteristic path length, and small-worldness are generally consistent across thresholds.

Modularity ($F_{(2,58)} = 7.083$, $p = .004$), clustering coefficient ($F_{(2,58)} = 6.408$, $p = .003$), and small worldness ($F_{(2,58)} = 8.313$, $p = .002$) all showed main effects of condition, similar to the results observed in the main analysis when only considering our *a priori* threshold of 10%. Some of these measures did show an interaction of condition and threshold, namely modularity ($F_{(8,232)} = 3.414$, $p < .001$), and small-worldness ($F_{(8,232)} = 2.385$, $p = .022$). However, as Figure 13 depicts, these metrics never appear to show a qualitatively different pattern of results across thresholds. We conclude from these results that our findings are generally robust across the range of thresholds tested.

Discussion

The results of the present study indicate that detection of task-irrelevant oddball images produces specific changes in global functional connectivity as assessed by graph theory. In particular, oddballs produced global changes in measures of functional segregation – namely modularity and clustering coefficient – as well as in small-worldness. The network and targeted attack analyses suggest that these changes were not specific to a subset of nodes or networks, but rather reflected truly global changes.

One of the two tested hypotheses, the "integration" account, predicted that the capture of attention by novel, task-irrelevant oddballs would cause increases in measures of global functional integration. The present results are inconsistent with this hypothesis. Indeed, rather than finding clear changes in integration metrics, we observed changes in measures that typically reflect properties of localized connectivity, with these effects culminating in significant changes in the brain's small-worldness, in direct support of the second hypothesis of global functional connectivity change with the capture of attention.

Small-world networks are often described as representing an efficient balance of functionally specialized and densely interconnected sub-networks with robust long-distance integration across these more specialized sub-components (Bullmore & Sporns, 2012; Rubinov & Sporns, 2010; Tononi et al., 1994; Watts & Strogatz, 1998). A wide variety of physical, technological, and social networks produce organizations that fall into this class of networks, a ubiquity that suggests these networks likely provide a unique efficiency in global information transfer – owing to a low average path length – without sacrificing the ability to perform more specialized processes by densely interconnected subcomponents with high clustering (Watts & Strogatz, 1998). Our results show that the changes associated with presentation of a task-irrelevant oddball specifically correlate with increases in clustering, a key component of small-worldness. It is possible that this increase in clustering reflects the panoply of specialized processes that are recruited with the capture of attention by a novel event (Pavlov, 1927; Sokolov, 1963), including primary sensory systems, executive control processes, and ultimately motor systems. According to this view, the increase in clustering coefficient is a result of widespread priming of functional neural connectivity in preparation for further processing of the oddball event, should it be advantageous for any given neural node or network to be involved in the swift processing of that event. Such large-scale increase in functional clustering may be brought about by activation of the locus coeruleus, as this region of the brainstem is known to briskly respond to novel events and to have widespread neuromodulatory effects in the cerebral cortex to facilitate rapid reorganizations of distributed functional networks (Sara & Bouret, 2012).

While we hypothesized that the capture of attention by a novel, salient event would trigger a surge in the brain's functional small-worldness, we did not expect that this change would be brought about by an increase in local functional connectivity rather than, or in addition to, a decrease in path length. A lack of effect on path length may be due to functional networks already being very efficient with near-optimal short-path lengths (Latora & Marchiori, 2001; van den Heuvel, Stam, Boersma, & Pol, 2008). Consistent with the present findings, a magnetoencephalography (MEG) study of the presentations of rare, deviant tones among standard tones reported that the deviant tones produced an increase in the clustering coefficient and other measures of localized processing (Nicol et al., 2012) rather than global increases in functional integration. What we show here is the profound functional implication that a change in clustering coefficient with the powerful capture of attention by oddballs can have on the brain's functional topology, and that is to increase its small-worldness, a fundamental and ubiquitous characteristic of efficient neural networks (Bullmore & Sporns, 2012; Humphries & Gurney, 2008; Watts & Strogatz, 1998).

An intriguing outcome of the present findings is how much they contrast with those obtained in Chapter II, in which conscious target perception was associated with large-scale increases in measures of functional integration and long-distance connectivity, consistent with global network models of awareness (Baars, 2005; Dehaene & Changeux, 2011; Ekman et al., 2012). Here we show that the powerful capture of attention by novel, task-irrelevant oddball stimuli produces changes in local connectivity patterns across the whole brain, similar to the profound and reliable pattern of effects observed in previous studies of capture by oddball stimuli (Asplund, Todd, Snyder, & Marois, 2010a; Braver et al., 2001; Marois et al., 2000; M. Stevens et al., 2005).

One possibility to account for these marked differences in functional dynamic topology may be the goal relevance of the attended or consciously perceived information. Greater long-range functional efficiency and functional integration has been shown to correlate with improved performance in goal-oriented working memory (Bassett et al., 2009; Kitzbichler et al., 2011) and visual discrimination tasks (Ekman et al., 2012). By contrast, the stimuli examined in the current study as well as in Nicol et al. (2011) can be considered distractors in the context of an ongoing goal-relevant task (though only the present data showed these distractors to interrupt ongoing goal-oriented behavior). It is therefore possible that functional integration is critical for the successful completion of task goals, whereas increased functional segregation emerges with the presentation of distractors that interrupt ongoing processing and capture attention from the goal-directed task sets. Functional integration and segregation may thus play specific and countervailing roles when it comes to processing of task-relevant and -irrelevant information. Alternatively, is also conceivable that rather than task relevance, it is the cognitive state that accounts for the differences in findings between Chapters II and III. That is, to the extent that attention and awareness can be dissociated (Koch & Tsuchiya, 2007; Lamme, 2003), it is possible that these two states are associated with distinct functional topologies. Given the significance that support for either of these two hypotheses would have in regards to understanding global brain functions, they merit further investigation.

In conclusion, the present findings indicate that attention capture by a novel, task-irrelevant event triggers a shift in the brain's small-world topology. Despite these changes occurring among local connectivity measures, they were observed across the entire graph system. It is perhaps through widespread increases in local connectivity that the brain is able to mobilize a wide variety of physiological and cognitive processes with attentional capture, thereby rendering the brain's functional topology a smaller world for swiftly and adaptively analyzing and responding to the attention capturing event.

CHAPTER IV

GLOBAL NETWORK PROPERTIES OF GOAL-RELEVANT AND -IRRELEVANT COMPETITION

Chapters II and III of this dissertation aimed to characterize awareness of percepts characterized as being respectively relevant (targets) and irrelevant (oddballs) for the ongoing cognitive tasks. Two qualitatively different patterns of global connectivity change correlated with each of these stimuli types. In everyday life we are constantly bombarded by stimuli, only a fraction of which will be relevant to our ongoing behavior. However, all of these stimuli, both relevant and irrelevant, will compete for representation and higher-level processing in the brain. When searching for a target among distractors that share target features, the amount of time it takes to find the target increases with the number of distractors (Neisser & Lazar, 1964; Wolfe, 2003) as well as with distractor similarity to the target (Duncan & Humphreys, 1989). When distractors are visually salient, they will attract attention away from goal-directed processes in a stimulus driven manner (Phelps, Ling, & Carrasco, 2006; Theeuwes, Kramer, & Kingstone, 2004; Yantis & Egeth, 1999). Additionally, when distractors share some stimulus feature(s) with targets, they will tend to capture attention via a hybrid of bottom-up and goal-driven processed referred to as contingent capture (Folk & Remington, 1992).

Both targets and nontargets are thought to contribute to a competition between stimuli fighting for access to the brain's information processing pathways. The biased competition model of selective attention (Desimone & Duncan, 1995; Duncan, Humphreys, & Ward, 1997) asserts that the processing of both goal-relevant and goal-irrelevant stimuli will be influenced by bottom-up and top-down biasing factor. These biases preferentially advantage each stimulus type to ultimately win this competition so that potentially relevant information reaches awareness. Competition appears to arise as multiple sources of information interact at common neural representations to mutually suppress one another (Beck & Kastner, 2009; Kelley & Lavie, 2011; Reynolds et al., 1999).

The sources of these biasing signals have often been linked to two separable cortical networks that control goal-driven and stimulus-driven orienting (Chiu & Yantis, 2009; Corbetta, 1998; Corbetta & Shulman, 2002). The goal-driven control of attention is localized to a dorsal attention work, comprised of bilateral human frontal eye fields (FEF) and intraparietal sulcus (IPS) (Corbetta & Shulman, 2002; Desimone & Duncan, 1995; Kastner et al., 1999). Stimulus-driven attention, recruited by salient or novel items in the environment, has been localized to the temporo-parietal junction (TPJ) and inferior frontal junction (IFJ) (Corbetta et al., 2008; Corbetta & Shulman, 2002; Marois et al., 2000) and is closely tied to a related network composed of the anterior cingulate cortex (ACC) and anterior insula (AI) tasked with indicating salience of objects in the environment(Han & Marois, 2014; V. Menon, 2011; V. Menon & Uddin, 2010).

The previous two chapters have described connectivity patterns that may reflect this biasing in favor of processing goal-relevant and goal-irrelevant stimuli, respectively. In Chapter II, awareness of the target produced a decrease in the modularity and increase in the functional participation of whole brain graphs relative to unperceived target trials. This pattern of results replicated across a comparison of hits to misses as well as a comparison of false alarms to correct rejections. Having perceived a target, regardless of whether one was actually presented, produced a consistent decrease in the functional separation between networks, brought about by increased cross-modular connections. Generally, these results are consistent with global theories of awareness including global workspace theories (Baars, 2005; Dehaene & Changeux, 2011) and information integration theory (Tononi, 2008; Tononi & Edelman, 1998). However, what this study could not address was whether information that was not relevant for an ongoing task would be distributed globally in a similar manner.

Chapter III assessed the graph properties of a task-irrelevant oddball stimulus to see whether information pertaining to distracting stimuli, rather than a task-relevant stimulus, entered a similar globally connected workspace. Rather than being the target of the ongoing task, these novel images decreased detection rates of the target, especially for the initial oddball presentations (Asplund, Todd, Snyder, & Marois, 2010a; Asplund, Todd, Snyder, Gilbert, & Marois, 2010b). Although these stimuli do not contain the goal-relevance as in the first study, the powerful orienting response they provoke does require the coordination of numerous distinct cognitive and physiological processes in order to orient to and evaluate novel stimuli (Han & Marois, 2014; Kahneman & Beatty, 1966; Sara & Bouret, 2012). Consequently, we might expect them to display similar widespread connectivity properties as were observed during detection of a task-relevant target. Ultimately, we found that global connectivity changes were associated with detection of a task-irrelevant stimulus, but that this pattern was qualitatively different than the one observed for target awareness: increased clustering drove an increase in the small-worldness of subjects' functional graphs during oddball detection, reflecting an increase in functional segregation. This increase in functional segregation was reflected in an increase of the modularity statistic, reflecting greater functional distinctiveness between modules.

Together, Chapters II and III suggest that the task-relevance of a stimulus (i.e., a stimulus that matches a target profile vs. an irrelevant or novel stimulus to be evaluated) may alter global patterns of connectivity, helping to communicate this information to appropriate brain structures for ongoing cognition. One alternative explanation for the source of these differences may come from a dichotomy between how task instructions are used to proactively bias information versus a more reactionary form of control brought about by the stimulus-driven properties of information at the moment it is presented (Braver, 2012; Braver, Paxton, Locke, & Barch, 2009). This "dual-mechanisms of control" (DMC) model predicts that expectations formed during cue presentation bias an individual to adopt either a more "proactive" form of control vs. a "reactive" control. This proactive bias is primarily described as primarily a goal-oriented form of control (Braver, 2012; E. K. Miller & Cohen, 2001). Alternatively, this model theorizes a second, reactive, form of control (Braver, 2012; Braver et al., 2009) that is more stimulus-driven in nature and tends to be adopted when expectations of distractor interference are low. Here, the dissociation between the effects of task-relevant and –irrelevant information is embodied primarily in whether the bias is cued in preparation for target presentation or whether stimulus-driven factors, such as distractor interference, necessitate more online control to filter out conflicting information.

In the present study, we aimed to test the relationship between functional topology and the processing of task-relevant and -irrelevant information. Awareness of task-relevant information may be crucially communicated and integrated throughout the brain by a highly interconnected set of global connections (Baars, 2005; Dehaene et al., 1998; Edelman, 2003; Tononi, 2008). These hypotheses were supported by the results of Chapter II in which detection of task-relevant stimuli (e.g. targets) produced decreases in functional modularity and increases in the functional participation of the whole brain. Consequently, for the current study this pattern of increased functional integration (indexed by decreases in modularity and increases in the participation coefficient) will serve a marker of task-relevant information processing. However, the presence of distracting information (e.g. oddball presentations) in Chapter III produced increases in functional segregation. We hypothesized that functional isolation (indexed by an increase in modularity, clustering coefficient, and small-worldness) may serve to prevent irrelevant information from out-competing goal-relevant information. Therefore, we would expect processing of task-irrelevant information would display these same patterns of functional segregation. A schematic of these hypothesized markers of task-relevant and task-irrelevant processing is shown in Figure 14.

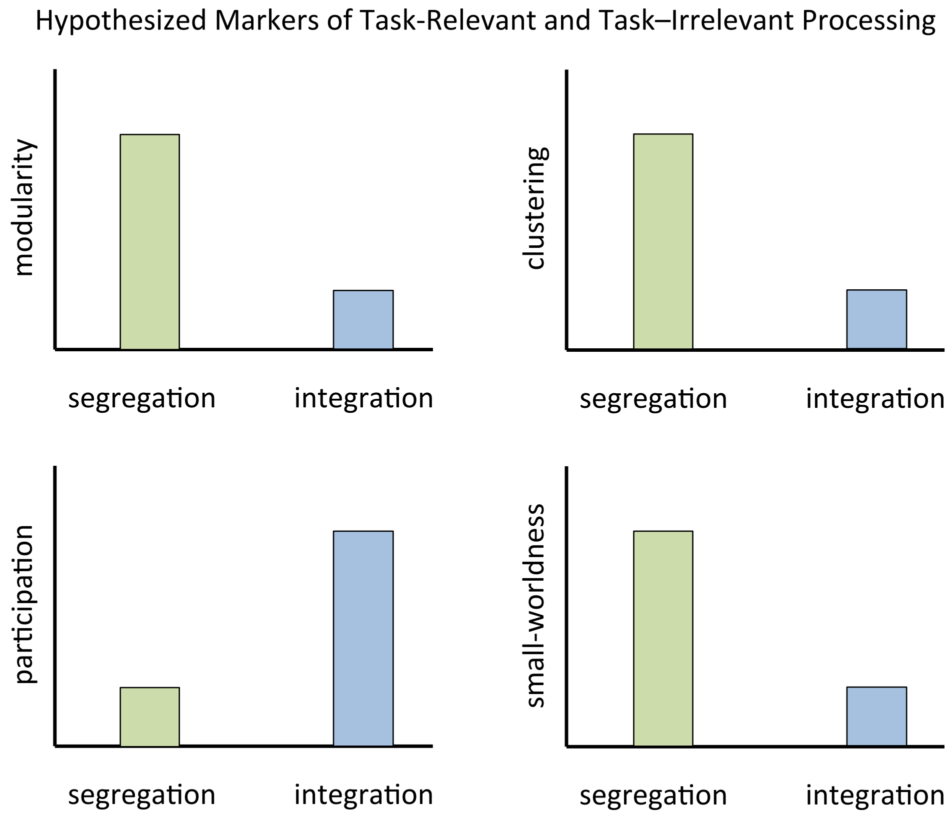

Figure 14. Hypothesized markers of task-relevant and task-irrelevant processing. Based on the results of Chapter II and III, the figure shows hypothesized markers of global functional connectivity associated with concurrent goal-relevant and goal-irrelevant information processing. No scale is given as these patters are intended to be relative changes in connectivity in comparison to other task periods. Goal-relevant processing, informed by the results of Chapter II, suggest that global functional integration (decreased modularity and increased participation) accompanies goal-relevant information. Chapter III showed increased global segregation (increased modularity, clustering, and small-worldness) associated with processing of task-irrelevant stimuli.

