A UNITARY FRAMEWORK DEFINING THE FUNCTIONAL SIGNIFICANCE OF NEURAL OSCILLATIONS IN THE ALPHA FREQUENCY

A Thesis
Submitted to the Graduate Faculty
of the
North Dakota State University
of Agriculture and Applied Science

By

Andrew Jordan Heinz

In Partial Fulfillment of the Requirements
for the Degree of
MASTER OF SCIENCE

Major Program:
Visual and Cognitive Neuroscience

June 2016

Fargo, North Dakota

North Dakota State University
Graduate School

Title

A UNITARY FRAMEWORK DEFINING THE FUNCTIONAL SIGNIFICANCE OF NEURAL OSCILLATIONS IN THE ALPHA FREQUENCY

By

Andrew Jordan Heinz

The Supervisory Committee certifies that this ***disquisition*** complies with North Dakota State University's regulations and meets the accepted standards for the degree of

MASTER OF SCIENCE

SUPERVISORY COMMITTEE:

Jeffrey Johnson
Chair

Benjamin Balas

Clayton Hilmert

Roger Green

Approved:

6/24/16	James Council
Date	Department Chair

ABSTRACT

Neural oscillations in the alpha frequency band (8-12 Hz) are among the most salient and ubiquitous phenomena observed in human electroencephalographic recordings. There have been various proposals regarding the functional significance of these oscillations, including the inhibition of task-irrelevant cortical areas, and the active retention of information in visual working memory. In an attempt to delineate between these two alternatives, I recorded EEG while participants performed two tasks requiring the short-term retention of visual working memory items. The results suggest that alpha-band oscillations reflect the implementation of an attentionally selective executive control process, inhibiting task-irrelevant processing, aiding the stability of active retention.

ACKNOWLEDGMENTS

First and foremost, I would like to thank my advisor, Dr. Jeffrey Johnson, for his guidance, mentorship and support these few past years. I would also like to thank my thesis committee, Dr. Benjamin Balas, Dr. Clayton Hilmert, and Dr. Roger Green, for all their contributions. Finally, I would like to thank my fellow lab members Dr. Amanda van Lamsweerde and Andrea Bocincova for all the assistance they provided in completing this project.

DEDICATION

I dedicate this manuscript to my parents, George and Kathleen Heinz, without whom I would be nowhere.

TABLE OF CONTENTS

LIST OF FIGURES

INTRODUCTION

"If you want to find the secrets of the universe, think in terms of energy, frequency and vibration." - Nikola Tesla

In 1929, Hans Berger reported a series of experiments demonstrating that electrical potentials generated by the brain could be recorded noninvasively from the human scalp (Berger, 1929). These early studies of the human electroencephalogram (EEG) revealed the presence of oscillatory rhythms embedded within neural activity, spanning a wide range of frequencies. Berger labeled the most prominent of these oscillations the alpha frequency, as this was the first distinct frequency band to be identified from scalp recordings. The alpha-band spans a frequency range from roughly 8-12 Hz. The amplitude of these oscillations was found to increase when an individual's eyes were closed and decrease when they were opened. Replications and extensions of these early observations led to the suggestion that increased alpha-band power (ABP) reflects an idling state within a given brain region. This is in contrast to higher frequency beta (~15-30 Hz) and gamma (~40-100 Hz) band oscillations associated with active processing (e.g.Pfurtscheller, Stancák, & Neuper, 1996) More recent evidence, however, has revealed ABP increases in a range of tasks involving attention and working memory, suggesting that alpha-band oscillations may play a more prominent role in cognitive processing than the idling account would suggest.

Research examining the functional role of alpha-band oscillations in cognition has led to two competing sets of proposals. The first associates increased power in the alpha-band with an inhibitory process that serves to mitigate the disruptive effects of distracting information during attention and working memory (WM) tasks (Dube, Payne, Sekuler, & Rotello, 2013; Fu et al.,

2001; Jensen, Gelfand, Kounios, & Lisman, 2002; Jensen & Mazaheri, 2010; Klimesch, Sauseng, & Hanslmayr, 2007; Payne, Guillory, & Sekuler, 2013). The second view, by contrast, suggests that increased ABP may instead reflect a critical component of the distributed network activity underlying selective attention and the maintenance of information in WM (Johnson, Sutterer, Acheson, Lewis-Peacock, & Postle, 2011; Mo, Schroeder, & Ding, 2011; S. Palva & Palva, 2007). By this view, increases in ABP reflect the active processes associated with increasing demands on attention and maintenance related neural systems, rather than inhibition. More recently, in an attempt to rectify these opposing viewpoints, it has been proposed that the specific functional contributions reflected by oscillatory activity in the alpha-band may vary between these two alternative accounts on the basis of the specific task's demands and/or the corresponding regions engaged (Jensen et al., 2002; J. M. Palva, Monto, Kulashekhar, & Palva, 2010; Sauseng et al., 2009).

In the sections that follow, I will review several lines of evidence that have been taken as support for each of these alternative accounts. On the basis of this review, I proposed a hypothesis: that elevated power in the alpha frequency band during the performance of attention and WM tasks, reflects a more general mechanism of cortical inhibition oriented towards potentially disruptive processes and enacted across levels of functional specificity within the brain. To provide further evidence for this interpretation I proposed two experiments, each testing specific predictions derived from the proposed account. Simultaneously, reframing two previously reported observations within the visual working memory (VWM) literature which are at first seemingly inconsistent within this inhibitory account.

Alpha-Band Power Modulations in Studies of Selective Attention

In this section I will briefly review experiments associated with the selective allocation of attention as it relates to visual stimulus processing. In so doing, I will provide evidence supporting the hypothesis that oscillations in the alpha frequency band reflect a mechanism suppressing task irrelevant processing in at least these cases.

Inter-sensory Attention

One variety of task in which an association between oscillations in the alpha-band and inhibition has been well documented are those that entail inter-sensory suppression. Inter-sensory suppression refers to instances in which individuals are required to orient attention toward stimuli within a single sensory modality (e.g., audition) to the exclusion of others (e.g., vision). Studies examining intersensory attention have revealed increases in ABP over task irrelevant sensory regions and a corresponding decrease within relevant sensory regions (Dube et al., 2013; Fu et al., 2001; Payne et al., 2013). For instance, in an experiment conducted by Payne and colleagues (2013), participants were asked to attend to auditory stimuli while ignoring visual stimuli. In this case, they observed increased ABP over regions thought to be involved in visual processing. Similarly, Dube and colleagues (2013), asked participants to ignore auditory stimuli while attending to visual features and observed an increase in ABP over regions associated with audition (see also Banerjee, Snyder, Molholm, & Foxe, 2011; Foxe, Simpson, & Ahlfors, 1998). Similar dynamics have also been observed in relation to motor tasks (Pfurtscheller, 1992).

If increased ABP truly reflects a mechanism of suppression engaged in supporting goal-oriented processing, one might expect to observe a relationship between an individual's ability to modulate ABP as a function of engagement in a given task and the performance of that task.

Performance here would presumably reflect variability in the efficacy of goal oriented processing. In keeping with this possibility, it has been shown consistently that the degree to which participants effectively modulate distributed alpha-band activity in accordance with task demands is positively correlated with measures of performance in various tasks (Babiloni, Vecchio, Bultrini, Luca Romani, & Rossini, 2006; Del Percio et al., 2007; Hanslmayr et al., 2007; Linkenkaer-Hansen, Nikulin, Palva, Ilmoniemi, & Palva, 2004; Manza, Hau, & Leung, 2014; Mazaheri, Nieuwenhuis, van Dijk, & Jensen, 2009) For example, Bollimunta and colleagues (2008) recorded local field potentials via linear multielectrode arrays during an experiment requiring non-human primates to attend to either visual or auditory stimuli. They found that increased ABP in early visual regions predicted faster auditory stimulus detection. This is significant, as faster stimulus detection is thought to be associated with attended, as opposed to unattended items (Gottlob, 2004; Soto & Blanco, 2004)

Spatial Attention

Another instance in which the observed patterns of alpha-band activity have been closely associated with shifts in regional excitability is within the domain of spatial attention. Following the reorientation of visual-spatial attention, the observed topography of ABP has been found to generally shift as a function of the specific retinotopic focus. For example, shifting attention to a visual stimulus on the left side of a computer screen produces a lateralized decrease in ABP over right occipital regions, and a related increase in ABP over left occipital regions (Jensen, Bonnefond, & VanRullen, 2012; Wolfgang Klimesch, 2012; Worden, Foxe, Wang, & Simpson, 2000). Such modulations of ABP topography can be quite specific. For example, Rihs and

colleagues (2007) were able to differentiate event-related increases in the topography of ABP for attention to as many as eight retinotopically distinct sub-regions.

Shifts in the topographic distribution of ABP such as these are thought to reflect goal directed, selective inhibition of activity in brain regions representing task-irrelevant spatial processing. The capacity for spatial attention to enhance early visual processing for items within specific spatial regions, while suppressing processing within irrelevant regions is well documented (Cepeda, Cave, Bichot, & Kim, 1998; Hillyard, Hink, Schwent, & Picton, 1973; Watson & Humphreys, 1997). Generally, the later suppressive case is reflected by a relative decrease in early event related potential (ERP) components (e.g. *P1, N1*) to stimuli appearing within irrelevant spatial locations. Note that this mirrors the task-dependent shifts observed in the topography of ABP. Thus, if oscillations in the alpha frequency band truly reflect an inhibitory mechanism this would be the pattern of suppression one would expect. There is some evidence to suggest that ERP suppression and ABP are correlated when participants must selectively allocate visual spatial attention (Huang & Sekuler, 2010; Kelly, Lalor, Reilly, & Foxe, 2006; Worden et al., 2000). Further evidence supporting this possibility will be reviewed below (see: *Cortical Excitability).*

Feature-Based Attention

Similar modulations of ABP have also been observed in tasks requiring selective processing of particular feature dimensions (e.g., color, direction of motion, etc.). As before, when observers are asked to attend to one feature of a visible stimulus (e.g., color) and ignore other features (e.g., direction of motion), one observes a relative decrease in local ABP over brain regions corresponding to the task-*relevant* dimension, and an opposing increase over task-

irrelevant regions. For instance, Snyder and Foxe (2010) cued subjects to attend to either the color or direction of motion of an upcoming dot field array. Motion processing is thought to occur in dorsal stream areas MT+ and MST, whereas color processing is thought to occur in more ventral areas of the occipital lobes (e.g., area V4, adjacent to the fusiform gyrus). As one might expect, they found ABP increased over ventral regions when *motion* was cued and an opposing ABP increase over dorsal regions when *color* was cued.

Neural Oscillations in the Alpha-Band and Cortical Excitability

If neural oscillations in the alpha-band are truly engaged in suppressing irrelevant processing, one would expect to observe co-variation between measures of stimulus processing and local oscillatory activity in the alpha band. Some of the more common metrics reflecting the degree of stimulus processing include perceptual reports, the amplitude ERP components, and evoked blood oxygen level dependent activity (BOLD). In line with expectation, several studies have revealed a relationship between alpha-band activity and measures of cortical excitability. For example, modulations of ABP have been correlated with variability in both baseline (Bollimunta, Mo, Schroeder, & Ding, 2011) and stimulus-evoked (Haegens, Nácher, Luna, Romo, & Jensen, 2011) firing rates in studies of multi-unit activity recorded via implanted electrode arrays. Further, others report a correlation between ABP and changes in ERPs following the presentation of otherwise identical sensory stimuli (Barry et al., 2004; Barry, de Pascalis, Hodder, Clarke, & Johnstone, 2003; Brandt & Jansen, 1991, 1991; Lakatos, Karmos, Mehta, Ulbert, & Schroeder, 2008; Maris, van Vugt, & Kahana, 2011; Mathewson, Gratton, Fabiani, Beck, & Ro, 2009; Vanrullen, Busch, Drewes, & Dubois, 2011; Voytek et al., 2010; Worden et al., 2000). For example, Zanto and Gazzaley (2009) presented dot field arrays largely

identical to those used by Foxe and Snyder (2010). Importantly, they attempted to separate the two feature domains, alternately presenting colored dot arrays and gray dot motion arrays. They found that attending to one feature domain or the other resulted in early ERP component amplitude suppression to stimuli corresponding to the unattended feature dimension. Interestingly, as observed above (see: Bollimunta, Chen, Schroeder, & Ding, 2008), ABP modulations predicted performance in this case, such that shifts in ABP were significantly smaller on low performance trials (Zanto & Gazzaley, 2009). Similarly, there is evidence to suggest that disrupting anticipatory oscillatory activity in the alpha frequency band via TMS is also associated with disrupting behavioral performance in attention tasks (Capotosto, Babiloni, Romani, & Corbetta, 2009). This is precisely what one would expect if oscillatory activity in the alpha-band truly reflects a mechanism of suppression engaged in supporting goal-oriented processing.

Features of local alpha-band activity, such as the instantaneous phase of the oscillations, have also been correlated with ongoing gamma frequency band amplitude fluctuations (Maris et al., 2011; Voytek et al., 2010), as well as the magnitude of the evoked BOLD response in early visual regions (Scheeringa, Mazaheri, Bojak, Norris, & Kleinschmidt, 2011). Both gamma activity and evoked BOLD activity are thought to reflect active cortical processing. These findings imply that fluctuations in excitability are correlated with variability in both the instantaneous phase and power of local alpha-band oscillations.

In several cases, researchers report that both the power and phase of oscillations in the alpha band appear to interact to predict the probability of perceiving a visual stimulus (Busch, Dubois, & VanRullen, 2009; Lange, Keil, Schnitzler, van Dijk, & Weisz, 2014; Mathewson et

al., 2009; Romei et al., 2008; Romei, Driver, Schyns, & Thut, 2011; Thut et al., 2011). This is consistent with the apparent interaction between phase and power. Further, this more directly links oscillations in the alpha-band with visual perception and cortical excitability. Lange and colleagues (2013), for example, found that reduced occipital ABP was correlated with the probability of incorrectly perceiving two successive visual or tactile stimuli when only one was present. Romei and colleagues (2008) report that both the power and phase of spontaneous oscillations in the alpha-band predict variability in an individual's threshold for perceiving transcranial magnetic stimulation (TMS) induced phosphenes—the perception of localized flashes of light produced by TMS of the primary visual cortex (see also, Dugue et al, 2011).

Finally, Romei and colleagues (2011) further attempted to investigate more directly the ability of oscillations in the alpha-band to bias the perception of stimuli within select spatial locations, in a manner analogous to the above visual-spatial attention case. To do this, they used TMS to artificially boost 8-12 Hz oscillations in the occipital cortex. TMS was applied to the scalp over either the right or the left hemisphere while subjects tried to detect near threshold visual stimuli. They found pre-stimulus TMS in the alpha-band exclusively impaired detection of targets contralateral to the site of stimulation.

Alpha-Band Power Modulations during Visual Working Memory Maintenance

Research reviewed in the previous sections support a view of alpha-band oscillations in which relatively higher power reflects increasingly greater selective inhibition of brain areas representing task-irrelevant processing (e.g., sensory modalities, spatial locations or feature-domains). In the present section, I review the VWM literature related to neural oscillations in the alpha frequency band.

ABP has been shown to be modulated as a function of working memory related processing demands across a variety of tasks. In keeping with the attention literature, some have proposed that delay period alpha band power (DPABP) modulations might reflect functional inhibition oriented toward many potential alternative sources of interference (Jensen et al., 2002; Jensen & Mazaheri, 2010; Klimesch et al., 2007). In contrast, others have proposed that increases may instead reflect the active processing and/or maintenance of task-relevant information (Johnson et al., 2011; Mo et al., 2011; S. Palva & Palva, 2007; von Stein, Chiang, & König, 2000), given observations that power scales parametrically as a function of WM load in some cases (Jensen et al., 2002; Scheeringa et al., 2009; Tuladhar et al., 2007). In the sections that follow, I will first review evidence supporting the inhibitory hypothesis, followed by a consideration of evidence supporting the active processing view. I will conclude by explaining how observations taken as support for each account could be reconciled within a unitary inhibitory framework, highlighting in particular, the logic underlying our approach to providing additional evidence supporting this reinterpretation.

The Alpha Inhibition Hypothesis

In addition to the studies of attention reviewed above, results from studies of WM have likewise provided support for the proposal that alpha-band oscillations track inhibitory processes. There are several alternative proposals regarding the specific focus of this inhibitory action, but all have in common the idea of suppressing sources of potentially disruptive processing, differing only with regards to what processing is suppressed and why (Bonnefond & Jensen, 2012; Haegens et al., 2011; Jensen et al., 2002; Jensen & Mazaheri, 2010; Klimesch, Doppelmayr, Schwaiger, Auinger, & Winkler, 1999; Sauseng et al., 2009). The most general form of the

inhibition account holds that increases in DPABP within task-irrelevant regions reflect inhibition oriented toward potentially disruptive processes within those regions (Jokisch and Jensen, 2007; Jensen and Mazaheri, 2010). Jokisch and Jensen (2007) provided evidence supporting this possibility using a task that required participants to maintain either the identities of target faces or their orientations across a brief delay. They argue that face identities are more of a ventral stream process, whereas orientation is more dorsal. If ABP reflects the same inhibitory process as in visual attention tasks, they should expect to observe similar results (see Snyder & Foxe, 2010). In line with this expectation, they observed increased DPABP over task-irrelevant dorsal stream regions during the more ventral stream face identity condition. They interpreted these results as reflecting inhibition of potentially disruptive processing of face orientation by the dorsal stream during retention of face identity, a putatively ventral stream process.

Paralleling observations in studies of spatial attention, shifts in the topography of posterior ABP have also been observed when maintaining spatial locations in WM. More specifically, DPABP decreases over task-relevant retinotopic locations and increases over task-irrelevant locations (Grimault et al., 2009; Poch, Campo, & Barnes, 2014; Sauseng et al., 2009; Van Der Werf, Jensen, Fries, & Medendorp, 2008) These findings are consistent with the proposal that increases in DPABP, in both attention and WM tasks, may reflect inhibition directed toward regions engaged in task-irrelevant processes.

In stark contrast to attention tasks, however, in several cases VWM tasks have also been associated with relative increases in alpha-band activity within task-relevant sensory regions (Bonnefond & Jensen, 2012; Cooper, Croft, Dominey, Burgess, & Gruzelier, 2003; Grimault et al., 2009; Jensen et al., 2002; Jensen & Mazaheri, 2010; Johnson et al., 2011; Leiberg,

Lutzenberger, & Kaiser, 2006; Lopes da Silva, 1991). If ABP increases track inhibition, why would they be present over task-relevant cortical regions?

One possibility, suggested by Klimesch et al. (1999), is that increased ABP over task relevant regions might serve to mitigate the effects of proactive interference by suppressing the representations of stimuli that were relevant on previous trials. In this case, ABP would be observed to increase over task-relevant regions, but its function would nonetheless be inhibitory. Another possibility, is that elevated ABP over task-relevant early sensory regions may reflect inhibitory gating. This is referred to as the sensory gating hypothesis.

Generally, this theory proposes that sensory inputs normally proceeding to later cortical areas presumed to be responsible for VWM maintenance, reflect a source of potential interference with respect to target representations. As a result, alpha power increases over visual sensory areas act to suppress processing within the earliest regions, protecting the target representation stored in VWM from such interference. In addition, by this account, early visual areas may not necessarily remain engaged (i.e., relevant) following initial encoding in VWM tasks, and therefore increased DPABP over these areas in effect represents another instance of ABP increases over a task-irrelevant region.

