THE UNIVERSITY OF CHICAGO

ALPHA-BAND OSCILLATIONS TRACK SPATIAL PRIORITY IN THE HUMAN BRAIN

A DISSERTATION SUBMITTED TO

THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES

IN CANDIDACY FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

DEPARTMENT OF PSYCHOLOGY

BY

JOSHUA JAMES FOSTER

CHICAGO, ILLINOIS

JUNE 2019

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF PUBLICATIONS

Sutterer, D.W.* Foster, J.J,* Adam K.C.S., Vogel, E.K., & Awh E. (2019). Item-specific delay activity demonstrates concurrent storage of multiple active neural representations in working memory. *PLOS Biology,* e3000239. doi:10.1371/journal.pbio.3000239

Foster, J.J. & Awh E. (2019). The role of alpha oscillations in spatial attention: limited evidence for a suppression account. *Current Opinion in Psychology*, *29*, 34-40. doi:10.1016/j.copsyc.2018.11.001 **This article is reproduced in part in Chapter 6.**

Sprague, T.C., Adam, K.C.S., Foster, J.J., Rahmati, M., Sutterer, D.W., & Vo, V.A. (2018). Inverted encoding models assay population level stimulus representations, not single-unit neural tuning. *eNeuro, 5*(3), 0098-18.2018. doi:10.1523/ENEURO.0098-18.2018

Van Moorselaar, D., Foster, J.J., Sutterer, D.W., Theeuwes, J., Olivers, C.N.L., & Awh, E. (2018). Spatially selective alpha oscillations reveal moment-by-moment trade-offs between working memory and attention. *Journal of Cognitive Neuroscience, 30*(2), 256-266. doi:10.1162/jocn_a_01198

Foster, J.J., Bsales, E.M., Jaffe, R.J., & Awh, E. (2017). Alpha-band activity reveals spontaneous representations of spatial position in visual working memory. *Current Biology,* 27(20), 31216-3223. doi:10.1016/j.cub.2017.09.031 **This article is reproduced in full as Chapter 4.**

Foster, J.J., Sutterer, D.W., Serences, J.T., Vogel, E.K., & Awh, E. (2017). Alpha-band oscillations enable spatially and temporally resolved tracking of covert spatial attention. *Psychological Science, 28*(7), 929-941. doi:10.1177/0956797617699167 **This article is reproduced in full as Chapter 2.**

Foster, J.J., Sutterer, D.W., Serences, J.T., Vogel, E.K., & Awh, E. (2016). The topography of alpha-band activity tracks the content of spatial working memory. *Journal of Neurophysiology, 115*(1), 168-177. doi:10.1152/jn.00860.2015 **This article is reproduced in full as Chapter 3.**

Grimshaw, G.M., Foster, J.J., & Corballis, P.M. (2014). Frontal and parietal asymmetries interact to predict attentional bias to threat. *Brain and Cognition, 90*, 76-86. doi:10.1016/j.bandc.2014.06.008

Foster, J.J.* & Adam, K.C.S.* (2014). Is feature-based attention spatially global during visual search? *Journal of Neuroscience, 34*(26), 8862-8864. doi:10.1523/JNEUROSCI.1616-14.2014

* Indicates that these authors contributed equally to the work

ACKNOWLEDGEMENTS

I am immensely grateful to everyone who has supported me on the road to completing this dissertation. Thank you to Gina Grimshaw, who inspired me to pursue a career in science, and whose mentorship ensured that I got the right start. To my graduate advisors, Ed Awh and Ed Vogel, I can't thank you enough for your guidance and support these past six years. I could not have hoped for a more rewarding and fun experience as a graduate student. I owe an especially big thank you to Ed Awh, my primary advisor, who has given me far more than my share of time and opportunities, and who has always challenged me to push the boundaries just a little further. I consider myself very lucky to be among your students. To my labmates, collaborators, conference buddies, and grad school friends (both in Eugene and Chicago): thank you for making me a better scientist and for making graduate school so much fun. I owe a special thanks to Dave Sutterer and Kirsten Adam, who were right there with me for almost all of grad school, including through a mid-Ph.D. move from University of Oregon to University of Chicago. I am truly lucky to be surrounded by such talented and supportive people. I also owe a big thank you also to the lab managers, research assistants, and other support staff for making this work possible. Thank you to my family for always being there for me. To my parents, Viv and Steve, thank you for giving me every opportunity, and supporting me in every pursuit, even when it takes me to the other side of the world. Last but not least, a heartfelt thank you to my wife, Natalie. Thank you for your unconditional love and support through this adventure abroad, and for all the sacrifices that you made to make it possible. I could not have done this without you. I look forward to many more adventures with you.

ABSTRACT

Spatial selection plays a central role in visual cognition, allowing us to prioritize processing at relevant locations. An emerging view is that alpha-band (8–12 Hz) oscillations play a key role in this core cognitive process. In this dissertation, I show that alpha-band oscillations track of spatial priority with remarkable spatial and temporal resolution, and across a range of contexts. To this end, I develop an encoding model approach to reconstruct spatially selective response profiles (called channel-tuning functions, or CTFs) from alpha-band power measured with EEG. These alpha-band CTFs reflect the spatial selectivity of the population-level activity that is measured with EEG. In Chapter 2, I show that alpha activity precisely tracks the locus of attention, and that the time course of this activity tracks trial-by-trial variations in the latency of covert orienting, establishing it as a powerful means of tracking the temporal dynamics of covert spatial attention. In Chapter 3, I show that alpha activity precisely encodes spatial positions held in working memory, consistent with the broad hypothesis that there is considerable functional overlap between attention and working memory. In Chapter 4, I show that alpha activity tracks the spatial position of non-spatial memoranda held in working memory, even when space is wholly irrelevant to the task, suggesting that spatial attention is also recruited during the maintenance of non-spatial features in working memory. Finally, having established a tight link between alpha oscillations and spatial selection, in Chapter 5 I examine the *consequences* of spatial attention. Here, I show that spatial attention increases the amplitude of population-level representations of stimulus position. Together, these studies establish that alpha-band oscillations enable spatially and temporally resolved tracking of spatial attention, and highlights the central role that spatial attention plays in a range of cognitive contexts.

CHAPTER 1. OVERVIEW

As we navigate our environment, we are confronted with a remarkable amount of visual information. However, our capacity to represent this information is limited (Rensink, Regan, & Clark, 2003; Tsubomi, Fukuda, Watanabe, & Vogel, 2013). Therefore, we must prioritize the processing of relevant information. Spatial attention plays a central role this effort, enabling us to prioritize information at relevant locations (Posner, 1980; Posner, Snyder, & Davidson, 1980). Often, we prioritize relevant locations *overtly*, by fixating the relevant location to bring it into high-resolution foveal vision (Posner, 1980). However, eye movements to relevant stimuli are not always possible: some stimuli come and go before the observer can initiate an eye movement, other times we must monitor multiple locations in the visual field at once (Awh & Pashler, 2000; Drew & Vogel, 2008; M. M. Müller, Malinowski, Gruber, & Hillyard, 2003). Situations like these require *covert* spatial attention – spatial selection in the absence of eye movements. There is considerable evidence that covert attention plays a central role in vision: covert spatial attention enhanced the fidelity and speed of processing at attended locations (Carrasco & McElree, 2001; Posner, 1980; for review, see Carrasco, 2011), and governs what stimuli we do and do not become aware of (Lavie, 2005; Mack & Rock, 1998).Thus, there is a strong motivation to understand the neural basis of this core cognitive process.

In this dissertation, I show that neural oscillations in the alpha frequency band (8–12 Hz) are intimately linked with spatial attention. I show that alpha activity precisely tracks *where* attention is deployed, and *when* it is deployed. Furthermore, I show that this selection mechanism is recruited in a range of contexts, including the selection of relevant locations in

perception, and the maintenance of information in working memory. Below, I outline the contributions of each of the chapters that follow.

Alpha-band oscillations precisely track covert spatial attention

Given the central role of covert attention in perception, there is strong motivation to understand the spatial and temporal dynamics of process (Egeth & Yantis, 1997). In service of this goal, one promising line of work has shown that oscillations in the alpha frequency band (8–12 Hz) are linked with covert spatial attention (for review, see Foxe & Snyder, 2011; Jensen & Mazaheri, 2010). Human electroencephalogram (EEG) studies have found that alpha oscillations track where attention is deployed. For example, when observers covertly attend one side of space, alpha-band power is reduced at posterior electrodes contralateral to the attended location (Kelly, Lalor, Reilly, & Foxe, 2006; Thut, Nietzel, Brandt, & Pascual-Leone, 2006; Worden, Foxe, Wang, & Simpson, 2000). More recent work suggests that the topographic distribution of alpha power tracks not just the attended hemifield but also covaries with the specific location in the visual field that is attended (Rihs, Michel, & Thut, 2007; Worden et al., 2000). In this dissertation, I build on this past work to show that alpha activity enables *spatially* and t*emporally* resolved tracking of covert spatial attention. To this end, I use the inverted encoding model (IEM) approach, which initially gained traction with functional MRI (fMRI) data (Brouwer & Heeger, 2009, 2011; Serences & Saproo, 2012; Sprague & Serences, 2013), to reconstruct spatially selective channel-tuning functions (CTFs) from the scalp distribution of alpha-band power. In Chapter 2, I use this approach to show that alpha-band oscillations not only precisely track where in the visual field attention is deployed, but also track trial-by-trial variations in the latency of covert orienting. This chapter establishes a powerful approach for tracking the spatial

and temporal dynamics of covert spatial attention, and provides new evidence for the tight link between alpha-band activity and spatial attention.

The role of alpha activity in the maintenance of online representations

Working memory (WM) is a system that allows us to hold relevant information in an online state .There is now substantial evidence for functional overlap between attention and WM (Awh & Jonides, 2001; Awh, Vogel, & Oh, 2006; Chun, 2011; Gazzaley & Nobre, 2012). For example, Awh and colleagues argued that covert spatial attention plays a key role in the maintenance of spatial locations in WM (Awh & Jonides, 2001). In support of this view, Awh and colleagues showed that holding locations in WM has the many of the same consequences as covertly attending those locations: stimuli that are presented at a remembered location are better processed (Awh, Jonides, & Reuter-Lorenz, 1998) and evoke a stronger neural response (Awh et al., 1999; Awh, Anllo-Vento, & Hillyard, 2000). Motivated by this view, in Chapter 3 I examine whether alpha activity tracks spatial positions held in WM as it does the locus of spatial attention. I show that alpha oscillations precisely track spatial positions maintained in WM. This finding provides new support for the functional overlap between attention and WM, and establishes spatially selective modulations of alpha-band activity as a common selection process that is recruited during selection of relevant locations in perception and during the maintenance of relevant locations in WM.

There is evidence that the contribution of spatial attention to WM is not limited to the maintenance of spatial information, but also support the maintenance of non-spatial information in WM. For example, in a color-WM task, Williams and colleagues (Williams, Pouget, Boucher, & Woodman, 2013) showed that observers tended to fixate the original locations of remembered

stimuli once the stimuli were no longer present, suggesting that observers spontaneously attend these locations even though they no longer contain relevant information (also see Kuo, Rao, Lepsien, & Nobre, 2009; Theeuwes, Kramer, & Irwin, 2011). Furthermore, Williams and colleagues found that diverting spatial attention away from the locations of the remembered stimuli disrupted memory for the remembered colors. In Chapter 4, I examine whether active representations of spatial position are maintained in non-spatial memory tasks in which spatial position was wholly irrelevant to the task. Here, too, I show that alpha activity tracks of the original stimulus position that persisted throughout the retention interval, even when spatial positon was not behaviorally relevant. This finding provides further evidence that spatial selection is also recruited during the maintenance of non-spatial information in WM. Thus, alpha-band oscillations reflect a general selection mechanism that operates in a range of cognitive contexts.

Spatial attention enhances population-level representations

Together, chapters 2-4 provide compelling evidence that alpha activity precisely tracks spatial selection in perception and during the maintenance of online representations. In Chapter 5, I address a distinct by related issue. In Chapter 5, I examine how spatial attention modulates population-level representations of visual stimuli (i.e., the *consequences* of spatial selection). Improved perception at attended locations is underpinned by improved neural coding of stimuli appearing at those locations (Maunsell, 2015; Sprague, Saproo, & Serences, 2015). Single-unit recordings in non-human primates have provided overwhelming evidence that spatial attention improves coding of stimulus properties at the level of individual neurons. For example, attention increases the amplitude of feature-selective tuning functions (e.g. McAdams & Maunsell, 1999).

However, perception is thought to be supported by the joint activity of populations of neurons (Saxena & Cunningham, 2019; Sprague et al., 2015), providing a strong motivation to understand how attention modulates population codes. In recent years, researchers have used the encoding model approach combined with fMRI data to examine how attention modulates population-level stimulus representations (e.g. Kay, Weiner, & Grill-Spector, 2015; Saproo & Serences, 2014; Sprague & Serences, 2013; Vo, Sprague, & Serences, 2017). However, with fMRI it has proven difficult to disentangle stimulus representations encoded by stimulus-driven neural activity from top-down selection signals that are independent of the stimulus-driven response (see Chapter 5 for further discussion). In Chapter 5, I apply the IEM approach to stimulus-evoked EEG signals to examine how attention modulates population-level representations of stimulus position. I show that attention enhances population-level representations of stimulus position by increasing the amplitude of these representations.