However, the data presented in Chapters II and III only measured functional changes concurrent with the presentation of targets and oddballs. The DMC theory suggests that when task-irrelevant information is likely, as indicated by an informative cue, proactive control is employed to optimally bias, in a sustained manner, the processing of goal-relevant information (Braver, 2012; Braver et al., 2009). When the likelihood of task-irrelevant information is low, reactive control only employs biasing procedures at the moment irrelevant information is detected (Braver, 2012; Braver et al., 2009). Therefore, from a DMC theory perspective, we would expect to observe differential patterns of global functional connectivity when distracting information is highly likely vs. when it is unlikely.

To test the differential predictions of the dual-mechanisms of control theory and predictions based only on immediate task-relevance, I reanalyzed a previously published experiment designed to probe preparatory influences on activity in visual cortex (Serences, Yantis, Culberson, & Awh, 2004b). Participants were asked to add two target numbers and make an immediate parity judgment. Targets were identified by spatially predictive cues that contained additional information about the likelihood of visual distractors during the task. Distractors were invalidly cued on a subset of trials allowing for the examination of cue-related preparatory activity as well the consequences of violations of expectations.

Methods

The data utilized in this study has been previously published as “Experiment 1” in Serences et al. (2004). The connectivity and graph theoretical methods described here are a novel analysis of this dataset aimed at addressing hypotheses formed based on the previous chapters of this dissertation. Additional details of the experimental design can be found in Serences et al. (2004) and Awh et al. (2003).

Participants

Twelve adults (10 female, aged 20-31) participated in this study. All participants provided informed consent and The Johns Hopkins University institutional review board approved the experimental design.

Behavioral Paradigm

The stimulus display consisted of white items overlaid on a black background. The trial design is shown in Figure 15. Each trial consisted of a cue period, target array, and pattern mask. The length of the cue period was drawn from a pseudo-exponential distribution ranging from 2500-8500 ms (mean 4,375 ms). Cues consisted of a center fixation square as well as two circle and two square stimuli presented at the corners of an invisible square, centered on the fixation square. Each pair of matching cue stimuli (circles, squares) was presented diagonally from one another. Half of all participants were instructed to attend to locations indicated by the square-cued locations, with the remaining participants instructed to attend to the circle-cued locations. The task required participants to add digits at the cued locations and provide an immediate odd/even judgment of the summed result. Digits appeared at all four locations on every trial, thus targets were defined by the appropriate cue location (and therefore cues were 100% valid with respect to target location).

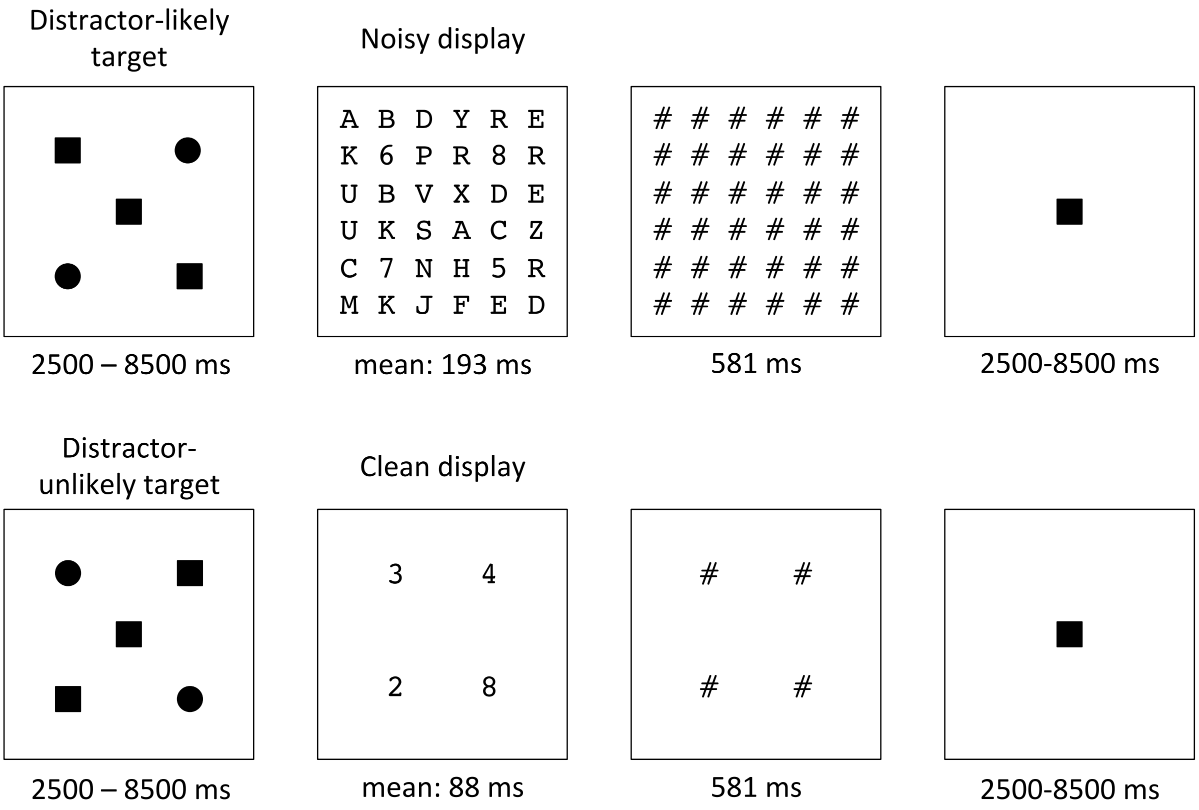

Figure 15. Task design depicting valid trial types. Participants were instructed to monitor two cued locations, indicated by either circle or square items in the cue display. The orientation of these cued locations indicated whether there were likely (top) or unlikely (bottom) to be distractors presented on that particular trial. Participants were instructed to add numbers in the target locations together and indicate whether the sum was even or odd with a button press. Figure adapted from Serences et al., (2004).

Participants were instructed that the location of the cue stimuli would indicate the likelihood of upcoming distractor stimuli to be presented during the target array. For half of participants, if the appropriate cue stimuli were presented in he upper left and lower right corners of the invisible square (unlikely-distractor cues), the target array was very likely to only consist of numbers presented at the four possible locations (clean target array). For the alternative orientation, lower left and upper right cues (likely-distractor cues) would indicate that the target array would likely be filled with irrelevant letters stimuli filling the screen at regular locations (noisy target array). Various aspects of cue meaning were counterbalanced across subjects. Distractor likelihood was cued with 80% validity; on 20% of unlikely-distractor trials participants would view a noisy target display during the target array period and on 20% of likely-distractor cues participants would view a clean target display during the target array period. Pattern masks consisting of "#" symbols at each target and distractor location were presented immediately following the target array and were displayed for 581 ms. No mask was presented at empty locations during clean target display trials.

Duration of the target array was determined prior to scanning by a staircasing procedure performed by each subject, independently for clean and noisy displays (only valid trials were presented during the staircasing procedure). This procedure aimed to equate accuracy across conditions, keep performance off of ceiling, and produce exposure durations too fast to successfully saccade between targets. The staircasing procedure titrated the exposure duration of the target to produce about 75% accuracy on both target array types. Displays were initially presented for each condition for 1,670 ms. The exposure duration decreased by 10% following correct responses and increased by 30% following incorrect response during this process. Eight blocks of the staircasing procedure, consisting of 40 trials each, were performed on the day prior to scanning. Average exposure durations during the final two blocks during this session were then used as a starting point for an additional three blocks of the staircasing procedure after the participant had entered the scanner, but prior to scanning. Exposure durations were manually adjusted during scanning between runs if participant performance deviated largely from around 75% accuracy. This procedure resulted in a mean exposure duration of 193 ms (41 ms standard deviation [SD]) during noisy displays. Clean displays, without distractors, resulted in a mean exposure duration of 88 ms (31 ms standard deviation) (Serences et al., 2004b).

fMRI Methods

The experiment was performed utilizing a Philips Intera 3T scanner at the F.M. Kirby Research Center for Functional Brain Imaging at the Kennedy Krieger Institute in Baltimore, MD. T1-weighted anatomical images were acquired with a 1mm isotropic voxel resolution. Functional scans utilized a SENSE (MRI Devices, Waukesha, WI) head coil with a SENSE factor of 2. Echoplanar functional image (EPI) acquisition parameters consisted of a repetition time (TR) of 1,830 ms, 25 ms echo time (TE), 70° flip angle, 240 mm FOV, 80 × 80 matrix. 26 axial slices were collected with a 3 mm thickness and a 1 mm gap. 10 functional runs were acquired, each consisting of 40 trials and 200 time points. In addition to the functional scans, two localizer scans were acquired for retinotopic mapping which utilized the same scanning parameters. These localizer scans are not considered in the present analysis.

Data preprocessing was performed using Brain Voyager QX 2.6 (Brain Innovation) and consisted of 3D motion correction, correction for slice scan timing, and linear trend removal. Anatomical and functional runs were co-registered and transformed into Talairach space (Talairach & Tournoux, 1988). Numerous sources of noise have been shown to negatively affect accurate estimation of functional connectivity (Power et al., 2012) including subject motion and cardiac/respiratory influences. Prior to estimation of functional connectivity measures, six estimates of motion, their first derivatives, as well as the timecourses from four regions of no interest (two white matter, two ventricular ROIs of similar size to spherically defined ROIs in the current study) were regressed out of the data to account for these spurious sources of variance.

White matter and ventricular regions of non-interest were manually defined on a per-subject basis.

Functional Connectivity Analysis

Estimation of functional connectivity was performed similarly to the experiments described in the previous chapter, utilizing 264 nodes of interest defined by Power and colleagues (2011) with known overlap with cognitive networks of interest. Six-millimeter radius spheres were drawn around each coordinate reported to form the 264 nodes with which to the whole brain graph for analysis. Connectivity estimates were derived from the condition-specific interaction regressors included in the GLM using the gPPI method.

Initial analysis of the data included four condition-specific regressors for cue periods and four condition-specific regressors for target periods. Cue specific conditions were defined by cue type (likely-distractors vs. unlikely-distractors) as well as what target array type was presented (clean target vs. noisy target). The original authors for this dataset were concerned that there was not adequate time between presentation of the cue and the target to distinguish the sources of these two signals (Serences et al., 2004b). Therefore, to test whether target-related activity affected estimates of cue activity, separate conditions for valid and invalid cues were included in the model. The authors predicted an interaction between cue type and target type would be indicative of potential cross-contamination of the two signals. We employ a similar test to assess these effects in our own analysis. Target-related conditions were similarly separated by the cue type and target-display type. Ultimately, this produced a 2 (distractor likely vs. unlikely) x 2 (clean vs. noisy target display) x 2 (cue vs. target periods) design. A secondary analysis was performed, combining cue and target conditions into a single condition beginning at the cue presentation period and ending with the target-array presentation. This secondary analysis constituted a single 2 (likely vs. unlikely cue) x 2 (clean vs. noisy target) design and is based on the assumption of signals from the cue and target related periods being confounded in connectivity analyses. All PPI analyses were conducted using BrainVoyager QX 2.6 and MATLAB custom scripts.

Graph Theoretical Analysis

Per-subject graphs were constructed for each of the conditions of interest based on parameter estimates of the PPI regressors. All analyses described below were performed on graphs proportionally thresholded at the top 10% of connections. All weights, as in analyses described in the previous chapters, were rescaled to the range [0,1] by dividing each connection by the maximum per-graph connectivity value. Each metric is normalized by calculating the ratio of the observed metric per-subject, per-condition to the average of 100 randomized null-graphs with equivalent degree distributions. Calculation of each metric as well as construction of reference null graphs was performed using the Brain Connectivity Toolbox for MATLAB (Rubinov & Sporns, 2010).

In the first-pass analysis, individual graphs were constructed for each cell of the initial 2 x 2 x 2 design. In the secondary analysis, a single cue-target period was defined and consequently no separation of cue and target related graphs were possible. Four graphs were composed for this secondary analysis, for each cell of the 2 x 2 design.

Several factors could determine the expected pattern of graph results for these conditions. Given that the results from Chapters II and III were measured concurrently with the presentation of task-relevant and irrelevant information, we might expect global patterns of connectivity to be dependent only on the current task relevance of the current visual display. This leads to the first set of predictions based on what I will refer to as my “task-relevance” hypothesis. Using the distinct patterns of the results observed for task-relevant and -irrelevant information processing in Chapters II and II as markers (Fig. 14), the task-relevance hypothesis would predict that currently relevant information produces a pattern of global functional integration. From a global workspace perspective, this relevant information would inhabit a centralized shared workspace, with efficient access to specialized processors in the brain such that relevant information could be easily processed, updated, and reprocessed based on going task demands (Baars, 2005; Dehaene et al., 1998). Without distraction, task-relevant information would easily win competition for central representation and therefore global functional integration would be an efficient means of communicating that information. Alternatively, task-irrelevant information, such as the presence of distracting information, would produce global functional segregation as a potential means of protecting ongoing goal-related information. Without the same efficiency of access to central representation and specialized cognitive processors that might be seen without distraction, task-irrelevant information may be more easily suppressed at the expense of ongoing task performance. This leads to the task-relevance prediction that cue-related information, which will always be task-relevant in this paradigm, to display characteristics of global functional integration. Noisy displays will then necessarily shift towards a more segregated pattern of topology, while clean displays will maintain a pattern of functional integration.

As an alternative to the task-relevance hypothesis, the dual-mechanisms of control model predicts a proactive form of control to be observed during cue-periods when distractor likelihood is high (Braver et al., 2009). If the patterns of connectivity observed in Chapter III, increased functional segregation, are in fact a means of protecting task-relevant information and suppressing distracting, task-irrelevant information, we would expect to see increased global segregation during the cue period on distractor-likely trials. Previous work has shown that in trials with greater expectation of distracting information, regions of the lateral prefrontal cortex show more sustained activity during cue-periods whereas decreased activity, time-locked to target and distractor presentation, occurs on trials when the distracting information is cued to be less likely (Burgess & Braver, 2010). Therefore, in the case of proactive control, no change in connectivity would be expected at the onset of the target/distractor array. I would then expect that on distractor-unlikely trials, according to this model, global connectivity patterns associated with task-irrelevant processing to occur coincident with the (invalid) presentation of a noisy display. This effect occurs due to the utilization of a reactive form of control, implemented in the absence of distractor expectation, detecting interference and providing an online shift in processing topology (Braver, 2012). The predictions of the task-relevance and dual-mechanisms of control hypotheses are summarized in Fig. 16.