The Active Processing Hypothesis

In contrast to the inhibitory view, proponents of the active processing view of alpha-band oscillations argue that in at least some cases, sustained increases in DPABP observed in WM tasks likely reflect a critical component of the distributed network activity underlying the selection and maintenance of objects in WM (Johnson et al., 2011; Mo et al., 2011; J. M. Palva et al., 2010; S. Palva & Palva, 2011). In particular, interregional phase interactions between

oscillations in the alpha frequency band are thought to play a role in coordinating neuronal processing within task-relevant regions across the brain. This coordination is potentially mediated by dynamic phase interactions between these regions and higher-level attentional and executive control regions.

There are many studies linking alpha-band activity with active processing during WM maintenance. For example, a number of researchers report positive cognitive or memory load dependent increases in ABP within frontoparietal regions during WM tasks (Mo et al., 2011; S. Palva, Kulashekhar, Hämäläinen, & Palva, 2011; Poch et al., 2014). In some instances, these increases in ABP are coincident with increased regional activation (Mo et al., 2011). Frontal and parietal regions are thought to underlie many central control processes in attention and VWM (Prabhakaran, Narayanan, Zhao, & Gabrieli, 2000; Rigotti et al., 2013; Rowe, Toni, Josephs, Frackowiak, & Passingham, 2000; Sakai, Rowe, & Passingham, 2002; Sreenivasan, Vytlacil, & D'Esposito, 2014; Warden & Miller, 2010). Thus, localization of load-dependent increases in DPABP in these regions has been suggested to reflect some component of these higher order control mechanisms (Palva et al., 2011). An emerging view argues that the cell assemblies representing items in WM may be coordinated via top-down inter-areal synchronization propagating from within these same regions (Jensen et al., 2002; J. M. Palva et al., 2010; Sauseng et al., 2009).

In keeping with this possibility, Palva and colleagues (2010) have shown that inter-areal synchronization between the intra-parietal sulcus and regions of the frontal and extrastriate visual cortical areas predicts WM capacity. These regions have similarly been associated with WM maintenance processes as well as many control processes (Todd & Marois, 2004; Vogel &

Machizawa, 2004; Xu & Chun, 2006). Another region commonly associated with the processing and short-term maintenance of visual object information is the inferior temporal cortex (IT). Increased DPABP within IT has been correlated with co-localized concurrent increases in multi-unit firing, and by extension, active processing (Mo et al., 2011).

Jensen and colleagues (2002) report similar parametric increases in posterior ABP as a function of VWM load over more posterior regions associated with visual processing (see also, Honkanen, Rouhinen, Wang, Palva, & Palva, 2014; Medendorp et al., 2007; Osipova et al., 2006; Tuladhar et al., 2007). Although the authors interpret these findings in light of the inhibition hypothesis, given the evidence reviewed within this section, it would seem at least equally likely that the load-dependent ABP increases observed within more posterior regions reflect active maintenance processes, similar to those associated with load-dependent increases in the aforementioned more anterior cases (i.e. Mo et al., 2011).

In keeping with the possibility that posterior load-dependent increases in DPABP reflect maintenance processes associated with active representations, there is evidence which suggests that DPABP also scales with the number of discrete task-relevant feature dimensions, even independent of the above load-dependent increases. A number of studies have observed relative increases in DPABP when participants were asked to maintain more complex VWM representations as compared to relatively simple ones (Hamidi, Slagter, Tononi, & Postle, 2009; Ikkai, Blacker, Lakshmanan, Ewen, & Courtney, 2014; Jensen et al., 2012; Johnson et al., 2011; Jokisch & Jensen, 2007; Klimesch et al., 1999; S. Palva et al., 2011; Sauseng et al., 2009). In effect, DPABP appears to increase only as a function of the relative degree of task relevant VWM processing. These results thus provide further support for the proposal that alpha-band

oscillations reflect a mechanism(s) supporting the retention of shape information and/or shape-location associations in WM, rather than, or in addition to, inhibition.

Reframing the Case for Alpha Inhibition

The above review highlights two lines of evidence that appear to support diametrically opposing roles for alpha-band activity during VWM tasks. In particular, there are several observations supporting versions of the active processing account, some, largely incompatible with most, if not all variants of the inhibition hypothesis as presented thus far. For example, evidence of ABP increases over putatively task-relevant areas (e.g., with increasing WM load), when no task-irrelevant features are present to be suppressed (see, e.g., Johnson et al., 2011), may present difficulties for a purely inhibitory account. However a pure form of the active processing view, in which ABP modulations exclusively reflect non-inhibitory attention- or maintenance-related activity, fails to adequately explain the ABP modulations observed within task-irrelevant regions. DPABP modulations observed in these cases, by contrast seem more or less exclusively consistent with an inhibitory view.

There are several possibilities for reconciling these views. For example, one approach would be to simply accept that the neural processes underlying increases in ABP play different functional roles under different circumstances. Although this remains a possibility, in the section that follows, I will articulate the alternative possibility that evidence taken as support for the active processing view is also consistent with a more general inhibitory framework. Specifically, I will attempt to address this discrepancy by reframing observations taken as support for the active processing view in terms of one of two alternative mechanisms of suppression. This framework could account for ABP increases over both task-relevant and irrelevant-regions alike.

The more general inhibition view that has been the focus of this review thus far holds that ABP increases reflect the suppression of potentially disruptive activity in cortical areas representing task-irrelevant information. Specific examples include the suppression of auditory cortex when attending to visual information in intersensory-attention tasks, and the suppression of dorsal stream areas processing item motion or orientation while attending to or remembering ventral stream features, such as color or face identity. In each of these cases, ABP-related suppression was directed towards sensory modalities or feature dimensions presumably involved in the processing of task-irrelevant aspects of the attended or remembered objects. As noted above, however, this view has difficulty accounting for cases in which ABP-modulations are observed either a) over putatively task-relevant areas (e.g., over occipital areas during the maintenance of VWM items), or b) increases appear to track the number of task-relevant items or features (e.g. scaling with VWM load increases).

In the context of WM, the first of these observations could be reconciled with an inhibitory account if, rather than being directed towards specific task-irrelevant properties of remembered objects, increased DPABP in some cases reflects a mechanism of sensory gating, aimed at protecting the contents of WM from potentially disruptive on-going sensory processing (i.e. bottom-up interference). For example, load-dependent DPABP increases observed over posterior visual areas in some studies of WM could reflect the increased gating of task-irrelevant sensory inputs to later cortical areas during maintenance, rather than maintenance-related activity per se. This would presumably be particularly important at higher loads, or in cases where memorized stimuli consist of more complex objects. As I point out above, this is the interpretation proposed by the sensory gating hypothesis, which accounts for load-dependent

increases observed over putatively task-relevant early visual areas in a manner consistent with the more general inhibitory view. Specifically, by this view, early sensory regions are recast as transitioning to a state of task-irrelevance during VWM delay periods. In this case, the load-dependent increases and associated suppression are by extension directed toward task-irrelevant processes. This possibility is also consistent with the observed negative relationship between ABP and regional excitability observed more generally across regions (Romei et al., 2008, 2011; Thut et al., 2011). Whether load-dependent increases in ABP over posterior visual areas specifically reflect a mechanism of sensory gating via modulations of visual cortex excitability will be examined in Experiment 1.

Although sensory gating recasts suppressed visual sensory regions as task-irrelevant, thus aligning such increases with the more general inhibitory account, DPABP has also been observed within other, task-relevant regions as well. The proposed inhibitory account could accommodate DPABP observed over task-relevant visual areas, if such increases were to reflect the suppression of specific features along a task-relevant dimension that happen to be task-irrelevant on a given trial. For example, in a typical WM task, on a given trial, an observer may be asked to remember only one or several values along a target dimension (e.g., 1-3 oriented bars), with all of the other possible values along that dimension being task irrelevant. In such cases, it may be advantageous for individuals to inhibit the currently-task-irrelevant features, while reducing inhibition over just those specific cortical sites representing the currently relevant feature(s). Such a mechanism of inhibition would result in increased DPABP over cortical regions sensitive to the remembered dimension(s) (i.e., over task-relevant areas), as has been observed in numerous studies.

One instance in which this more selective form of within-dimension suppression appears to occur is in spatial WM and attention tasks. Recall that, in these cases, increases in ABP are highly selective, appearing exclusively over early visual areas representing spatial locations that are not remembered on that trial, with an associated decrease over task-relevant regions. This relationship is so robust, in fact, that researchers are able to identify the specific target locations being maintained on any given trial based on the scalp topography of ABP alone (Meyer et al. 2015). Whether similar within-dimension feature suppression is also observed in non-spatial WM tasks is currently an open question.

One challenge in providing evidence for this possibility comes from the fact that the organization of feature-selective populations of neurons in the cortex occurs at a much finer scale than does spatial coding. For example, the spatial locations of objects are represented in a highly regular fashion across the cortex, with targets appearing on the left side of the screen being preferentially represented by cells in the contralateral (i.e., right) portion of the occipital lobes, and vice versa. Similarly, stimuli appearing in the upper visual field are represented by cells in more ventral portions of the occipital lobes, and vice versa. As a result, it can be relatively straightforward to determine from the pattern of ABP at posterior EEG sensors whether a remembered stimulus is on the left, right, upper or lower portion of the screen. Determining which of several possible orientations is being remembered on a given trial, however, could require more sensitive measures. This question, and the more specific question of whether within-dimension feature suppression is observed in non-spatial WM tasks, will be addressed in Experiment 2 using pattern classification techniques.

Conclusion

Evidence reviewed in the preceding sections suggests that ABP modulations observed in studies of attention most likely reflect the operation of an inhibitory mechanism that serves to suppress the processing of task-irrelevant visual information. Although similar modulations have also been observed in the context of WM, the specific functional role of such increases continues to be a matter of debate. On one hand, it has been argued that ABP increases observed in WM tasks reflect the operation of an inhibitory mechanism similar to that engaged during studies of attention. On the other hand, it has been argued that such increases may instead reflect the active processing or maintenance of task-relevant information in WM. In contrast to the latter view, I have argued that alpha-band activity observed in WM tasks may reflect an inhibitory mechanism, selective against non-target processes within task relevant and irrelevant regions alike. In each case, increased ABP serves to support task-relevant processing by inhibiting potentially disruptive processes that are not currently relevant to the task. Thus far, however, there is little evidence directly linking DPABP to cortical excitability in VWM, or to the suppression of non-target features within a task-relevant dimension. As such, an important first step toward differentiating the relative viability of the opposing interpretations, requires establishing a clear link between ABP and the distinct patterns of inhibition proposed in both the gating and within-dimension feature suppression cases of the proposed inhibition account.

To examine whether ongoing visual processing is modulated by VWM load-dependent increases in ABP, Experiment 1 will use a change detection task in which WM load is varied, and irrelevant probe stimuli are presented during the delay interval on a portion of trials. If load-dependent increases in DPABP reflect the progressive suppression of visual cortex, one would

expect to observe a correlation between posterior load-dependent increases in DPABP and the magnitude of the neural response elicited by the probe stimulus. To provide evidence for the within-dimension feature suppression account, participants in Experiment 2 will perform a precision delayed-estimation VWM task, in which the orientation of a single Gabor patch will be maintained in WM, and distractor stimuli varying in terms of orientation, relative to the remembered target, will be presented mid-way through the delay. If item identity can be decoded from distributed patterns of DPABP, and if the observed pattern enabling this decoding reflects suppression of non-target orientation features; then neural responses to distractors should be suppressed. Moreover, the degree of suppression should be correlated with the relative degree to which this pattern is tuned selectively to any given target item.

EXPERIMENT 1: ALPHA-BAND POWER AND CORTICAL EXCITABILITY

Studies exploring the role of neural oscillations in cognition have revealed sustained increases in ABP during the delay period of verbal and visual WM tasks which scale with increased cognitive load. There have been various proposals regarding the functional significance of such increases, including the inhibition of task-irrelevant processes, as well as the active retention of information in WM. In the present study, I explore these alternatives by examining the role of DPABP in mediating the effects of distractor processing during a VWM task.

To this end, EEG was recorded while subjects performed a change detection task requiring the retention of two or four novel shapes. Importantly, on a portion of trials, a task-irrelevant bilateral checkerboard probe was presented during the delay interval. If load-dependent increases in DPABP reflect an inhibitory mechanism gating ongoing visual processing, then they should be correlated with (1) modulations of the electrophysiological response evoked by the delay period probe and (2) the number of additional items individuals are able to retain in the probe condition.

Accordingly, analyses focused first on examining the relationship between set size dependent increases in DPABP and changes in the magnitude of the global response evoked by the probe (p-ER). Further, in attempt to provide evidence for the behavioral relevance of this inhibitory process, I explore the relationship between these increases and changes in VWM capacity across set sizes separately in each condition. Finally, oscillatory phase-based connectivity analysis was used in order to examine the possible source of ABP-based effects on distractor processing and behavior.

Method

Participants

Twenty-eight participants between the ages of 18 and 35 were recruited from the North Dakota State University student population. All participants reported normal or corrected to normal visual acuity, and provided written informed consent prior to participation. Participants received either course credit or monetary compensation ($10/hr) for their participation.

Stimuli and Procedure

During the experiment, participants were seated in a dimly lit, noise-controlled room. Stimulus presentation and response collection was controlled by a PC running Presentation software (Neurobehavioral Systems, Inc.). Stimuli were presented against a light grey background (RGB=[125,125,125]) on the surface of a 19" cathode ray tube monitor with a refresh rate of 100 Hz, at a viewing distance of 70 cm.

Participants were asked to perform a change detection task (Figure 1). Specifically, they were instructed to maintain either two or four target shapes in memory across a 1400-ms delay. Each trial began with the appearance of a centrally presented fixation cross. Participants were asked to maintain focus on this fixation cross for the duration of each trial. 300-500 ms after fixation onset, the memory display appeared. Target items were presented bilaterally around the fixation point, in one of eight possible locations; four per side. The memory display remained visible for 500 ms. Memory target items consisted of black abstract shapes (Attneave & Arnoult, 1956), drawn at random from among twelve such shapes.

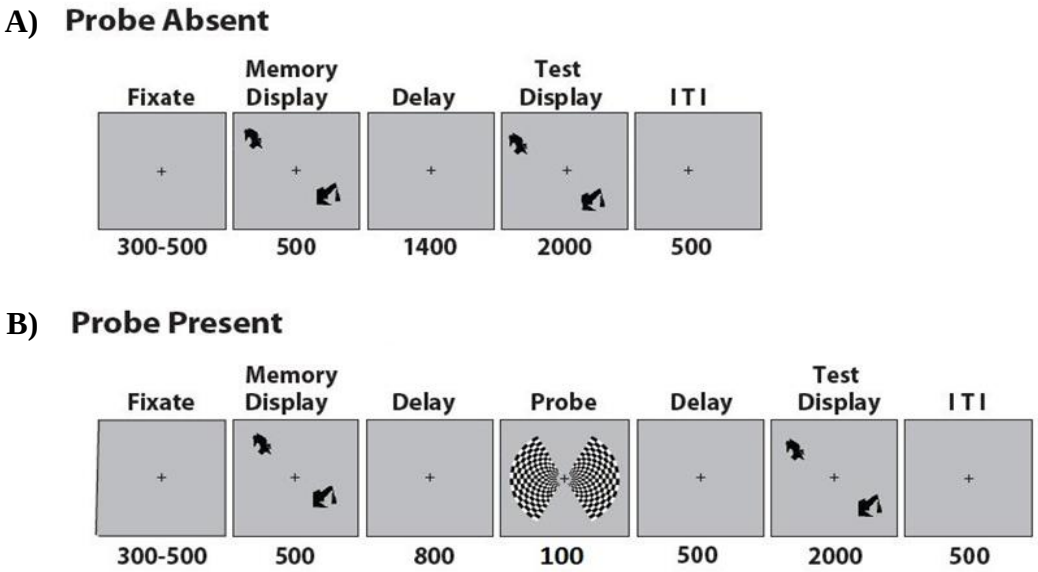

Figure 1. Experiment 1 Design. Experiment 1 (A) Probe absent condition and (B) Probe present condition.

Randomly, on two thirds of trials, a bilateral checkerboard probe stimulus was presented 800 ms into the delay period, for 100 ms. Participants were instructed to ignore this stimulus. The squares within the checkerboard alternated between white and light grey (RGB=[125,125,125]).

Following the delay period, participants reported whether or not any of the items within the test display had changed shape relative to those in the memory display. When a change occurred, one of the shapes was replaced by a different shape not present in the original memory display. Additionally, to constrain the magnitude of shape change across trials, the twelve shapes were grouped into four subsets of three items, based on experimenter estimates of subjective similarity (e.g. more elongated vs. more rectangular). This was done in order to minimize variability in the apparent magnitude of shape changes from one trial to the next. Participants responded using one of two response buttons. If none of the target items had changed, they made

a “no-change” response by pressing the button in their left hand. If there was a change, they made a “change” response by pressing the button in their right hand. The test display remained visible for 2 seconds or until a response was entered. Feedback was given on every trial and remained visible for 500 ms. For correct responses, the fixation cross briefly changed to a bold font. In contrast, when either an incorrect response was given or two seconds had passed with no response, the fixation cross changed to a minus sign. Participants completed a total of 36 trials per block and 20 blocks in all. This amounts to 480 probe trials and 240 no-probe trials.

To estimate participants’ WM capacity, a standard formula described by Cowan (2001) was used. Specifically, for each set-size (2, 4) and probe (present, absent) condition, WM capacity (K) was calculated using the formula K = set size × (hit rate – false alarm rate). K is a common metric utilized within the VWM literature that allows for comparisons across various set sizes (Fukuda et al, 2009).

EEG acquisition and preprocessing

EEG was recorded using active Ag/AgCl electrodes (BioSemi Active Two) positioned at the left and right mastoids and 64 scalp sites, according to the modified international 10-20 system (American Electroencephalographic Society, 1994). To detect eye movements and blinks, the electrooculogram (EOG) was recorded from electrodes placed at the outer canthi of each eye, as well as above and below each eye. All signals were recorded with a band-pass of .01-100 Hz and a sampling rate of 512 Hz.

Data was processed offline using the EEGLab (Delorme & Makeig, 2004) and ERPLab (Lopez-Calderon & Luck, 2014) open source Matlab-based toolboxes, and custom analysis scripts written in Matlab (MathWorks, Inc., Natic, MA). The EEG and EOG signals were band-

pass filtered using a non-causal Butterworth infinite impulse response filter with half-amplitude cut-offs of .1 and 30 Hz (12 dB/octave), and re-referenced to the average of the mastoid electrodes. The data were epoched into 2400-ms segments spanning the interval 500 ms before to 1900 ms after memory display onset (100 ms after test display onset).

The EOG signals were referenced into bipolar vertical and horizontal derivations and used in the detection of eyeblinks and saccades. Trials were automatically excluded if EEG amplitude exceeded 100 μV in any channel (or 70 μV in the vertical EOG channel) within a moving window of 200 ms. Trials were also rejected if a step function detected changes of more than 25 μV on the horizontal EOG channel, indicating the presence of lateral eye movements.