CHAPTER 2. ALPHA-BAND OSCILLATIONS ENABLE SPATIALLY AND TEMPORALLY RESOLVED TRACKING OF COVERT SPATIAL ATTENTION

Introduction

A typical visual scene contains more information than an observer can process at once. Therefore, the observer must focus limited processing resources on the most relevant aspects of the environment. Spatial attention plays a central role in this effort, enhancing the quality and speed of processing at attended locations (Carrasco & McElree, 2001; Eriksen & Hoffman, 1974; Posner, 1980; for review, see Carrasco, 2011). Because spatial attention is essential for normal perceptual function, there is great interest in understanding the neural basis of this process. One promising approach has been to examine the links between attentional states and rhythmic brain activity. A growing body of evidence suggests that oscillatory activity in the alpha frequency band (8–12 Hz) is integral to covert spatial attention. Measurements of the topographic distribution of alpha power across the scalp have revealed that alpha power is reduced contralateral to an attended location (e.g. Kelly et al., 2006; Sauseng et al., 2005; Thut et al., 2006). Further work has shown that the topography of alpha power tracks not just the hemifield but also the specific location that an observer is attending (e.g., Bahramisharif et al., 2010; Rihs et al., 2007; Worden et al., 2000). These findings suggest that spatially specific alpha-band activity directly tracks the deployment of spatial attention (Foxe & Snyder, 2011; Jensen & Mazaheri, 2010).

Nevertheless, the hypothesis that alpha-band activity is integral to spatial attention makes a clear prediction that remains untested: The topography of alpha-band activity should track not only the spatial locus of attention but also the time course of covert orienting. Extant studies

have not provided a rigorous analysis of the time course of spatially specific alpha activity or examined whether the time course of this activity tracks variations in the latency of covert spatial orienting. Thus, our goal was to determine whether dynamic changes in alpha-band activity provide a sensitive index of the speed of covert spatial orienting.

To this end, we used electroencephalography (EEG) recordings and an inverted encoding model (IEM; Brouwer & Heeger, 2011, 2009; Sprague & Serences, 2013; for review, see Sprague et al., 2015) to examine the time course of spatially specific alpha-band activity. This approach assumes that the topographic pattern of alpha power across the scalp reflects the activity of a number of underlying spatially tuned channels (or neuronal populations; Figure 2-1a). By first estimating the relative contributions of these channels to each electrode on the scalp (Figure 2-1b), the model can then be *inverted* so that the underlying response of these spatial channels can be estimated from the pattern of alpha power across the scalp (Figure 2-1c). The resulting profile of responses across the spatial channels (termed *channel-tuning functions*, or CTFs) reflects the spatial tuning of population-level alpha power, as measured with EEG. Thus, the IEM approach enables a straightforward quantification of spatially selective activity from a higher-dimensional pattern of alpha power on the scalp. By performing this analysis across separate points in time, we were able to examine the temporal dynamics of spatially selective alpha-band activity (also see J. J. Foster, Sutterer, Serences, Vogel, & Awh, 2016; Samaha, Sprague, & Postle, 2016).

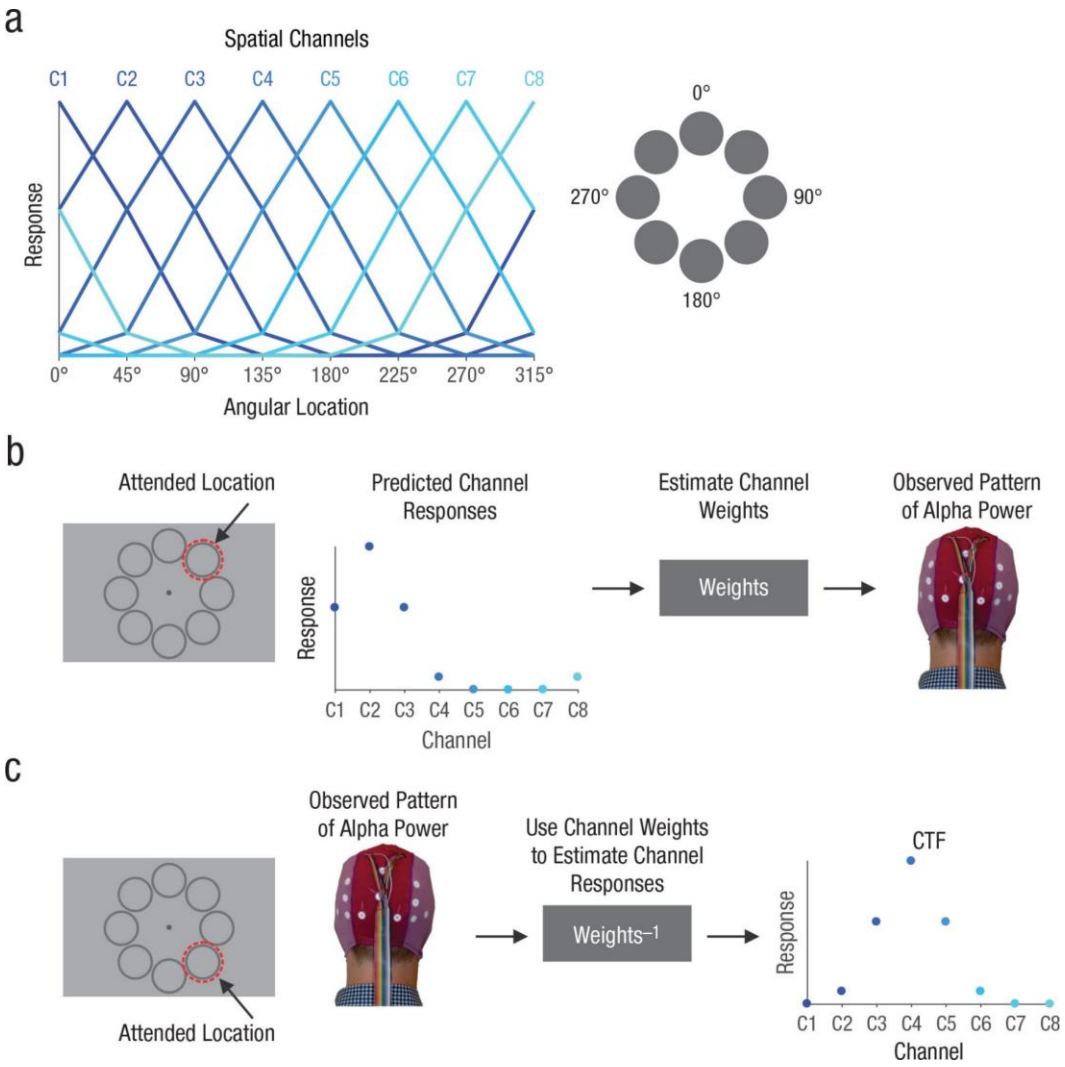

Figure 2-1. Illustration of the inverted encoding method for reconstructing spatial channel-tuning functions (CTFs) from the pattern of alpha-band power across the scalp.
We modeled alpha power measured at each electrode as the weighted sum of eight spatially tuned channels (C1–C8). Each curve in (a) shows the predicted response of one of the channels across eight possible attended angular locations (the locations shown here were used in Experiment 2-1; in Experiment 2-2, we used 22.5°, 67.5°, 112.5°, etc.). The gray circles on the right indicate how locations were labeled. In the training phase (b), we used predicted channel responses to estimate a set of channel weights that characterized the relative contribution of each of the spatial channels to the response measured at each of the scalp electrodes. The example shown here is for an attended location at 45°. In the test phase (c), using an independent set of data, we used the channel weights determined from the training data to estimate the channel responses from the observed pattern of alpha power on the scalp. The resulting CTF reflects the spatial selectivity of population-level alpha power, as measured by electroencephalography (EEG). The example shown here is for an attended location at 135°. For more details, see the Inverted Encoding Model section.

In two experiments, we tested whether the topographic distribution of alpha-band activity tracks the time course of covert orienting. In Experiment 2-1, subjects performed a spatial-cueing task, which allowed us to examine the time course of alpha-band CTFs as subjects shifted covert attention to the cued location following an attention-directing cue. In Experiment 2-2, subjects performed a visual search task, in which we manipulated the latency of covert orienting toward the target by varying search difficulty. This design allowed us to directly test whether the time course of alpha-band CTFs tracked differences in the latency of target selection across different levels of search difficulty and as a function of within-subject variations in orienting latency across trials, as indexed by reaction times.

General Method

Subjects

Fifty volunteers (20 in Experiment 2-1, 30 in Experiment 2-2[1]) participated in the experiments for monetary compensation ($10 per hr). Subjects were between 18 and 35 years old, reported normal or corrected-to-normal visual acuity, and provided informed consent according to procedures approved by the institutional review board at the University of Oregon.

In Experiment 2-1, we did not analyze data from subjects who provided fewer than 700 artifact-free trials (i.e., trials that were not contaminated by recording or ocular artifacts). This exclusion criterion was set during data collection and was chosen on the basis of our work using the IEM method to track locations stored in working memory (J. J. Foster et al., 2016). Two subjects were excluded because of excessive artifacts, and two subjects were excluded because

[1] Three additional volunteers completed the practice session for Experiment 2-2 (see Procedures) but did not return for the EEG session.

of an error with stimulus presentation. Thus, the final sample included a total of 16 subjects.[2] All subjects in the final sample provided data for at least 700 artifact-free trials (M = 1165, SD = 173).

In Experiment 2-2, we did not analyze data from subjects who provided fewer than 600 artifact-free trials (with correct responses) for each condition. We relaxed the exclusion criterion in Experiment 2-2 because we obtained fewer trials per condition, because Experiment 2-2 included two conditions rather than one. The exclusion criterion was determined during the course of data collection but before the data were analyzed. Seven subjects were excluded because of excessive artifacts, which left a total of 23 subjects.[3] All subjects in the final sample provided data for at least 600 trials per search condition (M = 772, SD = 79) after artifacts and incorrect responses were discarded.

Apparatus and Stimuli

We tested the subjects in a dimly lit, electrically shielded chamber. Stimuli were generated using MATLAB (The Mathworks, Natick, MA) and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) and were presented on a 17-in. CRT monitor (refresh rate = 60 Hz) at a viewing distance of approximately 100 cm. Stimuli were rendered in dark gray against a medium-gray background.

---

[2] For Experiment 2-1, our target sample size was 16 subjects, in keeping with our previous work using the method we used here to track locations held in spatial working memory (J. J. Foster et al., 2016).

[3] For Experiment 2-2, our target sample was a minimum of 20 subjects. Our target sample size was larger for Experiment 2-2 than for Experiment 2-1 because we had not run comparable tests for latency differences in previous work. Our lab was soon to relocate at the time of data collection. Thus, we continued data collection beyond our minimum sample until we no longer had access to the apparatus.

Procedure

After providing informed consent, the subjects were fitted with a cap embedded with 20 scalp electrodes before completing the experimental task. Including preparation time and experimental time, Experiment 2-1 took approximately 3 hr to complete, and Experiment 2-2 took approximately 3.5 hr to complete.

Experiment 2-1: spatial cueing task. Subjects in Experiment 2-1 completed a spatial-cuing task in which they were required to identify a target digit among distractor letters (Figure 2-2a). Subjects initiated each trial by pressing the space bar. Each trial began with a central fixation point (0.24° in diameter), surrounded by equally spaced placeholder rings (1.7° in diameter, with a border of 0.08°). Each placeholder was centered 2.4° from the fixation point. The exact angular position of the placeholders were jittered on each trial within a 45° wedge. Thus, the position of the first placeholder varied between −22.5° and 22.5°, the second varied between 22.5° and 67.5°, and so on. This jitter was not necessary for the IEM analysis.