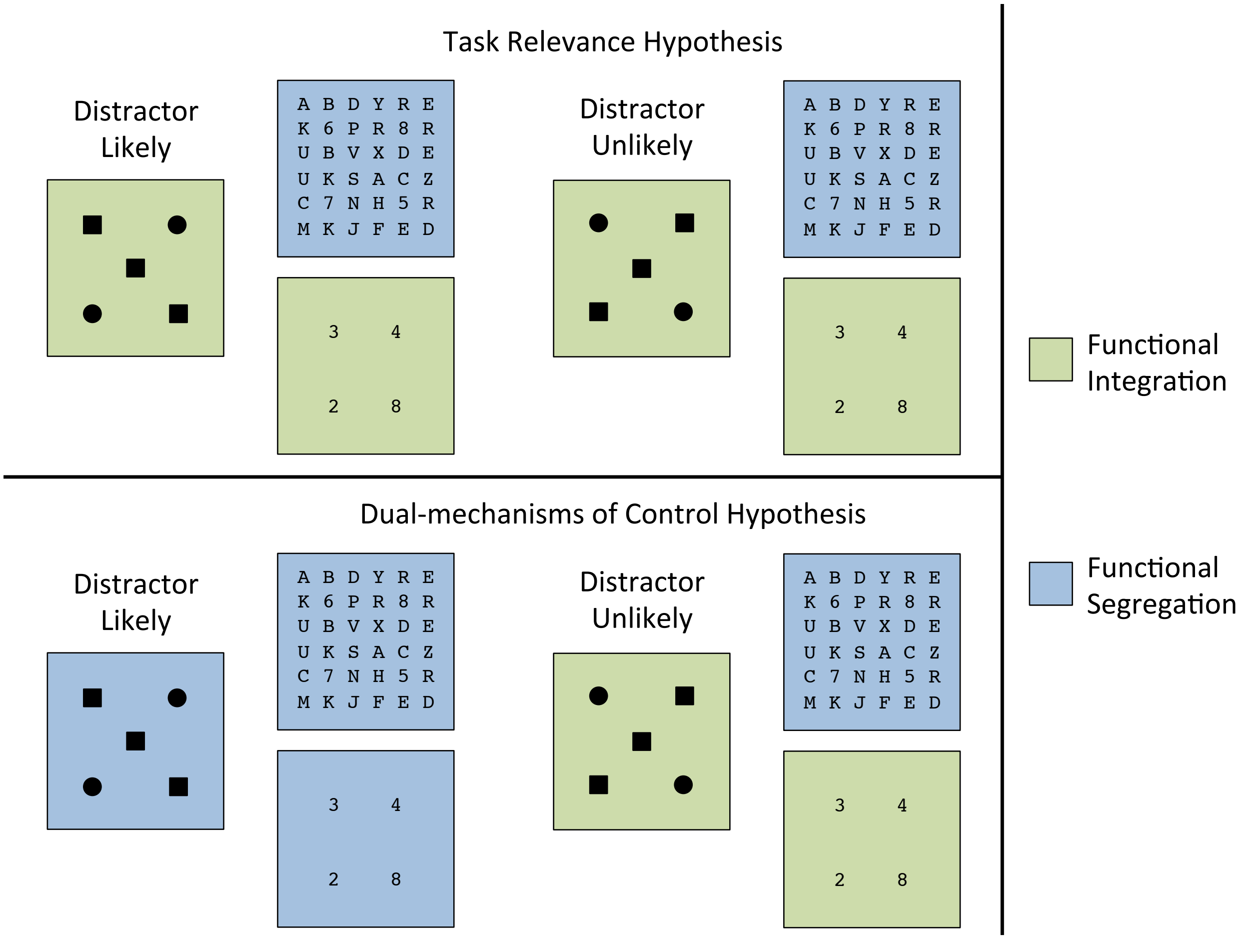

Figure 16. Patterns of topology predicted during cue and target display periods.

Results

Behavioral results are described in detail in Serences et al. (2004). Briefly, participants showed decreased accuracy on the parity judgment on distractor-present (noisy) trials when distractors were cued to be unlikely compared to all other trial types. No difference in performance was observed for distractor absent (clean) between distractor likely and unlikely cues. A significant interaction was observed between cue and target display type, with poor performance identified primarily in distractor-unlikely, noisy displays. Performance in validly and invalidly cued clean displays as well as validly cued noisy displays were roughly equivalent, all around performance determined by the staircasing procedure (~75% correct).

In the first-pass, 2 x 2 x 2 graph theoretical analysis, modularity for the cue-period showed no main effect of cue type (distractor likely vs. unlikely; $F_{(1,11)} = 1.322$, $p = .275$), nor of target display ($F_{(1,11)} = .297$; $p = .597$). However, modularity during the cue-period did show a marginal interaction between cue type and the target display type ($F_{(1,11)} = 4.635$; $p = .054$). This effect can be interpreted as a validity effect, with greater modularity seen for validly cued trials compared to invalidly cued trials. This marginal interaction effect is concerning for interpreting cue-period activity as it suggests that target related activity may be contaminating the cue epoch. No validity effect or cue by target interaction should be observed given that at the time of the cue the participant has seen no information about the target. The participation coefficient too showed no main effect of cue type ($F_{(1,11)} = .318$; $p = .584$) nor target type ($F_{(1,11)} = .054$; $p = .821$) while additionally showing a marginal interaction between cue and target types ($F_{(1,11)} = 4.151$; $p = .066$). Ultimately, only interactions between cue and target type were observed for metrics previously associated with a task-relevant pattern of global functional connectivity.

During the cue period, the clustering coefficient showed no main effect of cue type ($F_{(1,11)} = 1.153$; $p = .306$) and no main effect of target type ($F_{(1,11)} = 1.425$; $p = .258$). Again, however, the clustering coefficient displayed a marginal interaction between cue and target type ($F_{(1,11)} = 3.492$; $p = .089$). The characteristic path length showed a similar pattern of effects. No main effect of characteristic path length was observed during the cue-period ($F_{(1,11)} = .225$; $p = .645$) nor was a main effect of target observed ($F_{(1,11)} = .021$; $p = .888$). A significant interaction between cue and target type was observed during the cue-period for the characteristic path length ($F_{(1,11)} = 7.316$; $p = .02$). These interactions appear to be driven by a validity effect, with valid trials showing greater clustering as well as a larger characteristic path length. The small-worldness metric showed no main effects of cue type ($F_{(1,11)} = 1.165$; $p = .304$), target type ($F_{(1,11)} = 1.452$; $p = .254$), nor an interaction between cue and target ($F_{(1,11)} = 2.933$; $p = .115$). Again, for metrics previously associated with a task-irrelevant pattern of functional connectivity, we observed only interactions between cue and target types.

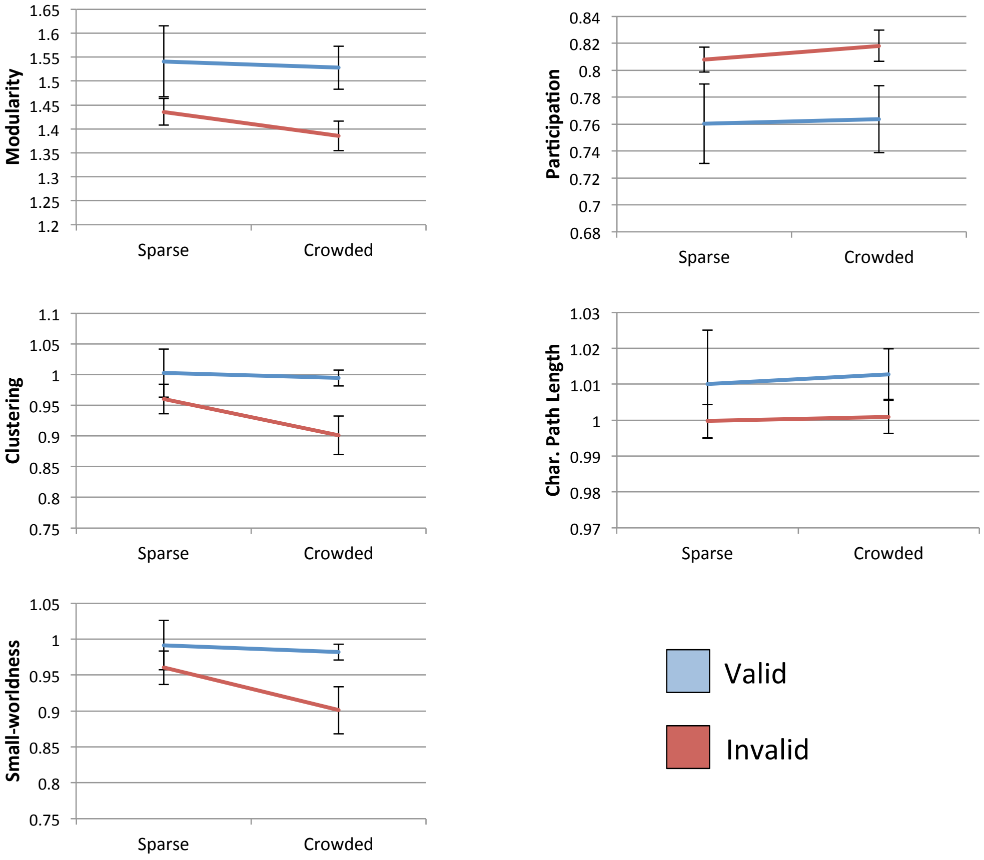

Figure 17. Cue period results for initial analysis for the top 10% of connections. Interactions between cue type and validity were seen for several metrics indicating contamination from target activity.

During the period following target-array presentation, no main effect of cue type ($F_{(1,11)} < .001$; $p = .996$) or target type ($F_{(1,11)} = 2.159$; $p = .17$) was observed in modularity, a metric predicted to be an indicator of goal-driven changes in functional connectivity. A significant interaction was observed between cue and target types ($F_{(1,11)} = 7.385$; $p = .02$), with the greatest modularity occurring for target-unlikely cues followed by noisy displays. Participation showed no main effect of cue type ($F_{(1,11)} = .247$; $p = .629$), but a marginal effect of target display type ($F_{(1,11)} = 3.832$; $p = .076$). Additionally, the participation coefficient showed a significant interaction between cue and target types ($F_{(1,11)} = 6.569$; $p = .026$), with the greatest functional participation occurring on trials with target-likely cues but clean displays. This trial type additionally showed the lowest values of modularity, reflecting the close relationship between an increase in cross modular communication and a decrease in measures of how separable those modules are.

The clustering coefficient showed no main effects of cue ($F_{(1,11)} = .826$; $p = .383$) or target type ($F_{(1,11)} = .053$; $p = .823$). However, similar to the goal-driven measures of functional connectivity, a significant interaction between cue and target type was observed ($F_{(1,11)} = 28.142$; $p < .001$). The characteristic path length showed no main effect of cue type ($F_{(1,11)} = .344$; $p = .57$), but a marginal effect of target type ($F_{(1,11)} = 3.989$; $p = .071$). No significant interaction effect was observed for the characteristic path length. The small-worldness metric mirrored the clustering coefficient, as expected given the lack of effects for the characteristic path length. No main effect of cue ($F_{(1,11)} = 1.27$; $p = .284$) or target ($F_{(1,11)} = .038$; $p = .849$) type was observed for small-worldness but a significant interaction ($F_{(1,11)} = 30.497$; $p < .001$) was observed. This interaction appears to underscore a crossover effect, with the greatest clustering values occurring for distractor-unlikely cues followed by noisy displays (an invalidly cued trial). The lowest clustering values occurred for distractor-likely cued trials and especially in the valid trials with noisy target displays.

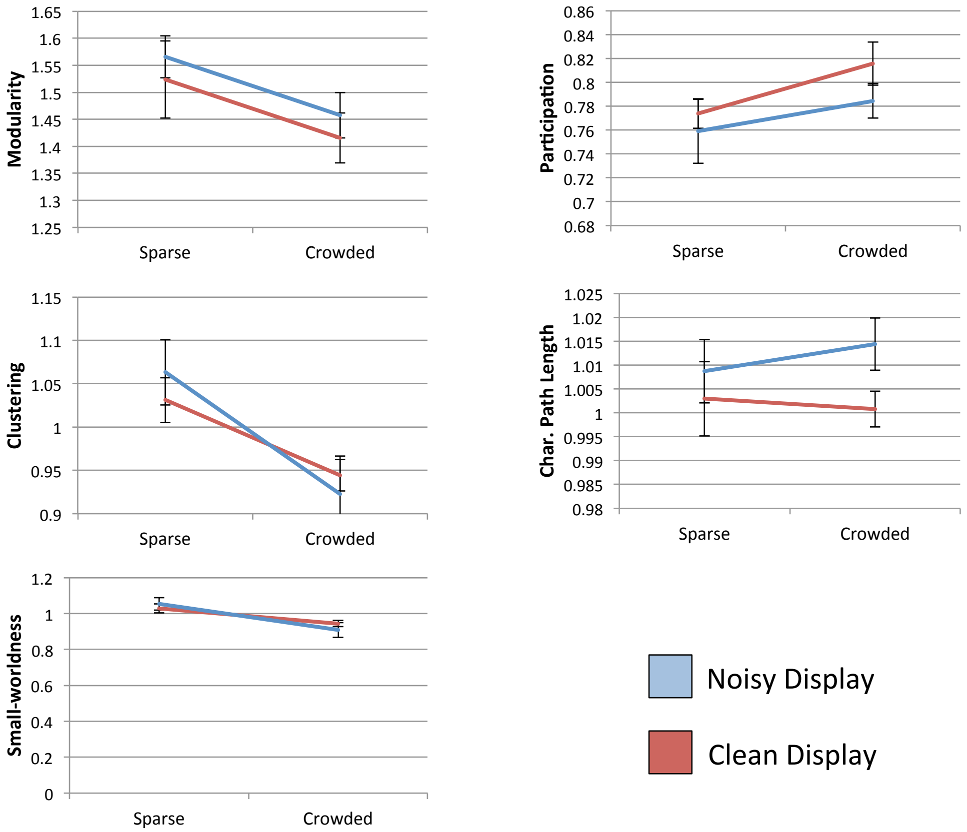

Figure 18. Target period results for first analysis at top 10% of connections.

The lack of interpretable cue effects during the cue period precludes any assessment of predictions made by the dual-mechanisms of control model or an task-relevance hypothesis. Additionally, evidence of a cue/target interaction in nearly all of the observed metrics indicates a high likelihood that the cue period was contaminated by target-related information in the current analysis. Although such an effect was not found in the original analysis of this manuscript (Serences et al., 2004b), those analyses tested peak differences in event related averages of BOLD-activation timecourses. These analyses somewhat bypass the influence of the target related activity apparent during the latter timepoints in the event related averages. Despite the lack of observed interaction, Serences et al. nevertheless retested their prediction in a slow-event related design specifically formulated to disentangle cue and target related activity, expressing concern over contamination of these measures (unfortunately, for technical reasons this data was not available for graph theory analysis). Given the current study's use of PPI methodology to estimate functional connectivity, a temporal dissociation between cue and target related activity may not be possible. Consequently, we performed a second analysis that combined cue and target periods into a single condition predictor, beginning with the cue presentation and ending with the target presentation timepoint.

This analysis (Fig. 19) was evaluated similarly to the above method, assessing graph metrics for main effects of target, cue, and interactions between these factors. Modularity showed no main effects of cue ($F_{(1,11)} = .92$; $p = .358$) or target ($F_{(1,11)} = 1.124$; $p = .312$). However, the interaction showed a significant effect ($F_{(1,11)} = 7.011$; $p = .023$), which appears to be driven by a decrease in modularity in the distractor likely but clean display trial type. The participation coefficient showed no main effects of cue ($F_{(1,11)} = .344$; $p = .569$) or target ($F_{(1,11)} = 1.252$; $p = .287$) type but also showed a significant interaction ($F_{(1,11)} = 8.969$; $p = .012$) between these factors.