EEG analyses

Time-frequency analysis of power. Time-frequency decomposition of EEG signals was performed using the Fieldtrip software package (Oostenveld et al., 2010), an open-source Matlab-based toolbox for the analysis of electrophysiological data, and custom scripts written in Matlab. Time frequency representations (TFRs) were obtained by convolving stimulus-locked single-trial data from all electrodes with complex Morlet wavelets: $e^{i2\pi ft}e^{-t^2}/(2\sigma^2)$, where *t* is time, *f* is frequency, which varied from 4-40 Hz in 37 logarithmically spaced steps, and σ is the width of each frequency band, defined as n/(2πf), where n is the number of wavelet cycles, which varied from 3 to 6 in logarithmically spaced steps to obtain comparable frequency precision at low and high frequencies. TFRs of power were estimated by squaring the complex convolution signal Z (*power* = real[z(t)]^2 + imag[z(t)]^2) and averaging across trials. Data was baseline corrected by subtracting mean frequency specific power from the -300 to -100ms pre-stimulus interval from the power values obtained at each time-frequency point.

Probe-evoked response analyses. To quantify the probe-evoked response (p-ER), artifact free trial epochs created during preprocessing were used to compute averaged ERP waveforms within a narrower time window. Specifically, for these analyses, the size of each trial epoch was reduced to encompass a 600-ms time window, ranging from 100 ms before to 500 ms after probe onset. Individual trials were baseline corrected using the 100 ms pre-probe interval. Time-domain analysis focused on comparing the mean amplitude of specific components of the p-ER, including the posterior P1 (75-125 ms) and N1 (125-175 ms) components, as well as a later anterior positive component (LAP) observed over fronto-central electrode sites during the 200-400 ms latency range. Specifically, for each load condition (2 or 4), mean amplitude was calculated at a subset of posterior (PO7/PO8) and frontal (FC1/FC2) electrodes for the P1/N1 and LAP respectively. To avoid biasing the results by using differences between conditions to select the time windows and electrodes of interest, selection was made on the basis of visual inspection of the p-ER waveforms collapsed across conditions. This revealed maximal p-ERs centered at PO7/O8, in the case of the P1/N1, and FC1/FC2, in the case of the later positive-going component.

Phase-based connectivity analyses. To examine load-specific differences in oscillatory coupling between frontal and posterior brain areas, I computed a measure of phase coherence known as inter-site phase clustering (ISPC). ISPC reflects the degree to which the phase angle differences between a given pair of electrodes are clustered in polar space, at a specific time point, across trials. This specific measure was chosen for two reasons. First, it provides relatively strong evidence for task-related modulations in connectivity of the kind I expected to observe across load conditions. Second, this method allows for a relatively high degree of temporal

precision as compared to other similar methods. The analysis was conducted using custom scripts written in Matlab, following the procedure described in Cohen (2014).

Prior to ISPC analysis, a surface Laplacian was applied to EEG data epochs (Perrin et al., 1989). The surface Laplacian is a spatial bandpass filter that attenuates low spatial frequencies, which helps to minimize spurious connectivity effects arising from volume conductance, and renders the data more appropriate for electrode-level connectivity analysis (Cohen, 2014).

Following application of the surface Laplacian, single-trial EEG epochs were decomposed into their constituent TFRs by convolving them with a set of Morlet wavelets, as described above, with frequencies ranging from 8 to 30 Hz, in 14 logarithmically spaced steps, and number of wavelet cycles varying from 4 to 8 in logarithmically spaced steps. Frequency-band specific ISPC was computed using the phase angle, φ_t = arctan(imag[z(t)]/real[z(t)]), of the complex convolution result. ISPC is defined as the trial-averaged phase angle difference between two electrodes j and k at each time-frequency point:

$$\left|\frac{1}{n}\sum_{t=1}^{n} e^{i(\varphi jt-\varphi kt)}\right|,$$

Where n is the trial count. The frequency-specific average of ISPC values over the –300 to –100 ms interval prior to the memory display was used for baseline correction.

Statistical Analyses

As I discuss above, previous research has revealed increases in DPABP as a function of WM load. These increases potentially reflect the engagement of an inhibitory mechanism that serves to suppress potentially distracting neural representations. This information could include not-currently-relevant sensory information (e.g., auditory information presented during the performance of a visual task), stimulus features (e.g., the orientation of a remembered face), or

spatial locations. However, it is unclear why DPABP would increase with load in the absence of specific task-irrelevant information that needs to be suppressed (e.g., during the unfilled delay interval of a WM task).

To reconcile this finding with the inhibitory view, I have proposed that, in addition to suppressing potentially disruptive representations, increased DPABP may serve a more general function, gating inputs from earlier visual areas to higher order regions engaged in WM maintenance. Such inputs could be due to the appearance of stimuli in the task space, or simply a as a consequence of neural noise. To provide support for this possibility, I conducted the following analyses investigating the relationship between ABP, ERP measures and behavior.

Behavioral data

To examine load- and probe-related changes in capacity (K) I conducted a two-way repeated measures ANOVA with factors of set size (2, 4) and probe (present, absent). Additionally, to examine the relationship between observed changes in capacity across conditions and changes in spectral properties of the EEG (e.g., DPABP), correlation analysis was conducted using Pearson's r (described further below).

EEG data

A critical assumption of the present study is that DPABP will increase as a function of WM load (set size). Thus, the first step in the EEG analysis was to determine whether this was in fact the case. To do this, estimates of DPABP, obtained using the wavelet method described above, were compared across load conditions. Next, load-dependent changes in the p-ER were assessed via three separate paired sample t-tests comparing the mean amplitude of the P1, N1, and LAP across load conditions.

To examine the relationship between DPABP and the p-ER, correlation analysis (Pearson's r) was used to determine whether observed load-dependent differences in DPABP were predictive of load-dependent changes in the p-ER. If increases in DPABP reflect a VWM maintenance process, as some have proposed, we would not expect to observe a correlation between increases in DPABP and suppression of specific p-ER components. If oscillatory activity in the alpha frequency band reflects a form of gating, however, I would expect load-dependent increases to be associated with a reduction in one or more components of the p-ER.

More specifically, if DPABP reflects the inhibition of earlier sensory processing, as the gating hypothesis proposes, one would expect to observe a correlation between these increases and suppression of earlier ERP components (P1, N1) as well as potentially the LAP. An alternative possibility, is that these increases might reflect the gating of visual processing at a later point in the visual hierarchy. For example, increases might reflect the inhibition of visual features or specific objects stored beyond the primary sensory cortices, which represent potential sources of disruption with respect to the specific target items being maintained on any given trial. If this were the case, I would expect load-dependent changes in DPABP to be correlated with suppression of only the later ERP component (LAP).

In addition to these hypotheses, I further expected load-dependent changes in DPABP to be related in systematic ways to load- and probe-related changes in WM capacity, as estimated by K. If load-dependent changes in DPABP reflect an increase in neural activity related to the active maintenance of information in WM, these changes should be related to increases in capacity estimates across set sizes observed in the No Probe condition. To examine this possibility, for each participant, I first calculated the change in capacity as a function of set size

($K_{SS4} - K_{SS2}$), and used these change values to sort participants into separate High and Low capacity groups. I then performed an independent samples t-test to assess differences in observed load-dependent changes in DPABP across these groups.

Mirroring the above p-ER analysis, I additionally wanted to investigate the alternative possibility that set-size dependent increases in DPABP reflect inhibition gating either relatively early or later on-going visual processing. If increases in DPABP reflect a more general gating mechanism, suppressing early sensory processing as the gating hypothesis proposes, then I would expect increases to be correlated with the difference in capacity between the probe and no-probe condition at SS4, for which significant probe-disruption effects were observed. To get at this, I conducted a median split based on the differences in K observed across probe conditions at SS4. I then performed an independent samples t-test comparing DPABP during the pre-probe interval for individuals showing a relatively small decrease in capacity across probe-conditions against those with comparatively larger decreases. Finally, I conducted a third analysis to investigate the possibility, discussed above with respect to the p-ER, that increases in DPABP reflect a process insulating additional WM items from distraction within a putatively later visual region. If this were the case, I would expect those individuals able to retain relatively more items at SS4 (SS4-SS2) in the probe condition to exhibit a larger increase in DPABP across set sizes. To get at this, I again performed a median split, sorting participants into separate High and Low K groups based on observed load-dependent changes in K, this time in the Probe condition alone. I then performed an independent samples t-test comparing DPABP during the pre-probe interval for individuals showing a relatively small increase in capacity across set-sizes against those with comparatively larger increases.

Finally, to assess whether observed load-dependent increases in DPABP were a consequence of an attentionally selective frontal control process, Pearson's r was calculated to assess the correlation between set-size dependent increases in ISPC and similar modulations of DPABP. If set-size dependent increases in DPABP are the consequence of an executive control mechanism, I expected to observe a positive correlation between increases in DPABP and ISPC prior to the probe. Additional exploratory connectivity analyses were also conducted to further understand the dynamic inter-regional interactions underlying DPABP.

Results

Behavioral results

The behavioral results, highlighted in Figure 2, imply that on average participants were able to maintain the same number of items at SS2 in both the probe and no-probe conditions. At SS4, however, participants exhibit an increase in VWM capacity in the no-probe condition, whereas no such increase is observed in the probe condition. This would appear to suggest that participants were holding more information in VWM at SS4, as we would expect, but maintenance may be disrupted by the probe stimulus at higher loads, resulting in lower capacity estimates. A two-way analysis of variance exploring these effects yielded a main effect of probe condition, $F(1, 17) = 11.30$, $p < .05$, such that the average K was significantly higher in the no-probe condition ($M = 1.61$, $SD = .37$) than for probe condition ($M = 1.49$ $SD = .34$). The main effect of set-size approached but did not achieve significance, $F(1, 17) = 4.21$, $p=.056$, as did the interaction effect, $F(1, 17) = 4.01$, $p =.062$.

Given the results of the ANOVA, I conducted follow up t-tests comparing each specific load- and probe-related change in capacity (K) across conditions. As shown in Figure 2, a t-test

revealed that on average subjects did not experience a significant disruption as a result of the probe at SS2 $t(17) = -1.257$, $p = .112$. At SS4, however, the probe significantly disrupted performance on average $t(17) = -2.92$, $p = .004$. This suggests that participants were generally more susceptible to interference at higher set-sizes as the figure shows. Further, participants were able to maintain significantly more target items in the SS4 as compared to SS2 condition $t(17) = 2.42$, $p = 0.019$, confirming that individuals were actually exhibiting a significant change in VWM load on average across set-sizes when no probe was present. By contrast, participants did not retain significantly more items at SS2 as compared to SS4 in the probe condition $t(17) = .302$, $p = 0.383$.

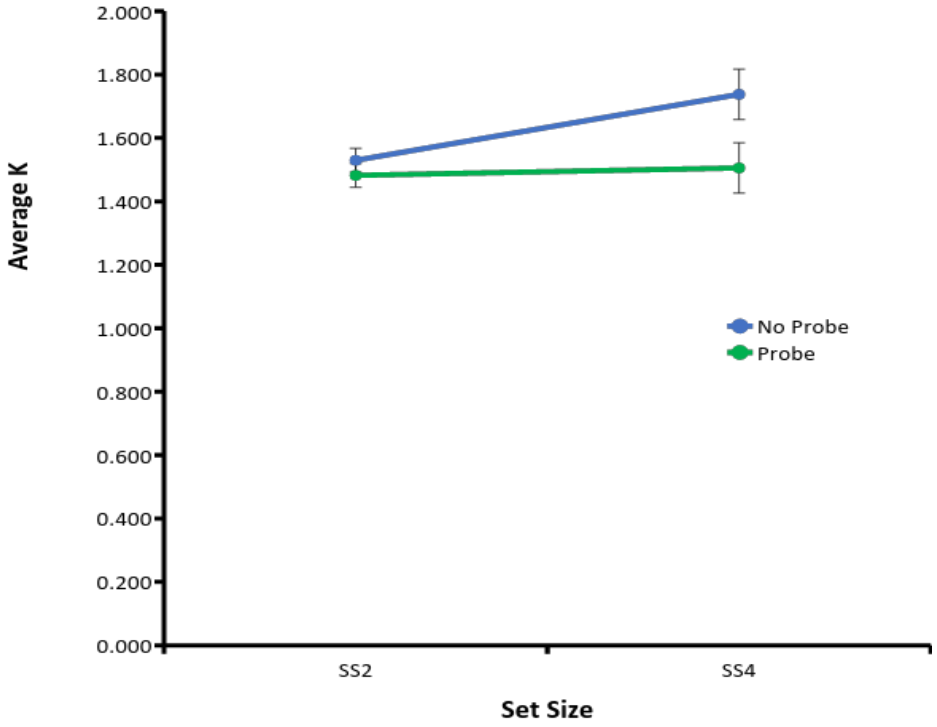

Figure 2. Working Memory Capacity (K) Differences across Probe Conditions. Error bars reflect the standard error of the mean.

EEG results

Delay-period alpha-band power differences A t-test confirms that during the 100ms prior to the probe onset there was a significant increase in DPABP across set sizes $t(17) = 3.38$, $p = .004$ at electrodes POZ and PZ. Looking at Figure 3A and 3B, we can clearly see this increase in power across set sizes in the probe condition. This increase is primarily focused within the alpha and beta frequency bands, beginning just prior to the onset of the probe. Because I was interested in the effect of DPABP on biasing ongoing visual processing, I focused on the average power in the 8-12 Hz range during the 100ms interval proceeding probe onset. Figure 3B reflects the topographic distribution of power averaged across this time and frequency range, here we see that DPABP is greatest at these two posterior electrodes specifically and increasing across loads. Figure 3C reflects the change in ABP at these two electrodes across time across all probe conditions separately. Here, following memory display offset we observe an initial desynchronization, average power then gradually increases across conditions. At around 1000ms the lines begin to diverge and we observe a load dependent difference in DPABP in both the probe and no probe condition (solid versus dashed lines, respectively). In both cases, the increase in power peaks roughly coincident with the onset of the expected probe stimulus. In the no-probe condition, DPABP then appears to trail off gradually over the remainder of the delay interval. This pattern seems consistent with what one might expect if these increases were specifically related to modulating probe processing, as I propose. By comparison, in the probe condition, DPABP drops off sharply following the probe. The sharp decrease in this latter case may be the result of the probe disrupting the VWM processes underlying the increases, or the process may be actively discontinued in both conditions as a probe is no longer expected to appear.

Probe-evoked response differences. Figure 4B shows the difference in the scalp topography of the evoked response comparing SS2 and SS4. Both bilateral pairs of electrodes [FC1, FC2 and PO7, PO8] are marked in black. A comparison of the amplitude of the evoked-response across set sizes revealed a significant set-size dependent change in the P1, $t(17) = 3.24$, $p = .004$, and N1, $t(17) = 6.43$, $p < .001$ (Figure4C). Similarly, analysis revealed a significant difference in the amplitude of the LAP component across set sizes $t\ (17) = 2.54$, $p = 0.021$ (Figure4A).

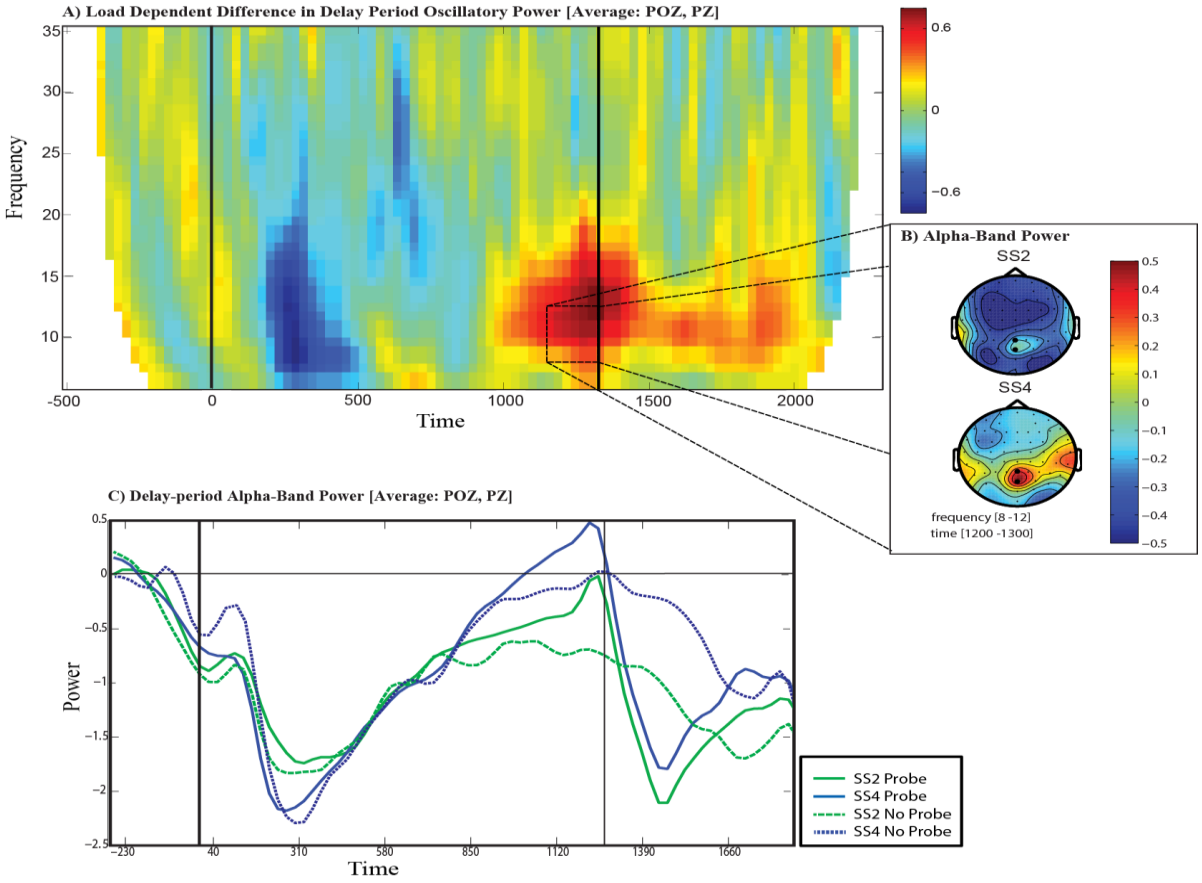

Figure 3. Delay Period Alpha Band Power across Conditions. (A) Time frequency plot of the load dependent change in DPABP. The boxed region reflects the 100ms prior to probe onset (B) Topographic Distribution of DPABP at set size 2 and set size 4. (C) DPABP across time.