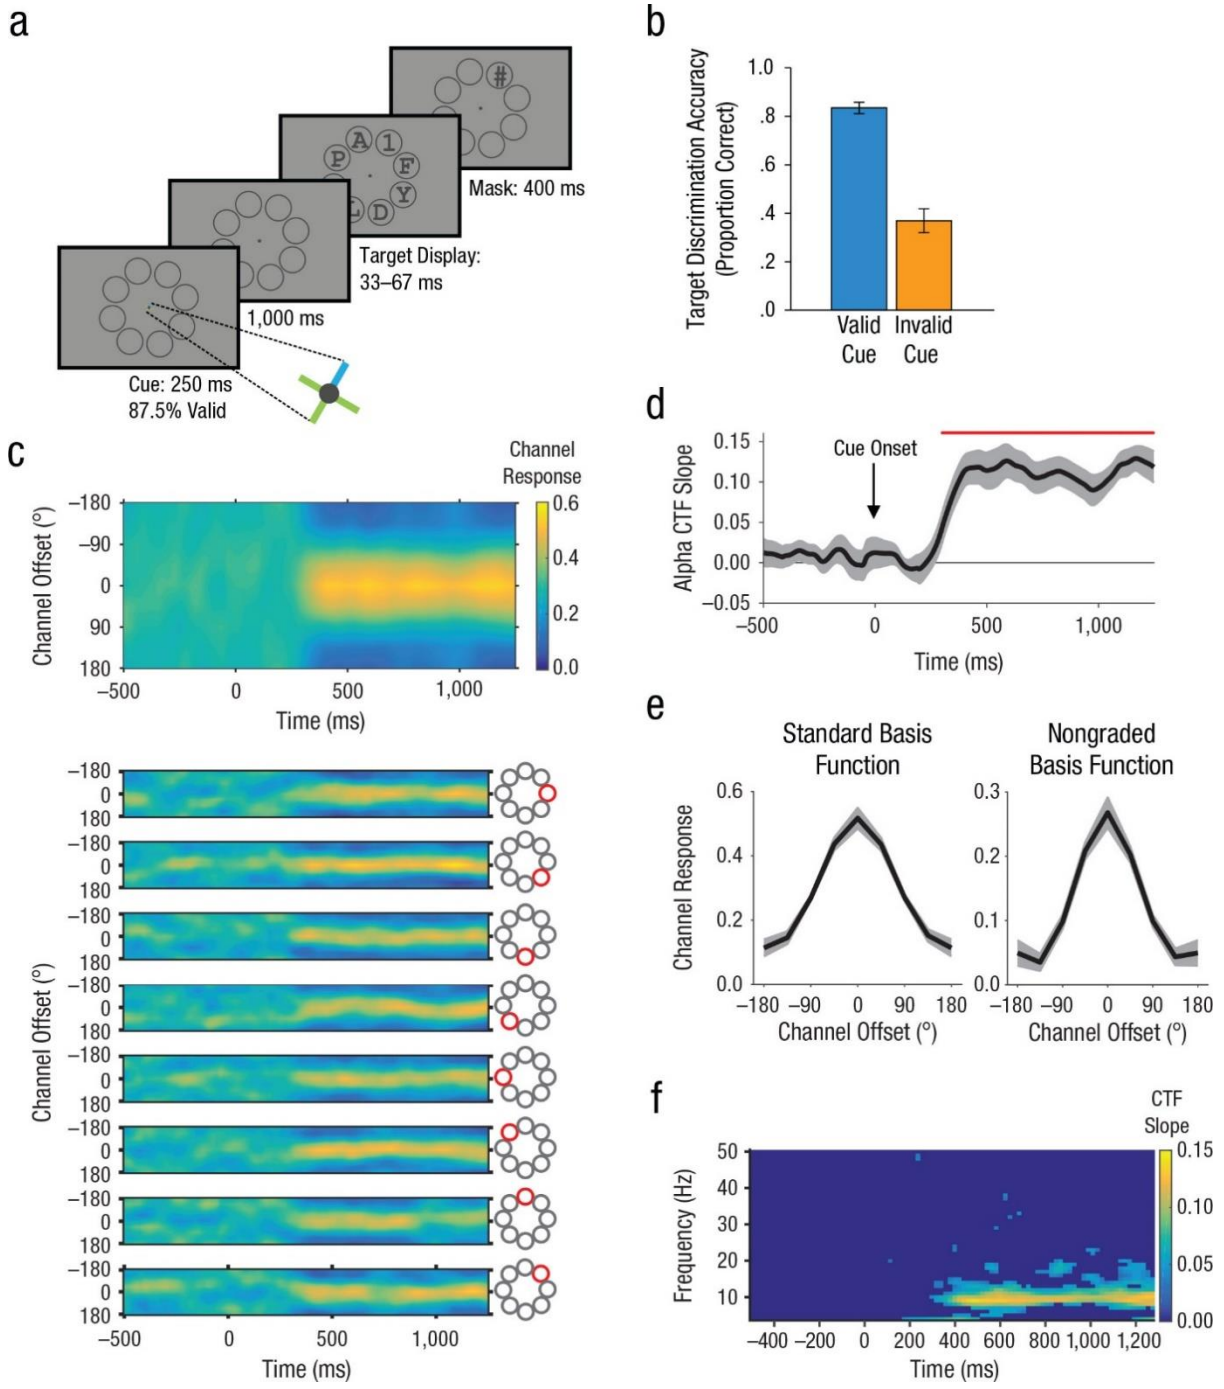

Figure 2-2. Task and results from Experiment 2-1.
In the spatial-cuing task (a), a central cross (cue) with three arms of one color and one arm of a different color directed the subjects to attend one of eight placeholders. After a 1,000-ms delay, the target digit was displayed among distractor letters and then masked with a pound sign. The graph in (b) shows target-discrimination accuracy (proportion correct) separately for trials with valid and invalid cues. The top graph in (c) shows the average alpha-band channel-tuning

Figure 2-2, continued..
function (CTF) across time for all eight locations. The yellow band shows the peak channel response. The eight graphs below that show the average CTF for each of the eight cued locations separately. In (d), the graph shows the average slope of the alpha-band CTF as a function of time. The red marker indicates the area of reliable CTF selectivity. The shaded areas indicate ±1 bootstrapped SEM. The graphs in (e) show the channel-response profile recovered using the standard graded basis function (left) and a nongraded basis function (right). The shaded areas indicate ±1 bootstrapped SEM. The slopes of the CTFs reconstructed from the topographic distribution of oscillatory power across a broad range of frequencies (4–50 Hz, in increments of 1 Hz) are plotted in (f). Points at which the CTF slope was not reliably above zero as determined by a permutation test are set to zero (dark blue).

After a variable interval between 800 and 1,500 ms, a central cue (87.5% valid), presented for 250 ms, indicated the likely location of a subsequent target. The cue was a small cross (0.6° wide; arms were 0.08° thick) with three green arms and one blue arm (or vice versa, counterbalanced across subjects). The uniquely colored arm of the cue pointed toward the cued placeholder. In a target display presented 1,250 ms after cue onset, each placeholder was occupied by a letter or digit. The display included one target (a digit between 1 and 9) among distractors (uppercase letters). Digits and letters were approximately 0.9° tall and 0.8° wide. The distractor letters were randomly selected without replacement from all possible letters (except for "I," "S," and "Z" because of their similarity to the digits "1," "5," and "2," respectively). The target was backward-masked with a pound symbol ("#") presented for 400 ms. Following the mask, the subjects reported the target identity using the number pad on a standard keyboard; there was no time limit. The reported digit appeared ~1° above the fixation point, and the subjects could correct their response if they pressed a wrong key. Finally, the subjects confirmed their response by pressing the space bar.

To encourage the subjects to attend the cued location in advance of the target display, we adjusted the duration of the target display for each subject using a staircase procedure. Subjects

completed one or two blocks (72 trials per block) of this procedure at the start of the session. During the staircase procedure, the cue was valid on all trials, and the subjects were instructed to attend the cued location. Exposure duration was decreased by 16.7 ms (i.e., one refresh cycle at 60 Hz) when the subjects made a correct response or increased by 33.3 ms (i.e., two refresh cycles at 60 Hz) when the subjects made an incorrect response, until performance reached an asymptote. The resulting duration of the target display varied between 33.3 ms and 66.7 ms across subjects. This staircase procedure was somewhat coarse because changing exposure duration by 16.7 ms had a considerable effect on task difficulty. Nevertheless, this procedure ensured that target identification was adequately difficult for all the subjects: Target identification accuracy ranged be-tween 65.7% and 97.0% (M = 83.5%, SD = 9.5) on trials with valid cues, and all subjects showed a large spatial-cuing effect, which ranged between 26.6% and 69.8% (M = 46.6%, SD = 13.0).

After the staircase procedure, the subjects completed as many blocks of the spatial-cuing task as time permitted, but the goal was 20 blocks. Each block contained 72 trials, and all subjects completed at least 13 blocks. Each of the eight placeholders was cued equally often within each block of trials.

Experiment 2-2: visual search task. Subjects in Experiment 2-2 performed a visual search task in which they searched for a target (a vertical or horizontal bar) among distractors (Figure 2-3a). Each item in the search display consisted of a dark gray bar (1.5° × 0.2°) superimposed on a gray circle (2.1° in diameter). The items were equally spaced in a circle around a dark gray fixation point (0.2° in diameter). Each item was centered 3° from the fixation point. Stimulus positions were not jittered in Experiment 2-2.

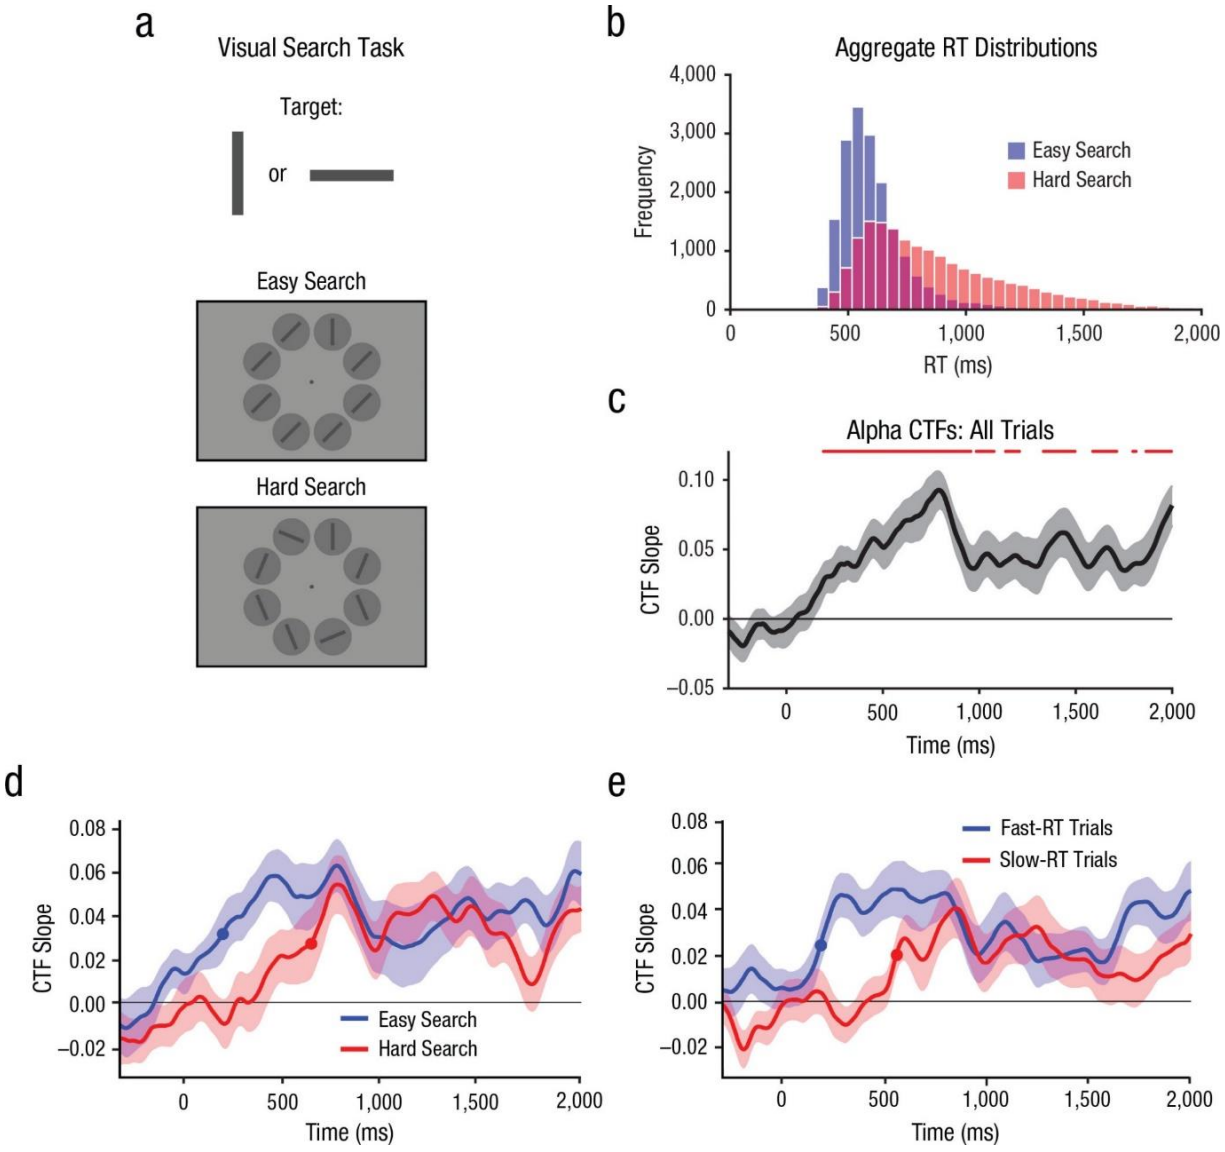

Figure 2-3. Task and results from Experiment 2-2.
Subjects searched for a target (a vertical or horizontal bar) among distractors and reported the target's orientation (a). The histogram in (b) shows the distributions of response times (RT) for easy and hard search. In (c), the plot shows the selectivity of the average target-related alpha-band channel-tuning function (CTF) across time, collapsed across the search conditions (easy and hard). Time points at which CTF selectivity was reliable are indicated by the red marker. The plots in (d) and (e) show the average target-related CTF slope across time for (d) easy and hard search and (e) for fast- and slow-RT trials (regardless of search condition). The filled circles mark the points at which the target-related CTF reached the onset criterion (50% of maximum amplitude). The shaded areas in (c), (d), and (e) represent ±1 bootstrapped SEM.