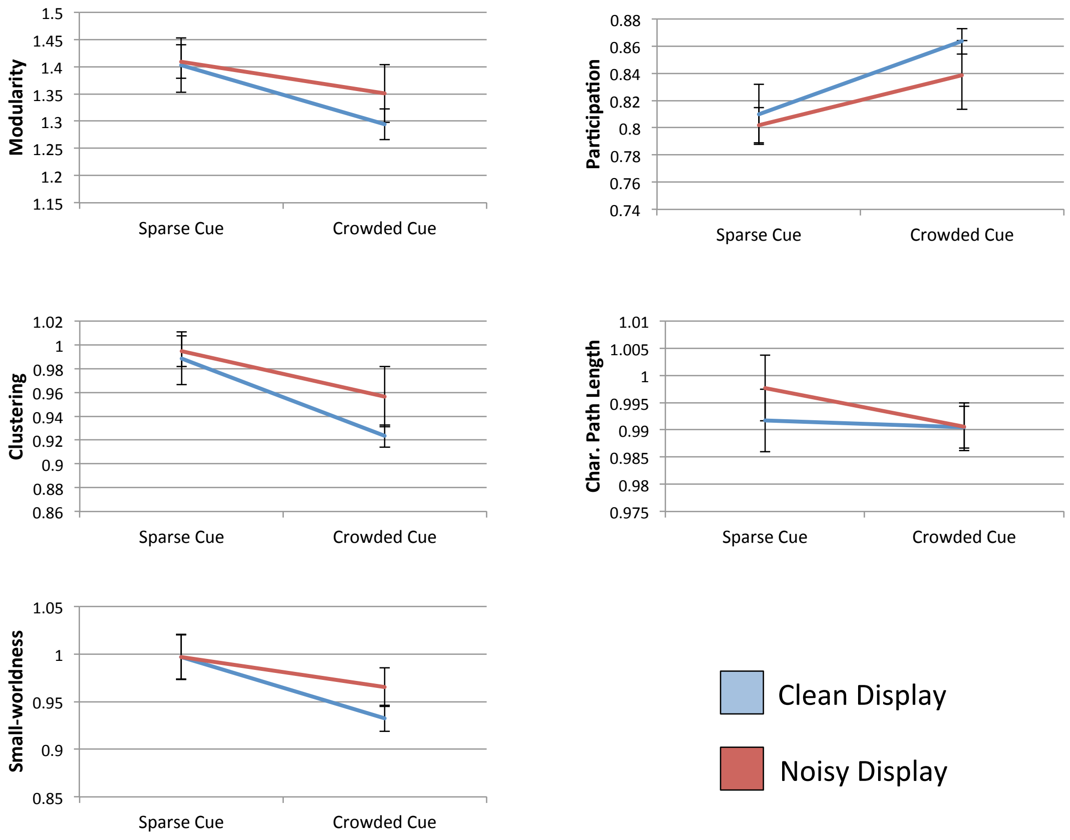

Figure 19. Results from secondary analysis combining cue and target presentations for the top 10% of connections.

The clustering coefficient showed no main effect of cue ($F_{(1,11)} = .879$; $p = .369$) or target ($F_{(1,11)} = .756$; $p = .403$) type. A marginal interaction effect was observed for the clustering coefficient ($F_{(1,11)} = 4.513$; $p = .057$). The characteristic path length showed no effects of cue ($F_{(1,11)} = .387$; $p = .547$), target ($F_{(1,11)} = .366$; $p = .558$), or an interaction ($F_{(1,11)} = 2.062$; $p = .179$) between cue and target. Small-worldness, tracking the clustering coefficient, showed no main effect of target ($F_{(1,11)} = 1.275$; $p = .283$) or cue ($F_{(1,11)} = .563$; $p = .469$) type and only a marginal interaction ($F_{(1,11)} = 3.483$; $p = .089$) between target and cue.

The primary finding of this second analysis appears to be a pattern of interaction between cue and target, with significant interactions found for functional modularity and participation coefficient. Marginal interactions were apparent in the clustering coefficient and small-worldness of the graph. In each of these interactions, the effect appears to be largely driven by one condition in particular: likely-distractor/clean target.

As with the previous chapters, we aimed to evaluate whether the global pattern of results was primarily driven by a small number of networks we might a priori have believed to be drivers of connectivity change from a network-level perspective. The original Serences et al. (2004) manuscript found activation differences in visual cortex as a result of cue differences (distractor likely vs. unlikely). We thus explored the specific changes in functional connectivity that might be found in visual cortex. Examining only the nodes in a visual network defined by Power et al., (2011), we probed the participation and clustering coefficients, metrics designed to measure inter-modular and intra-modular connectivity changes respectively. The participation coefficient of the visual network showed no main effect of cue ($F_{(1,11)} = 2.625$; $p = .133$) nor a main effect of target type ($F_{(1,11)} = .646$; $p = .439$), but did display an interaction between cue and target type ($F_{(1,11)} = 11.524$; $p = .006$) that is driven by a decrease in participation in the distractor-likely/noisy display condition. This pattern of results in the functional participation of the visual network mirrors the pattern of significance seen in the global statistics, suggesting the visual network connectivity likely tracks the global pattern of results. The clustering coefficient, again similar to the global pattern of changes, showed no main effect of cue type ($F_{(1,11)} = .501$; $p = .494$), no main effect of target type ($F_{(1,11)} = 2.211$; $p = .165$), and no interaction between cue and target ($F_{(1,11)} = 1.535$; $p = .241$). The lack of target display effects raises concerns about the sensitivity of our analysis with the current dataset as we would generally expect noisy and clean displays to elicit strong differences in visual cortex.

Discussion

Several findings complicate the overall interpretation of the current results. First and foremost, in our initial analysis many graph theoretical metrics displayed an interaction between cue and target type during the cue-period. This interaction suggests that activity due to target and distractor presentation contaminated cue-period measures of functional connectivity. This contamination likely stems from a fundamental step of the PPI analysis, construction of the interaction regressor. Correlations in activity are measured over the time period during which a typical hemodynamic response might occur (Gitelman et al., 2003; McLaren et al., 2012). Given the relatively short period between cue and target in the current dataset, we might not expect to be able to differentiate these signals, as the cue interaction period will overlap with the subsequent target. A typical BOLD activation analysis would typically deal with this confound in signal estimation by two means: a variable delay between cue and target and/or a deconvolution analysis (Burock, Buckner, Woldorff, Rosen, & Dale, 1998; Friston, Zarahn, Josephs, Henson, & Dale, 1999; Serences, 2004). Although the current dataset utilized a variable delay, the signal of interest in a PPI study does not conform to the typical hemodynamic response function (HRF; and this signal is specifically disregarded) and may not directly benefit from the increased ability to disentangle the HRF due to cue and target. Given this limitation, examination of the primary analysis in the current study is likely un-interpretable. Our secondary analysis aimed to combine signals from cue and target into a single experimental predictor, as this analysis should allow us to interpret results by including the conflation of cue and target presentation signals in our model. The findings from this analysis were solely an interaction between cue and target types for the modularity and participation coefficient, with a marginal effect in clustering and small-worldness. These effects appear to be driven by the interaction between distractor-likely cues paired with a clean display. However, because our ultimate findings were primarily driven by the results in a single condition, I am ultimately unable to interpret the full range of predictions made by the dual-mechanisms of control and task-relevance models with these findings.

Finally, the lack of target display effects in the visual network promotes additional concern that the current analysis has a limited ability to discern potential changes in connectivity effects in the context of the current analysis. If any stimulus difference were to produce effects in visual areas, it would be the comparison of a crowded visual display vs. a relatively clean display. The lack of main effect due to target could be due to the confluence of top-down effects on these visual areas due to cue combined with any feed-forward effects, complicating expectations of activity in these areas.

Given the limitations of interpretation in the current study, an ideally designed follow-up would include additional time between cue and target to allow for an adequate estimation of connectivity for cues and targets. Experiment 2 in Serences et al. (2004) provides additional time between cue and target to better estimate activity of these two task periods independently. This data was, unfortunately, unavailable for analysis. A second alternative to the difficulties faced by the current study would be adopt a class of task-concurrent connectivity analyses based on correlating concatenated time-series after regressing out task-signal (Cole et al., 2014; Friston, 2011a; Norman-Haignere, McCarthy, Chun, & Turk-Browne, 2011). These analysis methods do not typically separate the intrinsic connectivity structure from task-evoked perturbations, and consequently display modularity highly correlated with the resting state (Cole et al., 2014). These analyses were not pursued for two reasons. First, the variance associated with intrinsic connectivity has been shown to be much larger than the variance associated with task-induced changes in connectivity (Cole et al., 2014). Consequently, examination of graph metrics would be largely driven by the organization of the intrinsic functional connectivity of the brain, rather than by the task evoked signal. Next, in order to compare between conditions such that the intrinsic connectivity does not dominate the graph organization, we would take the difference between conditions to “cancel out” the intrinsic signal. This step then requires interpretation of graph theoretical results of difference scores between conditions rather than direct estimations of the task-induced connectivity. These methods do, however, have the advantage of being more flexible in terms of the time-periods during which connectivity can be correlated compared to the PPI style analysis utilized in the current study. Further study with well-controlled design and analysis will be necessary to best answer the questions posed in this chapter.

CHAPTER V

CONCLUSIONS

Modern theories of brain function contend that specific cognitive operations can be localized to functionally and anatomically specialized sites in the brain. However, coordination across multiple regions is necessary for the production of coherent behavior (Bullmore & Sporns, 2009; Friston, 2005; John, Easton, & Isenhart, 1997; McIntosh, 2000; Mesulam, 1990; Posner et al., 1988; Varela, Lachaux, Rodriguez, & Martinerie, 2001). Network science and graph theory provide a flexible and generalizable vantage from which to address questions about the coordination of activity across neural systems. A primary goal of the studies presented in this dissertation was to find evidence for and characterize whole brain changes in connectivity patterns associated with attention and awareness.

Chapter II addressed the theoretical debate between neurobiological theories of conscious awareness. Focal, network-level, and global theories provided distinct and testable predictions as to the expected patterns of functional connectivity. Focal and network-level theories would predict that the changes associated with conscious awareness of a target would be restricted to individual nodes of the global network or to a circumscribed set of fronto-parietal attention networks, respectively. However, detection of a simple target stimulus produced a much more diffuse pattern of effects than would have been expected from the standpoint of these theories. Functional connections spanning the entire brain appeared to participate in a global decrease in the functional modularity of the brain, measuring of how well one can disentangle the various specialized networks in the brain. Concurrent with, and likely contributing to, these changes in modularity were increases in global measures of participation, or cross-modular connections. Consciousness has been previously correlated with increases in long-range oscillatory synchrony (Gross et al., 2004; King et al., 2013; Melloni et al., 2007; Schröter et al., 2012), a mechanism that appears to fit well as a potential source of the signal indexed by changes in participation.

The observed changes associated with conscious detection of a target (or the percept of detection of a target, in the case of false alarms) tightly align with the predictions of global models of conscious awareness. Rather than requiring specialized processing in a restricted set of regions, these theories argue that large-scale coordination across multiple modalities converge on a unified representation of currently accessible information. Formalized, the Global Workspace hypothesis (Baars, 2005; Baars, 1983; 2002; Dehaene et al., 1998; Dehaene & Changeux, 2011), Dynamic Core hypothesis (Edelman, 2003; Tononi & Edelman, 1998) and the Information Integration Theory (Tononi, 2008; Tononi & Sporns, 2003) each require information to be integrated via whole brain, recurrent connections as a prerequisite for awareness of that information. Of course, these results do not discount the findings of numerous task activation studies that found fronto-parietal activations in contrasts of seen versus unseen stimuli. Rather, the current results compliment these findings and highlight the need to examine large-scale functional dynamics as coordination across the entire brain is necessary for the production of a coherent cognitive experience (Bullmore & Sporns, 2009; Friston, 2005; John et al., 1997; McIntosh, 2000; Mesulam, 1990; Posner et al., 1988; Varela et al., 2001). These global changes in connectivity may work in tandem with the specific recruitment of fronto-parietal networks, facilitating efficient connectivity between these regions, sensory areas, and control regions necessary for flexibly responding behaviorally to incoming relevant stimuli.

Although the results in this study were global, no differences were found in the characteristic path length or functional clustering of whole-brain graphs in aware compared to unaware conditions. These two measures are component characteristics of the brain's small-worldness, a holistic property of the brain that it shares with many other organized systems (Bullmore & Sporns, 2012; Humphries & Gurney, 2008; Watts & Strogatz, 1998). This functional balance between computational specialization and integration of information across domains can be found at numerous levels of analysis including structural (He et al., 2007; Sporns & Zwi, 2004; Supekar, Musen, & Menon, 2009; Watts & Strogatz, 1998), resting state (Achard, Salvador, Whitcher, Suckling, & Bullmore, 2006; Liao et al., 2011; Spoormaker et al., 2010; Uehara et al., 2014; van den Heuvel et al., 2008), and task-concurrent studies (Bassett et al., 2009; Cole et al., 2014; Yu et al., 2011). Several factors may have contributed to the modularity and participation effects in the absence of clustering and characteristic path lengths effects, including the high levels of functional efficiency already inherent in the network (Estrada & Hatano, 2008), insensitivity of clustering and characteristic path length measures to specifically cross-modular connections (Humphries & Gurney, 2008), and the lack of distracting information, allowing an unimpeded shift benefitting functional integration.

Chapter III describes the connectivity changes associated with detection of a task-irrelevant oddball that interrupts ongoing goal-directed behavior. Novel, salient stimuli bring about a host of both physiological (Beatty, 1982; Kahneman & Beatty, 1966; Tracy et al., 2000) and cognitive changes (Han & Marois, 2014; Kahneman, 1973; Sara & Bouret, 2012) required to produce an effective response to unexpected information. By presenting this information during the engagement of an otherwise (goal-driven) attentionally demanding task, this study assessed how the stimulus-driven recruitment of attention altered the functional human connectome.

Salient information may (or may not) be important for both immediate goals (like completing an experimental task) and higher-order goals (like survival) and consequently must be evaluated before such information is allowed to influence future behavior. If task-relevant information is globally communicated or represented, as shown in Chapter II, detection of a potentially relevant (but ultimately irrelevant) stimulus could produce complex changes to this pattern of connectivity. Detecting a novel stimulus requires evaluation of that stimulus while at the same time maintaining your current goals.

Largely mirroring the patterns of behavioral effects with repeated oddball presentations (Asplund, Todd, Snyder, & Marois, 2010a; Asplund, Todd, Snyder, Gilbert, & Marois, 2010b), initial presentations of a salient oddball caused significant changes in global functional connections. These initial oddballs significantly altered functional modularity, this time via an increase in the local clustering causing networks to ultimately be more functionally segregated from each other. Subsequent presentations of oddballs showed decreased measures of functional segregation (decreased modularity, clustering, small-worldness) relative to initial oddballs, but greater than those observed during no-target trials. Unlike the changes with target awareness described in Chapter II, detection of an oddball stimulus (or at least the initial presentations of an oddball) did produce changes in the functional small-world structure of brain. An increase in the small-worldness suggests that detection of this oddball produces a more functionally segregated connectome, perhaps reflecting recruitment of many specialized processes necessary to evaluate novel information that might necessitate behavioral change (Han & Marois, 2014; Pavlov, 1927; Sara & Bouret, 2012; Sokolov, 1963).

These results stand in contrast to a hypothetical “integration” hypothesis which, based on the results from Chapter II, hypothesized that information – regardless of its content – would necessitate functional integration for global broadcasting (Baars, 2005; Dehaene & Changeux, 2011; Edelman, 2003; Tononi & Edelman, 1998). Instead, we found a pattern of results more in line with a "segregation" hypothesis that predicted increases in measures of small-worldness perhaps because of the multitude of specialized processes brought online in order to evaluate novel stimuli (Pavlov, 1927; Sokolov, 1963). Of course, participants do consciously perceive oddball stimuli, and thus one could have expected that there would be high degree of integration in the global network as well. The current results represent a potential shift in the prioritization of the two driving influences on small-world organization: integration (for information sharing between networks) and segregation (for specialized processing).

The discrepancies in results found in Chapters II and III presented an interesting question: does task-relevant and task-irrelevant information processing produce differential patterns of global connectivity? Chapter II found evidence supporting global theories of awareness, which argue that in order for information to be consciously perceived, it needs be globally broadcast throughout the brain favoring the demands of functional integration. Similar patterns of global integration have been found for other goal-driven behaviors (Ekman et al., 2012; Kitzbichler et al., 2011). Chapter III found a contrasting pattern of results, suggesting that when task-irrelevant, novel information (e.g. an oddball) interrupts goal-oriented behavior, it produces a more functionally segregated brain (also see Nicol et al., 2011). Perhaps the behavioral relevance of a stimulus has a direct influence on the global functional connectome in order to effectively facilitate complex demands of the environment. As noted above, distracting or irrelevant stimuli must be specifically suppressed while attempting to maintain performance levels dependent on the processing of goal-relevant information.