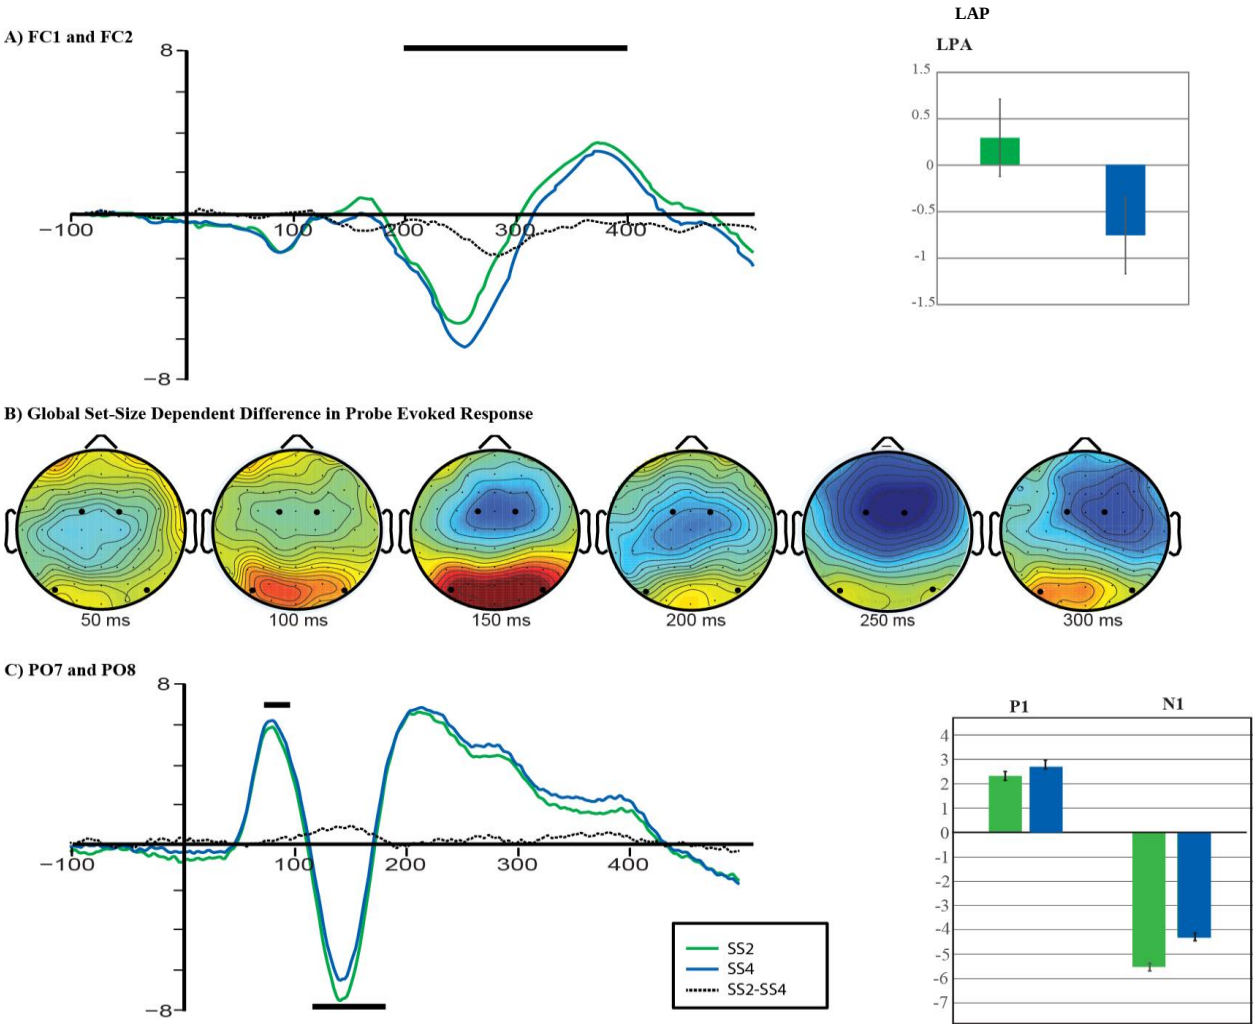

Figure 4. Evoked Response Differences across Set-Sizes. (A) LAP wave form set size two and four. The black line reflects the interval over which activity was averaged for the bar plots (B) Difference in the probe evoked response across electrodes [SS2-SS4] (C) P1 and N1 wave form across set sizes. As before, the black line reflects the interval over which activity was averaged for the bar plots

Relationship between delay-period alpha-band power and p-ER. A Pearson's r confirms that set-size dependent increases in DPABP are correlated with changes in the amplitude of the N1 evoked response across set-sizes $r(16) = -0.567$, $p = .014$ (Figure 5B) but not the P1 $r(16) = -.088$, $p = .719$ (Figure 5A). Despite the fact that individuals on average show a significant

reduction in the N1 amplitude, the correlation in this case appears to be opposite from what I predicted. Larger increases in DPABP across set-sizes were instead correlated with a smaller change in the amplitude of the N1. The LAP was also significantly correlated with the set-size dependent increase in DPABP $r(16) = -0.507$, $p = .030$ (Figure 5C) In this case, however, this correlation matched our prediction. Larger increases in DPABP across set-sizes were correlated with larger reductions in the amplitude of the LAP at SS4. The latency of this modulation may suggest that DPABP reflects the inhibition of ongoing visual processing at later rather than earlier stages.

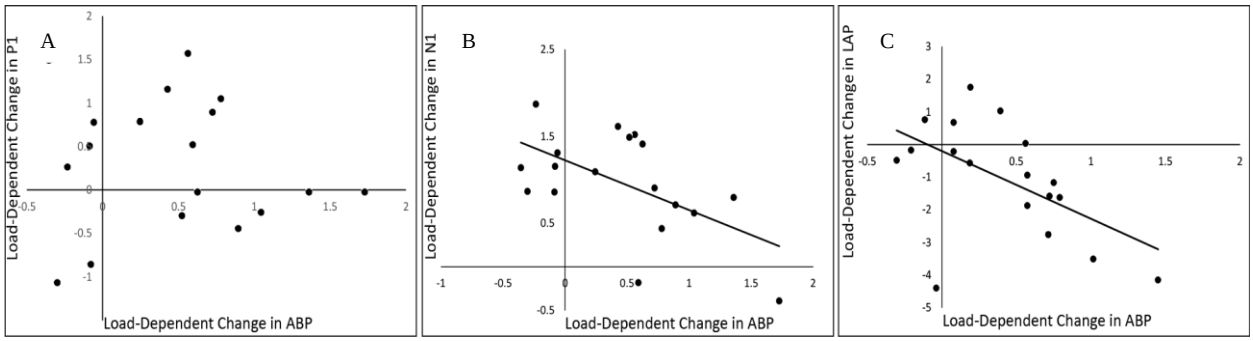

Figure 5. Set-size Dependent change in ABP and ERP Amplitudes. Relationship between load-dependent change in DPABP and load-dependent change in amplitude of P1 (A), N1 (B), and LAP (C) components of the probe-evoked response.

Relationship between delay-period alpha-band power and changes in K. Three separate tests were conducted to further clarify the relationship between DPABP and maintenance versus inhibitory processes. The first contrast tested the hypothesis, derived from the active processing view of alpha, that those individuals exhibiting the largest change in K with load, would also exhibit the largest change in DPABP. Although DPABP was slightly larger for individuals exhibiting a large change in K (see Figure 6A), this difference was not significant, $t(17) = -.753$,

p = .478. This result suggests that DPABP is not directly tracking the amount of information that is stored in WM, contrary to the active processing view.

The second contrast tested the hypothesis, derived from the sensory gating view, that individuals showing the smallest probe-interference effect at SS4 ($IE=K_{no\ probe} - K_{probe}$), would exhibit greater change in DPABP across set sizes. As can be seen in Figure 6B, there was a small tendency in the opposite direction; individuals exhibiting the largest interference effect had slightly higher DPABP. However, this difference was also not significant, $t(17) = -0.96$, p =.448. This result suggests that DPABP is also not directly tracking a general gating process working at an early level of processing to prevent interference with items held in WM.

The third and final contrast tested a further hypothesis derived from the sensory gating view: that the set-size dependent increase in DPABP should be greater for those individuals who are able to store more versus less information in the Probe condition at SS4 versus SS2. As can be seen in Figure 6C, DPABP was considerably higher for those subjects exhibiting a large versus a small increase in K from SS2 to SS4 ($K_{ss4} - K_{ss2}$)—i.e., those who were able to store additional items in spite of the probe presentation. Confirming this, a t-test revealed significantly greater DPABP for those individuals showing the greatest increase in the number of items that were successfully stored in the SS4 versus SS2 conditions with a probe present, $t(17) = 2.37$, p = .03. This finding suggests a more specific form of gating, in which DPABP reflects the selective insulation of particular items from probe-related disruption, which, I propose, is likely to occur at later stages of visual processing, potentially reflected in the LAP.

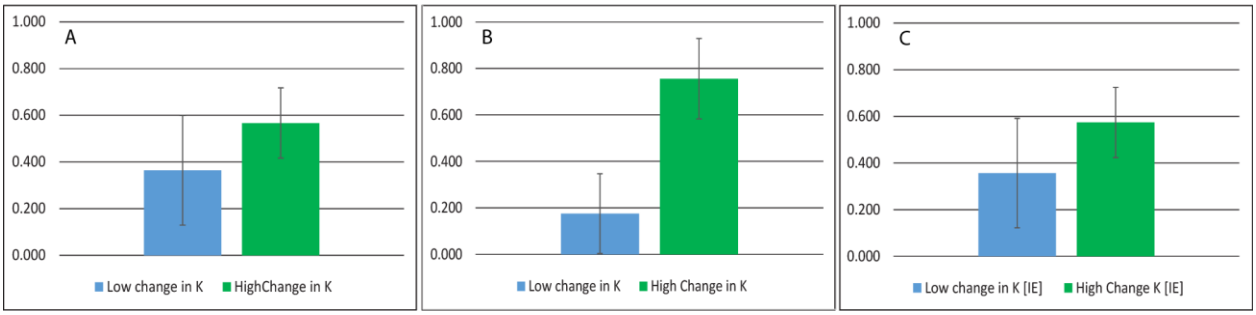

Figure 6. Median Split Analysis. Difference in DPABP for subjects with low versus high K change across set sizes in the no-probe (A), and probe (B) conditions, and across probe conditions at SS4 (C).

Inter-site phase clustering. Looking at Figure 7, we see that ISPC in the probe condition exhibits a pattern very similar to ABP across the same trials (see Figure 3C), as we would expect if ISPC was reflecting a process underlying these increases. In this case, however, the set-size dependent difference emerges far earlier, increasing still further later in the delay, at around the same time point as the initial difference in DPABP is observed. Further mirroring the observed modulations in DPABP, this increase appears to peak around the expected onset of the probe stimulus, falling off abruptly thereafter. Though ISPC estimates become highly unstable with lower trial numbers, it is worth noting that, in the no probe condition, ISPC again exhibits an analogous time course. Following the expected onset of the probe, ISPC in the no probe condition remains stable, trailing off gradually later in the delay, as observed for DPABP across these same trials.

Looking at the same 100ms interval prior to probe onset in the probe conditions, a t-test confirmed that ISPC increased significantly from SS2 to SS4, $t(17) = 2.91$, $p = .004$. Pearson's r comparing the change in ISPC with the change in DPABP across this pre-probe interval provides further evidence supporting the possibility that the increases are related. The set-size dependent increase in DPABP was correlated at an individual subject level with the set-size dependent

increase in ISPC during the same time interval $r(16) = .527$, $p=.024$ (Figure 8). Further, Pearson's r suggests the presence of a significant correlation between disrupted ISPC following the probe and the probe interference effect observed at SS4, $r(16) = .637$, $p=.004$. Interestingly, the disruption in ISPC following the probe is also correlated with observed modulations of both the N1, $r(16) = -.682$, $p=.001$, and LAP, $r(16) = -.647$, $p=.003$, components of the p-ER.

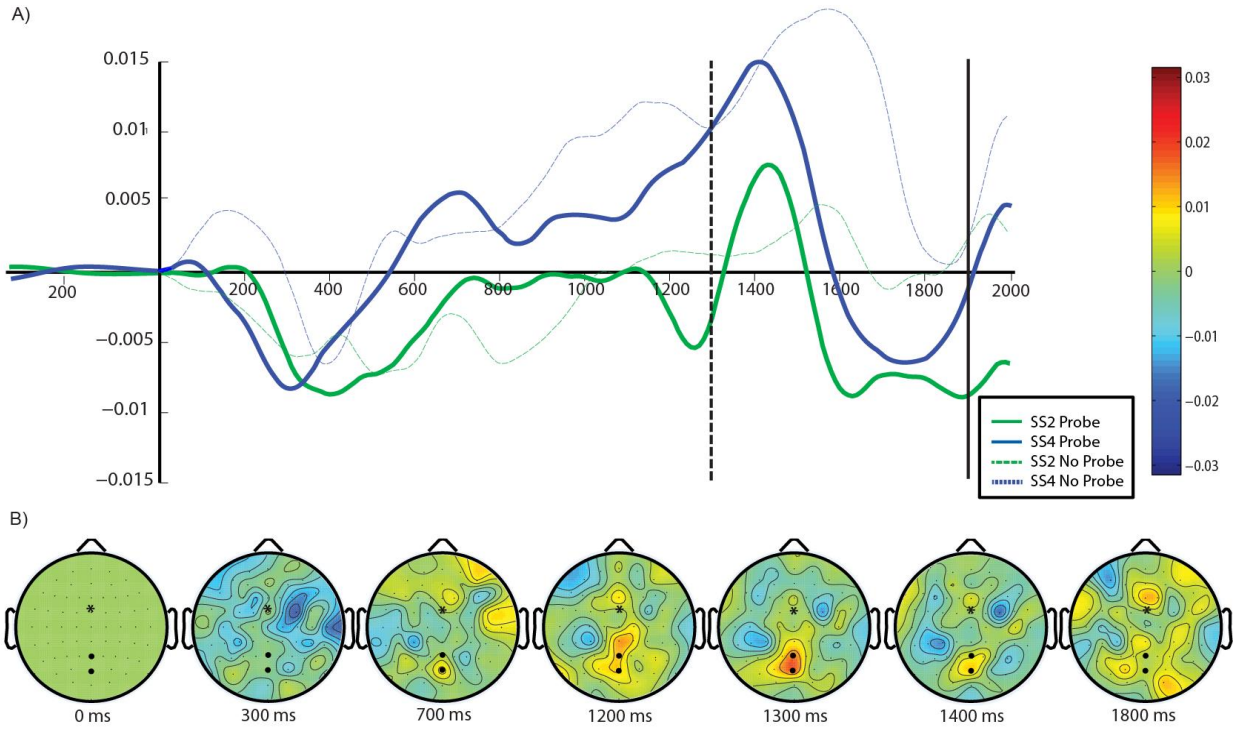

Figure 7. Inter-site Phase Clustering. (A) Inter-Site phase clustering averaged between FC1 and POZ/PZ. The vertical axis reflects memory display onset. The horizontal axis reflects the degree of phase clustering in polar space. The dotted line reflects probe onset. (B) The difference across set sizes between the topographic distribution of ISPC between FC1 and all

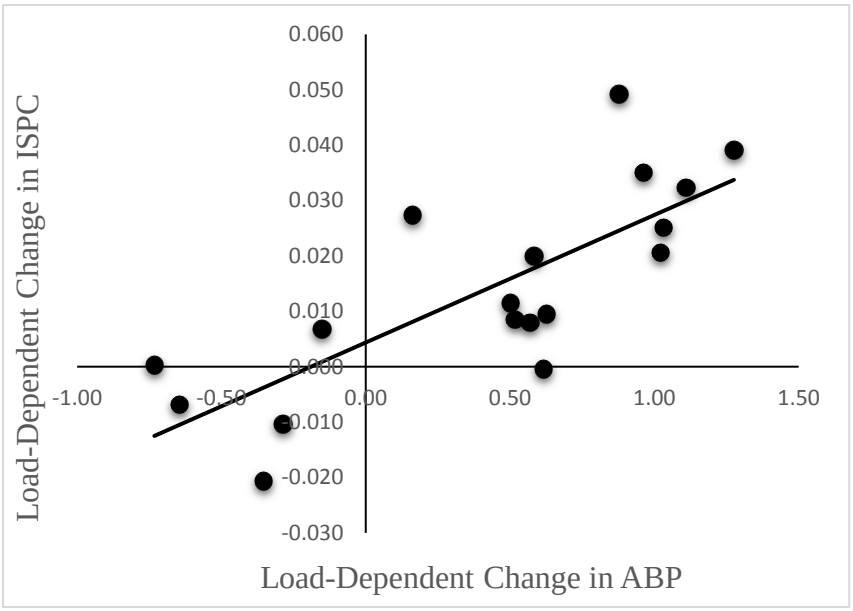

Figure 8. Set-Size Dependent change in ABP and ISPC Magnitude (prior to probe onset)

Discussion

There have been various proposals regarding the functional significance of increases in ABP during the delay period of verbal and visual WM tasks which scale with increased cognitive load, including the inhibition of task-irrelevant processes, as well as the active retention of information in WM. In the present study, I explored these alternatives by examining the role of DPABP in mediating the effects of distractor processing during a VWM task. The evidence suggests that, probe processing disrupts target maintenance to the degree that DPABP is unsuccessful in filtering visual inputs on the basis of goal-dependent patterns of suppression. This gating appears to scale with VWM load, but there is reason to think that the increases reflect the progressive inhibition of visual representations inconsistent with maintained target items.

The results further suggest that selection and the resulting inhibition of task-irrelevant processing, is potentially implemented via executive top-down control. This mechanism appears

to enact this protective measure predictively within the several hundred millisecond interval surrounding the expected probe, with a peak centered roughly coincident with onset.

Disrupting maintenance related activity, reflected by this connectivity, is correlated with the degree of disruption observed with respect to VWM performance on this task. In line with a gating interpretation, this disruption may arise as a function of the degree to which the probe processing is insufficiently suppressed and thus allowed to persist beyond the earliest visual regions. As I proposed it is likely through this gating that neural oscillations in alpha-band contribute to the maintenance and stability of VWM representations.

Further research is necessary to clarify whether this, or other interpretations of these results is the most plausible. For instance, more direct evidence for a causal link between disrupted frontal-posterior connectivity and increased bottom-up interference, as well as the possibility that the pattern of suppression occurs as a function of the specific features of the visual representations being maintained on a given trial. This will be the focus of experiment 2.

EXPERIMENT 2: DECODING ITEM IDENTITY

Studies exploring the role of neural oscillations in the alpha frequency band during VWM tasks have revealed increases within task-relevant cortical regions. This pattern suggests a possible role for such oscillations in suppressing competing feature representations in support of target maintenance. However, there have been alternative proposals that potentially explain the observed changes as well. As in Experiment 1, I attempted to explore the relative viability of these alternatives by examining the role of DPABP in mediating the processing of, and interference associated with, visual distractors. In this case, specifically, those differing from the target with respect to the task-relevant feature dimension.

In pursuit of this goal, I asked participants to complete a precision delayed recall task while I recorded concurrent EEG activity. As before, on each trial, an irrelevant distractor was presented during the delay. Importantly, distractors in this case varied in terms of their similarity relative to the target along the task-relevant feature dimension, i.e. orientation. The resulting alpha-band activity and distractor evoked response (d-ER) estimates were then used in combination with an inverted encoding model of orientation selectivity in order to establish whether any correlation exists between the target-dependent topography of posterior ABP and the selective suppression of target-inconsistent stimulus features, from within the task relevant feature dimension.

The specific goals of the analysis described in the following sections are twofold. The primary goal was to determine whether the pattern of oscillatory activity in the alpha frequency band across the scalp contains information specific to the memory target's orientation, and further, whether the specificity of this distributed pattern can be used to predict variability in the

d-ER. Second, I examined the degree to which target-related information is maintained following distractor presentation and whether or not this is predictive with regards to the accuracy of participant's response at test.