We varied the difficulty of visual search by manipulating both distractor variability (i.e., distractor orientation was uniform or varied) and target-distractor similarity (i.e., the extent to which the distractors resembled the target). In the easy-search condition, all distractors were identical and were rotated 45° clockwise or counter clockwise from the possible target orientations (Figure 2-3a). Thus, distractor variability and target-distractor similarity were low. In the hard-search condition, the distractors were heterogeneous and were rotated 22.5° from the possible target orientations (Figure 2-3a). Thus, distractor variability and target-distractor similarity were higher than in the easy-search condition, resulting in a more difficult search (Duncan & Humphreys, 1989).

Each search array was presented for 2 s, separated by a variable intertrial interval between 1.8 and 2.3 s, during which only the fixation point remained visible. Subjects reported whether the target was vertical or horizontal by pressing the "z" key (left index finger) or "/" key (right index finger), respectively. Subjects were instructed to respond as quickly and as accurately as possible. Feedback (mean response time, or RT, and accuracy) was provided at the end of each block of trials. To minimize artifacts during the stimulus display and a 300-ms prestimulus baseline period, the subjects were instructed to maintain fixation throughout each block of the search task and to blink (if necessary) immediately after the offset of the search array.

The subjects completed 30 blocks of 64 trials each. The search conditions (easy or hard) were blocked, and the blocks alternated between conditions. The order of the conditions (easy first or hard first) was counterbalanced across subjects. Before beginning the session, the subjects completed two blocks of practice trials (easy search followed by hard search).

The subjects also completed a short practice session the day before the EEG session to familiarize themselves with the visual search task. During this session, the subjects completed three blocks of the easy-search condition followed by three blocks of the hard-search condition.

Electrophysiology

EEG was recorded using 20 tin electrodes mounted in an elastic cap (Electro-Cap International, Eaton, OH). We recorded from International 10/20 sites F3, FZ, F4, T3, C3, CZ, C4, T4, P3, PZ, P4, T5, T6, O1, and O2, along with five nonstandard sites: OL midway between T5 and O1, OR midway between T6 and O2, PO3 midway between P3 and OL, PO4 midway between P4 and OR, and POz midway between PO3 and PO4. All sites were recorded with the left mastoid as a reference; they were re-referenced off-line to the algebraic average of the left and right mastoids. To detect horizontal eye movements, we used horizontal electrooculography (EOG) recorded from electrodes placed approximately 1 cm from the external canthus of each eye. To detect blinks and vertical eye movements, we recorded vertical EOG from an electrode placed below the right eye and referenced to the left mastoid. The EEG and EOG were amplified using an SA Instrumentation (San Diego, CA) amplifier with a band-pass filter of 0.01 to 80 Hz and were digitized at 250 Hz using LabVIEW 6.1 (National Instruments, Austin, TX) running on a PC. Trials were visually inspected for artifacts, and we discarded trials (both EEG and behavioral data) contaminated by blocking (i.e., amplifier saturation), blinks, detectable eye movements, excessive muscle noise, or skin potentials. Removal of ocular artifacts was effective: Variation in the grand-averaged horizontal EOG waveforms by cued and target locations was less than 3 μV. Given that eye movements of about 1° of visual angle produce a deflection in the horizontal EOG of approximately 16 μV (Lins, Picton, Berg, & Scherg, 1993a),

the residual variation in the average horizontal EOG corresponds to variations in eye position of less than 0.2° of visual angle (i.e., smaller than the size of the fixation point).

Time-Frequency Analysis

Time-frequency analyses were performed using the Signal Processing toolbox and EEGLAB toolbox (Delorme & Makeig, 2004) for MATLAB (The Mathworks, Natick, MA). To isolate frequency-specific activity, we band-pass-filtered the raw EEG signal using a two-way least-squares finite-impulse-response filter ("eegfilt.m" from EEGLAB Toolbox; Delorme & Makeig, 2004). This filtering method uses a zero-phase forward and reverse operation, which ensures that phase values are not distorted, as can occur with forward-only filtering methods. A Hilbert Transform (Matlab Signal Processing Toolbox) was applied to the bandpass-filtered data, producing the complex analytic signal, $z(t)$, of the filtered EEG, $f(t)$:

$$z(t) = f(t) + i\tilde{f}(t)$$

where $\tilde{f}(t)$ is the Hilbert Transform of $f(t)$, and $i = \sqrt{-1}$. The complex analytic signal was extracted for each electrode using the following Matlab syntax:

hilbert(eegfilt(data,F,f1,f2)')')

In this syntax, data is a 2-D matrix of raw EEG (number of trials × number of samples), *F* is the sampling frequency (250 Hz), *f*1 is the lower bound of the filtered frequency band, and *f*2 is the upper bound of the filtered frequency band. For alpha-band analyses, we used an 8- to 12-Hz band-pass filter; thus, *f*1 and *f*2 were 8 and 12, respectively. For the time-frequency analysis, we searched a broad range of frequencies (4–50 Hz, in increments of 1 Hz with a 1-Hz band pass). For these analyses, *f*1 and *f*2 were 4 and 5 to isolate 4- to 5-Hz activity, 5 and 6 to isolate 5- to 6-

Hz activity, and so forth. Instantaneous power was computed by squaring the complex magnitude of the complex analytic signal.

Inverted Encoding Model

In keeping with our previous work on spatial working memory (J. J. Foster et al., 2016), we used an IEM to reconstruct location-selective CTFs from the topographic distribution of oscillatory power across electrodes. We assumed that power measured at each electrode reflected the weighted sum of eight spatial channels (i.e., neuronal populations), each tuned for a different angular location (see Brouwer & Heeger, 2009; Sprague & Serences, 2013). We modeled the response profile of each spatial channel across angular locations as a half sinusoid raised to the seventh power:

$$R = sin(0.5\theta)^7$$

where θ is angular location (ranging from 0° to 359°) and R is the response of the spatial channel in arbitrary units. This response profile was shifted circularly for each channel such that the peak response of each spatial channel was centered over one of the eight locations (corresponding to the cued locations 0°, 45°, 90°, etc., for Experiment 2-1 and the target locations 22.5°, 67.5°, 112.5°, etc., for Experiment 2-2; see Figure 2-1a).

An IEM routine was applied to each time point in the alpha-band analyses and each time-frequency point in the time-frequency analysis.[4] We partitioned our data into independent sets of training data and test data (for details, see the Training and Test Data section). The routine

[4] For the time-frequency analysis, the IEM was applied across many frequency bands. To reduce computation time, we down- sampled power values from 250 Hz to 50 Hz (i.e., one sample every 20 ms). We down-sampled power values after filtering and applying the Hilbert transform so that down-sampling did not affect how power values were obtained.

proceeded in two stages (training and test). In the training stage (Figure 2-1b), the training data (*B1*) were used to estimate weights that approximated the relative contributions of the eight spatial channels to the observed response (i.e., oscillatory power) measured at each electrode. We define *B1* (*m* electrodes × *n1* measurements) as a matrix of the power at each electrode for each measurement in the training set, *C1* (*k* channels × *n1* measurements) as a matrix of the predicted response of each spatial channel (specified by the basis function for that channel) for each measurement, and *W* (*m* electrodes × *k* channels) as a weight matrix that characterizes a linear mapping from channel space to electrode space. The relationships among *B1*, *C1*, and *W* can be described by a general linear model of the following form:

$$B_1 = WC_1$$

The weight matrix was obtained via least squares estimation was as follows:

$$\widehat{W} = B_1 {C_1}^T \left(C_1 {C_1}^T\right)^{-1}$$

In the test stage (Figure 2-1c), we inverted the model to transform the test data, *B2* (*m* electrodes × *n2* measurements), into estimated channel responses, $\widehat{C_2}$ (*k* channels × *n2* measurements), using the estimated weight matrix, $\widehat{W}$, that we obtained in the training phase:

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

Each estimated channel-response function was circularly shifted to a common center, so that the center channel was the channel tuned for the cued or target location (i.e., 0° on the channel offset axes of Figure. 2-2c and Figure 2-2e). We then averaged these shifted channel-response functions to obtain the CTF averaged across the eight cued (or target) locations. The IEM routine was performed separately for each time point.

Finally, because the exact contributions of the spatial channels to electrode responses (i.e., the channel weights, *W*) were expected to vary by subject, the IEM routine was applied separately for each subject, and statistical analyses were performed on the reconstructed CTFs. This approach allowed us to disregard differences in how location-selective activity was mapped to scalp-distributed patterns of power across subjects and instead focus on the profile of activity in the common stimulus, or information, space (J. J. Foster et al., 2016; Sprague et al., 2015).

Training and Test Data

For the IEM procedure, we partitioned artifact-free trials for each subject into independent sets of training data and test data. Specifically, we divided the trials into three sets. For each of these sets, we averaged power across trials for each cued location (Experiment 2-1) or target location (Experiment 2-2), which resulted in three 20 (electrodes) × 8 (locations) matrices of power values, one for each set. We used a leave-one-out cross-validation routine such that two of these matrices served as the training data (*B1*, 20 electrodes × 16 measurements), and the remaining matrix served as the test data (*B2*, 20 electrodes × 8 measurements). Because no trial belonged to more than one of the three sets, the training and test data were always independent. We applied the IEM routine using each of the three matrices as the test data, and the remaining two matrices as the training data. The resulting CTFs were averaged across the three test sets.

When we partitioned the trials into three sets, we constrained the assignment of trials to the sets so that all eight locations in all three sets had the same number of trials. To that end, we determined the minimum number of trials per subject for any location, *n*, and assigned *n*/3 trials for each location to each set. For example, if *n* was 100, we assigned 33 trials for each location to

each set. Because of this constraint, some excess trials did not belong to any block. In Experiment 2-2, we compared CTFs across search conditions (easy or hard) and across trials with fast and slow RTs. To obtain a CTF for each condition separately, we partitioned each condition into three sets of data as described, which resulted in six sets in total (three for each condition). We constrained the assignment of trials as in Experiment 2-1. Note that this ensured that the same number of trials was used for each of the conditions.

We used an iterative approach to make use of all available trials. For each iteration, we randomly partitioned the trials into three sets (as just described) and performed the IEM routine on the resulting training and test data. We repeated this process of partitioning trials into sets multiple times (5 times for the full time-frequency analyses, 10 times for the alpha-band analyses, and 100 times for the latency analyses in Experiment 2-2). For each iteration, the subset of trials that were assigned to blocks was randomly selected. Therefore, the trials that were not included in any block were different for each iteration. We averaged the resulting channel-response profiles across iterations. This iterative approach reduced noise in the resulting CTFs by minimizing the influence of idiosyncrasies that were specific to any given assignment of trials to blocks.

Statistical Analysis

Quantifying CTF selectivity. To quantify the location selectivity of CTFs, we used linear regression to estimate CTF slope (i.e., slope of channel response as a function of location channels after collapsing across channels that were equidistant from the channel tuned to the location of the evoking stimulus). Higher CTF slope indicates greater location selectivity.

Permutation test. In Experiment 2-1, to test whether CTF selectivity was reliably above chance, we tested whether CTF slope was greater than zero using a one-sample *t* test. Because mean CTF slope may not be normally distributed under the null hypothesis, we used a Monte Carlo randomization procedure to empirically approximate the null distribution of the *t* statistic. Specifically, we implemented the IEM as described earlier but randomized the location labels within each block so that the labels were random with respect to the observed responses in each electrode. This randomization procedure was repeated 1,000 times to obtain a null distribution of *t* statistics. To test whether the observed CTF selectivity was reliably above chance, we calculated the probability of obtaining (from the surrogate null distribution) a *t* statistic that was greater than or equal to the observed *t* statistic (i.e., the probability of a Type I error). Our permutation test was therefore a one-tailed test. CTF selectivity was deemed reliably above chance if the probability of a Type I error was less than .01. This testing procedure was applied to each time-frequency point in the time-frequency analyses and to each time point in the alpha-band analyses.

Jackknife test for latency differences. In Experiment 2-2, we tested for differences in CTF onset latency between easy-search and hard-search trials and between trials with fast RTs and trials with slow RTs. We used a jackknife-based procedure (Miller, Patterson, & Ulrich, 1998) to test for latency difference in CTF onset. CTF onset latency was measured as the earliest time at which CTF slope reached 50% of its maximum amplitude. The latency difference between conditions, *D*, was measured as the difference in onset latency between conditions in the time courses of the CTF slopes averaged across subjects. We used a jackknife procedure (Miller et al., 1998) to estimate the standard error of the latency difference, SE_D, from the latency

differences obtained for subsamples that included all but one subject. Specifically, the latency differences, D_{-i} (for $i = 1, \ldots, N$, where N is the sample size), were calculated where D_{-i} was the latency difference for the sample with all subjects except for subject i. The jackknife estimate of the was calculated as

$$SE_D = \sqrt{\frac{N-1}{N}\sum_{i=1}^{N}(D_{-i} - \bar{J})^2}$$

where $\bar{J}$ is the mean of the differences obtained for all subsamples (i.e., $\bar{J} = \sum D_{-i}/N$).