Analyses presented in Chapter IV were intended to disentangle the potential effects of task-relevant changes in the functional connectome from task-irrelevant effects. In particular, either task-relevant information (targets) was expected to produce a pattern of goal-driven global functional integration, to be identified by changes in functional modularity and participation. Complementary, we hypothesized an increase in small-worldness and measures of functional segregation with presentations of task-irrelevant distractors, particularly when these stimuli were unexpected. Participants were cued to the locations of target digits along with information as to the likelihood that distracting, task-irrelevant letters would simultaneously be displayed (Awh, Matsukura, & Serences, 2003; Serences et al., 2004b). On a subset of trials, cues invalidly informed participants as to the upcoming presence of distractors, providing a mismatch between the expectations and control exerted during the cue period versus the need to reactively adapt behavior to the presence of unforeseen task-irrelevant information (Braver, 2012; Braver et al., 2009). Ideally, the design would have allowed comparison of distractor-likely and unlikely cues separately from distractor-present and absent trials.

Results of these analyses suggested that the task design was insufficient to accurately disentangle connectivity differences in cue and target-display related activity. Specifically, cue-period graph theoretical measures showed an interaction between cued likelihood and distractor presence in the subsequent target epoch, suggesting that activity during the target presentation period contaminated connectivity measurements during the cue period. Consequently, further analyses were unable to disentangle whether observed differences in graph theoretical measures were due to cue or target related activity, interpretations with divergent expectations given our a priori task-relevance and dual-mechanisms of control hypotheses. Several possible alternate routes for testing these hypotheses in the context of a similar attentional task were proposed in Chapter IV.

Ultimately, despite a complicated pattern of results in Chapter IV, Chapters II and III have highlighted novel neurobiological characteristics of well-studied psychological phenomena. Despite well-characterized network-level activations in awareness and attention research, the results presented here suggest that there are other dimensions of these processes still to study. Although the current results have focused on one kind of graph theoretical result (global changes in connectivity), there are several other avenues of study available to expand upon typical activation based research. For example, the study of hub regions, or nodes vitally important for the efficient transfer of information, may aid in the identification and explanation of information processing bottlenecks.

Multitasking costs are ubiquitous and profound despite the frequency with which humans engage in multiple concurrent tasks (M. C. Smith, 1967; Welford, 1952). The central bottleneck hypothesis presumes that each task recruits a common, structural mechanism responsible for the response selection phase of task performance. This structure is only available to perform one mapping at a time, resulting in serial performance of each task (Marois & Ivanoff, 2005; Pashler, 1994; M. C. Smith, 1967). Imaging research into the identity of potential bottlenecks has identified a few likely candidates including: lateral prefrontal cortex (lPFC), inferior frontal junction (IFJ), dorsal premotor cortex (FEF), anterior medial prefrontal cortex, anterior insula, and intraparietal sulcus (Dux & Marois, 2009; Dux, Ivanoff, Asplund, & Marois, 2006; Jiang & Kanwisher, 2003; Marois & Ivanoff, 2005; Marois, Larson, Chun, & Shima, 2005; Sigman & Dehaene, 2008; Tombu et al., 2011). These studies typically involve identifying foci revealed by the contrast of dual vs. single task or long vs. short stimulus onset asynchrony (SOA) in a psychological refractory period (PRP) design. These foci make sense as potential central bottlenecks of information processing given their purported involvement in various aspects of cognition including cognitive control (Derrfuss, Brass, & Yves von Cramon, 2004; Dosenbach et al., 2007; 2008), attention (Corbetta et al., 2008; Corbetta & Shulman, 2002), and are generally recruited across a wide range of cognitively demanding tasks (Duncan, 2010; Duncan & Owen, 2000; Duncan, Emslie, Williams, Johnson, & Freer, 1996).

Although performance of dual tasks produces recognizable deficits, performance of specific tasks over an extended period of time shows improvement through practice (Jonides, 2004; Schumacher et al., 2001; Tombu & Jolicoeur, 2004). Dux et al. (2009) find that practice effects appear consistent with an “improved efficiency” hypothesis in which connection strengths between regions involved in task performance are improved and progressively routed away from speed-limiting bottlenecks (Poldrack, 2005). Consequently, through practice the very bottleneck regions identified by multitasking deficits appear to be routed around as a means of speeding processing. The network perspective afforded by graph theory offers an opportunity to examine the degree to which whole-brain, network, and individual hub node efficiency improves with practice of a task compared to early performance and in particular, test theories as to whether processing becomes more efficient in these nodes vs. new routes of information transfer are formed (Dux et al., 2009).

Additionally, examination of hub regions may play a future role in better elucidating the interactions between individual networks, such as the dorsal and ventral attention networks. Corbetta and Shulman (2002; 2008) have hypothesized that the ventral attention network acts as a “circuit breaker,” responsible for interrupting current task goals when salient stimuli occur outside the current goal-directed focus of attention. Diffusion tensor imaging (DTI) has shown several anatomical connections between the dorsal and ventral attention networks (de Schotten et al., 2011) suggesting a plausible role of ventral attention regions responsible for pivoting attention away from the current focus via interactions with the dorsal network. This interaction was hypothesized to occur between the TPJ and nearby intraparietal regions (Astafiev et al., 2006; Corbetta et al., 2008; Corbetta & Shulman, 2002). In fact, posterior parietal damage resulting in neglect symptoms produces breakdown of resting-state functional connectivity in both dorsal and ventral attention regions suggesting an important role of this region for connections between both attention networks (Corbetta & Shulman, 2011; He et al., 2007). However, other work implicates IFJ as a linking node between dorsal and ventral attention networks. Asplund et al. (2010) showed a switch in patterns of IFJ activity during capture of attention by a surprising, irrelevant stimulus and goal-directed search. During goal-directed search, IFJ showed similar activity to dorsal attention regions but coactivates with TPJ during capture by the surprising stimulus (Asplund, Todd, Snyder, & Marois, 2010a). Other work has shown that during covert shifts of spatial attention the ventral frontal cortex showed increased psychophysiological interaction (PPI) connectivity with dorsal attention regions for validly cued trials, again suggesting a coordinating role for IFJ (Weissman & Prado, 2012). If either of these regions (IFJ or TPJ) is a crucial, pivot points for the "circuit-breaking" process, we would expect to see a decrease in the average path length of these regions (suggesting efficient communication relies on these areas) and potentially an increase in the clustering coefficient (implying that these areas become more interconnected overall) in comparisons of oddball vs. target detection.

Additional probes of task-relevant and task-irrelevant information as well as differences between top-down and bottom-up attention will be necessary to determine the accuracy of these hypotheses. However, towards developing these hypotheses, these studies provided a novel characterization of attention and awareness from the perspective of changes in global functional connectivity. These studies have, hopefully, highlighted the utility of graph theory as a means of succinctly describing the complex behaviors of the brain's functional connectivity. Further work will be necessary to map the relationship between complex network structures of the brain and the specialized neural circuitry required to produce multifaceted human behavior. Ideally, future work would incorporate a more holistic, potentially graph theoretical, examination of the whole-brain connectivity data to account for the global changes observed in this dissertation.

REFERENCES

Achard, S., Salvador, R., Whitcher, B., Suckling, J., & Bullmore, E. (2006). A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs. *Journal of Neuroscience*, *26*(1), 63–72. http://doi.org/10.1523/JNEUROSCI.3874-05.2006

Ahmadlou, M., Rostami, R., & Sadeghi, V. (2012). Neuroscience Letters. *Neuroscience Letters*, *516*(1), 156–160. http://doi.org/10.1016/j.neulet.2012.03.087

Arnold, D. H., Clifford, C. W., & Wenderoth, P. (2001). Asynchronous processing in vision: color leads motion. *Current Biology*, *11*(8), 596–600.

Arrington, C. M., Carr, T. H., Mayer, A. R., & Rao, S. M. (2000). Neural mechanisms of visual attention: object-based selection of a region in space. *Journal of Cognitive Neuroscience*, *12 Suppl 2*, 106–117. http://doi.org/10.1162/089892900563975

Asplund, C. L., Todd, J. J., Snyder, A. P., & Marois, R. (2010a). A central role for the lateral prefrontal cortex in goal-directed and stimulus-driven attention. *Nature Neuroscience*, *13*(4), 507–512. http://doi.org/10.1038/nn.2509

Asplund, C. L., Todd, J. J., Snyder, A. P., Gilbert, C. M., & Marois, R. (2010b). Surprise-induced blindness: A stimulus-driven attentional limit to conscious perception. *Journal of Experimental Psychology: Human Perception and Performance*, *36*(6), 1372–1381. http://doi.org/10.1037/a0020551

Astafiev, S. V., Shulman, G. L., & Corbetta, M. (2006). Visuospatial reorienting signals in the human temporo-parietal junction are independent of response selection. *The European Journal of Neuroscience*, *23*(2), 591–596. http://doi.org/10.1111/j.1460-9568.2005.04573.x

Awh, E., Matsukura, M., & Serences, J. T. (2003). Top-down control over biased competition during covert spatial orienting. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(1), 52–63. http://doi.org/10.1037/0096-1523.29.1.52

Baars, B. (2005). Global workspace theory of consciousness: toward a cognitive neuroscience of human experience. *Boundaries of Consciousness: Neurobiology and Neuropathology*, *150*, 45–53. http://doi.org/10.1016/S0079-6123(05)50004-9

Baars, B. J. (1983). Conscious Contents Provide the Nervous System with Coherent, Global Information. *Consciousness and Self-Regulation*, 1–21.

Baars, B. J. (2002). The conscious access hypothesis: origins and recent evidence. *Trends in Cognitive Sciences*, *6*(1), 47–52.

Bachiller, A., Poza, J., Gómez, C., Molina, V., Suazo, V., Díez, A., & Hornero, R. (2014). Graph-Theoretical Analysis in Schizophrenia Performing an Auditory Oddball Task. In *IFMBE Proceedings* (Vol. 41, pp. 799–802). Cham: Springer International Publishing. http://doi.org/10.1007/978-3-319-00846-2_198

Bassett, D. S., Bullmore, E. T., Meyer-Lindenberg, A., Apud, J. A., Weinberger, D. R., Coppola, R., & Kopell, N. J. (2009). Cognitive Fitness of Cost-Efficient Brain Functional Networks. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(28), 11747–11752.

Bassett, D. S., Nelson, B. G., Mueller, B. A., Camchong, J., & Lim, K. O. (2012). Altered resting state complexity in schizophrenia. *NeuroImage*, *59*(3), 2196–2207. http://doi.org/10.1016/j.neuroimage.2011.10.002

Bassett, D. S., Wymbs, N. F., Porter, M. A., Mucha, P. J., Carlson, J. M., & Grafton, S. T. (2011). Dynamic reconfiguration of human brain networks during learning. *Proceedings of the National Academy of Sciences*, *108*(18), 7641–7646. http://doi.org/10.1073/pnas.1018985108

Beatty, J. (1982). Task-evoked pupillary responses, processing load, and the structure of processing resources. *Psychological Bulletin*, *91*(2), 276–292.

Beck, D. M., & Kastner, S. (2009). Vision Research. *Vision Research*, *49*(10), 1154–1165. http://doi.org/10.1016/j.visres.2008.07.012

Beck, D. M., Rees, G., Frith, C. D., & Lavie, N. (2001). Neural correlates of change detection and change blindness. *Nature Neuroscience*, *4*(6), 645–650. http://doi.org/10.1038/88477

Bledowski, C., Prvulovic, D., Goebel, R., Zanella, F. E., & Linden, D. E. J. (2004). Attentional systems in target and distractor processing: a combined ERP and fMRI study. *NeuroImage*, *22*(2), 530–540. http://doi.org/10.1016/j.neuroimage.2003.12.034

Boly, M., Balteau, E., Schnakers, C., Degueldre, C., Moonen, G., Luxen, A., et al. (2007). Baseline brain activity fluctuations predict somatosensory perception in humans. *Proceedings of the National Academy of Sciences of the United States of America*, *104*(29), 12187–12192. http://doi.org/10.1073/pnas.0611404104

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*(4), 433–436.

Braver, T. S. (2012). The variable nature of cognitive control: a dual mechanisms framework. *Trends in Cognitive Sciences*, *16*(2), 105–112. http://doi.org/10.1016/j.tics.2011.12.010

Braver, T. S., Barch, D. M., Gray, J. R., Molfese, D. L., & Snyder, A. (2001). Anterior cingulate cortex and response conflict: effects of frequency, inhibition and errors. *Cerebral Cortex (New York, NY : 1991)*, *11*(9), 825–836.

Braver, T. S., Paxton, J. L., Locke, H. S., & Barch, D. M. (2009). Flexible neural mechanisms of cognitive control within human prefrontal cortex. *Proceedings of the National Academy of Sciences*, *106*(18), 7351–7356. http://doi.org/10.1073/pnas.0808187106

Broadbent, D. E. (1981). Selective and control processes. *Cognition*, *10*(1-3), 53–58.

Bullmore, E. T., & Bassett, D. S. (2011). Brain graphs: graphical models of the human brain connectome. *Annual Review of Clinical Psychology*, *7*, 113–140. http://doi.org/10.1146/annurev-clinpsy-040510-143934

Bullmore, E., & Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, *10*(3), 186–198. http://doi.org/10.1038/nrn2575

Bullmore, E., & Sporns, O. (2012). The economy of brain network organization. *Nature Reviews Neuroscience*, *13*(5), 336–349. http://doi.org/10.1038/nrn3214

Burgess, G. C., & Braver, T. S. (2010). Neural mechanisms of interference control in working memory: effects of interference expectancy and fluid intelligence. *PLoS One*, *5*(9), e12861. http://doi.org/10.1371/journal.pone.0012861

Burock, M. A., Buckner, R. L., Woldorff, M. G., Rosen, B. R., & Dale, A. M. (1998). Randomized event-related experimental designs allow for extremely rapid presentation rates using functional MRI. *Neuroreport*, *9*(16), 3735–3739.

Caeyenberghs, K., Leemans, A., Heitger, M. H., Leunissen, I., Dhollander, T., Sunaert, S., et al. (2012). Graph analysis of functional brain networks for cognitive control of action in traumatic brain injury. *Brain : a Journal of Neurology*, *135*(4), 1293–1307. http://doi.org/10.1093/brain/aws048

Cao, H., Plichta, M. M., Schäfer, A., Haddad, L., Grimm, O., Schneider, M., et al. (2014). Test–retest reliability of fMRI-based graph theoretical properties during working memory, emotion processing, and resting state. *NeuroImage*, *84*, 888–900. http://doi.org/10.1016/j.neuroimage.2013.09.013

Chawla, D., Rees, G., & Friston, K. J. (1999). The physiological basis of attentional modulation in extrastriate visual areas. *Nature Neuroscience*, *2*(7), 671–676. http://doi.org/10.1038/10230

Chica, A. B., Valero-Cabre, A., Paz-Alonso, P. M., & Bartolomeo, P. (2014). Causal Contributions of the Left Frontal Eye Field to Conscious Perception. *Cerebral Cortex (New York, NY : 1991)*, *24*(3), 745–753. http://doi.org/10.1093/cercor/bhs357

Chiu, Y.-C., & Yantis, S. (2009). A domain-independent source of cognitive control for task sets: shifting spatial attention and switching categorization rules. *Journal of Neuroscience*, *29*(12), 3930–3938. http://doi.org/10.1523/JNEUROSCI.5737-08.2009

Cisler, J. M., Bush, K., & Steele, J. S. (2014). A comparison of statistical methods for detecting context-modulated functional connectivity in fMRI. *NeuroImage*, *84*, 1042–1052. http://doi.org/10.1016/j.neuroimage.2013.09.018

Cole, M. W., & Schneider, W. (2007). The cognitive control network: Integrated cortical regions with dissociable functions. *NeuroImage*, *37*(1), 343–360. http://doi.org/10.1016/j.neuroimage.2007.03.071

Cole, M. W., Bassett, D. S., Power, J. D., Braver, T. S., & Petersen, S. E. (2014). Intrinsic and Task-Evoked Network Architectures of the Human Brain. *Neuron*, *83*(1), 238–251. http://doi.org/10.1016/j.neuron.2014.05.014

Cole, M. W., Pathak, S., & Schneider, W. (2010). Identifying the brain's most globally connected regions. *NeuroImage*, *49*(4), 3132–3148. http://doi.org/10.1016/j.neuroimage.2009.11.001

Cole, M. W., Reynolds, J. R., Power, J. D., Repovs, G., Anticevic, A., & Braver, T. S. (2013). Multi-task connectivity reveals flexible hubs for adaptive task control. *Nature Neuroscience*, *16*(9), 1348–1355. http://doi.org/10.1038/nn.3470

Corbetta, M. (1998). Frontoparietal cortical networks for directing attention and the eye to visual locations: identical, independent, or overlapping neural systems? *Proceedings of the National Academy of Sciences of the United States of America*, *95*(3), 831–838.