Method

Participants

Seventeen participants between the ages of 18 and 35 were recruited from the North Dakota State University student population. All participants had normal or corrected to normal visual acuity, and provided written informed consent prior to participation. Participants received either course credit or monetary compensation ($15/hr) for their participation.

Materials and Stimuli

Stimulus presentation and response recording were controlled by a PC running Matlab (Mathworks, Inc.) with Psychophysics Toolbox extensions (Brainard, 1997; Pelli, 1997). Stimuli were presented against a light grey background (RGB=[125,125,125]) on the surface of a 19" cathode ray tube monitor with a refresh rate of 100 Hz at a viewing distance of 70 cm. Participants were asked to maintain fixation on a centrally presented black fixation cross throughout each trial.

Participants were required to remember the orientation of a memory target and to ignore a subsequently presented distractor stimulus (Figure 9). The memory target and distractor stimulus presented on each trial consisted of a single oriented Gabor patch. In each case, the diameter of the Gabor was 7.1° of visual angle. The phase and frequency of the Gabor patch was varied between memory and test display. This was done in order to encourage participants to maintain

the actual orientation of the target stimulus rather than a specific spatial location along the perimeter of the Gabor patch.

On each trial, the orientation of the memory target was drawn at random from within one of six orientation bins, with the constraint that each bin was sampled from equally often within and across blocks. Each bin spanned a 7.5° range of possible orientations from 0° to 180°. The boundaries between any two adjacent target orientation bins was separated by 15°. Each target bin corresponded to the following orientations: [Bin1 = 15°:22.5°; Bin2 = +/- 37.5°:45°; Bin3= 60°:67.5°; Bin4 = 105°:112.5°; Bin5 = 127.5°:135°; Bin6=165°-172.5°]. The orientation of the delay period distractor on a given trial was chosen from one of three possible bins that varied in terms of their similarity to the target [Bin1 = +/- 22.5 °; Bin2= +/- 45°; Bin3 = +/- 67.5]. The selection among distractor bins was balanced within each target bin for a given block and randomized across blocks. Thus, all possible combinations of target bin and relative distractor orientation bins were probed equally often. Finally, at the end of each trial, participants were presented with a test stimulus, a Gabor patch identical to the target, differing only in terms of its phase, frequency and starting orientation. Again, this was done in order to encourage participants to maintain the actual orientation of the target rather than a specific spatial location along the perimeter. Participants estimated the orientation of the memory target by using a computer mouse to adjust the angle of the test stimulus to match the remembered orientation.

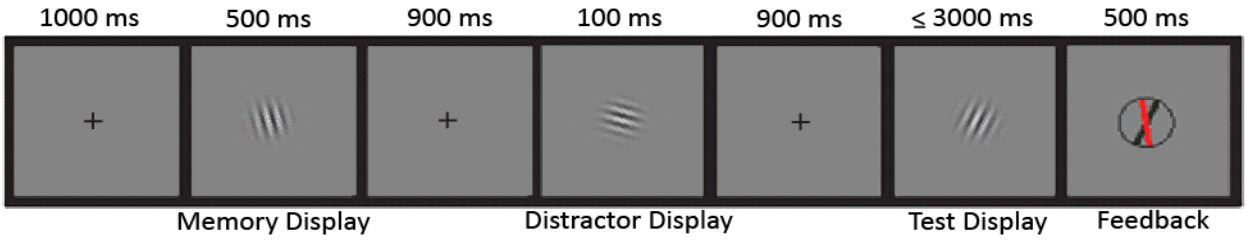

Figure 9. Experiment 2 Design.

Procedure

During the experiment, participants were seated in a dimly lit, noise-controlled and electromagnetically shielded room. Each trial began with the presentation of a memory target for 500ms (see Figure 9). The memory target was followed by a 900ms inter-stimulus interval (ISI-1) and the 100-ms presentation of a distractor stimulus. The distractor was then followed by another 900ms ISI (ISI-2) and the eventual appearance the test display probing the participant's memory for the orientation of the target VWM item. The test display remained on screen until the participant either made a response or 5 seconds had elapsed, whichever occurred first. Participants were instructed to respond as accurately as possible. Response speed was not emphasized. Following the test display, participants received feedback in the form of a red line, whose orientation corresponded to the true orientation of the target, and a black line reflecting the orientation of their response. If the response was perfect, participants would see only the red line. The feedback screen was followed by an inter-trial interval varying in length from 950-1000ms.

EEG Acquisition and preprocessing

EEG acquisition was identical to that described for Experiment 1. EEG preprocessing followed the procedures outlined in Experiment 1 Methods, with the exception that, in this instance, the data was not filtered prior to artifact detection, but instead was filtered between 8-12 Hz following artifact rejection while creating data epochs (see ***EEG analyses*** section for further details). Using the pre-processed data, individual EEG epochs were created within a 3000-ms time window, beginning 300-ms before the onset of the memory target and ending 500-ms after the appearance of the test array. Individual trial epochs were baseline corrected by

subtracting the mean amplitude in the 300-ms pre-stimulus time window from each trial time point. Artifact identification and rejection procedures were identical to those used in Experiment 1.

Statistical Analyses

Modeling response error distributions

To investigate any possible systematic response distortions as a function of distractor similarity, for each trial, a response error value was calculated, reflecting the difference between the participant's response and the orientation of the target stimulus (ranging from -90 to 90°). Maximum likelihood estimation was used to fit a Von Mises distribution to the distribution of response error values, using methods described by Zhang & Luck (2008). This method was used to estimate two parameters of the Von Mises distribution: the mean (μ), which gives an indication of how similar recall responses are, on average, to the target orientation, and the standard deviation (SD), which indicates the relative precision of the stored orientation, with smaller SD values corresponding to higher mnemonic precision.

EEG Analyses

The analysis described in the following sections will have a number of related components. The primary goal will be to determine whether one can decode the identity of the target memory item based on the distributed topography of ABP during the first and second inter-stimulus intervals. If the topographic distribution of DPABP reflects a unique pattern of inhibition dependent upon the specific target item being remembered, then the distributed pattern of activity should be uniquely tuned to this item. As a result, I should be able to decode its identity on the basis of this pattern of activity. Additionally, if this pattern reflects the

suppression of non-target features along the target feature dimension, as I propose, then during ISI-1 it should be associated with the degree to which such features are suppressed. Accordingly, the second goal of this analysis will be to assess whether the relative specificity of the distributed pattern of DPABP is significantly less consistent across trials on which the electrophysiological response evoked by the distractor items are relatively pronounced.

Global mean field power analysis

To examine differences in the electrophysiological response to the distractor stimulus, as well as how they may relate to the distributed pattern of DPABP, I computed global mean field power (GMFP, Lehman & Skrandies, 1980). GMFP reflects the spatial standard deviation of activity across all 64 electrodes, averaged across a given time interval, and provides a global measure of the electrophysiological response evoked by the distractor stimulus.

To determine whether distractor processing is influenced by the distributed pattern of DPABP, artifact free trial epochs created during preprocessing were used to compute GMFP across the interval from 300 ms prior to the distractor onset up through 900 ms following offset. The 300ms pre-distractor interval served as a baseline. I then visually inspected the results and focused on the average across the peak magnitude of this response, the time interval from 100 – 400ms following distractor onset. The resulting GMFP estimates were used to examine whether variations in the magnitude of the GMFP between these trial sets were correlated with the changes in the distributed pattern of DPABP.

Inverted Encoding Model Based Analysis

The proposed theory suggests that the distributed pattern of DPABP will vary as a function of task demands. In the present study I was interested in how they may vary as a

function of the cognitive demands imposed by maintaining a specific task-relevant visual feature. Because all other task demands remain more or less constant across trials, each target representation should exhibit a relatively consistent, corresponding pattern of alpha activity, differing only as a function of its unique orientation. If the distributed pattern of alpha activity reflects the suppression of a differential subset of non-target orientations, as I propose, and the specific subset is dependent upon the target items maintained on any given trial, then I should be able to decode the target item on the basis of this pattern of activity. Further, the relative specificity of this pattern should be related to the evoked response, as measured by GMFP, produced by a task-irrelevant distractor stimulus with non-target orientation features.

In an attempt to show that the pattern of alpha activity generally reflects a target dependent pattern of inhibition, I used an inverted encoding model (IEM) of orientation selectivity. The IEM assumes that the pattern of responses across each electrode samples from underlying neural populations tuned to varying degrees to the full range of orientation features. The response of any given electrode then is proportional to the summed response of all of these underlying sub-populations, each differentially responding to a given target (Foster et al, 2016). If DPABP reflects the inhibition of such orientation tuned subpopulations, then the IEM, which itself assumes activation tuned differentially on this basis, should successfully decode the identity of the target item. The IEM should however not successfully decode the target item if the pattern of activity is not selectively tuned in this manner, making it a more robust test of our hypothesis as compared with other less constrained decoding methods.

When conducting this analysis, I followed a procedure similar to that outlined by Foster et al. (2016). I began by band-pass filtering the EEG signals within a 3100-ms time window

surrounding the memory target (see *EEG acquisition and preprocessing* section) into five frequency-specific bands ranging from 8-12 Hz in 1 Hz steps using a two-way least squares finite impulse response filter (modified version of eegfilt.m from EEGLAB toolbox) (Delorme & Makeig, 2004). A Hilbert transform was then applied to the resulting frequency-specific filtered waveforms, resulting in a complex analytic signal for all 64 electrodes. The frequency specific instantaneous amplitude was then extracted from this complex analytic signal for each of eleven posterior electrodes (Figure 10). These electrodes were chosen as a result of their relative proximity to regions believed to be responsible for processing orientation features (i.e., the occipital cortex).

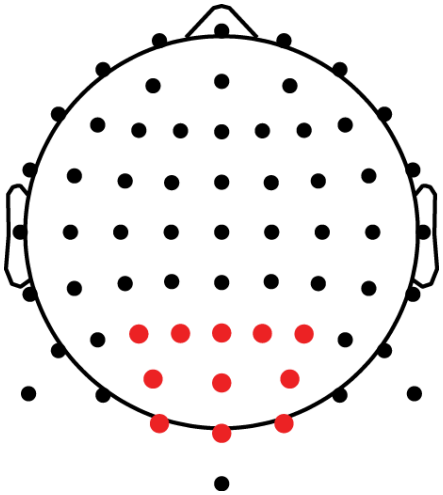

Figure 10. Electrodes used in Decoding Analysis (highlighted in red).

The frequency specific instantaneous amplitude for each electrode was used in order to extract single-trial waveforms of induced oscillatory power. Induced power was estimated by squaring the magnitude of the complex analytic signal obtained from the Hilbert transform, for each trial separately. These values were then averaged across all trials within a given block, for

each of the six orientation bins, again separately. This means of estimating power emphasizes signals that are not necessarily phase locked to stimulus onset. Induced power for each frequency was estimated for each block and time point across the selected set of posterior electrodes. Blocks were derived by first identifying for each participant the target bin retaining the fewest trials following artifact rejection. These trials were then evenly distributed, at random, across ten training blocks. The data for all trials among the five remaining target bins, were then also distributed evenly across these same ten training blocks such that the number of trials assigned to each training block was equivalent across all bins. Consequently, all target bins had the same number of trials within as well as across the ten blocks. Induced power was then averaged across each target bin, at each time point within a given block.

After dividing trial data among the 10 blocks, the blocks were further divided into separate training and test sets. Each training set consisted of nine blocks. The tenth block was used as a test set. The procedure, detailed below, was carried out for each frequency and time point of interest. Further, to minimize the effect of any issues arising from idiosyncrasies in the random distribution of trials across blocks and to assure that no trial was left out of the analysis all together, I repeated the random block assignment and the following decoding procedure 15 times per subject, and used the average of these 15 iterations for all subsequent analyses.

Next, we used an inverted encoding model to reconstruct orientation-selective channel tuning functions (CTFs) from the topographic distribution of oscillatory power across electrodes. In line with the 6 target bins, a basis set which assumed that the power measured across individual electrodes was the weighted sum of at least 6 orientation tuned neural sub-populations

was used. This basis set was a matrix reflecting the hypothetical channel response from these populations as given by:

$$\mathbf{R = \sin(0.5\ 0)^5}$$

This response profile was circularly shifted for each channel such that the peak response was centered over one of the 6 orientation bins. Data from the training set were mapped onto channel outputs from the matrix reflecting the hypothetical channel response (i.e. the basis set) via a weight matrix W [*electrodes* × *observations*]. This matrix was estimated with a general linear model (*B1 = WC1*) where *B1* is a training set, *W* is a weight matrix and *C1* reflects the derived hypothetical channel response for a given training set. This matrix *W* reflects the estimated weights on the hypothetical orientation channels for each electrode. This channel weight matrix was then derived via repeated least squares estimation comparing training and test data across iterations using the following formula:

$$\widehat{\boldsymbol{W}} = \mathbf{B_1\ C_1^T\ (C_1\ C_1^T)^{-1}}$$

The training data B1 was used to estimate these weights. W reflects a set of weights mapping the linear transformation between the hypothetical idealized response of the basis set and the observed activity of B1. The model was then inverted in order to transform the observed data B_2 into estimated channel responses C_2 as follows:

$$\mathbf{C_2 = (\widehat{W}^T\ \widehat{W})^{-1}\ \widehat{W}^T\ B_2}$$

I then used a "leave-one-out" cross validation procedure in which nine of the ten blocks were used as the training set (B1) in the process of estimating the weight matrix ($\widehat{W}$). The remaining block always served as the test set (B_2) used to estimate the channel response function

(C_2). This process was repeated until all ten blocks, had been used as the test set. As noted above, this process was iterated 15 times to account for idiosyncrasies in random block assignment.

The resulting orientation-specific channel response functions were then circularly shifted to a common center and averaged together, yielding the overall CTF. The location of the CTF's peak along the x-axis reflects the degree to which the oscillatory response is tuned to the target orientation across trials. An offset relative to 0 implies activity tuned to orientations other than the target item. These CTF's were averaged across all ten iterations of the leave one out procedure and then the fifteen iterations of the random block assignment.

Channel Tuning Function Analysis

CTFs were used to determine whether the distributed pattern of oscillatory activity across electrodes exhibit activity tuned to the target orientation. This was done by investigating whether patterns of alpha-band activity across trials are distinct enough between orientation bins to enable decoding via the IEM across the entire delay interval.

The efficacy of decoding was quantified by comparing the slopes of derived CTFs across time against a hypothetical null distribution of CTF slopes. This null distribution was generated using a randomized trial-bin label permutation procedure. This entails generating CTFs by randomizing target-bin labels across trials and then generating CTF's from the resulting scrambled datasets. This process was iterated a thousand times. On each iteration the CTF was derived using the same procedure detailed above (see: Encoding Model Based Decoding Analysis). The slopes for each of the resulting CTFs were used to generate a hypothetical null distribution meant to approximate the probability of achieving the observed level of target-orientation selectivity from data for which, inherently, no significant relationship should exist

between the target bin and the distributed pattern of activity. Here, above chance classification accuracy will support the proposal that the distributed pattern of DPABP is dependent upon the specific target item.

A further question was whether or not I observed a disruption with respect to the CTF during ISI-2 following the distractor, as compared to ISI-1. Here I separately averaged the CTF's observed across ISI-1 and ISI-2 to see if they differed significantly in terms of orientation selectivity.

In order to quantify changes in CTF orientation selectivity, I performed a two-step curve fitting procedure. This entailed fitting a Gaussian function to the CTFs derived for each participant. With respect to this function, u=mean, a=the amplitude scaling factor, b = the baseline scaling factors, and ∂ = dispersion. These parameters were derived via a MATLAB routine that attempts to optimize the fit of a Gaussian function to the CTF curve. In terms of the resulting parameters, a will be an estimate of the response amplitude for a given stimulus, while ∂ will be an estimate of the dispersion within the resulting CTF, and by extension the degree to which the distributed pattern of ABP is tuned to the target orientation.

Channel Tuning Function

Much of the remaining analysis will focus on both the amplitude and dispersion parameters derived from this procedure across various CTFs. This will be done in order to investigate a number of questions regarding the degree to which adherence to the aforementioned target-dependent alpha distribution predicts both the variability in the evoked responses as well as cognitive and behavioral factors that I am attempting to associate with such variations.

First and foremost, while decoding of the target VWM item is consistent with the proposed view, this merely implies that the distributed pattern of DPABP is selectively tuned to the specific orientation being held in memory. This result alone would not differentiate between my proposed inhibitory interpretation and the alternative, whereby the orientation selective activity is representing the maintenance of the target item per se.

If, as I propose, alpha reflects a mechanism whereby each target item is selectively suppressing a unique subset of non-target orientation features, then the distributed pattern of DPABP would, under ideal circumstances, exhibit a selective topography corresponding to this suppression and by extension the associated target item. Across trials for which the distributed pattern of ABP is relatively less consistent with this ideal, the CTFs, which rely on this pattern, should exhibit less aggregate tuning to the target item and consequently higher dispersion. If as I predict, the selectivity of this topography reflects the suppression of non-target orientations, then it stands to reason that the trials exhibiting on average higher dispersion and/or lower amplitude CTF's should exhibit, in aggregate, less consistent suppression of non-target orientation stimuli, i.e. the distractors. This finding would support the view that alpha reflects a process mediating this stimulus specific pattern of suppression.

To examine this possibility, I conducted a median split analysis to determine whether variability in the global electrophysiological response evoked by the distractor is associated with variability in CTF dispersion and/or amplitude during the 100ms prior to onset. I median split trials based on the magnitude of GMFP. Using the same methods as described above, I then generated CTFs for each subset of trials and compared the resulting CTFs in terms of the associated amplitude and dispersion parameters. This was done in order to see if the pattern of

oscillatory alpha activity corresponding to each of these CTF's is shown to vary more on average across trials during which the proposed corresponding inhibition is also less consistent. Here, I expected to observe a negative relationship, with increased CTF dispersion (i.e., higher variability in stimulus-specific patterns of alpha) predicting a larger d-ER (i.e., reduced distractor suppression).

As I explained, CTF dispersion and to a lesser degree amplitude, arguably reflect coarse measures of the degree to which the stimulus specific patterns of alpha are consistent across trials. I argue that this less selective pattern of alpha reflects inadequate inhibition of non-target orientation feature processing, and thus an associated failure to mitigate the disruptive effects of the distractor. This would presumably lead to a degraded target VWM representation. If alpha-activity reflects processes protecting target representations, then it is reasonable to expect that the degree to which this suppression is successful might be meaningfully correlated with disrupting the target representation. This would likely result in the associated target-dependent inhibitory activity becoming less consistent following the presentation of various distractor orientations. Accordingly, I wanted to investigate whether I observed any disruption with respect to channel tuning across ISI's in cases in which disruptive distractor processing is less suppressed, i.e. following a larger distractor evoked response. To explore this possibility, I looked at whether CTFs generated during intervals following large distractor evoked responses show a significant increase in dispersion between ISI1 and ISI2; as compared to those following relatively suppressed evoked responses. In particular, I compared the CTF's generated following the GMFP based median split across these intervals.

Further, if the evoked response predicts differences in CTF's during ISI-2 and this reflects a disruption of processes supporting target maintenance, then it is possible that relatively higher dispersion and/or lower amplitude CTF's here could predict an increase in the magnitude of response errors across those trials. I thus generated CTFs from trials median split on the basis of response error magnitude. Then compared both dispersion and amplitude of the resulting CTFs during ISI-2 to see if they differed significantly between conditions. Again, if I observed a significant difference across the subsets of trial this would suggest that the distributed pattern of alpha activity is meaningfully related to the ongoing stability of working memory representations.