A jackknifed t statistic, t_j, was then calculated as

$$t_j = \frac{D}{SE_D}$$

which follows an approximate t distribution with $N - 1$ degrees of freedom under the null hypothesis. Our jackknife tests for latency differences in CTF onsets were one-tailed because we had clear directional hypotheses: that CTF onset would be delayed for hard search compared with easy search and for trials with slow RTs compared with trials with fast RTs. The jackknife approach to testing for latency differences between conditions circumvents the need to calculate latency differences for individual subjects, which are often noisy because of the low signal-to-noise ratio of EEG data (Miller et al., 1998).

Results

Experiment 2-1

Subjects in Experiment 2-1 performed a spatial-cuing task in which they identified a target digit among distractor letters (Figure 2-2a). A central cue indicated the likely location of the target (valid on 87.5% of trials). We observed a robust spatial-cuing effect in the accuracy of target discrimination (Figure 2-2b): Target-discrimination accuracy was higher on trials with a

valid cue (M = 83.5%, SD = 9.5) than on those with an invalid cue (M = 36.9%, SD = 18.2), $t(15)$ = 14.33, $p < .001$, Cohen's d_z = 3.58.

Having established that the subjects attended the cued location, we tested whether the topography of alpha-band activity tracked shifts of covert attention to the cued location (collapsing across trials with valid and invalid cues). Using an IEM, we reconstructed spatial CTFs from the scalp distribution of alpha power. A spatially selective CTF emerged several hundred milliseconds after cue onset and was sustained until the search array was presented (Figure 2-2c, top). Figure 2-2d shows CTF selectivity across time (quantified as CTF slope; see General Method). A permutation test revealed that CTF selectivity was reliably above chance beginning 304 ms after cue onset. Therefore, covert attention must have been shifted to the cued location by this time at the latest. The time course of the attention-related CTF dovetails with past behavioral work, which has shown that endogenously cued shifts of attention typically take 200 to 400 milliseconds to execute (Cheal & Lyon, 1991; Eriksen & Collins, 1969; Liu, Stevens, & Carrasco, 2007; H. J. Müller & Rabbitt, 1989; Nakayama & Mackeben, 1989; for review, see Egeth & Yantis, 1997).

Next, we examined whether alpha-band activity tracked the specific location that was attended. The alpha-band CTFs that we have reported so far reflected channel response profiles that were averaged across all possible cued locations. We did observe reliable spatial selectivity in the averaged CTF. Nevertheless, it is possible that the spatial selectivity of the averaged CTF reflects selectivity for some locations but not others, which leads to reliable spatial selectivity on average (J. J. Foster et al., 2016). Thus, we inspected the alpha-band CTFs for each cued location separately (Figure 2-2c, bottom). For each location, the CTF peaked at the channel tuned for the

cued location (i.e., a channel offset of 0°) starting approximately 300 ms after cue onset, which demonstrated that time-resolved alpha-band CTFs tracked which of the eight locations was attended. Thus, alpha-band activity tracks the locus of covert attention in a spatially precise fashion.

Alpha-band CTFs showed a graded response profile: The strongest response was in the channel tuned for the cued location, and responses steadily decreased across channels tuned for other locations (Figure 2-2e, left). However, our standard set of basis functions (the basis set) specified a graded channel-response profile across locations. Therefore, the graded profile of alpha-band CTFs might be imposed by the graded basis function rather than reflecting truly graded spatially selective activity. To test this possibility, we reconstructed CTFs with the IEM again, with a basis set of Kronecker delta functions (stick functions; J. J. Foster et al., 2016). These basis functions do not specify a graded channel-response profile. Thus, a graded CTF profile seen when using this modified basis set necessarily reflects graded activity in the data itself rather than a pattern imposed by the basis function. Alternatively, if spatially selective alpha activity does not follow a graded format, then we should recover a peak in the channel tuned for the attended location and uniform responses across the other channels. Using this modified basis set, we found that alpha-band CTFs (averaged from 300 through 1,250 ms) showed a graded profile across channels (Figure 2-2e, right), which demonstrated that the graded profile of alpha-band CTFs reflects the underlying spatial selectivity of covert spatial attention.

Having established that the topographic distribution of alpha power tracked the spatial locus of covert attention, we tested whether such spatially selective activity was specific to the alpha band (8–12 Hz). We used the IEM to search a range of frequencies (4–50 Hz, in

increments of 1 Hz) across time to identify the frequency bands in which the topographic distribution of power carried information about the attended location (Figure 2-2f). We found that spatially selective oscillatory activity was largely restricted to the alpha band.

Experiment 2-2

In Experiment 2-1, we showed that spatially selective CTFs can be reconstructed from the topographic distribution of alpha power after a central, attention-directing cue. This spatially selective activity emerged several hundred milliseconds after cue onset, so that it dovetailed with behavioral estimates of the time course of endogenous shifts of spatial attention (e.g., H. J. Müller & Rabbitt, 1989). Although this finding suggests that alpha-band CTFs track spatial attention in a temporally resolved fashion, a direct test requires a manipulation of covert orienting speed. Thus, in Experiment 2-2, we manipulated the speed of target selection during visual search. Subjects performed a visual search task in which they searched for a target (a horizontal or vertical bar) among distractors (Figure 2-3a). We varied the difficulty of search by manipulating both distractor variability and target-distractor similarity (Duncan & Humphreys, 1989), and we measured reaction time to obtain a trial-by-trial estimate of the time taken to attend the target. This approach allowed us to test whether the time course of alpha-based CTFs tracked differences in the latency of target selection across different levels of search difficulty and as a function of within-subject differences in orienting latency across trials.

Figure 2-3b shows the aggregate RT distributions for easy and hard search. Our manipulation of search difficulty was effective: Median RTs were slower for hard search (M = 829 ms, SD = 153) than for easy search (M = 593 ms, SD = 71), $t(22) = 12.31$, $p < .001$, Cohen's d_z = 2.57, and accuracy was lower for hard search (M = 91.7%, SD = 4.4) than for easy search

(M = 97.0%, SD = 2.4), $t(22) = 6.09$, $p < .001$, Cohen's $d_z = 1.27$. We first tested whether alpha-band CTFs tracked orienting to the target location during the visual search task. As in Experiment 2-1, we quantified the spatial selectivity of alpha-band CTFs as CTF slope (see General Method). A spatially selective alpha-band CTF emerged soon after onset of the search array (Figure 2-3c). A permutation test revealed that CTF selectivity was reliably above chance starting 196 ms after cue onset. Thus, alpha-band CTFs tracked covert orienting to the target's location during visual search. To test whether alpha-band CTFs track the latency of covert orienting to the search target, we compared the onset latency of target-related CTFs between the easy- and hard-search conditions. To measure the difference in onset latency of target-related CTF between the search conditions, we used a jackknife-based procedure with a 50% maximum amplitude criterion (Miller, Patterson, & Ulrich, 1998; see General Method). The filled circles in Figure 3d mark the CTF onset estimates during easy and hard search. We found that the target-related CTF onset was 440 ms later for hard-search trials than for easy-search trials, $t(22) = 2.48$, $p = .011$, Cohen's $d_z = 0.52$ (one-tailed test). Thus, the onset latency of the target-related CTF was delayed in hard search compared with easy search, which demonstrated that alpha-based CTFs revealed the difference in the latency of orienting attention to the target between the search conditions.

Although RTs were generally slower for the hard-search condition than for the easy-search condition, there was also considerable overlap in RTs between the easy- and hard-search conditions (Figure 2-3b). This overlap was expected because target selection should sometimes occur very quickly during hard search, when the target happens to be one of the initial items to be selected. Given the overlap in RTs between conditions, we examined the onset latency of

target-related CTFs split by RT, comparing the 50% of trials with the fastest RTs (fast trials) with the 50% of trials with the slowest RTs (slow trials), regardless of search condition (Figure 2-3e). If alpha-band CTFs reveal the latency of target selection, CTF latency should covary with trial-by-trial reaction times. Indeed, we found that the target-related CTF onset was 372 ms later for slow trials than for fast trials, $t(22) = 7.22$, $p < .001$, Cohen's $d_z = 1.51$ (one-tailed test), which provides further evidence that the onset of target-related CTFs tracks the latency of covert orienting to the target item during visual search.

Discussion

The central role of covert spatial attention in visual cognition has motivated a sustained effort to elucidate the neural mechanisms that underpin this process. One productive avenue has been to examine the links between oscillatory alpha-band activity and spatial attention. A growing body of evidence has shown that the topographic distribution of alpha power tracks the locus of spatial attention (e.g., Kelly et al., 2006; Rihs et al., 2007; Thut et al., 2006; Worden et al., 2000), which suggests that alpha oscillations play a role in biasing visual processing toward attended locations (Foxe & Snyder, 2011; Jensen & Mazaheri, 2010). According to this view, the topographic distribution of alpha power should track the temporal dynamics of covert spatial attention. However, this prediction has not previously been subjected to a rigorous test.

Our findings provide a compelling confirmation of this prediction. In Experiment 2-1, we showed that variations in scalp distribution of alpha power enabled the reconstruction of spatially specific response profiles (i.e., CTFs) that track the endogenous orienting of spatial attention after a central cue. These alpha-band CTFs precisely discriminated the attended position beginning approximately 300 ms after the onset of the central cue, consistent with past estimates

of the time taken to endogenously shift attention (e.g., H. J. Müller & Rabbitt, 1989; Nakayama & Mackeben, 1989; for review, see Egeth & Yantis, 1997). Critically, Experiment 2-2 extended this finding by showing that dynamic changes in alpha topography tracked the latency of covert orienting during visual search. The onset of alpha-band CTFs was delayed during difficult search compared with easy search and for trials with slow responses compared with trials with fast responses. Together these findings demonstrate that moment-by-moment changes in the topography of alpha-band activity track the temporal dynamics of covert spatial attention, which closes a significant gap in the evidence linking alpha activity with covert spatial orienting.

Experiment 2-2 also provides important evidence that spatially specific alpha-band activity plays a role in covert spatial attention in a range of paradigms. It has long been thought that covert spatial orienting plays a central role in visual search (e.g., Kim & Cave, 1995; Luck, Fan, & Hillyard, 1993). However, evidence linking alpha-band activity to covert orienting during visual search has been lacking because studies that have linked alpha-band activity with spatial attention have relied almost exclusively on spatial-cuing tasks (e.g., Thut et al., 2006; Worden et al., 2000). Our finding that alpha-band CTFs tracked the latency of orienting to a target during visual search provides clear evidence for the role of alpha-band activity in visual search. Thus, spatially specific alpha-band activity plays a general role in covert orienting in a range of paradigms.

Our findings also have important methodological implications for a field that has had a long-standing interest in the spatial and temporal dynamics of covert orienting (Egeth & Yantis, 1997). Early work relied on overt behavioral responses to probe these dynamics (e.g., Downing, 1988; H. J. Müller & Rabbitt, 1989). More recently, however, neural signals that track the

allocation of attention have played a central role in this endeavor, in part because they circumvent the need for overt behavioral responses. Functional magnetic resonance imaging (fMRI) precisely tracks the spatial locus of covert attention (e.g., Brefczynski & DeYoe, 1999; Tootell et al., 1998) but provides little information about the time course of attention because of the slow hemodynamic response. Thus, researchers have relied on electrophysiological signals to examine the temporal dynamics of attention (e.g. Garcia, Srinivasan, & Serences, 2013; M. M. Müller, Teder-Sälejärvi, & Hillyard, 1998).

One productive approach has been to measure the consequences of spatial attention by examining stimulus-evoked potentials rather than directly measuring endogenous, attention-related activity. Sensory components that are amplified by spatial attention (e.g., the P1 component; Hillyard, Vogel, & Luck, 1998) have allowed researchers to probe the spatial allocation of attention. For example, Hopfinger and Mangun (1998) showed that the P1 response evoked by a probe stimulus was amplified after an exogenous spatial cue, which means that exogenous orienting shapes early stages of visual processing. However, although this approach has provided important insights into how and when attention modulates evoked visual responses, it does not reveal the time course of covert orienting before the evoking stimulus. To overcome this limitation, some studies have focused on the rhythmic brain activity—called a steady-state visual evoked potential—evoked by a flickering stimulus; this activity is amplified by spatial attention (Morgan, Hansen, & Hillyard, 1996). Thus, by examining the time course of amplitude modulations, it has been possible to measure the latency of orienting toward a flickering target (e.g. M. M. Müller et al., 1998). Nevertheless, stimulus-evoked approaches are not without limitations. Note that because these approaches rely on stimulus-evoked activity, they cannot be

used to track attention to empty locations, which restricts the kinds of questions that can be addressed with these methods. Thus, there is much to be gained from a temporally resolved signal that tracks spatial attention in the absence stimulus-evoked activity.