Corbetta, M., & Shulman, G. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*(3), 201–215.

Corbetta, M., & Shulman, G. L. (2011). Spatial Neglect and Attention Networks. *Annual Review of Neuroscience*, *34*(1), 569–599. http://doi.org/10.1146/annurev-neuro-061010-113731

Corbetta, M., Kincade, J. M., Ollinger, J. M., McAvoy, M. P., & Shulman, G. L. (2000). Voluntary orienting is dissociated from target detection in human posterior parietal cortex. *Nature Neuroscience*, *3*(3), 292–297. http://doi.org/10.1038/73009

Corbetta, M., Patel, G., & Shulman, G. L. (2008). The reorienting system of the human brain: from environment to theory of mind. *Neuron*, *58*(3), 306–324. http://doi.org/10.1016/j.neuron.2008.04.017

Crossley, N. A., Mechelli, A., Vértes, P. E., Winton-Brown, T. T., Patel, A. X., Ginestet, C. E., et al. (2013). Cognitive relevance of the community structure of the human brain functional coactivation network. *Proceedings of the National Academy of Sciences*, *110*(28), 11583–11588. http://doi.org/10.1073/pnas.1220826110

de Haan, W., van der Flier, W. M., Koene, T., Smits, L. L., Scheltens, P., & Stam, C. J. (2012). Disrupted modular brain dynamics reflect cognitive dysfunction in Alzheimer's disease. *NeuroImage*, *59*(4), 3085–3093. http://doi.org/10.1016/j.neuroimage.2011.11.055

de Schotten, M. T., Dell'Acqua, F., Forkel, S. J., Simmons, A., Vergani, F., Murphy, D. G. M., & Catani, M. (2011). A lateralized brain network for visuospatial attention. *Nature Neuroscience*, *14*(10), 1245–1246. http://doi.org/10.1038/nn.2905

De Vico Fallani, F., Astolfi, L., Cincotti, F., Mattia, D., Marciani, M. G., Gao, S., et al. (2008). Structure of the cortical networks during successful memory encoding in TV commercials. *Clinical Neurophysiology : Official Journal of the International Federation of Clinical Neurophysiology*, *119*(10), 2231–2237. http://doi.org/10.1016/j.clinph.2008.06.018

Dehaene, S., & Changeux, J.-P. (2011). Experimental and Theoretical Approaches to Conscious Processing. *Neuron*, *70*(2), 200–227. http://doi.org/10.1016/j.neuron.2011.03.018

Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences*, *10*(5), 204–211. http://doi.org/10.1016/j.tics.2006.03.007

Dehaene, S., Kerszberg, M., & Changeux, J. P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences of the United States of America*, *95*(24), 14529–14534.

Dehaene, S., Naccache, L., Cohen, L., Le Bihan, D., Mangin, J., Poline, J., & Riviere, D. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nature Neuroscience*, *4*(7), 752–758.

Derrfuss, J., Brass, M., & Yves von Cramon, D. (2004). Cognitive control in the posterior frontolateral cortex: evidence from common activations in task coordination, interference control, and working memory. *NeuroImage*, *23*(2), 604–612. http://doi.org/10.1016/j.neuroimage.2004.06.007

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*(1), 193–222.

Dosenbach, N. U. F., Fair, D. A., Cohen, A. L., Schlaggar, B. L., & Petersen, S. E. (2008). A dual-networks architecture of top-down control. *Trends in Cognitive Sciences*, *12*(3), 99–105. http://doi.org/10.1016/j.tics.2008.01.001

Dosenbach, N. U. F., Fair, D. A., Miezin, F. M., Cohen, A. L., Wenger, K. K., Dosenbach, R. A. T., et al. (2007). Distinct brain networks for adaptive and stable task control in humans. *Proceedings of the National Academy of Sciences of the United States of America*, *104*(26), 11073–11078. http://doi.org/10.1073/pnas.0704320104

Downar, J., Crawley, A. P., Mikulis, D. J., & Davis, K. D. (2000). A multimodal cortical network for the detection of changes in the sensory environment. *Nature Publishing Group*, *3*(3), 277–283. http://doi.org/10.1038/72991

Downar, J., Crawley, A. P., Mikulis, D. J., & Davis, K. D. (2002). A cortical network sensitive to stimulus salience in a neutral behavioral context across multiple sensory modalities. *Journal of Neurophysiology*, *87*(1), 615–620.

Drakesmith, M., Caeyenberghs, K., Dutt, A., Lewis, G., David, A. S., & Jones, D. K. (2015). Overcoming the effects of false positives and threshold bias in graph theoretical analyses of neuroimaging data. *NeuroImage*, *118*(C), 313–333. http://doi.org/10.1016/j.neuroimage.2015.05.011

Duncan, J. (1980). The locus of interference in the perception of simultaneous stimuli. *Psychological Review*, *87*(3), 272–300.

Duncan, J. (2006). EPS Mid-Career Award 2004: brain mechanisms of attention. *Quarterly Journal of Experimental Psychology (2006)*, *59*(1), 2–27. http://doi.org/10.1080/17470210500260674

Duncan, J. (2010). The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. *Trends in Cognitive Sciences*, *14*(4), 172–179. http://doi.org/10.1016/j.tics.2010.01.004

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, *96*(3), 433–458.

Duncan, J., & Owen, A. M. (2000). Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends in Neurosciences*, *23*(10), 475–483.

Duncan, J., Emslie, H., Williams, P., Johnson, R., & Freer, C. (1996). Intelligence and the frontal lobe: the organization of goal-directed behavior. *Cognitive Psychology*, *30*(3), 257–303. http://doi.org/10.1006/cogp.1996.0008

Duncan, J., Humphreys, G., & Ward, R. (1997). Competitive brain activity in visual attention. *Current Opinion in Neurobiology*, *7*(2), 255–261.

Dux, P. E., & Marois, R. (2009). The attentional blink: a review of data and theory. *Attention, Perception & Psychophysics*, *71*(8), 1683–1700. http://doi.org/10.3758/APP.71.8.1683

Dux, P. E., Ivanoff, J., Asplund, C. L., & Marois, R. (2006). Isolation of a Central Bottleneck of Information Processing with Time-Resolved fMRI. *Neuron*, *52*(December21), 1109–1120. http://doi.org/10.1016/j.neuron.2006.11.009

Dux, P. E., Tombu, M. N., Harrison, S., Rogers, B. P., Tong, F., & Marois, R. (2009). Training improves multitasking performance by increasing the speed of information processing in human prefrontal cortex. *Neuron*, *63*(1), 127–138. http://doi.org/10.1016/j.neuron.2009.06.005

Edelman, G. M. (2003). Naturalizing consciousness: a theoretical framework. *Proceedings of the National Academy of Sciences of the United States of America*, *100*(9), 5520–5524. http://doi.org/10.1073/pnas.0931349100

Egeth, H. E., Virzi, R. A., & Garbart, H. (1984). Searching for conjunctively defined targets. *Journal of Experimental Psychology: Human Perception and Performance*, *10*(1), 32–39.

Ekman, M., Derrfuss, J., Tittgemeyer, M., & Fiebach, C. J. (2012). Predicting errors from reconfiguration patterns in human brain networks. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(41), 16714–16719. http://doi.org/10.1073/pnas.1207523109/-/DCSupplemental

Estrada, E., & Hatano, N. (2008). Communicability in complex networks. *Physical Review. E, Statistical, Nonlinear, and Soft Matter Physics*, *77*(3 Pt 2), 036111.

Feinstein, J. S., Stein, M. B., Castillo, G. N., & Paulus, M. P. (2004). From sensory processes to conscious perception. *Consciousness and Cognition*, *13*(2), 323–335. http://doi.org/10.1016/j.concog.2003.10.004

Ferri, R., Rundo, F., Bruni, O., Terzano, M. G., & Stam, C. J. (2008). The functional connectivity of different EEG bands moves towards small-world network organization during sleep. *Clinical Neurophysiology*, *119*(9), 2026–2036. http://doi.org/10.1016/j.clinph.2008.04.294

Folk, C., & Remington, R. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of …*.

Forman, S. D., Cohen, J. D., Fitzgerald, M., Eddy, W. F., Mintun, M. A., & Noll, D. C. (1995). Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): use of a cluster-size threshold. *Magnetic Resonance in Medicine*, *33*(5), 636–647.

Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences of the United States of America*, *102*(27), 9673–9678. http://doi.org/10.1073/pnas.0504136102

Fox, M., Corbetta, M., Snyder, A., Vincent, J., & Raichle, M. (2006). Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. *Proceedings of the National Academy of Sciences*, *103*(26), 10046–10051.

Fries, P., Reynolds, J. H., Rorie, A. E., & Desimone, R. (2001). Modulation of Oscillatory Neuronal Synchronization by Selective Visual Attention. *Science*, *291*(5508), 1560–1563. http://doi.org/10.1126/science.1055465

Friston, K. (2011a). Dynamic causal modeling and Granger causality Comments on: The identification of interacting networks in the brain using fMRI: Model selection, causality and deconvolution. *NeuroImage*, *58*(2), 303–305. http://doi.org/10.1016/j.neuroimage.2009.09.031

Friston, K. J. (2005). Models of Brain Function in Neuroimaging. *Annu. Rev. Psychol.*, *56*(1), 57–87. http://doi.org/10.1146/annurev.psych.56.091103.070311

Friston, K. J. (2011b). Functional and effective connectivity: a review. *Brain Connectivity*, *1*(1), 13–36. http://doi.org/10.1089/brain.2011.0008

Friston, K. J., Frith, C. D., Liddle, P. F., & Frackowiak, R. S. J. (1993). Functional Connectivity: The Principal-Component Analysis of Large (PET) Data Sets. *Journal of Cerebral Blood Flow and Metabolism : Official Journal of the International Society of Cerebral Blood Flow and Metabolism*, *13*(1), 5–14. http://doi.org/10.1038/jcbfm.1993.4

Friston, K. J., Zarahn, E., Josephs, O., Henson, R. N., & Dale, A. M. (1999). Stochastic designs in event-related fMRI. *NeuroImage*, *10*(5), 607–619. http://doi.org/10.1006/nimg.1999.0498

Friston, K., Moran, R., & Seth, A. K. (2013). Analysing connectivity with Granger causality and dynamic causal modelling. *Current Opinion in Neurobiology*, *23*(2), 172–178. http://doi.org/10.1016/j.conb.2012.11.010

Gati, J. S., Menon, R. S., Ugurbil, K., & Rutt, B. K. (1997). Experimental determination of the BOLD field strength dependence in vessels and tissue. *Magnetic Resonance in Medicine : Official Journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, *38*(2), 296–302.

Ginestet, C. E., & Simmons, A. (2011). Statistical parametric network analysis of functional connectivity dynamics during a working memory task. *NeuroImage*, *55*(2), 688–704. http://doi.org/10.1016/j.neuroimage.2010.11.030

Gitelman, D. R., Penny, W. D., Ashburner, J., & Friston, K. J. (2003). Modeling regional and psychophysiologic interactions in fMRI: the importance of hemodynamic deconvolution. *NeuroImage*, *19*(1), 200–207. http://doi.org/10.1016/S1053-8119(03)00058-2

Godwin, D., Barry, R. L., & Marois, R. (2015). Breakdown of the brain's functional network modularity with awareness. *Proceedings of the National Academy of Sciences*, 201414466. http://doi.org/10.1073/pnas.1414466112

Grissom, R. J., & Kim, J. J. (2012). Effect sizes for research: Univariate and multivariate applications.

Gross, J., Schmitz, F., Schnitzler, I., Kessler, K., Shapiro, K., Hommel, B., & Schnitzler, A. (2004). Modulation of long-range neural synchrony reflects temporal limitations of visual attention in humans. *Proceedings of the National Academy of Sciences of the United States of America*, *101*(35), 13050–13055. http://doi.org/10.1073/pnas.0404944101

Hahn, B., Ross, T. J., Yang, Y., Kim, I., Huestis, M. A., & Stein, E. A. (2007). Nicotine enhances visuospatial attention by deactivating areas of the resting brain default network. *Journal of Neuroscience*, *27*(13), 3477–3489. http://doi.org/10.1523/JNEUROSCI.5129-06.2007

Han, S. W., & Marois, R. (2014). Functional Fractionation of the Stimulus-Driven Attention Network. *Journal of Neuroscience*, *34*(20), 6958–6969. http://doi.org/10.1523/JNEUROSCI.4975-13.2014

Haynes, J.-D., Driver, J., & Rees, G. (2005). Visibility Reflects Dynamic Changes of Effective Connectivity between V1 and Fusiform Cortex. *Neuron*, *46*(5), 811–821. http://doi.org/10.1016/j.neuron.2005.05.012

He, Y., Chen, Z. J., & Evans, A. C. (2007). Small-world anatomical networks in the human brain revealed by cortical thickness from MRI. *Cerebral Cortex (New York, NY : 1991)*, *17*(10), 2407–2419. http://doi.org/10.1093/cercor/bhl149

Hermundstad, A. M., Bassett, D. S., Brown, K. S., Aminoff, E. M., Clewett, D., Freeman, S., et al. (2013). Structural foundations of resting-state and task-based functional connectivity in the human brain. *Proceedings of the National Academy of Sciences*, *110*(15), 6169–6174. http://doi.org/10.1073/pnas.1219562110

Hon, N., Epstein, R. A., Owen, A. M., & Duncan, J. (2006). Frontoparietal activity with minimal decision and control. *Journal of Neuroscience*, *26*(38), 9805–9809. http://doi.org/10.1523/JNEUROSCI.3165-06.2006

Hopfinger, J. B., Buonocore, M. H., & Mangun, G. R. (2000). The neural mechanisms of top-down attentional control. *Nature Neuroscience*, *3*(3), 284–291. http://doi.org/10.1038/72999

Humphries, M. D., & Gurney, K. (2008). Network "small-world-ness": a quantitative method for determining canonical network equivalence. *PLoS One*, *3*(4), e0002051. http://doi.org/10.1371/journal.pone.0002051

Jiang, Y., & Kanwisher, N. (2003). Common neural substrates for response selection across modalities and mapping paradigms. *Journal of Cognitive Neuroscience*, *15*(8), 1080–1094. http://doi.org/10.1162/089892903322598067

John, E. R., Easton, P., & Isenhart, R. (1997). Consciousness and cognition may be mediated by multiple independent coherent ensembles. *Consciousness and Cognition*.

Jonides, J. (1981). Voluntary versus automatic control over the mind"s eye"s movement. Attention and performance IX.