Results

Channel Tuning Function

I first sought to establish that I was able to decode the content of VWM based on the topographic distribution of induced DPABP. To this end, I used an inverted encoding model to reconstruct orientation-selective CTF's. If the multivariate pattern of power across electrodes was tuned to the specific target item, then the IEM should reveal a graded tuning function with a clear peak in the channel tuned to the remembered orientation. With respect to the circularly shifted and averaged CTFs, this would be reflected by a peak centered over 0 along the horizontal axis. If DPABP was not selectively tuned to a specific orientation stimulus than I would expect to observe flat or generally distorted CTF, reflecting this fact.

Figure 11 clearly shows that throughout the delay I was able to decode the specific target orientation held in memory. This result was compared against the null distribution I derived via a

permutation analysis, the results of which are depicted in Figure 11. Here, intervals during which the slope of the tuning curve was significantly above chance are marked by a black line.

A further question was whether or not I observed a disruption with respect to the CTF during ISI-2 following the distractor, as compared to ISI-1. To do this, I calculated the average CTF observed across both ISI-1 and ISI-2 (see Figure 12) and determined whether they differed in terms of their orientation selectivity using a t test. Although the CTF is not as peaked in the ISI-2 versus ISI-1, the results suggest that there is no significant difference in either the amplitude, $t(14) = -.968$, $p=.174$, or dispersion $t(14) = .747$, $p=.233$, of the observed CTFs.

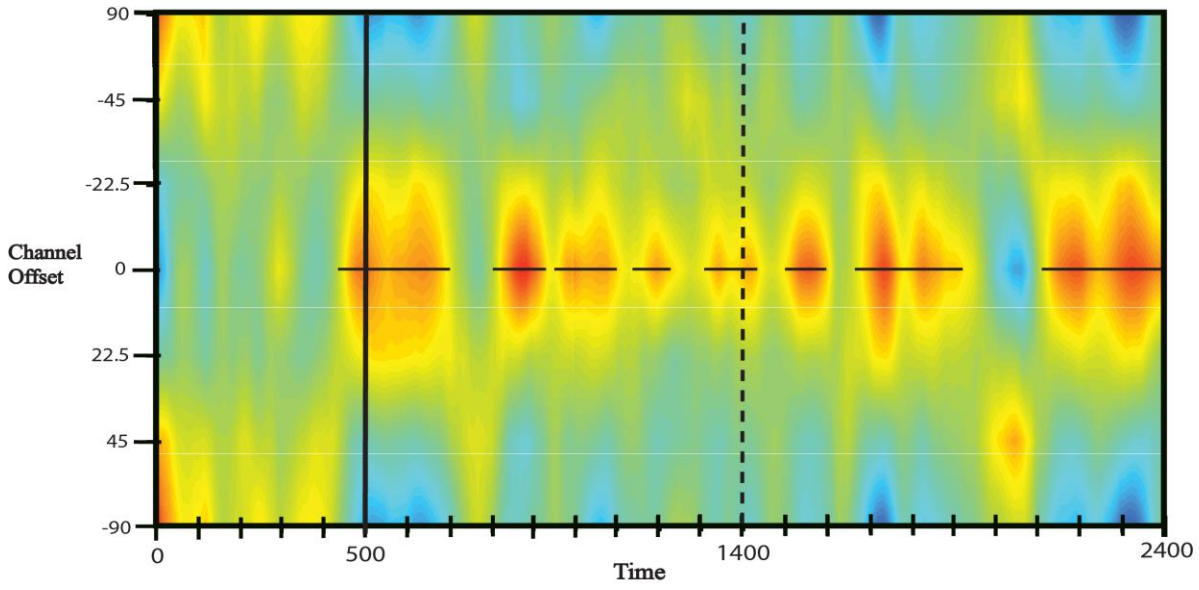

Figure 11. Channel Tuning Across Time.

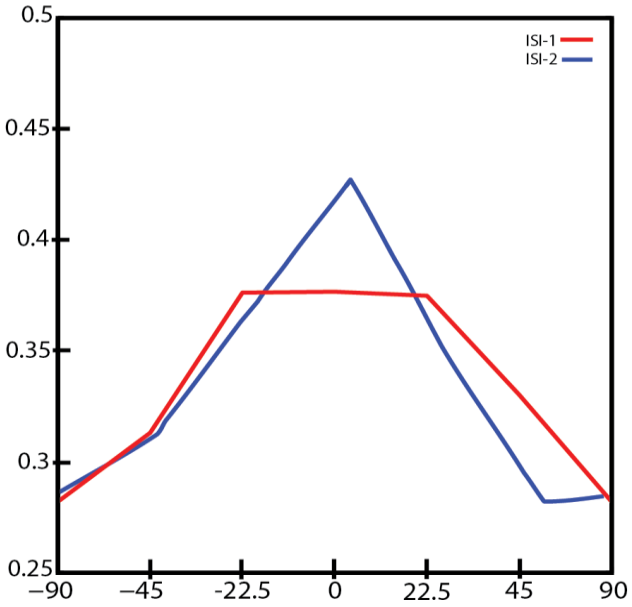

Figure 12. CTF Differences ISI-1 vs ISI-2. CTF's averaged across ISI-1 (red line) and ISI-2 (blue line)

Global Mean Field Power based Median Split Analyses

In order to provide evidence differentiating between the two possible explanations for the observed pattern of DPABP, I performed a median split of trials based on the response evoked by the probe, specifically in terms of GMFP during the interval from 100-300ms (Figure 13). This interval was chosen as it encompasses the peak magnitude of the distractor-evoked response. I then generated CTF's for each of these subsets of trials separately for each subject. Finally, I took the average of these CTF's during the interval 100ms prior to probe onset (Figure 14) and the 200ms following the peak response (Figure 15). While the differences between the functions across each subset of trials are striking visually, I aimed to quantify and compare these differences statistically. Accordingly, I again decomposed the observed channel response profiles into discrete metrics of tuning amplitude and dispersion by fitting each participant's CTF

response profile with a Gaussian function. I then used a t-test to compare both dispersion and amplitude across median split conditions, both before and after the distractor evoked response, separately.

The analysis revealed that there were significant differences in both the amplitude and dispersion of the resulting channel tuning functions for the pre-distractor interval. Specifically, on trials during which the d-ER was smaller, the CTF's exhibit significantly greater amplitude, $t(14) = -2.97$, $p=0.005$, as well as lower dispersion, $t(14) = 1.85$, $p=.042$, as compared to those trials on which the d-ER was larger. This is consistent with what one would expect given my proposal. A similar pattern held for ISI-2. In this case, however, only dispersion was significantly lower, $t(14) = 2.14$, $p=.024$, for trials on which the d-ER was smaller, whereas amplitude was only marginally significant, $t(14) = -1.71$, $p=.054$. Thus, participants exhibit greater CTF distortions throughout the both ISI's period on trials during which the probe evoked a larger response.

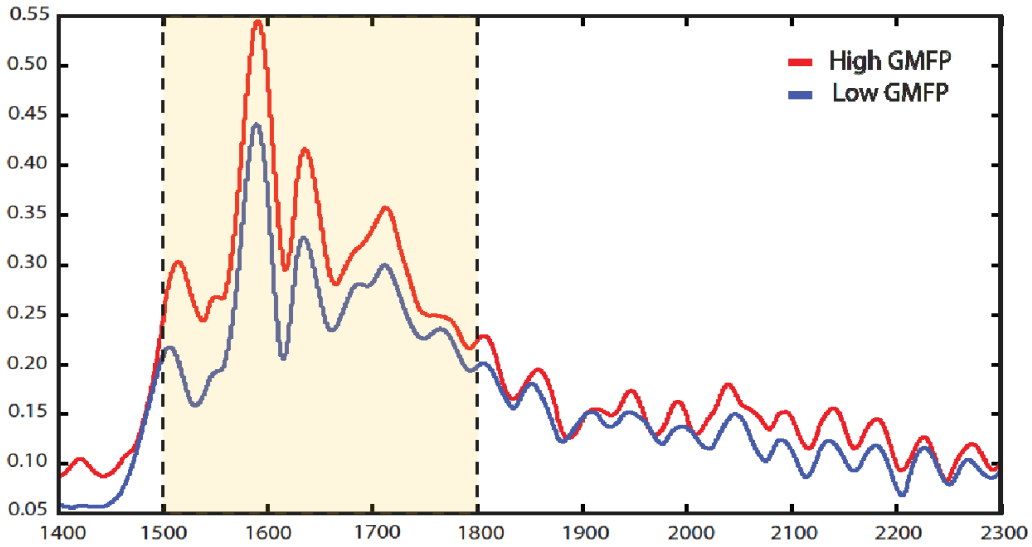

Figure 13. GMFP Median Split. Average GMFP across subjects following the probe. The region boxed in by the red line reflects the interval across which GMFP magnitude was averaged in order to perform the median split.

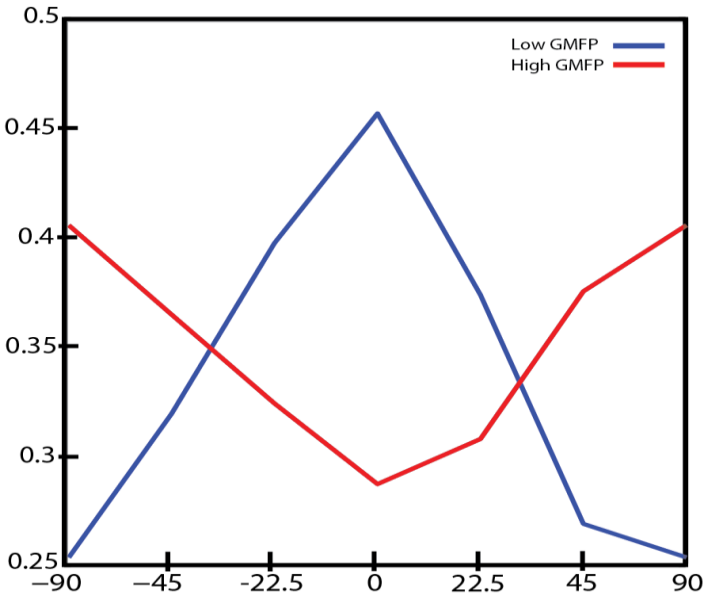

Figure 14. ISI-1 GMFP Median Split CTF. GMFP median split CTF for pre-probe interval

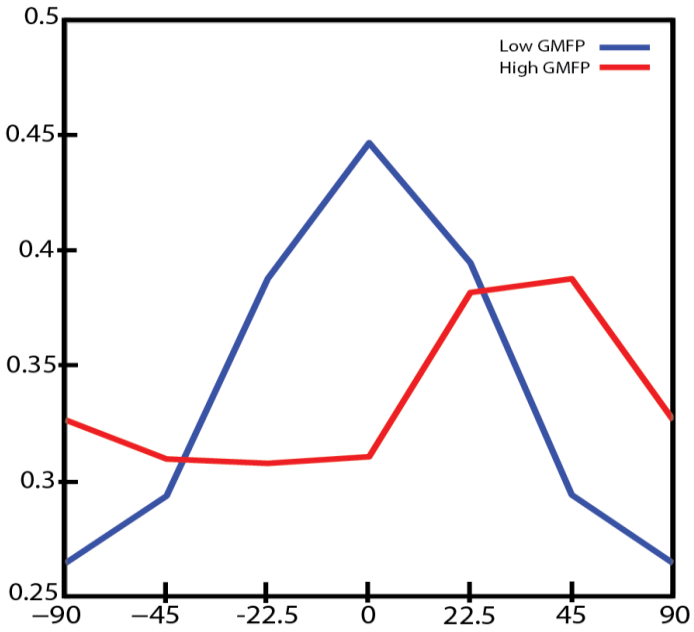

Figure 15. ISI-2 GMFP Median Split CTF. GMFP median split CTF from interval following

Response Errors

Looking to Figure16A, there does not appear to be any difference in SD for Von Mises distribution across distractor bins, there does however appear to be an attraction effect for the mean (Figure16B), in which responses are biased in the direction of the distractor stimulus. One-sample t-tests performed across distractor bins reveal that this directional bias is significant for all bins at $p<.001$, with the exception of distractor bin 1 (-67.5) $t(14) = 1.42$ $p = .176$ and bin 3 (-22.5) $t(14) = .350$ $p = .731$. Consistent with this observation, an ANOVA confirms that distractor bin did not appear to significantly influence standard deviation of the distribution, i.e. precision, $F(1, 14)= 2.29$ $p = .055$, but there was a main effect of distractor bin with respect to mean error, i.e. accuracy, $F(1, 14)= 15.903$ $p < .001$.

The results of the GMFP based median split of trials suggest that following a larger d-ER, individuals did not exhibit significantly greater reductions in either the accuracy, $t(14) = .932$, $p=.367$ (Figure17A), or precision of responses, $t(14) = -.956$, $p=.355$ (Figure17B). These results are not what we would predict if distractor processing is the primary source of interference with respect to target maintenance.

The results further suggest that the distributed pattern of DPABP is tracking response accuracy to some degree, that is to say the absolute difference between the target orientation and response orientation. In particular, when a median split is performed on the basis of the magnitude of response errors (large vs. small) across trials (Figure 18), the results show that larger response errors are associated with significantly lower amplitude CTFs, $t(14) = 1.77$, $p=.049$, and significantly higher dispersion $t(14) = -2.96$, $p=.005$.

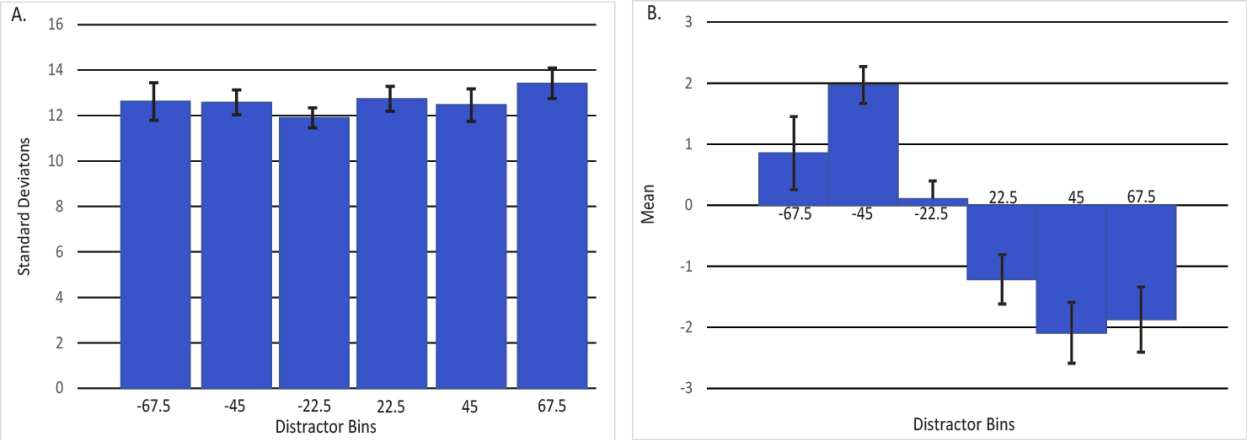

Figure 16. Differences in the Standard Deviation (A) and Mean (B) of Von Mises Distribution across Distractor Bins.

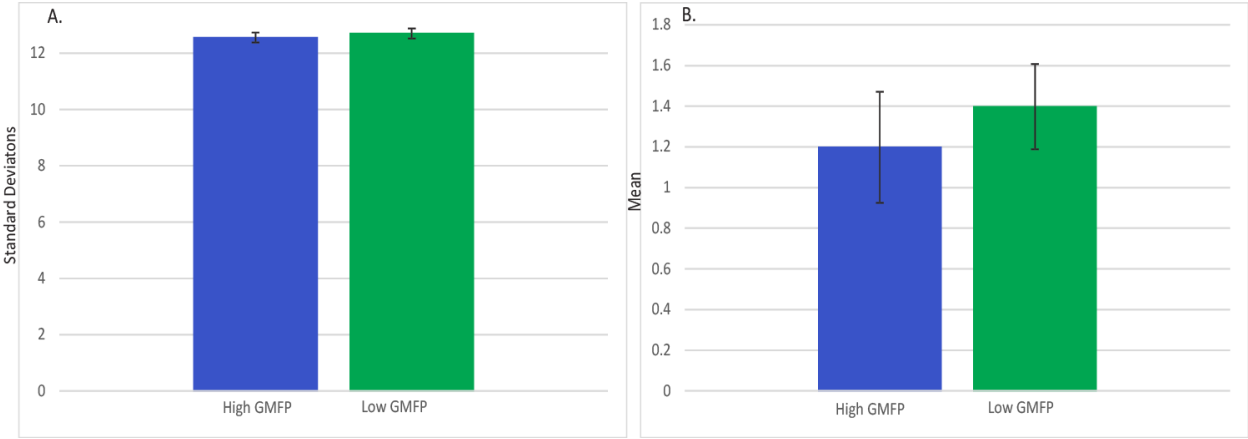

Figure 17. GMFP Median Split Differences in Standard Deviation (A) and Mean (B) of Von Mises Distribution.

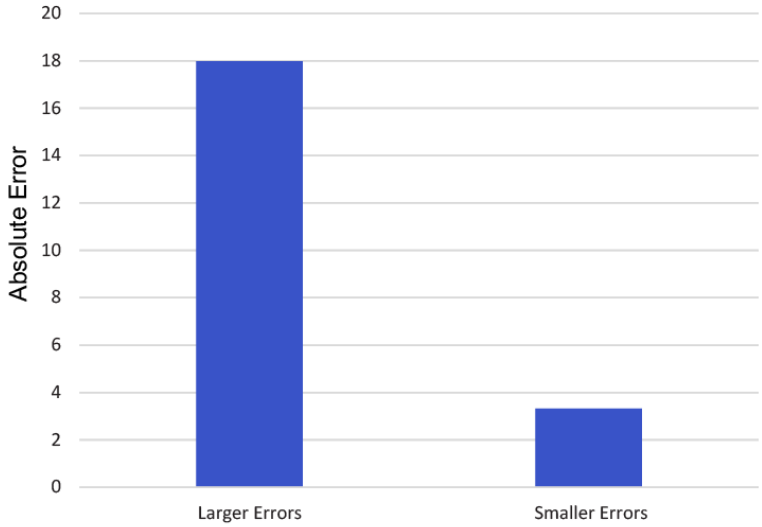

Figure 18. Median Split of Response Errors. Average offset of responses in degrees (absolute difference between the target orientation and response orientation) following a median split of trials based on magnitude of these response errors across trials.