Our findings suggest that spatially specific alpha-band activity provides such an opportunity. Alpha-band CTFs tracked the locus of covert spatial attention in balanced visual displays and in the absence of transient evoked activity, suggesting that spatially specific alpha-band activity reflects endogenous shifts of spatial attention rather than stimulus-evoked activity. Thus, given its spatial and temporal precision, this method provides a promising approach for obtaining a moment-by-moment index of the locus of covert attention across a broad range of paradigms.

Conclusions

In the current study, we showed that the topographic distribution of alpha-band activity tracked the spatial locus of covert attention after attention-directing cues and during visual search. These results demonstrate that alpha-band activity plays a central role in covert orienting in a range of paradigms. Critically, the time course of spatially specific alpha activity tracked trial-by-trial variations in the speed of covert orienting during visual search. Our results provide critical new evidence for the link between alpha-band activity and covert spatial attention.

CHAPTER 3. THE TOPOGRAPHY OF ALPHA-BAND ACTIVITY TRACKS THE CONTENT OF SPATIAL WORKING MEMORY

Introduction

A range of evidence suggests that neuronal oscillations in the alpha band (8–12 Hz) play a central role in the selection and storage of information in the brain (Canolty & Knight, 2010; Fries, 2005; Klimesch, 2012). For instance, many studies have shown that alpha-band activity covaries with the deployment of spatial attention, such that posterior alpha power is reduced contralateral to attended locations (e.g., Gould, Rushworth, & Nobre, 2011; Kelly, Lalor, Reilly, & Foxe, 2006; Thut, Nietzel, Brandt, & Pascual-Leone, 2006). Indeed, the topographic distribution of alpha power not only tracks the attended visual hemifield, but also the specific retinotopic coordinates that are attended (Rihs et al., 2007; Worden et al., 2000). For example, Rihs and colleagues cued participants to attend one of eight placeholder locations around a central fixation point and found that the topography of alpha power systematically varied with the cued location, with more similar topographies associated with adjacent locations (Rihs et al., 2007). Extending these findings, others have decoded both horizontal and vertical shifts of attention from patterns of alpha power (Bahramisharif et al., 2010; van Gerven & Jensen, 2009), and shown that lateralized modulations of alpha power are sensitive to the eccentricity of the attended location (Bahramisharif, Heskes, Jensen, & van Gerven, 2011). Thus, alpha-band activity enables tracking of the locus of spatial selective attention.

Here, motivated by past work positing a strong functional overlap between spatial attention and spatial working memory (WM; Awh & Jonides, 2001; Awh, Vogel, & Oh, 2006; Gazzaley & Nobre, 2012), we examined whether alpha-band activity tracks spatial

representations held in WM. Consistent with this possibility, past work has already shown that alpha power is reduced contralateral to locations held in spatial WM (Medendorp et al., 2007; Van Der Werf, Jensen, Fries, & Medendorp, 2008; van Dijk, Van Der Werf, Mazaheri, Medendorp, & Jensen, 2010). However, these contralateral modulations do not establish whether alpha-band activity – as in the case of spatial selective attention – tracks the specific locations that are stored. Instead, these modulations might reflect a lateralized memory operation (e.g., a top-down control signal) that tracks the visual hemifield of the remembered location but not the exact location that is stored. Here, we tested whether alpha-band activity relates to the *coding* of the precise positions stored in WM, rather than a lateralized memory process than is not sensitive the specific content of spatial WM.

In three experiments, participants performed spatial WM tasks that required them to remember the precise angular location of a stimulus, sampled from a 360° space. Using an inverted spatial encoding model (IEM; Brouwer & Heeger, 2011, 2009; for review, see Sprague et al., 2015), we identified the oscillatory frequency bands in which the topographic distribution of power – measured using EEG – carried location-specific information. We found that the topographic distribution of both evoked (i.e., activity phase-locked to stimulus onset) and total (i.e., activity regardless of phase) power across a range of low-frequency bands transiently tracked stimulus location. However, only the topography of total power in the alpha-band tracked the content of spatial WM throughout the delay period, indicating that alpha-band activity is directly related to the coding of spatial representations held in WM. The IEM allowed us to reconstruct spatially specific response profiles (termed channel-tuning functions, or CTFs) that tracked the stored location during both encoding and delay periods of WM tasks. Thus,

alpha-band activity enables time-resolved tracking of spatial representations held in WM. These findings underscore the central role that alpha-band activity plays in coding the content of spatial WM, and provide new evidence for the functional overlap between spatial WM and spatial attention.

Materials and Methods

Subjects

Fifteen volunteers took part in each experiment for monetary compensation ($10 per hour). The participants in each experiment were non-overlapping, with the exception of one participant who took part in both Experiments 3-1 and 3-3. Participants reported normal or corrected-to-normal vision, were between 18 and 35 years old, and provided informed consent according to procedures approved by the University of Oregon Institutional Review Board. Participants were replaced if more than 25% of trials were lost due to recording or ocular artifacts, and/or if the participants did not complete all trials during the session (five in Experiment 3-1, seven in Experiment 3-2, and three in Experiment 3-3).

Stimulus Displays

Stimuli were generated in Matlab (Mathworks, Inc.) using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997), and were presented on a 17-inch CRT monitor (refresh rate: 60 Hz) at a viewing distance of ~100 cm. All stimuli were rendered in dark gray against a medium gray background.

The spatial WM tasks required participants to remember the angular location of a sample stimulus. The stimulus was a circle (1.6° in diameter) centered 4° of visual angle from the central fixation point (0.2° in diameter). For each trial, the angular location of the stimulus was sampled

from one of eight location bins spanning 0-315° (in angular location), in steps of 45°, with jitter added to cover all 360° of possible locations to prevent categorical coding of stimulus location. In all experiments, the location of the sample stimulus was drawn from each bin equally often, and in a random order, within each block of trials.

Procedures

After providing informed consent, participants were fitted with a cap embedded with 20 scalp electrodes before completing a spatial WM task (see below for details). Testing took place in a dimly lit, electrically shielded chamber. In each experiment, the spatial WM task comprised 15 blocks of 64 trials, and took approximately 2-2.5 hours to complete.

Delayed estimation task. Participants in Experiments 3-1 and 3-2 performed a spatial delayed-estimation task (Wilken & Ma, 2004; Zhang & Luck, 2008) (see Figure 3-1). These experiments differed only in trial timing. Participants began each trial by pressing the spacebar. The trial began with a fixation display lasting between 600-1500 ms. A sample stimulus was then presented (250 ms in Experiment 3-1; 1000 ms in Experiment 3-2), followed by a delay period during which only the fixation point remained visible (1750 ms in Experiment 3-1; 1000 ms in Experiment 3-2). After the delay period, participants used a mouse to click on the perimeter of a probe ring (8° in diameter, 0.2° thick) to report the remembered location of the sample disc as precisely as possible. Before starting the task, participants completed a brief set of practice trials to ensure they understood the task.

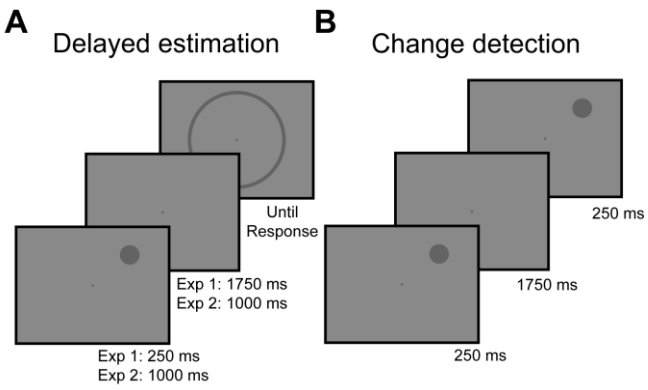

Figure 3-1. Spatial working memory tasks used in Chapter 3.
(a) In experiments 3-1 and 3-2, participants performed a delayed estimation task. Participants maintained fixation while they remembered the angular location of a sample stimulus, which they reported after a delay period by clicking on the perimeter of a rim. Experiments 1and 2 differed only in trial timing. (b) In experiment 3-3, participants performed a change detection task. Rather than report the location of the sample stimulus, a test stimulus was presented after the delay period and participants reported whether or not the location of the stimulus had changed. Stimuli not perfectly to scale.

Spatial change detection task. Participants in Experiment 3-3 performed a spatial change detection task (see Figure 3-1). Rather than reporting the remembered location after the delay period with a mouse click, a test stimulus (identical to the sample stimulus) was presented for 250 ms. On half of trials, the test stimulus was presented in the same location as the sample stimulus (*no-change trials*), while on the other half of trials the test stimulus was shifted 20° clockwise or anticlockwise from the sample location (*change trials*). Participants indicated with a key press whether or not the location of the stimulus had changed. The timing of the task was identical to Experiment 3-1 (i.e., 250 ms sample stimulus, 1750 ms delay period). Before starting the task, participants completed a brief set of practice trials with feedback to ensure they knew the size of the change to expect for change of trials.

Modeling Response Error Distributions

In the delayed estimation experiments (Experiments 3-1 and 3-2), response error, the angular difference between the reported and presented locations, could range from -180° to 180°. The response error distribution for each participant was modeled as the mixture of a von Mises distribution and a uniform distribution, corresponding to trials in which the sample stimulus was successfully or unsuccessfully stored, respectively (see Zhang & Luck, 2008). Maximum likelihood estimates were obtained for three parameters: (1) the mean of the von Mises component (μ), corresponding to response bias; (2) the dispersion of the von Mises distribution (*s.d.*), corresponding to mnemonic precision; and (3) the height of the uniform distribution (P_f), corresponding to the probability of forgetting the sample stimulus. Parameter estimates were obtained using the 'MemFit.m' function of MemToolbox (Suchow, Brady, Fougnie, & Alvarez, 2013).

EEG Acquisition

We recorded EEG from 20 tin electrodes mounted in an elastic cap (Electro-Cap International, Eaton, OH). We recorded from International 10/20 sites F3, FZ, F4, T3, C3, CZ, C4, T4, P3, PZ, P4, T5, T6, O1, and O2, along with five nonstandard sites: OL midway between T5 and O1, OR midway between T6 and O2, PO3 midway between P3 and OL, PO4 midway between P4 and OR, and POz midway between PO3 and PO4. All sites were recorded with a left-mastoid reference, and were re-referenced offline to the algebraic average of the left and right mastoids. To detect horizontal eye movements, horizontal electrooculogram (EOG) was recorded from electrodes placed ~1 cm from the external canthus of each eye. To detect blinks and vertical eye movements, vertical EOG was recorded from an electrode placed below the

right eye and referenced to the left mastoid. The EEG and EOG were amplified with an SA Instrumentation amplifier with a bandpass of 0.01 to 80 Hz and were digitized at 250 Hz using LabVIEW 6.1 running on a PC. Trials were visually inspected for artifacts, and we discarded trials contaminated by blocking, blinks, detectable eye movements, excessive muscle noise, or skin potentials. An average of 12.2% (*SD* = 5.6%) of trials were rejected per participant across all three experiments.

Time-Frequency Analysis

Time-frequency analyses were performed using Matlab in conjunction with the Signal Processing and EEGLAB toolbox (Delorme & Makeig, 2004). To isolate frequency-specific activity, the raw EEG signal was bandpass filtered using a two-way least-squares finite impulse response filter (Delorme & Makeig, 2004). This filtering method uses a zero-phase forward and reverse operation, which ensures that phase values are not distorted, as can occur with forward-only filtering methods. A Hilbert Transform (Matlab Signal Processing Toolbox) was applied to the bandpass-filtered data, producing the complex analytic signal, $z(t)$, of the filtered EEG, $f(t)$, where $z(t) = f(t) + i\tilde{f}(t) = A(t)e^{i\phi(t)}$, from which instantaneous amplitude, $A(t)$, was extracted; $\tilde{f}(t)$ is the Hilbert Transform of $f(t)$, and $i = \sqrt{-1}$. The complex analytic signal was extracted for each electrode using the following Matlab syntax:

hilbert(eegfilt(data,F,f1,f2)')')

where data is a 2D matrix of raw EEG (number of trials × number of samples), *F* is the sampling frequency (250 Hz), *f1* is the lower bound of the filtered frequency band, and *f2* is the upper bound of the filtered frequency band. For alpha-band analyses, we used an 8-12 Hz bandpass filter, thus *f1* and *f2* were eight and twelve, respectively. For the time-frequency analysis, we

searched a broad range of frequencies (4-50 Hz, in increments of 1 Hz with a 1 Hz bandpass). For these analyses *f1* and *f2* were four and five to isolate 4-5 Hz activity, five and six to isolate 5-6 Hz activity, and so on.

Total power was computed by squaring the complex magnitude of the complex analytic signal, and then averaging across trials. Thus, total power reflects ongoing activity irrespective of its phase-relationship to onset of the sample stimulus. In contrast, evoked power was calculated by first averaging the complex analytic signal across trials, and then squaring the complex magnitude of the averaged analytic signal. Evoked power reflects activity phase-locked to stimulus onset because only activity with consistent phase across trials remains after averaging the complex analytic signal.