Jonides, J. (2004). How does practice makes perfect? *Nature Neuroscience*, *7*(1), 10–11. http://doi.org/10.1038/nn0104-10

Kahneman, D. (1973). Attention and effort. Prentice Hall.

Kahneman, D., & Beatty, J. (1966). Pupil diameter and load on memory. *Science*, *154*(3756), 1583–1585.

Kastner, S., De Weerd, P., Desimone, R., & Ungerleider, L. G. (1998). Mechanisms of directed attention in the human extrastriate cortex as revealed by functional MRI. *Science*, *282*(5386), 108–111. http://doi.org/10.1126/science.282.5386.108

Kastner, S., Pinsk, M. A., De Weerd, P., Desimone, R., & Ungerleider, L. G. (1999). Increased activity in human visual cortex during directed attention in the absence of visual stimulation. *Neuron*, *22*(4), 751–761.

Kelley, T. A., & Lavie, N. (2011). Working memory load modulates distractor competition in primary visual cortex. *Cerebral Cortex (New York, NY : 1991)*, *21*(3), 659–665. http://doi.org/10.1093/cercor/bhq139

Kim, C.-Y., & Blake, R. (2005). Psychophysical magic: rendering the visible 'invisible'. *Trends in Cognitive Sciences*, *9*(8), 381–388. http://doi.org/10.1016/j.tics.2005.06.012

Kim, J., & Horwitz, B. (2008). Investigating the neural basis for fMRI-based functional connectivity in a blocked design: application to interregional correlations and psycho-physiological interactions. *Magnetic Resonance Imaging*, *26*(5), 583–593. http://doi.org/10.1016/j.mri.2007.10.011

Kincade, J. M., Abrams, R. A., Astafiev, S. V., Shulman, G. L., & Corbetta, M. (2005). An event-related functional magnetic resonance imaging study of voluntary and stimulus-driven orienting of attention. *Journal of Neuroscience*, *25*(18), 4593–4604. http://doi.org/10.1523/JNEUROSCI.0236-05.2005

King, J.-R., Sitt, J. D., Faugeras, F., Rohaut, B., Karoui, El, I., Cohen, L., et al. (2013). Information sharing in the brain indexes consciousness in noncommunicative patients. *Current Biology*, *23*(19), 1914–1919. http://doi.org/10.1016/j.cub.2013.07.075

Kitzbichler, M. G., Henson, R. N. A., Smith, M. L., Nathan, P. J., & Bullmore, E. T. (2011). Cognitive Effort Drives Workspace Configuration of Human Brain Functional Networks. *Journal of Neuroscience*, *31*(22), 8259–8270. http://doi.org/10.1523/JNEUROSCI.0440-11.2011

Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3. *Perception*, *36*.

Kleinschmidt, A., Büchel, C., Zeki, S., & Frackowiak, R. S. (1998). Human brain activity during spontaneously reversing perception of ambiguous figures. *Proceedings. Biological Sciences / the Royal Society*, *265*(1413), 2427–2433. http://doi.org/10.1098/rspb.1998.0594

Klugkist, I., Laudy, O., & Hoijtink, H. (2005). Inequality Constrained Analysis of Variance: A Bayesian Approach. *Psychological Methods*, *10*(4), 477–493. http://doi.org/10.1037/1082-989X.10.4.477

Koch, C., & Tsuchiya, N. (2007). Attention and consciousness: two distinct brain processes. *Trends in Cognitive Sciences*, *11*(1), 16–22. http://doi.org/10.1016/j.tics.2006.10.012

Kranczioch, C., Debener, S., Schwarzbach, J., Goebel, R., & Engel, A. K. (2005). Neural correlates of conscious perception in the attentional blink. *NeuroImage*, *24*(3), 704–714. http://doi.org/10.1016/j.neuroimage.2004.09.024

Kunimoto, C., Miller, J., & Pashler, H. (2001). Confidence and Accuracy of Near-Threshold Discrimination Responses. *Consciousness and Cognition*, *10*(3), 294–340. http://doi.org/10.1006/ccog.2000.0494

Lamme, V. A. F. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences*, *7*(1), 12–18.

Latora, V., & Marchiori, M. (2001). Efficient Behavior of Small-World Networks. *Physical Review Letters*, *87*(19), 198701. http://doi.org/10.1103/PhysRevLett.87.198701

Lau, E. F., Phillips, C., & Poeppel, D. (2008). A cortical network for semantics: (de)constructing the N400. *Nature Reviews Neuroscience*, *9*(12), 920–933. http://doi.org/10.1038/nrn2532

Lau, H. C., & Passingham, R. E. (2006). Relative blindsight in normal observers and the neural correlate of visual consciousness. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(49), 18763–18768. http://doi.org/10.1073/pnas.0607716103

Lefton, L. A., & Newman, Y. (1976). Metacontrast and paracontrast: Both photopic and scotopic luminance levels yield monotones. *Bulletin of the Psychonomic Society*.

Liao, W., Ding, J., Marinazzo, D., Xu, Q., Wang, Z., Yuan, C., et al. (2011). Small-world directed networks in the human brain: multivariate Granger causality analysis of resting-state fMRI. *NeuroImage*, *54*(4), 2683–2694. http://doi.org/10.1016/j.neuroimage.2010.11.007

Lord, L.-D., Allen, P., Expert, P., Howes, O., Lambiotte, R., McGuire, P., et al. (2011). Characterization of the anterior cingulate's role in the at-risk mental state using graph theory. *NeuroImage*, *56*(3), 1531–1539. http://doi.org/10.1016/j.neuroimage.2011.02.012

Lumer, E. D., & Rees, G. (1999). Covariation of activity in visual and prefrontal cortex associated with subjective visual perception. *Proceedings of the National Academy of Sciences of the United States of America*, *96*(4), 1669–1673.

Lumer, E. D., Friston, K. J., & Rees, G. (1998). Neural correlates of perceptual rivalry in the human brain. *Science*, *280*(5371), 1930–1934.

Lynall, M.-E., Bassett, D. S., Kerwin, R., McKenna, P. J., Kitzbichler, M., Muller, U., & Bullmore, E. (2010). Functional connectivity and brain networks in schizophrenia. *Journal of Neuroscience*, *30*(28), 9477–9487. http://doi.org/10.1523/JNEUROSCI.0333-10.2010

Marois, R., & Ivanoff, J. (2005). Capacity limits of information

processing in the brain. *Trends in Cognitive Sciences*, *9*(6), 1–10. http://doi.org/10.1016/j.tics.2005.04.010

Marois, R., Larson, J. M., Chun, M. M., & Shima, D. (2005). Response-specific sources of dual-task interference in human pre-motor cortex. *Psychological Research*, *70*(6), 436–447. http://doi.org/10.1007/s00426-005-0022-6

Marois, R., Leung, H. C., & Gore, J. C. (2000). A stimulus-driven approach to object identity and location processing in the human brain. *Neuron*, *25*(3), 717–728.

Marois, R., Yi, D.-J., & Chun, M. M. (2004). The neural fate of consciously perceived and missed events in the attentional blink. *Neuron*, *41*(3), 465–472.

McIntosh, A. R. (2000). Towards a network theory of cognition. *Neural Networks : the Official Journal of the International Neural Network Society*, *13*(8-9), 861–870.

McLaren, D. G., Ries, M. L., Xu, G., & Johnson, S. C. (2012). A generalized form of context-dependent psychophysiological interactions (gPPI): a comparison to standard approaches. *NeuroImage*, *61*(4), 1277–1286. http://doi.org/10.1016/j.neuroimage.2012.03.068

Melloni, L., Molina, C., Pena, M., Torres, D., Singer, W., & Rodriguez, E. (2007). Synchronization of neural activity across cortical areas correlates with conscious perception. *Journal of Neuroscience*, *27*(11), 2858–2865. http://doi.org/10.1523/JNEUROSCI.4623-06.2007

Menon, V. (2011). Large-scale brain networks and psychopathology: a unifying triple network model. *Trends in Cognitive Sciences*, *15*(10), 483–506. http://doi.org/10.1016/j.tics.2011.08.003

Menon, V., & Uddin, L. Q. (2010). Saliency, switching, attention and control: a network model of insula function. *Brain Structure and Function*, *214*(5-6), 655–667. http://doi.org/10.1007/s00429-010-0262-0

Mesulam, M. (1990). Large-scale neurocognitive networks and distributed processing for attention, language, and memory. *Annals of Neurology*, *28*(5), 597–613.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*(1), 167–202. http://doi.org/10.1146/annurev.neuro.24.1.167

Moore, T. (2003). Microstimulation of the Frontal Eye Field and Its Effects on Covert Spatial Attention. *Journal of Neurophysiology*, *91*(1), 152–162. http://doi.org/10.1152/jn.00741.2002

Moutoussis, K., & Zeki, S. (2002). The relationship between cortical activation and perception investigated with invisible stimuli. *Proceedings of the National Academy of Sciences*, *99*(14), 9527–9532. http://doi.org/10.1073/pnas.142305699

Mulder, J., Klugkist, I., van de Schoot, R., Meeus, W. H. J., Selfhout, M., & Hoijtink, H. (2009). Bayesian model selection of informative hypotheses for repeated measurements. *Journal of Mathematical Psychology*, *53*(6), 530–546. http://doi.org/10.1016/j.jmp.2009.09.003

Naghavi, H. R., & Nyberg, L. (2005). Common fronto-parietal activity in attention, memory, and consciousness: shared demands on integration? *Consciousness and Cognition*, *14*(2), 390–425. http://doi.org/10.1016/j.concog.2004.10.003

Neisser, U., & Lazar, R. (1964). Searching for novel targets. *Perceptual and Motor Skills*. http://doi.org/10.2466/pms.1964.19.2.427

Newman, M. (2004). Analysis of weighted networks. *Physical Review E*, *70*(5), 056131. http://doi.org/10.1103/PhysRevE.70.056131

Nicol, R. M., Chapman, S. C., Vertes, P. E., Nathan, P. J., Smith, M. L., Shtyrov, Y., & Bullmore, E. T. (2012). Fast reconfiguration of high-frequency brain networks in response to surprising changes in auditory input. *Journal of Neurophysiology*, *107*(5), 1421–1430. http://doi.org/10.1152/jn.00817.2011

Nomura, E. M., Gratton, C., Visser, R. M., Kayser, A., Perez, F., & D'Esposito, M. (2010). Double dissociation of two cognitive control networks in patients with focal brain lesions. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(26), 12017–12022. http://doi.org/10.1073/pnas.1002431107/-/DCSupplemental

Norman-Haignere, S. V., McCarthy, G., Chun, M. M., & Turk-Browne, N. B. (2011). Category-Selective Background Connectivity in Ventral Visual Cortex. *Cerebral Cortex (New York, NY : 1991)*. http://doi.org/10.1093/cercor/bhr118

Ogawa, S., Menon, R. S., Tank, D. W., Kim, S. G., Merkle, H., Ellermann, J. M., & Ugurbil, K. (1993). Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging. A comparison of signal characteristics with a biophysical model.

Biophysical Journal, *64*(3), 803–812. http://doi.org/10.1016/S0006-3495(93)81441-3

Palva, S., & Palva, J. M. (2012). Discovering oscillatory interaction networks with M/EEG: challenges and breakthroughs. *Trends in Cognitive Sciences*, *16*(4), 219–230. http://doi.org/10.1016/j.tics.2012.02.004

Palva, S., Monto, S., & Palva, J. M. (2010). Graph properties of synchronized cortical networks during visual working memory maintenance. *NeuroImage*, *49*(4), 3257–3268. http://doi.org/10.1016/j.neuroimage.2009.11.031

Park, H. J., & Friston, K. (2013). Structural and Functional Brain Networks: From Connections to Cognition. *Science*, *342*(6158), 1238411–1238411. http://doi.org/10.1126/science.1238411

Pashler, H. (1994). Dual-task interference in simple tasks: data and theory. *Psychological Bulletin*, *116*(2), 220–244.

Pavlov, I. (1927). Conditioned reflexes: an investigation of the physiological processes of the cerebral cortex. London: Oxford UP. http://doi.org/10.1111/j.1365-2869.2010.00842.x/full

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442.

Petersen, S. E., & Posner, M. I. (2012). The attention system of the human brain: 20 years after. *Annual Review of Neuroscience*, *35*, 73–89. http://doi.org/10.1146/annurev-neuro-062111-150525

Pettersson-Yeo, W., Allen, P., Benetti, S., McGuire, P., & Mechelli, A. (2011). Neuroscience and Biobehavioral Reviews. *Neuroscience and Biobehavioral Reviews*, *35*(5), 1110–1124. http://doi.org/10.1016/j.neubiorev.2010.11.004

Phelps, E. A., Ling, S., & Carrasco, M. (2006). Emotion facilitates perception and potentiates the perceptual benefits of attention. *Psychological Science*, *17*(4), 292–299. http://doi.org/10.1111/j.1467-9280.2006.01701.x

Poldrack, R. A. (2005). The Neural Correlates of Motor Skill Automaticity. *Journal of Neuroscience*, *25*(22), 5356–5364. http://doi.org/10.1523/JNEUROSCI.3880-04.2005

Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. *Annual Review of Neuroscience*, *13*, 25–42. http://doi.org/10.1146/annurev.ne.13.030190.000325

Posner, M. I., Petersen, S. E., Fox, P. T., & Raichle, M. E. (1988). Localization of cognitive operations in the human brain. *Science*, *240*(4859), 1627–1631.

Power, J. D., Barnes, K. A., Snyder, A. Z., Schlaggar, B. L., & Petersen, S. E. (2012). Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *NeuroImage*, *59*(3), 2142–2154. http://doi.org/10.1016/j.neuroimage.2011.10.018

Power, J. D., Cohen, A. L., Nelson, S. M., Wig, G. S., Barnes, K. A., Church, J. A., et al. (2011). Functional Network Organization of the Human Brain. *Neuron*, *72*(4), 665–678. http://doi.org/10.1016/j.neuron.2011.09.006

Rees, G., Kreiman, G., & Koch, C. (2002). Neural correlates of consciousness in humans. *Nature Reviews Neuroscience*, *3*(4), 261–270. http://doi.org/10.1038/nrn783

Ress, D., & Heeger, D. J. (2003). Neuronal correlates of perception in early visual cortex. *Nature Neuroscience*, *6*(4), 414–420. http://doi.org/10.1038/nn1024

Reynolds, J. H., & Chelazzi, L. (2004). Attentional modulation of visual processing. *Annual Review of Neuroscience*, *27*(1), 611–647. http://doi.org/10.1146/annurev.neuro.26.041002.131039

Reynolds, J. H., Chelazzi, L., & Desimone, R. (1999). Competitive Mechanisms Subserve Attention in Macaque Areas V2 and V4. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *19*(5), 1736–1753. http://doi.org/10.1146/annurev.ne.18.030195.001205

Rogers, B. P., Morgan, V. L., Newton, A. T., & Gore, J. C. (2007). Assessing functional connectivity in the human brain by fMRI. *Magnetic Resonance Imaging*, *25*(10), 1347–1357. http://doi.org/10.1016/j.mri.2007.03.007

Rosenberg, M. D., Finn, E. S., Scheinost, D., Papademetris, X., Shen, X., Constable, R. T., & Chun, M. M. (2015). A neuromarker of sustained attention from whole-brain functional connectivity. *Nature Neuroscience*, *19*(1), 165–171. http://doi.org/10.1038/nn.4179

Rubinov, M., & Sporns, O. (2010). Complex network measures of brain connectivity: uses and interpretations. *NeuroImage*, *52*(3), 1059–1069. http://doi.org/10.1016/j.neuroimage.2009.10.003

Rubinov, M., & Sporns, O. (2011). Weight-conserving characterization of complex functional brain networks. *NeuroImage*, *56*(4), 2068–2079. http://doi.org/10.1016/j.neuroimage.2011.03.069

Rubinov, M., Knock, S. A., Stam, C. J., Micheloyannis, S., Harris, A. W. F., Williams, L. M., & Breakspear, M. (2009). Small-world properties of nonlinear brain activity in schizophrenia. *Human Brain Mapping*, *30*(2), 403–416. http://doi.org/10.1002/hbm.20517

Sadaghiani, S., Hesselmann, G., & Kleinschmidt, A. (2009). Distributed and Antagonistic Contributions of Ongoing Activity Fluctuations to Auditory Stimulus Detection. *Journal of Neuroscience*, *29*(42), 13410–13417. http://doi.org/10.1523/JNEUROSCI.2592-09.2009

Sara, S. J., & Bouret, S. (2012). Orienting and Reorienting: The Locus Coeruleus Mediates Cognition through Arousal. *Neuron*, *76*(1), 130–141. http://doi.org/10.1016/j.neuron.2012.09.011

Schall, J. D. (2002). The neural selection and control of saccades by the frontal eye field. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *357*(1424), 1073–1082. http://doi.org/10.1098/rstb.2002.1098

Schiller, P. H., & Smith, M. C. (1966). Detection in metacontrast. *Journal of Experimental Psychology*, *71*(1), 32–39.