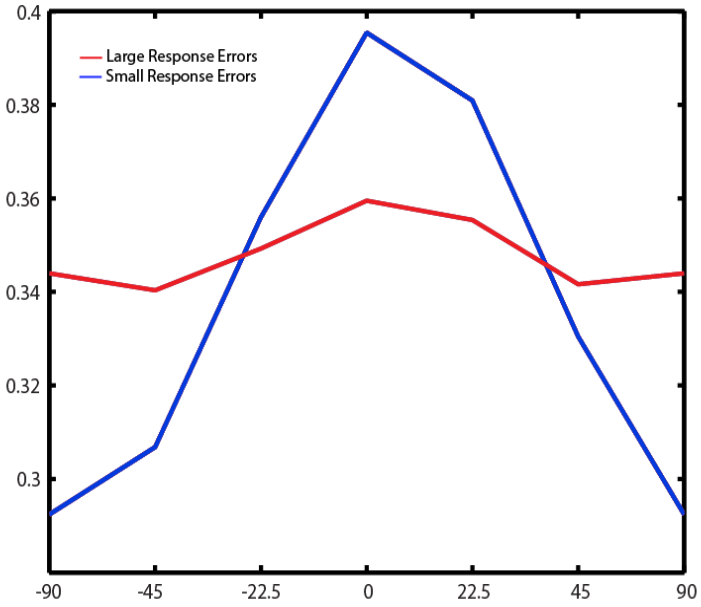

Figure 19. CTF Differences High vs Low Response Errors: ISI-2. CTFs generated following median split of trials based on the magnitude of response errors.

Discussion

The pattern of results I observed in Experiment 2 are consistent with the proposed view that the distributed pattern of DPABP observed over task-relevant regions reflect the inhibition of task-irrelevant representations along the target feature dimensions. Specifically, CTF's derived via the IEM suggest that DPABP across posterior electrodes, exhibits a distributed pattern of activity that is uniquely tuned to the specific target orientation across trials. The results of the permutation analysis confirm that this selectivity remains stable across almost the entire delay, despite the intervening presentation of a salient distractor. Again, if the observed pattern of DPABP reflects a mechanism by which each target item is inhibiting a unique subset of non-target orientation tuned neural populations, as I proposed, this is what one would expect. This result alone, however, does not rule out non-inhibitory explanations for this activity; it remains possible that DPABP in this case might reflect the activity of the populations tuned to the target orientation itself.

To provide further support for the inhibitory interpretation, I wanted to see how the relative specificity of this distributed pattern of oscillatory activity was related to the degree to which non-target orientations were suppressed across trials. Following a median split of trials based on the magnitude of the distractor evoked GMFP, the two resulting CTF's exhibited significant differences in terms of both amplitude and dispersion. As expected, on those trials during which the distractor evoked a relatively suppressed response, the resulting CTF's exhibit a distributed pattern of activity tuned selectively to the target item, with significantly higher amplitude and lower dispersion as compared to the CTF generated from trials on which the response was relatively larger.

As an extension of this GMFP-based median split analysis I was interested in investigating the possibility that a larger distractor evoked response might disrupt VWM representations to varying degrees across trials, resulting in a relatively greater CTF distortion during ISI-2. The results of this analysis were consistent with this possibility. Further, those trials on which I observed greater response errors also resulted in more distorted CTF's, consistent with the idea that these patterns of activity are dependent upon the actual representation being maintained across time. Surprisingly, however, GMFP magnitude was not significantly correlated with the precision or accuracy of VWM recall at test as one might expect if this were the primary source of interference.

GENERAL DISCUSSION

For the present thesis I conducted a pair of experiments attempting to provide evidence for the proposal that elevated power in the alpha frequency band, during the performance of attention and working memory tasks, reflects a mechanism of attentionally selective cortical inhibition. This suppression is oriented toward potentially disruptive task-irrelevant processes as a function of concurrent task demands. As such, the specific pattern of inhibition enacted is dependent upon the task-relevant processes engaged, as well as any foreseeable interference with respect to these processes given the context. This mechanism, I argue, is enacted across levels of functional specificity within the brain, from simple features to entire functionally distinct sub-regions; ultimately, gating, filtering and otherwise guiding distributed patterns of neural activation across a variety of situations.

The majority of evidence throughout the attention and working memory literature is consistent with this inhibitory view. In the context of the WM literature, however, it has been argued that in some cases increases may instead reflect the distributed network activity underlying the allocation of attention and the maintenance of information in WM. In these cases, it is thought that increases in ABP may reflect the active processes associated with increasing demands on attention and maintenance related neural systems, rather than inhibition. I have argued that these observations could instead be just as easily explained by the aforementioned inhibitory mechanism.

Experiment 1 focused on the exploring the functional significance of increases in DPABP which scale with VWM load. Because ABP appears to scale with the number of relevant items held in mind, it is possible that it could reflect processes supporting maintenance of these

additional items. I hypothesized that these increases may instead reflect a mechanism that serves to suppress on-going visual processing, insulating target representations from sources of bottom-up interference. To examine this possibility, this experiment focused on how increases in DPABP are associated with the maintenance of additional items in VWM, efficacious suppression of ongoing visual processing.

As hypothesized, the pattern of results were consistent with this second perspective. The behavioral evidence suggests that the increases in DPABP were not correlated with the actual number of additional items participants were holding in memory across set-sizes, as one would expect if this activity was supporting maintenance per se. Instead, increases in DPABP were correlated with the number of additional items an individual was successfully holding in memory during only the probe-present trials. That is, while on average most individual's exhibit no significant difference in working memory load across set sizes in the probe condition, the individuals showing the greatest increases exhibit a significantly greater increase in DPABP. Consistent with the claim that these oscillations reflect processes gating additional target items against interference.

If increases reflect an individual's attempt to insulate the maintained items against interference, then when no probe is presented, the degree to which one has implemented this strategy is irrelevant. Assuming all else is equal, an individual could completely fail to gate target processing and still maintain a number of items close to capacity nearly as effectively as subjects who had. By comparison, if one attempts to maintain items near capacity in the presence of a significantly disruptive probe and fails to inhibit processing, they would be more susceptible to interference and thus show an associated drop in estimated capacity. This is why one would

expect increases in DPABP to be correlated with the number of VWM targets successfully insulated against interference, rather than only changes in load in the no-probe condition.

Given the way in which DPABP and capacity are estimated, it is unclear whether this reflects the fact that individuals are protecting all the maintained items on more trials but are completely failing on others, or if they are able to protect only a certain number of items against interference on any given trial. If the former possibility were the case, I should have observed a correlation between increases in DPABP and the interference effect at set-size 4. The lack of a correlation in this case lends some further credence to the proposed interpretation.

In order to provide further evidence that increases in these oscillations reflect an inhibitory process, I looked to see if they were correlated with any set-size dependent suppression of either early and/or late components of the p-ER. The results of this analysis further bolster support for the proposed interpretation. The increases in DPABP were significantly correlated with a reduction in the amplitude of the later component of the p-ER (LAP). The results of my analysis were not however consistent with the possibility that increases in DPABP are engaged in suppressing relatively early visual sensory processing, as the sensory gating hypothesis proposes. If alpha-band oscillations did reflect sensory suppression in the current task, one would have expected to observe a correlation between increases in DPABP and suppression of the P1 and N1 components as well as potentially a reduction in the interference effect. The results of both experiments however do allow for the possibility that gating/filtering is occurring at various, possibly later, points along visual processing hierarchy.

Evidence from the visual attention literature suggests that ERP's temporally coincident with the LAP are associated with filtering items on the basis of relevant features and task

demands. Specifically, the P2 component may be associated with selectively filtering visual stimuli on the basis of task relevant features (Evans and Federmeier, 2009; Kaan and Carlisle, 2014), while a larger P300 amplitude has been associated with conscious perception of visual stimuli as well as later recall and task-relevance of sensory stimuli (Rutiku, Aru, Bachmann, 2016; Donchin, 1981). Modulations of the LAP, which bridges the interval during which these ERP's are observed, thus may reflect processes related to the degree to which the probe is filtered on the basis of expectations, in particular various visual features.

Feature and shape processing are thought to occur within successively more anterior regions in the brain. Prior evidence implicates these regions, initially engaged in processing shape/feature representations, in the processes maintaining the corresponding representations in VWM. Sensory inputs normally proceeding to later cortical areas, could then reflect a source of interference with respect to maintained target representations. It would follow logically that the inhibitory process underlying increases in DPABP could be selectively filtering sensory inputs by suppressing task-irrelevant processing at, and/or before, these more anterior cortical regions, protecting the target representation(s), the findings appear to corroborate this interpretation

Specifically, in Experiment 1, the later anterior positive component may be suppressed to the degree that this gating is successful in preventing the flow of feed forward probe processing from reaching these more anterior cortical regions, representing the abstract target shapes. This interpretation squares with the observed latency of the LAP modulations and the more anterior topography of its peak amplitude. The correlation between LAP amplitude and the later disruption in top-down connectivity, DPABP and the performance deficits associated with the magnitude of this disruption, follow logically from this perspective. This is consistent with the

idea that increases in DPABP reflect the progressive implementation of an attentionally selective top-down mechanism of inhibition, which serves to suppress sources of potential interference.

It is generally accepted that in attention and working memory tasks, these frontal and parietal regions underlie many of the same attentionally selective control processes attributed to oscillations in the alpha frequency band. Localization of increased DPABP in these regions has been suggested to reflect some component of these higher order control mechanisms coordinating cell assemblies and selectively maintaining items in VWM via top-down inter-areal synchronization. The increases in ISPC observed between frontal electrode sites and those regions exhibiting the load-dependent increases in DPABP just prior to the onset of the probe stimulus in Experiment 1, I argue reflect the implementation of this type of top down control. The significant correlation between set-size dependent increases in DPABP and ISPC just prior to probe onset supports this view.

The temporal dynamics of each of these components across the delay lend further support to this possibility (see Figure 3 and Figure 7). Recall, ISPC reflects the degree to which the phase of oscillations are aligned at a given time point across trials. So, increases in ISPC just prior to the probe suggest that, on a greater proportion of trials, the control process this connectivity reflects, has been initiated or sustained at or before that point in time. In this way, both the magnitude of ISPC and the increases in DPABP can be thought of, not as gradually ramping up, but instead as initiated at various points during the delay period, with a dramatic increase in frequency and/or retention just prior to the probe. The onset latency, with respect to the point at which I observe increases in each clearly lends itself to the interpretation that ISPC is driving increases in DPABP. In addition, the fact that the frequency with which it is initiated increases

just prior to the expected onset of the probe, supports the view that they are possibly preparatory in nature. Given the predictable onset of the probe, the processes underlying ISPC may be specifically initiated because the probe was expected to occur around that point in time.

With regards to this line of reasoning, it is worth highlighting both DPABP and ISPC as observed in the no probe condition (see Figure 3 and Figure 7). In each case, they clearly peak roughly around probe onset and trail off gradually following the interval during which the probe should have appeared. This suggests that connectivity is initiated and maintained with what is effectively a probability distribution centered more or less over the expected onset of the probe. Further corroborating the idea that these changes are related specifically to the temporal expectation of the impending probe. Low trial numbers in the no-probe condition resulted in estimates of ISPC that exhibit significantly more noise, as compared to the probe condition. This number of trials is not as impactful with respect to estimates of ABP. That being said, both the onset latency and set-size dependent increases remain evident with respect to both.

This apparent trailing off of DPABP and ISPC, following the expected onset of the probe, is noteworthy for another reason. By comparison, in the probe condition, both ISPC and DPABP drop off abruptly following the actual probe display. As I explain above, ISPC reflects the proportion of trials upon which the phase between frontal and posterior electrodes sites are aligned. This sudden drop in ISPC and DPABP lend further, bi-directional support to the relationship between the two phenomena, i.e. ISPC possibly driving increases in DPABP. In addition, this suggests that on a majority of trials, for most subjects, something about the probe is disrupting the phase alignment and thus connectivity observed between frontal and posterior electrodes.

This supports the view that the functional contribution underlying ISPC, and the associated rise in DPABP, are likely directly related to modulating processes associated with this impending probe. Additionally, the set-size dependence and pattern of behavioral disruptions associated with modulations in ISPC suggests that they are intimately related to the effectual implementation of a protective VWM maintenance process. Such predictive biasing of sensory processing is a hallmark of the functional contribution made on the part of executive control processes, which I propose underlie implementation of this mechanism and are reflected by ISPC.

Clearly, experiment one provides a number of insights supporting the proposed functional significance of set-size dependent increases in DPABP, the temporal dynamics of these changes as well as their potential sources. In particular, providing evidence that alpha reflects the selective suppression of on-going visual processing. This suppression is potentially driven by top-down control as reflected by interregional phase clustering. Arguably this facilitates maintenance, in that, the degree to which processing is effectively suppressed, is ultimately relevant to behavior at test.

Experiment 2 builds on these findings, providing additional support for the ubiquity of the proposed inhibitory mechanism, investigating more directly the potential functional significance of increases observed over task-relevant regions. The ability to decode the content of VWM in Experiment 2, implies that these increases likely reflect a processes supporting the maintenance of target items. It is difficult to explain why else the distributed pattern of activity would be specifically and differentially tuned to the orientation of target items across trials.

However, this observation alone does not necessarily differentiate between the maintenance account and the more general inhibitory view proposed here.

By the latter view, the increases in DPABP in this case are supporting maintenance by inhibiting non-target orientation features, thus insulating target representations from these potential sources of lateral interference. In this way they may help to stabilize the maintained target orientation throughout the delay interval. The results of the median split analysis were consistent with this view in that trials on which the target dependent subset of non-target orientations features are less consistently suppressed, are those trials on which the pattern of activity, proposedly reflecting this suppression, is less consistently tuned to the target item. This pattern clearly parallels the shifts in these oscillations across most other attention and working memory tasks. For example, the aforementioned increases observed over task-irrelevant spatial regions within early retinotopic cortex, which similarly allow one to decode the target location being held in VWM on the basis of inhibitory activity. These results are thus consistent with the view that, in addition to the more regional suppression previously observed in attention and working memory tasks, alpha-band activity may be related to suppressing task irrelevant processes within relatively more circumscribed neural networks during VWM maintenance.

Conclusion

Together the two experiments provide multiple lines of evidence supporting the proposed inhibitory view and challenging the alternative interpretation for both load-dependent increases in DPABP as well as increases observed within putatively task-relevant cortical areas. Elevated power in the alpha frequency band, during the performance of attention and working memory tasks, likely reflect attentionally selective cortical inhibition, oriented toward potentially

disruptive task-irrelevant processes. For example, in Experiment 1 and 2, gating sources of bottom up interference which might otherwise disrupt target maintenance. The observed temporal and spatial dynamics of this activity further support the fact that this occurs predictively in response to expected sources of interference in a given context.

Finally, previous evidence has consistently shown that this inhibitory process is likely enacted across various regions along the visual hierarchy, but the current results further support the idea that this mechanism is potentially enacted across levels of specificity within as well as across these regions. In Experiment 2, specifically, potentially suppressing non-target orientation features along the task relevant dimension.

The sum total of these results provide a variety of evidence supporting the ubiquity of the proposed inhibitory mechanism. While there are alternative interpretations for some of these results, in light of the preexisting literature, it seems increasingly unreasonable to consider these as comparably viable. While these alternatives seem reasonable in isolation, the fact that such observations are just as reasonably explained by the inhibitory account, supported by so much of the literature, seems to beg the question: Why should we assume that in these cases oscillations reflect an entirely distinct mechanism as compared with the majority of others, when the relatively more parsimonious alternative is at least as plausible?

REFERENCES

Anderson, D. E., Serences, J. T., Vogel, E. K., & Awh, E. (2014). Induced Alpha Rhythms Track the Content and Quality of Visual Working Memory Representations with High Temporal Precision. The Journal of Neuroscience, 34(22), 7587–7599. http://doi.org/10.1523/JNEUROSCI.0293-14.2014

Attneave, F., & Arnoult, M. D. (1956). The quantitative study of shape and pattern perception. Psychological Bulletin, 53(6), 452–471. http://doi.org/10.1037/h0044049

Babiloni, C., Vecchio, F., Bultrini, A., Luca Romani, G., & Rossini, P. M. (2006). Pre- and poststimulus alpha rhythms are related to conscious visual perception: a high-resolution EEG study. Cerebral Cortex (New York, N.Y.: 1991), 16(12), 1690–1700. http://doi.org/10.1093/cercor/bhj104

Banerjee, S., Snyder, A. C., Molholm, S., & Foxe, J. J. (2011). Oscillatory alpha-band mechanisms and the deployment of spatial attention to anticipated auditory and visual target locations: supramodal or sensory-specific control mechanisms? The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 31(27), 9923–9932. http://doi.org/10.1523/JNEUROSCI.4660-10.2011

Barry, R. J., de Pascalis, V., Hodder, D., Clarke, A. R., & Johnstone, S. J. (2003). Preferred EEG brain states at stimulus onset in a fixed interstimulus interval auditory oddball task, and their effects on ERP components. International Journal of Psychophysiology: Official Journal of the International Organization of Psychophysiology, 47(3), 187–198.

Barry, R. J., Rushby, J. A., Johnstone, S. J., Clarke, A. R., Croft, R. J., & Lawrence, C. A. (2004). Event-related potentials in the auditory oddball as a function of EEG alpha phase

at stimulus onset. Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology, 115(11), 2593–2601. http://doi.org/10.1016/j.clinph.2004.06.004

Berger, H. (1929). Über das Elektrenkephalogramm des Menschen. Archiv für Psychiatrie und Nervenkrankheiten, 87(1), 527–570. http://doi.org/10.1007/BF01797193

Bollimunta, A., Chen, Y., Schroeder, C. E., & Ding, M. (2008). Neuronal Mechanisms of Cortical Alpha Oscillations in Awake-Behaving Macaques. The Journal of Neuroscience, 28(40), 9976–9988. http://doi.org/10.1523/JNEUROSCI.2699-08.2008

Bollimunta, A., Mo, J., Schroeder, C. E., & Ding, M. (2011). Neuronal mechanisms and attentional modulation of corticothalamic α oscillations. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 31(13), 4935–4943. http://doi.org/10.1523/JNEUROSCI.5580-10.2011

Bonnefond, M., & Jensen, O. (2012). Alpha Oscillations Serve to Protect Working Memory Maintenance against Anticipated Distracters. Current Biology, 22(20), 1969–1974. http://doi.org/10.1016/j.cub.2012.08.029

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

Brandt, M. E., & Jansen, B. H. (1991). The Relationship Between Prestimulus Alpha Amplitude and Visual Evoked Potential Amplitude. International Journal of Neuroscience, 61(3-4), 261–268. http://doi.org/10.3109/00207459108990744

Busch, N. A., Dubois, J., & VanRullen, R. (2009). The phase of ongoing EEG oscillations predicts visual perception. The Journal of Neuroscience: The Official Journal of the

Society for Neuroscience, 29(24), 7869–7876. http://doi.org/10.1523/JNEUROSCI.0113-09.2009

Capotosto, P., Babiloni, C., Romani, G. L., & Corbetta, M. (2009). Frontoparietal Cortex Controls Spatial Attention through Modulation of Anticipatory Alpha Rhythms. The Journal of Neuroscience, 29(18), 5863–5872. http://doi.org/10.1523/JNEUROSCI.0539-09.2009

Cepeda, N. J., Cave, K. R., Bichot, N. P., & Kim, M.-S. (1998). Spatial selection via feature-driven inhibition of distractor locations. Perception & Psychophysics, 60(5), 727–746. http://doi.org/10.3758/BF03206059

Cooper, N. R., Croft, R. J., Dominey, S. J. J., Burgess, A. P., & Gruzelier, J. H. (2003). Paradox lost? Exploring the role of alpha oscillations during externally vs. internally directed attention and the implications for idling and inhibition hypotheses. International Journal of Psychophysiology: Official Journal of the International Organization of Psychophysiology, 47(1), 65–74.