Because calculating evoked power requires averaging across trials, artifact-free trials were partitioned into three blocks. To prevent bias in our analysis, we equated the number of observations across location bins within each block. To this end, we calculated the minimum number of trials for any given location bin *n* for each participant, and assigned *n*/3 many trials for each location bin to each of the three blocks. Importantly, the blocks were independent (i.e., no trial was repeated across blocks) to prevent circularity in the cross-validation procedures used for the IEM routine (see Inverted Encoding Model). Evoked and total power was then calculated for each location bin for each block, resulting in an $l*b \times m \times s$ matrix of both evoked and total power values, where l is the number of location bins, b is the number of blocks, m is the number of electrodes, and s is the number of time samples. For the analysis in which the IEM is applied across many frequency bands, we down-sampled the data matrix to a sample rate of 50 Hz (i.e.,

one sample every 20 ms) to reduce computation time. The data matrix was not down-sampled for analyses restricted to the alpha-band.

Finally, because we equated the number of trials across location bins within blocks, a random subset of trials were not included in any block. Thus, we randomly generated multiple block assignments (five for the full time-frequency analyses, and ten for the alpha-band analyses), each resulting in an $l*b \times m \times s$ power matrix. The IEM routine (see Inverted Encoding Model) was applied to the matrices of power values for each block assignment, and their outputs (i.e., channel response profiles) were averaged. This approach better utilized the complete data set for each participant and minimizes the influence of idiosyncrasies in estimate of evoked and total power specific to certain assignments of trials to blocks.

Inverted Encoding Model

We used an IEM to reconstruct location-selective CTFs from the topographic distribution of oscillatory power across electrodes. We assumed that power measured at each electrode reflects the weighted sum of eight spatial channels (i.e., neuronal populations), each tuned for a different angular location (c.f. Brouwer & Heeger, 2011, 2009; Sprague, Ester, & Serences, 2014; Sprague & Serences, 2013). We modeled the response profile of each spatial channel across angular locations as a half sinusoid raised to the seventh power, given by:

$$R = sin(0.5\theta)^7,$$

where θ is angular location (ranging from 0° to 359°), and R is the response of the spatial channel in arbitrary units. This response profile was circularly shifted for each channel such that the peak response of each spatial channel was centered over one of the eight location bins (i.e.,

0°, 45°, 90° etc.). The predicted channel responses for each location bin were derived from these basis functions (calculated using the angular location at the center of each bin).

An IEM routine was applied to each time-frequency point in the time-frequency analyses, and to each time point in the alpha-band analyses. This routine proceeded in two stages (train and test). In the training stage, training data B_1 were used to estimate weights that approximate the relative contribution of the eight spatial channels to the observed response measured at each electrode. Let B_1 (m electrodes × n_1 observations) be the power at each electrode for each measurement in the training set, C_1 (k channels × n_1 observations) be the predicted response of each spatial channel (determined by the basis functions) for each measurement, and W (m electrodes × k channels) be a weight matrix that characterizes a linear mapping from "channel space" to "electrode space". The relationship between B_1, C_1, and W can be described by a general linear model of the form:

$$B_1 = WC_1$$

The weight matrix was obtained via least-squares estimation as follows:

$$\widehat{W} = B_1{C_1}^T\left(C_1{C_1}^T\right)^{-1}$$

In the test stage, with the weights in hand, we inverted the model to transform the observed test data B_2 (m electrodes × n_2 observations) into estimated channel responses, C_2 (k channels x n_2 observations):

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

Each estimated channel response function was circularly shifted to a common center (i.e., 0° on the "Channel Offset" axes of the figures) by aligning the estimated channel responses to the channel tuned for the stimulus bin to yield CTFs. The IEM routine was performed separately for

each sample point from 500 ms prior to stimulus onset through to the end of the delay period (2000 ms).

Importantly, we used a "leave-one-out" cross validation routine such that two blocks of estimated power values (see Time-Frequency Analysis) served as B_1 and were used to estimate , and the remaining block served as B_2 and was used to estimate C_2. Thus, the training and test data were always independent. This process was repeated until each of the three blocks were held out as the test set, and the resulting CTFs were averaged across each test block.

Finally, because the exact contributions of each spatial channel to each electrode (i.e., the channel weights, *W*) will likely vary by subject, the IEM routine is applied separately for each subject, and statistical analyses were performed on the reconstructed CTFs. This approach allowed us to disregard differences in the how location-selectivity is mapped to scalp-distributed distributed patterns of power across subjects, and instead focus on the profile of activity in the common stimulus or "information" space (Sprague & Serences, 2015).

Statistical Analysis

To quantify the location selectivity of CTFs, we used linear regression to estimate CTF slope (i.e., slope of channel response as a function of location channels after collapsing across channels that were equidistant from the channel tuned to the location of the evoking stimulus), where higher CTF slope indicates greater location selectivity. To test whether CTF selectivity was reliably above chance, we tested whether CTF slope was greater than zero using a one-sample *t* test. Because mean CTF slope may not be normally distributed under the null hypothesis, we employed a Monte Carlo randomization procedure to empirically approximate the null distribution of the *t* statistic. Specifically, we implemented the IEM as described above but

randomized the location labels within each block so that the labels were random with respect to the observed responses in each electrode. This randomization procedure was repeated 1000 times to obtain a null distribution of *t* statistics. To test whether the observed CTF selectivity was reliably above chance, we calculated the probability of obtaining a *t* statistic from the surrogate null distribution greater than or equal to the observed *t* statistic (i.e., the probability of a Type 1 Error). Our permutation test was therefore a one-tailed test. CTF selectivity was deemed reliably above chance if the probability of a Type 1 Error was less than .01. This permutation testing procedure was applied to each time-frequency point in the time-frequency analyses, and to each time point in the alpha-band analyses.

Quantifying Biases in Eye Position

Although we discarded trials with detectable eye movements, small but systematic biases in eye position towards the remembered location may still exist. Indeed, we found evidence for very small but reliable biases in horizontal EOG amplitude that tracked stimulus location (see Results). If eye position is biases towards stimulus location, then there should be a linear relationship between the horizontal position of the stimulus and horizontal EOG amplitude. Thus, we used linear regression to calculate an eye bias score to quantify the extent to which eye position covaried with stimulus location. The eye bias score was calculate as the slope of the best fitting linear function describing horizontal EOG amplitude (in μV) as a function of the horizontal location of the stimulus (in degrees of visual angle). Higher eye bias scores therefore reflect greater changes in eye position as a function of stimulus location. This eye bias score was calculated across time for each participant separately.

Results

Behavior

Task performance confirmed that participants were engaged in the spatial WM tasks in all three experiments. In the delayed-estimation experiments (Experiments 3-1 and 3-2), mnemonic precision computed from a mixture model (Suchow et al., 2013; Zhang & Luck, 2008) was high, indicated by low *s.d.* values (Experiment 3-1: $M = 6.6°$, $SD = 1.8°$; Experiment 3-2: $M = 5.2°$, $SD = 1.6°$), and the probability that the stimulus was forgotten, P_f, was extremely low (Experiment 3-1: $M = .002$, $SD = .002$; Experiment 3-2: $M = .002$, $SD = .002$). In the spatial change detection task (Experiment 3-3), change detection accuracy was high ($M = 90.1\%$, $SD = 4.7\%$).

Experiment 3-1

Identifying frequency bands that track the content of spatial working memory. We first sought to identify the frequency bands in which the topographic distribution of oscillatory power tracked the content of spatial WM. To this end, we used an IEM to reconstruct location-selective CTFs (see Materials and Methods). Using the IEM, we searched a broad range of frequencies (4-50 Hz, in increments of 1 Hz) across time to identify the frequency bands in which the topographic distribution of evoked and total power tracked the location of the sample stimulus. If the multivariate pattern of power across electrodes carries information about stimulus location, then the IEM should reveal a graded CTF profile, with a clear peak in the channel tuned for the remembered location. On the other hand, if the multivariate pattern of power does not carry information about stimulus location, then the IEM should produce a flat CTF profile, indicating no location tuning. Figure 3-2 shows location selectivity of reconstructed

CTFs (as measured by CTF slope) as a function of time and frequency for both evoked and total power. We performed a permutation test at each time-frequency point to identify the points at which CTF slope was reliably above zero (see Materials and Methods). We found the topographic distribution of both evoked and total power transiently tracked the location of the sample stimulus across a range of low frequencies (~4-20 Hz; Figure 3-2a). In contrast, only total alpha power (~8-12 Hz) enabled sustained tracking of the stored location throughout the blank delay (Figure 3-2b). Thus, the topographic distribution of total alpha power tracked spatial representations held in WM.

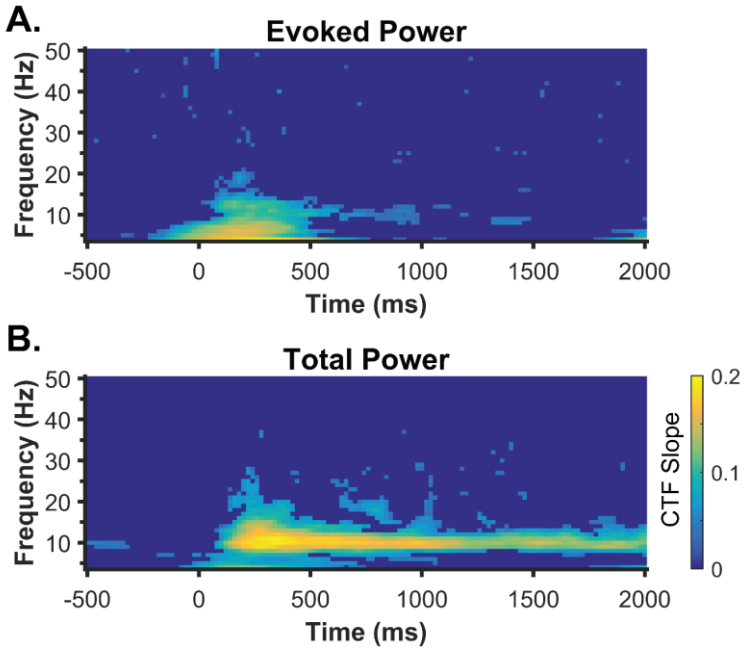

Figure 3-2. Identifying the frequency bands that track the content of spatial working memory in Experiment 3-1.
An IEM was used to reconstruct location-selective channel-tuning functions (CTFs) from the topographic distribution of evoked (a) or total (b) power across a broad range of frequencies (4–50 Hz, in increments of 1 Hz) and time. Evoked and total power transiently tracked stimulus location after stimulus onset across a broad range of frequencies (4 to ~20 Hz). However, only total alpha power tracked the content of spatial WM throughout the delay period. Color represents CTF slope, a measure of CTF selectivity that quantifies the location-specific activity in the topographic distribution of power. Points at which CTFs slope was not reliably above zero as determined by a permutation test are set to zero (dark blue).

Alpha power tracks the content of spatial WM with fine-grained spatial resolution. Previous work has demonstrated that alpha power decreases contralateral to a location held in spatial WM (Medendorp et al., 2007; Van Der Werf et al., 2008; van Dijk et al., 2010). It is possible that location-selectivity in the alpha-band, measured using CTF slope, simply reflects sensitivity of alpha power to the hemifield (or quadrant) of the remembered location. Such coarse location tuning could give rise to a graded CTF when shifted and averaged across location bins. To test this possibility, we examined the channel responses profiles reconstructed from total alpha (8-12 Hz) power in Experiment 3-1 for each of the eight location bins separately. If the topographic distribution of alpha power tracks the precise location held in WM, then the channel response profiles should differ for each of the eight location bins. Specifically, for each location bin, we should observe a graded tuning profile, with a peak response in the channel tuned for the remembered location. Figure 3-3 shows the channel responses (averaged across time from 0-2000 ms) for each of the eight location bins. Indeed, we found that the peak channel response was always seen in the channel tuned for the remembered location, providing clear evidence that the topography of alpha power tracked the specific angular location held in WM.

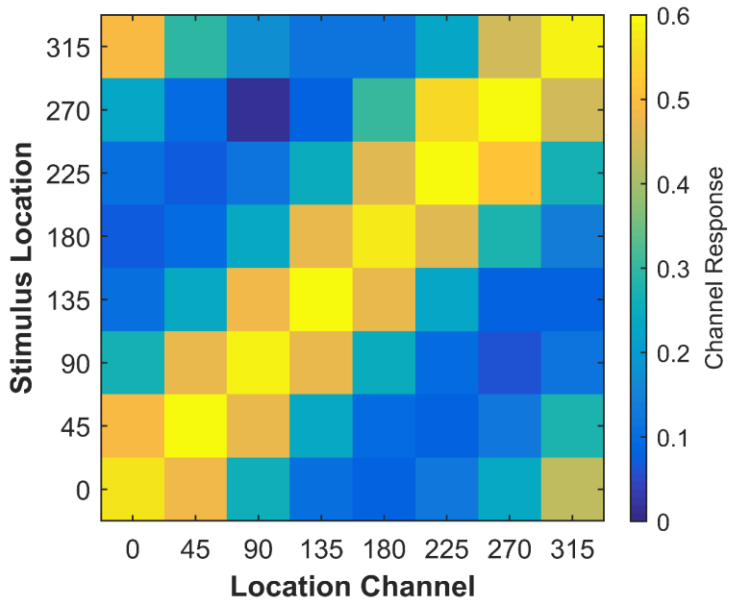

Figure 3-3. CTFs precisely track the location of the remembered stimulus. Unshifted channel responses reconstructed from total alpha power after stimulus onset (averaged from 0–2,000 ms) shown for each of the eight stimulus location bins in Experiment 3-1. For all location bins, the peak response is seen in the channel tuned for that stimulus location, confirming that the topography of alpha power tracked the precise location held in spatial WM.