Schröter, M. S., Spoormaker, V. I., Schorer, A., Wohlschläger, A., Czisch, M., Kochs, E. F., et al. (2012). Spatiotemporal reconfiguration of large-scale brain functional networks during propofol-induced loss of consciousness. *Journal of Neuroscience*, *32*(37), 12832–12840. http://doi.org/10.1523/JNEUROSCI.6046-11.2012

Schumacher, E. H., Seymour, T. L., Glass, J. M., Fencsik, D. E., Lauber, E. J., Kieras, D. E., & Meyer, D. E. (2001). Virtually Perfect Time Sharing in Dual-Task Performance: Uncorking the Central Cognitive Bottleneck. *Psychological Science*, *12*(2), 101–108. http://doi.org/10.1111/1467-9280.00318

Seeley, W. W., Menon, V., Schatzberg, A. F., Keller, J., Glover, G. H., Kenna, H., et al. (2007). Dissociable Intrinsic Connectivity Networks for Salience Processing and Executive Control. *Journal of Neuroscience*, *27*(9), 2349–2356. http://doi.org/10.1523/JNEUROSCI.5587-06.2007

Serences, J. T. (2004). A comparison of methods for characterizing the event-related BOLD timeseries in rapid fMRI. *NeuroImage*, *21*(4), 1690–1700. http://doi.org/10.1016/j.neuroimage.2003.12.021

Serences, J. T., Schwarzbach, J., Courtney, S. M., Golay, X., & Yantis, S. (2004a). Control of object-based attention in human cortex. *Cerebral Cortex (New York, NY : 1991)*, *14*(12), 1346–1357. http://doi.org/10.1093/cercor/bhh095

Serences, J. T., Shomstein, S., Leber, A. B., Golay, X., Egeth, H. E., & Yantis, S. (2005). Coordination of voluntary and stimulus-driven attentional control in human cortex. *Psychological Science*, *16*(2), 114–122. http://doi.org/10.1111/j.0956-7976.2005.00791.x

Serences, J. T., Yantis, S., Culberson, A., & Awh, E. (2004b). Preparatory activity in visual cortex indexes distractor suppression during covert spatial orienting. *Journal of Neurophysiology*, *92*(6), 3538–3545. http://doi.org/10.1152/jn.00435.2004

Seth, A. K., Dienes, Z., Cleeremans, A., Overgaard, M., & Pessoa, L. (2008). Measuring consciousness: relating behavioural and neurophysiological approaches. *Trends in Cognitive Sciences*, *12*(8), 314–321. http://doi.org/10.1016/j.tics.2008.04.008

Seth, A. K., Izhikevich, E., Reeke, G. N., & Edelman, G. M. (2006). Theories and measures of consciousness: an extended framework. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(28), 10799–10804. http://doi.org/10.1073/pnas.0604347103

Sheffield, J. M., Repovs, G., Harms, M. P., Carter, C. S., Gold, J. M., MacDonald, A. W., Daniel Ragland, J., et al. (2015a). Fronto-parietal and cingulo-opercular network integrity and cognition in health and schizophrenia. *Neuropsychologia*, *73*, 82–93. http://doi.org/10.1016/j.neuropsychologia.2015.05.006

Sheffield, J. M., Repovs, G., Harms, M. P., Carter, C. S., Gold, J. M., MacDonald, A. W., Ragland, J. D., et al. (2015b). Evidence for Accelerated Decline of Functional Brain Network Efficiency in Schizophrenia. *Schizophrenia Bulletin*, sbv148. http://doi.org/10.1093/schbul/sbv148

Shulman, G. L., Ollinger, J. M., Akbudak, E., Conturo, T. E., Snyder, A. Z., Petersen, S. E., & Corbetta, M. (1999). Areas involved in encoding and applying directional expectations to moving objects. *Journal of Neuroscience*, *19*(21), 9480–9496.

Sigman, M., & Dehaene, S. (2008). Brain mechanisms of serial and parallel processing during dual-task performance. *Journal of Neuroscience*, *28*(30), 7585–7598. http://doi.org/10.1523/JNEUROSCI.0948-08.2008

Smith, M. C. (1967). Theories of the psychological refractory period. *Psychological Bulletin*, *67*(3), 202–213.

Sokolov, E. N. (1963). Higher nervous functions: The orienting reflex. *Annual Review of Physiology*.

Spoormaker, V. I., Schröter, M. S., Gleiser, P. M., Andrade, K. C., Dresler, M., Wehrle, R., et al. (2010). Development of a large-scale functional brain network during human non-rapid eye movement sleep. *Journal of Neuroscience*, *30*(34), 11379–11387. http://doi.org/10.1523/JNEUROSCI.2015-10.2010

Sporns, O. (2011). The human connectome: a complex network. *Annals of the New York Academy of Sciences*, *1224*, 109–125. http://doi.org/10.1111/j.1749-6632.2010.05888.x

Sporns, O., & Betzel, R. F. (2015). Modular Brain Networks. *Annu. Rev. Psychol.*, *67*(1), annurev–psych–122414–033634–29. http://doi.org/10.1146/annurev-psych-122414-033634

Sporns, O., & Zwi, J. D. (2004). The small world of the cerebral cortex. *Neuroinformatics*, *2*(2), 145–162. http://doi.org/10.1385/NI:2:2:145

Sporns, O., Honey, C. J., & Kötter, R. (2007). Identification and classification of hubs in brain networks. *PLoS One*, *2*(10), e1049. http://doi.org/10.1371/journal.pone.0001049

Sterzer, P., Kleinschmidt, A., & Rees, G. (2009). The neural bases of multistable perception. *Trends in Cognitive Sciences*, *13*(7), 310–318. http://doi.org/10.1016/j.tics.2009.04.006

Stevens, A. A., Tappon, S. C., Garg, A., & Fair, D. A. (2012). Functional Brain Network Modularity Captures Inter- and Intra-Individual Variation in Working Memory Capacity. *PLoS One*, *7*(1), e30468. http://doi.org/10.1371/journal.pone.0030468.s004

Stevens, M. C., Calhoun, V. D., & Kiehl, K. A. (2005). fMRI in an oddball task: Effects of target-to-target interval. *Psychophysiology*, *42*(6), 636–642. http://doi.org/10.1111/j.1469-8986.2005.00368.x

Supekar, K., Musen, M., & Menon, V. (2009). Development of large-scale functional brain networks in children. *Plos Biology*, *7*(7), e1000157. http://doi.org/10.1371/journal.pbio.1000157

Talairach, J., & Tournoux, P. (1988). Co-planar stereotaxic atlas of the human brain. 3-Dimensional proportional system: an approach to cerebral imaging. *(1988)*.

Theeuwes, J., Kramer, A. F., & Kingstone, A. (2004). Attentional capture modulates perceptual sensitivity. *Psychonomic Bulletin & Review*, *11*(3), 551–554.

Thompson, K. G., Biscoe, K. L., & Sato, T. R. (2005). Neuronal basis of covert spatial attention in the frontal eye field. *Journal of Neuroscience*, *25*(41), 9479–9487. http://doi.org/10.1523/JNEUROSCI.0741-05.2005

Thompson, R., & Duncan, J. (2009). Attentional modulation of stimulus representation in human fronto-parietal cortex. *NeuroImage*, *48*(2), 436–448. http://doi.org/10.1016/j.neuroimage.2009.06.066

Tombu, M. N., Asplund, C. L., Dux, P. E., Godwin, D., Martin, J. W., & Marois, R. (2011). A Unified attentional bottleneck in the human brain. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(33), 13426–13431.

Tombu, M., & Jolicoeur, P. (2004). Virtually no evidence for virtually perfect time-sharing. *Journal of Experimental Psychology: Human Perception and Performance*, *30*(5), 795–810. http://doi.org/10.1037/0096-1523.30.5.795

Tononi, G. (2008). Consciousness as integrated information: a provisional manifesto. *The Biological Bulletin*, *215*(3), 216–242.

Tononi, G., & Edelman, G. M. (1998). Consciousness and complexity. *Science*, *282*(5395), 1846–1851.

Tononi, G., & Sporns, O. (2003). Measuring information integration. *BMC Neuroscience*, *4*, 31. http://doi.org/10.1186/1471-2202-4-31

Tononi, G., Sporns, O., & Edelman, G. M. (1994). A measure for brain complexity: relating functional segregation and integration in the nervous system. *Proceedings of the National Academy of Sciences of the United States of America*, *91*(11), 5033–5037.

Tracy, J. I., Mohamed, F., Faro, S., Tiver, R., Pinus, A., Bloomer, C., et al. (2000). The effect of autonomic arousal on attentional focus. *Neuroreport*, *11*(18), 4037–4042.

Triesman, A. M. (1964). Selective attention in man. British Medical Bulletin.

Tse, P. U., Martinez-Conde, S., Schlegel, A. A., & Macknik, S. L. (2005). Visibility, visual awareness, and visual masking of simple unattended targets are confined to areas in the occipital cortex beyond human V1/V2. *Proceedings of the National Academy of Sciences of the United States of America*, *102*(47), 17178–17183. http://doi.org/10.1073/pnas.0508010102

Uddin, L. Q., Supekar, K., Amin, H., Rykhlevskaia, E., Nguyen, D. A., Greicius, M. D., & Menon, V. (2010). Dissociable connectivity within human angular gyrus and intraparietal sulcus: evidence from functional and structural connectivity. *Cerebral Cortex (New York, NY : 1991)*, *20*(11), 2636–2646. http://doi.org/10.1093/cercor/bhq011

Uehara, T., Yamasaki, T., Okamoto, T., Koike, T., Kan, S., Miyauchi, S., et al. (2014). Efficiency of a "Small-World" Brain Network Depends on Consciousness Level: A Resting-State fMRI Study. *Cerebral Cortex (New York, NY : 1991)*, *24*(6), 1529–1539. http://doi.org/10.1093/cercor/bht004

van den Heuvel, M. P., & Sporns, O. (2013). Network hubs in the human brain. *Trends in Cognitive Sciences*, *17*(12), 683–696. http://doi.org/10.1016/j.tics.2013.09.012

van den Heuvel, M. P., Stam, C. J., Boersma, M., & Pol, H. E. H. (2008). Small-world and scale-free organization of voxel-based resting-state functional connectivity in the human brain. *NeuroImage*, *43*(3), 528–539. http://doi.org/10.1016/j.neuroimage.2008.08.010

Van Dijk, K. R. A., Sabuncu, M. R., & Buckner, R. L. (2012). The influence of head motion on intrinsic functional connectivity MRI. *NeuroImage*, *59*(1), 431–438. http://doi.org/10.1016/j.neuroimage.2011.07.044

van Gaal, S., & Lamme, V. A. F. (2012). Unconscious high-level information processing: implication for neurobiological theories of consciousness. *The Neuroscientist*, *18*(3), 287–301. http://doi.org/10.1177/1073858411404079

van Wijk, B. C. M., Stam, C. J., & Daffertshofer, A. (2010). Comparing brain networks of different size and connectivity density using graph theory. *PLoS One*, *5*(10), e13701. http://doi.org/10.1371/journal.pone.0013701

Varela, F., Lachaux, J. P., Rodriguez, E., & Martinerie, J. (2001). The brainweb: phase synchronization and large-scale integration. *Nature Reviews Neuroscience*, *2*(4), 229–239. http://doi.org/10.1038/35067550

Vincent, J. L., Kahn, I., Snyder, A. Z., Raichle, M. E., & Buckner, R. L. (2008). Evidence for a Frontoparietal Control System Revealed by Intrinsic Functional Connectivity. *Journal of Neurophysiology*, *100*(6), 3328–3342. http://doi.org/10.1152/jn.90355.2008

Vossel, S., Thiel, C. M., & Fink, G. R. (2006). Cue validity modulates the neural correlates of covert endogenous orienting of attention in parietal and frontal cortex. *NeuroImage*, *32*(3), 1257–1264. http://doi.org/10.1016/j.neuroimage.2006.05.019

Wager, T. D., Jonides, J., & Reading, S. (2004). Neuroimaging studies of shifting attention: a meta-analysis. *NeuroImage*, *22*(4), 1679–1693. http://doi.org/10.1016/j.neuroimage.2004.03.052

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of “small-world” networks. *Nature*, *393*(6684), 440–442. http://doi.org/10.1038/30918

Weissman, D. H., & Prado, J. (2012). Heightened activity in a key region of the ventral attention network is linked to reduced activity in a key region of the dorsal attention network during unexpected shifts of covert visual spatial attention. *NeuroImage*, *61*(4), 798–804. http://doi.org/10.1016/j.neuroimage.2012.03.032

Weisz, N., Wuhle, A., Monittola, G., Demarchi, G., Frey, J., Popov, T., & Braun, C. (2014). Prestimulus oscillatory power and connectivity patterns predispose conscious somatosensory perception. *Proceedings of the National Academy of Sciences*, *111*(4), E417–E425. http://doi.org/10.1073/pnas.1317267111

Welford, A. T. (1952). The “Psychological Refractory Period” and the Timing of High-Speed Performance - A Review and a Theory. *British Journal of Psychology General Section*, *43*(1), 2–19. http://doi.org/10.1111/j.2044-8295.1952.tb00322.x

Wolfe, J. M. (2003). Moving towards solutions to some enduring controversies in visual search. *Trends in Cognitive Sciences*, *7*(2), 70–76.

Womelsdorf, T., Fries, P., Mitra, P. P., & Desimone, R. (2005). Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature*, *439*(7077), 733–736. http://doi.org/10.1038/nature04258

Yantis, S., & Egeth, H. E. (1999). On the distinction between visual salience and stimulus-driven attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, *25*(3), 661–676.

Yantis, S., & Jonides, J. (1990). Abrupt visual onsets and selective attention: voluntary versus automatic allocation. *Journal of Experimental Psychology: Human Perception and Performance*, *16*(1), 121–134.

Yantis, S., Schwarzbach, J., Serences, J. T., Carlson, R. L., Steinmetz, M. A., Pekar, J. J., & Courtney, S. M. (2002). Transient neural activity in human parietal cortex during spatial attention shifts. *Nature Publishing Group*, *5*(10), 995–1002. http://doi.org/10.1038/nn921

Yu, Q., Sui, J., Rachakonda, S., He, H., Pearlson, G., & Calhoun, V. D. (2011). Altered small-world brain networks in temporal lobe in patients with schizophrenia performing an auditory oddball task. *Frontiers in Systems Neuroscience*, *5*, 7. http://doi.org/10.3389/fnsys.2011.00007

Zeki, S. (2007). The disunity of consciousness. *Boundaries of Consciousness: Neurobiology and Neuropathology*, *168*, 11–268. http://doi.org/10.1016/S0079-6123(07)68002-9