Cowan, N. (2001). The magical number 4 in short-term memory: a reconsideration of mental storage capacity. The Behavioral and Brain Sciences, 24(1), 87–114; discussion 114–185.

Cohen, M.X., and Cavanagh, J.F. (2011). Single-trial regression elucidates the role of prefrontal theta oscillations in response conflict. Front.Psychol. 2:30. doi: 10.3389/fpsyg.2011.00030

Delorme, A., & Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. Journal of Neuroscience Methods, 134(1), 9–21. http://doi.org/10.1016/j.jneumeth.2003.10.009

Del Percio, C., Marzano, N., Tilgher, S., Fiore, A., Di Ciolo, E., Aschieri, P., … Eusebi, F. (2007). Pre-stimulus alpha rhythms are correlated with post-stimulus sensorimotor performance in athletes and non-athletes: a high-resolution EEG study. Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology, 118(8), 1711–1720. http://doi.org/10.1016/j.clinph.2007.04.029

Dube, C., Payne, L., Sekuler, R., & Rotello, C. M. (2013). Paying attention to attention in recognition memory: Insights from models and electrophysiology. Psychological Science, 24(12), 2398–2408. http://doi.org/10.1177/0956797613492426

Foxe, J. J., Simpson, G. V., & Ahlfors, S. P. (1998). Parieto-occipital approximately 10 Hz activity reflects anticipatory state of visual attention mechanisms. Neuroreport, 9(17), 3929–3933.

Fu, K. M., Foxe, J. J., Murray, M. M., Higgins, B. A., Javitt, D. C., & Schroeder, C. E. (2001). Attention-dependent suppression of distracter visual input can be cross-modally cued as indexed by anticipatory parieto-occipital alpha-band oscillations. Brain Research. Cognitive Brain Research, 12(1), 145–152.

Gottlob, L. R. (2004). Location cuing and response time distributions in visual attention. Perception & Psychophysics, 66(8), 1293–1302.

Grimault, S., Robitaille, N., Grova, C., Lina, J.-M., Dubarry, A.-S., & Jolicoeur, P. (2009). Oscillatory activity in parietal and dorsolateral prefrontal cortex during retention in visual short-term memory: additive effects of spatial attention and memory load. Human Brain Mapping, 30(10), 3378–3392. http://doi.org/10.1002/hbm.20759

Haegens, S., Nácher, V., Luna, R., Romo, R., & Jensen, O. (2011). α-Oscillations in the monkey sensorimotor network influence discrimination performance by rhythmical inhibition of neuronal spiking. Proceedings of the National Academy of Sciences, 108(48), 19377–19382. http://doi.org/10.1073/pnas.1117190108

Hamidi, M., Slagter, H. A., Tononi, G., & Postle, B. R. (2009). Repetitive Transcranial Magnetic Stimulation Affects behavior by Biasing Endogenous Cortical Oscillations. Frontiers in Integrative Neuroscience, 3, 14. http://doi.org/10.3389/neuro.07.014.2009

Hanslmayr, S., Aslan, A., Staudigl, T., Klimesch, W., Herrmann, C. S., & Bäuml, K.-H. (2007). Prestimulus oscillations predict visual perception performance between and within subjects. NeuroImage, 37(4), 1465–1473. http://doi.org/10.1016/j.neuroimage.2007.07.011

Hillyard, S. A., Hink, R. F., Schwent, V. L., & Picton, T. W. (1973). Electrical signs of selective attention in the human brain. Science (New York, N.Y.), 182(4108), 177–180.

Honkanen, R., Rouhinen, S., Wang, S. H., Palva, J. M., & Palva, S. (2014). Gamma Oscillations Underlie the Maintenance of Feature-Specific Information and the Contents of Visual Working Memory. Cerebral Cortex (New York, N.Y.: 1991). http://doi.org/10.1093/cercor/bhu263

Huang, J., & Sekuler, R. (2010). Attention protects the fidelity of visual memory: behavioral and electrophysiological evidence. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 30(40), 13461–13471. http://doi.org/10.1523/JNEUROSCI.2560-10.2010

Ikkai, A., Blacker, K. J., Lakshmanan, B. M., Ewen, J. B., & Courtney, S. M. (2014). Maintenance of Relational Information in Working Memory Leads to Suppression of the Sensory Cortex. Journal of Neurophysiology, jn.00134.2014. http://doi.org/10.1152/jn.00134.2014

Jensen, O., Bonnefond, M., & VanRullen, R. (2012). An oscillatory mechanism for prioritizing salient unattended stimuli. Trends in Cognitive Sciences, 16(4), 200–206. http://doi.org/10.1016/j.tics.2012.03.002

Jensen, O., Gelfand, J., Kounios, J., & Lisman, J. E. (2002). Oscillations in the Alpha Band (9–12 Hz) Increase with Memory Load during Retention in a Short-term Memory Task. Cerebral Cortex, 12(8), 877–882. http://doi.org/10.1093/cercor/12.8.877

Jensen, O., & Mazaheri, A. (2010). Shaping Functional Architecture by Oscillatory Alpha Activity: Gating by Inhibition. Frontiers in Human Neuroscience, 4. http://doi.org/10.3389/fnhum.2010.00186

Johnson, J. S., Sutterer, D. W., Acheson, D. J., Lewis-Peacock, J. A., & Postle, B. R. (2011). Increased Alpha-Band Power during the Retention of Shapes and Shape-Location Associations in Visual Short-Term Memory. Frontiers in Psychology, 2. http://doi.org/10.3389/fpsyg.2011.00128

Jokisch, D., & Jensen, O. (2007). Modulation of Gamma and Alpha Activity during a Working Memory Task Engaging the Dorsal or Ventral Stream. The Journal of Neuroscience, 27(12), 3244–3251. http://doi.org/10.1523/JNEUROSCI.5399-06.2007

Kelly, S. P., Lalor, E. C., Reilly, R. B., & Foxe, J. J. (2006). Increases in alpha oscillatory power reflect an active retinotopic mechanism for distracter suppression during sustained

visuospatial attention. Journal of Neurophysiology, 95(6), 3844–3851. http://doi.org/10.1152/jn.01234.2005

Klimesch, W. (2012). α-band oscillations, attention, and controlled access to stored information. Trends in Cognitive Sciences, 16(12), 606–617. http://doi.org/10.1016/j.tics.2012.10.007

Klimesch, W., Doppelmayr, M., Schwaiger, J., Auinger, P., & Winkler, T. (1999). "Paradoxical" alpha synchronization in a memory task. Brain Research. Cognitive Brain Research, 7(4), 493–501.

Klimesch, W., Sauseng, P., & Hanslmayr, S. (2007). EEG alpha oscillations: the inhibition-timing hypothesis. Brain Research Reviews, 53(1), 63–88. http://doi.org/10.1016/j.brainresrev.2006.06.003

Lakatos, P., Karmos, G., Mehta, A. D., Ulbert, I., & Schroeder, C. E. (2008). Entrainment of Neuronal Oscillations as a Mechanism of Attentional Selection. Science, 320(5872), 110–113. http://doi.org/10.1126/science.1154735

Lange, J., Keil, J., Schnitzler, A., van Dijk, H., & Weisz, N. (2014). The role of alpha oscillations for illusory perception. Behavioural Brain Research, 271, 294–301. http://doi.org/10.1016/j.bbr.2014.06.015

Leiberg, S., Lutzenberger, W., & Kaiser, J. (2006). Effects of memory load on cortical oscillatory activity during auditory pattern working memory. Brain Research, 1120(1), 131–140. http://doi.org/10.1016/j.brainres.2006.08.066

Linkenkaer-Hansen, K., Nikulin, V. V., Palva, S., Ilmoniemi, R. J., & Palva, J. M. (2004). Prestimulus oscillations enhance psychophysical performance in humans. The Journal of

Neuroscience: The Official Journal of the Society for Neuroscience, 24(45), 10186–10190. http://doi.org/10.1523/JNEUROSCI.2584-04.2004

Lopes da Silva, F. (1991). Neural mechanisms underlying brain waves: from neural membranes to networks. Electroencephalography and Clinical Neurophysiology, 79(2), 81–93.

Lopez-Calderon, J., & Luck, S. J. (2014). ERPLAB: an open-source toolbox for the analysis of event-related potentials. Frontiers in Human Neuroscience, 8. http://doi.org/10.3389/fnhum.2014.00213

Manza, P., Hau, C. L. V., & Leung, H.-C. (2014). Alpha power gates relevant information during working memory updating. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 34(17), 5998–6002. http://doi.org/10.1523/JNEUROSCI.4641-13.2014

Maris, E., van Vugt, M., & Kahana, M. (2011). Spatially distributed patterns of oscillatory coupling between high-frequency amplitudes and low-frequency phases in human iEEG. NeuroImage, 54(2), 836–850. http://doi.org/10.1016/j.neuroimage.2010.09.029

Mathewson, K. E., Gratton, G., Fabiani, M., Beck, D. M., & Ro, T. (2009). To see or not to see: prestimulus alpha phase predicts visual awareness. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 29(9), 2725–2732. http://doi.org/10.1523/JNEUROSCI.3963-08.2009

Mazaheri, A., Nieuwenhuis, I. L. C., van Dijk, H., & Jensen, O. (2009). Prestimulus alpha and mu activity predicts failure to inhibit motor responses. Human Brain Mapping, 30(6), 1791–1800. http://doi.org/10.1002/hbm.20763

Medendorp, W. P., Kramer, G. F. I., Jensen, O., Oostenveld, R., Schoffelen, J.-M., & Fries, P. (2007). Oscillatory activity in human parietal and occipital cortex shows hemispheric lateralization and memory effects in a delayed double-step saccade task. Cerebral Cortex (New York, N.Y.: 1991), 17(10), 2364–2374. http://doi.org/10.1093/cercor/bhl145

Mo, J., Schroeder, C. E., & Ding, M. (2011). Attentional Modulation of Alpha Oscillations in Macaque Inferotemporal Cortex. The Journal of Neuroscience, 31(3), 878–882. http://doi.org/10.1523/JNEUROSCI.5295-10.2011

Osipova, D., Takashima, A., Oostenveld, R., Fernández, G., Maris, E., & Jensen, O. (2006). Theta and gamma oscillations predict encoding and retrieval of declarative memory. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 26(28), 7523–7531. http://doi.org/10.1523/JNEUROSCI.1948-06.2006

Oostenveld, R., Fries, P., Maris, E., and Schoffelen, J-M. (2011). FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data. Computational Intelligence and Neuroscience, vol. 2011, Article ID 156869, 9 pages. doi:10.1155/2011/156869

Palva, J. M., Monto, S., Kulashekhar, S., & Palva, S. (2010). Neuronal synchrony reveals working memory networks and predicts individual memory capacity. Proceedings of the National Academy of Sciences, 107(16), 7580–7585. http://doi.org/10.1073/pnas.0913113107

Palva, S., Kulashekhar, S., Hämäläinen, M., & Palva, J. M. (2011). Localization of cortical phase and amplitude dynamics during visual working memory encoding and retention. The

Journal of Neuroscience : The Official Journal of the Society for Neuroscience, 31(13), 5013–5025. http://doi.org/10.1523/JNEUROSCI.5592-10.2011

Palva, S., & Palva, J. M. (2007). New vistas for alpha-frequency band oscillations. Trends in Neurosciences, 30(4), 150–158. http://doi.org/10.1016/j.tins.2007.02.001

Palva, S., & Palva, J. M. (2011). Functional Roles of Alpha-Band Phase Synchronization in Local and Large-Scale Cortical Networks. Frontiers in Psychology, 2. http://doi.org/10.3389/fpsyg.2011.00204

Payne, L., Guillory, S., & Sekuler, R. (2013). Attention-modulated alpha-band oscillations protect against intrusion of irrelevant information. Journal of Cognitive Neuroscience, 25(9), 1463–1476. http://doi.org/10.1162/jocn_a_00395

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

Pfurtscheller, G. (1992). Event-related synchronization (ERS): an electrophysiological correlate of cortical areas at rest. Electroencephalography and Clinical Neurophysiology, 83(1), 62–69.

Pfurtscheller, G., Stancák Jr., A., & Neuper, C. (1996). Event-related synchronization (ERS) in the alpha band — an electrophysiological correlate of cortical idling: A review. International Journal of Psychophysiology, 24(1–2), 39–46. http://doi.org/10.1016/S0167-8760(96)00066-9

Poch, C., Campo, P., & Barnes, G. R. (2014). Modulation of alpha and gamma oscillations related to retrospectively orienting attention within working memory. The European Journal of Neuroscience, 40(2), 2399–2405. http://doi.org/10.1111/ejn.12589

Prabhakaran, V., Narayanan, K., Zhao, Z., & Gabrieli, J. D. (2000). Integration of diverse information in working memory within the frontal lobe. Nature Neuroscience, 3(1), 85–90. http://doi.org/10.1038/71156

Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. Nature, 497(7451), 585–590. http://doi.org/10.1038/nature12160

Rihs, T. A., Michel, C. M., & Thut, G. (2007). Mechanisms of selective inhibition in visual spatial attention are indexed by α-band EEG synchronization. European Journal of Neuroscience, 25(2), 603–610. http://doi.org/10.1111/j.1460-9568.2007.05278.x

Romei, V., Brodbeck, V., Michel, C., Amedi, A., Pascual-Leone, A., & Thut, G. (2008). Spontaneous fluctuations in posterior alpha-band EEG activity reflect variability in excitability of human visual areas. Cerebral Cortex (New York, N.Y.: 1991), 18(9), 2010–2018. http://doi.org/10.1093/cercor/bhm229

Romei, V., Driver, J., Schyns, P. G., & Thut, G. (2011). Rhythmic TMS over parietal cortex links distinct brain frequencies to global versus local visual processing. Current Biology: CB, 21(4), 334–337. http://doi.org/10.1016/j.cub.2011.01.035

Rowe, J. B., Toni, I., Josephs, O., Frackowiak, R. S., & Passingham, R. E. (2000). The prefrontal cortex: response selection or maintenance within working memory? Science (New York, N.Y.), 288(5471), 1656–1660.

Sakai, K., Rowe, J. B., & Passingham, R. E. (2002). Active maintenance in prefrontal area 46 creates distractor-resistant memory. Nature Neuroscience, 5(5), 479–484. http://doi.org/10.1038/nn846

Sauseng, P., Klimesch, W., Heise, K. F., Gruber, W. R., Holz, E., Karim, A. A., … Hummel, F. C. (2009). Brain oscillatory substrates of visual short-term memory capacity. Current Biology: CB, 19(21), 1846–1852. http://doi.org/10.1016/j.cub.2009.08.062

Scheeringa, R., Mazaheri, A., Bojak, I., Norris, D. G., & Kleinschmidt, A. (2011). Modulation of visually evoked cortical FMRI responses by phase of ongoing occipital alpha oscillations. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 31(10), 3813–3820. http://doi.org/10.1523/JNEUROSCI.4697-10.2011

Scheeringa, R., Petersson, K. M., Oostenveld, R., Norris, D. G., Hagoort, P., & Bastiaansen, M. C. M. (2009). Trial-by-trial coupling between EEG and BOLD identifies networks related to alpha and theta EEG power increases during working memory maintenance. NeuroImage, 44(3), 1224–1238. http://doi.org/10.1016/j.neuroimage.2008.08.041

Snyder, A. C., & Foxe, J. J. (2010). Anticipatory attentional suppression of visual features indexed by oscillatory alpha-band power increases: a high-density electrical mapping study. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 30(11), 4024–4032. http://doi.org/10.1523/JNEUROSCI.5684-09.2010

Soto, D., & Blanco, M. J. (2004). Spatial attention and object-based attention: a comparison within a single task. Vision Research, 44(1), 69–81. http://doi.org/10.1016/j.visres.2003.08.013

Srinivasan,R., Winter,W.R., Ding,J., and Nunez,P.L. (2007). EEG and MEG coherence: measures of functional connectivity at distinct spatial scales of neo-cortical dynamics. J. Neurosci. Methods 166, 41–52. doi:10.1016/j.jneumeth.2007. 06.026

Sreenivasan, K. K., Vytlacil, J., & D'Esposito, M. (2014). Distributed and dynamic storage of working memory stimulus information in extrastriate cortex. Journal of Cognitive Neuroscience, 26(5), 1141–1153. http://doi.org/10.1162/jocn_a_00556

Thut, G., Veniero, D., Romei, V., Miniussi, C., Schyns, P., & Gross, J. (2011). Rhythmic TMS causes local entrainment of natural oscillatory signatures. Current Biology: CB, 21(14), 1176–1185. http://doi.org/10.1016/j.cub.2011.05.049

Todd, J. J., & Marois, R. (2004). Capacity limit of visual short-term memory in human posterior parietal cortex. Nature, 428(6984), 751–754. http://doi.org/10.1038/nature02466

Tuladhar, A. M., ter Huurne, N., Schoffelen, J.-M., Maris, E., Oostenveld, R., & Jensen, O. (2007). Parieto-occipital sources account for the increase in alpha activity with working memory load. Human Brain Mapping, 28(8), 785–792. http://doi.org/10.1002/hbm.20306

Van Der Werf, J., Jensen, O., Fries, P., & Medendorp, W. P. (2008). Gamma-band activity in human posterior parietal cortex encodes the motor goal during delayed prosaccades and antisaccades. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 28(34), 8397–8405. http://doi.org/10.1523/JNEUROSCI.0630-08.2008

Vanrullen, R., Busch, N. A., Drewes, J., & Dubois, J. (2011). Ongoing EEG Phase as a Trial-by-Trial Predictor of Perceptual and Attentional Variability. Frontiers in Psychology, 2, 60. http://doi.org/10.3389/fpsyg.2011.00060

Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. Nature, 428(6984), 748–751. http://doi.org/10.1038/nature02447

Von Stein, A., Chiang, C., & König, P. (2000). Top-down processing mediated by interareal synchronization. Proceedings of the National Academy of Sciences of the United States of America, 97(26), 14748–14753. http://doi.org/10.1073/pnas.97.26.14748

Voytek, B., Canolty, R. T., Shestyuk, A., Crone, N. E., Parvizi, J., & Knight, R. T. (2010). Shifts in gamma phase-amplitude coupling frequency from theta to alpha over posterior cortex during visual tasks. Frontiers in Human Neuroscience, 4, 191. http://doi.org/10.3389/fnhum.2010.00191

Warden, M. R., & Miller, E. K. (2010). Task-Dependent Changes in Short-Term Memory in the Prefrontal Cortex. Journal of Neuroscience, 30(47), 15801–15810. http://doi.org/10.1523/JNEUROSCI.1569-10.2010

Watson, D. G., & Humphreys, G. W. (1997). Visual marking: prioritizing selection for new objects by top-down attentional inhibition of old objects. Psychological Review, 104(1), 90–122.

Worden, M. S., Foxe, J. J., Wang, N., & Simpson, G. V. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 20(6), RC63.

Xu, Y., & Chun, M. M. (2006). Dissociable neural mechanisms supporting visual short-term memory for objects. Nature, 440(7080), 91–95. http://doi.org/10.1038/nature04262

Zanto, T. P., & Gazzaley, A. (2009). Neural suppression of irrelevant information underlies optimal working memory performance. The Journal of Neuroscience: The Official

Journal of the Society for Neuroscience, 29(10), 3059–3066. http://doi.org/10.1523/JNEUROSCI.4621-08.2009

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. Nature, 453(7192), 233–235. http://doi.org/10.1038/nature06860