Examining the format of spatial representations tracked by alpha power. Next we examined the *format* of the spatial representations tracked by the topographic distribution of alpha power. Our standard basis set specified a "graded" channel response profile, with the peak response in the channel tuned for the remembered location, and with gradually diminishing response for channels tuned for other locations. This graded response profile across feature-selective cells is the hallmark of population coding of sensory variables (Pouget et al., 2000), rather than more abstracted (e.g., categorical) representations. We found that CTFs reconstructed from alpha-band activity showed this graded profile, suggesting that that modulations of alpha power follow a graded, sensory format (Figure 3-4a). However, using an IEM, it is always possible that the graded CTF profile reflects the graded basis function itself rather than truly

graded location-selective activity (Ester, Sprague, & Serences, 2015; Saproo & Serences, 2014). To rule out this possibility, we reconstructed CTFs with the IEM, this time with a modified basis set of eight orthogonal Kronecker delta functions, each centered on one of the eight location bins. These functions do not specify a graded profile of responses across channels. Thus, if a graded profile is observed using this basis function it must reflect a graded pattern in the data itself. In contrast, if the spatial representations tracked by alpha-band power does not follow this graded format, we would expect to recover a peak in the channel tuned for the remembered, and uniform responses across the rest of the channels. Using this orthogonal basis set, reconstructed CTFs had a graded profile (Figure 3-4b), confirming that graded location tuning is an intrinsic property of alpha activity and is not imposed by the graded basis function. Figure 3-4c shows that this graded profile was consistent across time. Therefore, alpha-band activity follows the expected format of a sensory code.

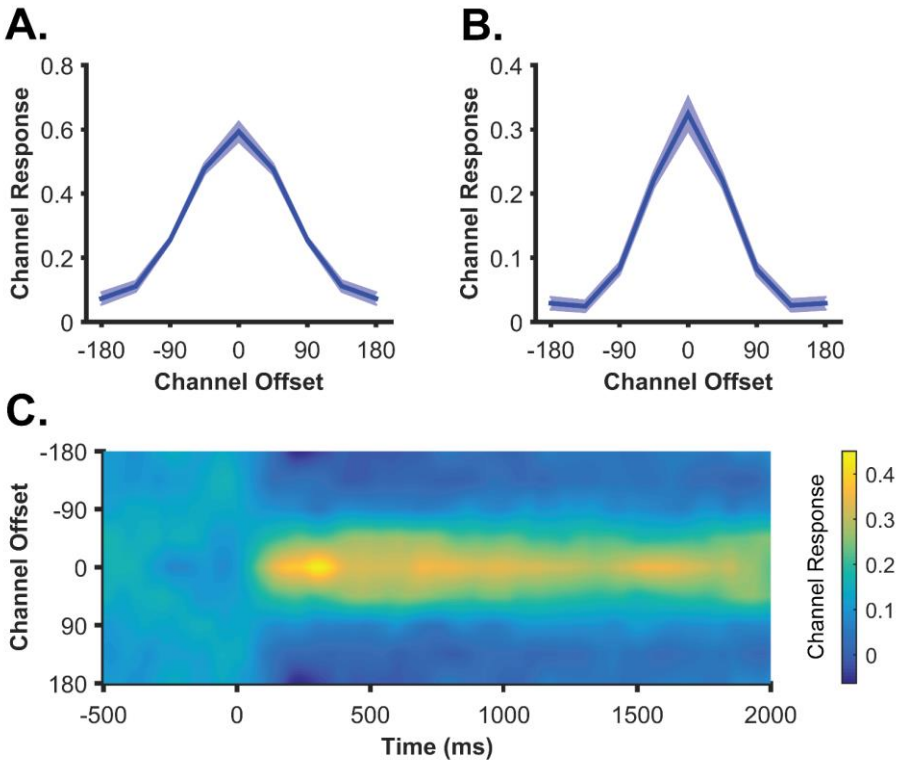

Figure 3-4. Examining the format of spatial representations tracked by alpha power. CTFs reconstructed from total alpha (8–12 Hz) power after stimulus onset (averaged from 0–2,000 ms) using the standard, "graded" basis set (a) and using a basis set of eight orthogonal Kronecker delta functions (b) for Experiment 3-1. CTFs reconstructed using the orthogonal basis set still showed a smooth, graded profile, demonstrating that the intrinsic tuning properties of alpha power are well-described by the graded basis function of the spatial encoding model. Shaded error bars reflect bootstrapped standard error of the mean. (c) Time-resolved CTF reconstructed using the orthogonal basis set, showing that the graded CTF profile was consistent across time points.

Ruling out biases in eye position. Eye movements generate electrical potentials that affect EEG recordings. We instructed participants to maintain fixation during the spatial WM task, and discarded trials with detectable eye movements. Nevertheless, small but systematic biases in eye position towards the remembered location may still exist. To examine this possibility, we inspected horizontal electrooculogram as a function of stimulus location. We found a reliable bias in eye position that tracked stimulus position. However, this bias was

remarkably small (< 2 μV), corresponding to shifts from fixation of less than 0.15° of visual angle on average (Lins, Picton, Berg, & Scherg, 1993b; Lins et al., 1993a). We found that CTF slope in the alpha-band is greatest early in the trial, and decreased towards the end of the delay period (Figure 3-5a). In contrast, bias in eye position, quantified using an eye bias metric (see Materials and Methods), increased gradually as the delay period progressed (Figure 3-5b). Therefore, subtle biases in eye position cannot explain the link between alpha activity and the content of spatial WM.

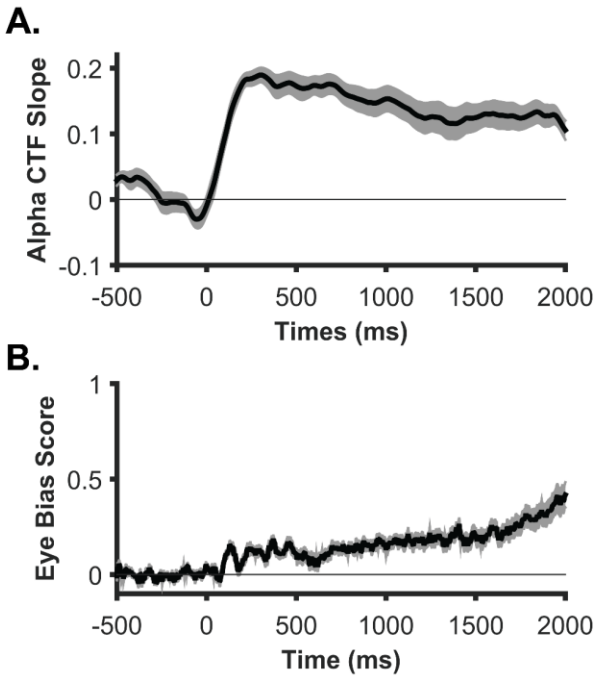

Figure 3-5. Ruling out biases in eye position as a source of spatially selective alpha power. Alpha CTF selectivity and small biases in eye position toward the remembered location show different time courses. (a) Location-selectivity of CTFs reconstructed from total alpha (8–12 Hz) power across time. Although CTF selectivity is robust throughout the delay period, it gradually declines across time. (b) Eye bias score quantifying bias in eye position toward (positive values) or away (negative values) from the remembered location. This score reflects the change in HEOG amplitude (in μV) seen when the horizontal stimulus position changes by one degree of visual angle. In contrast to CTF selectivity, bias in eye position toward the remembered location increases across time, demonstrating that location-selective patterns of alpha power are not accounted for by small but reliable biases in eye position. Shaded error bars reflect bootstrapped standard error of the mean.

Experiment 3-2: Alpha activity tracks spatial representations when the stimulus remains in view

In Experiment 3-1, we demonstrated that that topographic patterns of alpha power track the content of spatial WM during the delay period of a WM task. Recent work suggests that WM is important not only for maintaining representations of stimuli that are no longer externally available but also for representing stimuli that remain within view (Chun, 2011; Tsubomi et al., 2013). In Experiment 3-2, we used an extended encoding period (1000 ms) to examine whether the topographic distribution of alpha power track the spatial representation of a stimulus that remains in view. We observed location-selective CTFs reconstructed from total alpha activity (Figure 3-6b). Importantly, robust CTFs were seen throughout both the stimulus (0-1000 ms) and delay (1000-2000 ms) periods, demonstrating that alpha activity tracks the location of a to-be-remembered stimulus, even when the stimulus remains in view. Thus, total alpha power tracks spatial representations both in the presence and absence of visual input.

In Experiment 3-2, we again found that evoked alpha power generated reliable CTFs following stimulus onset (Figure 3-6a). However, weak but reliable CTF selectivity was also seen following stimulus offset (at 1000 ms). This second burst of location-specific evoked activity may reflect resynchronization of low-frequency activity caused by abrupt visual transient of stimulus offset (Gruber, Klimesch, Sauseng, & Doppelmayr, 2005; Hanslmayr et al., 2007). However, further work is necessary to test this possibility.

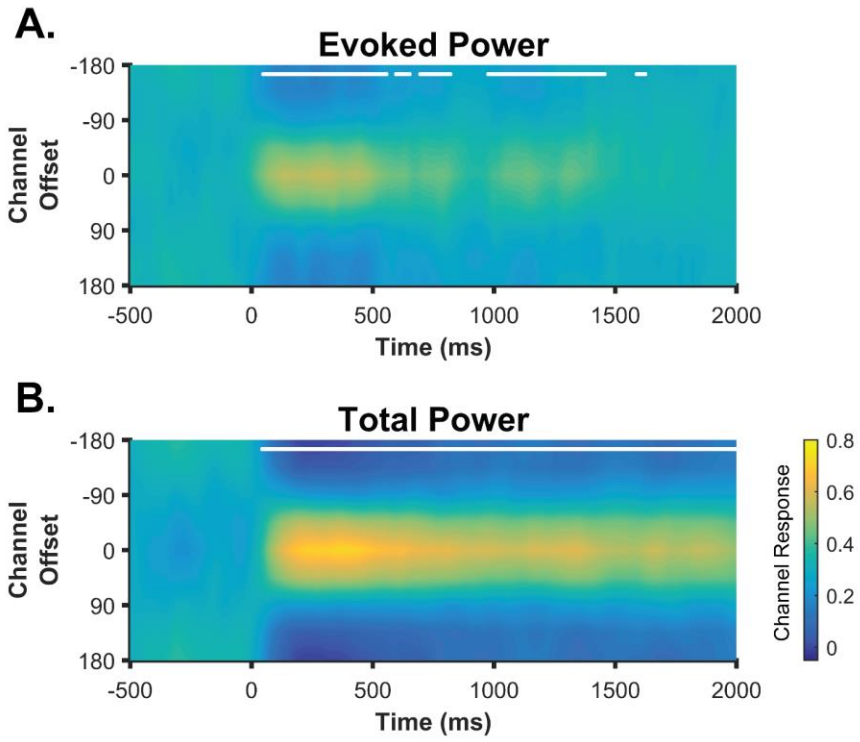

Figure 3-6. Total alpha activity tracks stimulus location when the stimulus remains in view. Location CTFs reconstructed from evoked (a) and total (b) alpha (8–12 Hz) activity during the spatial WM task with a long encoding period (1,000 ms; Experiment 3-2). Despite the stimulus remaining in view from 0 to 1,000 ms, total alpha power tracked stimulus location during this period, demonstrating that total alpha power tracks spatial representations both in the presence and absence of visual input. Evoked alpha power showed reliable but weaker location selectivity during the delay period. White markers along the top of the panels indicate the points at which CTF slope was reliably above chance as determined by a permutation test.

Experiment 3-3: Ruling out response confounds

In our first two experiments, we used a delayed-estimation task in which participants used a mouse to click on a ring around fixation to report the remembered location (Figure 3-1a). Consequently, the remembered location covaried with the required response. In Experiment 3-3, we sought to rule out the possible contribution of preparatory motor activity to the location-selective CTFs that we observed in the delayed-estimation task. To this end, participants performed a spatial change detection task, in which they remembered the precise location of a

sample stimulus, and reported whether the location of a test stimulus – presented after the delay period – differed from that of the sample stimulus (Figure 3-1b). Critically, participants could not plan their response ("change" vs. "no change") until the test stimulus was presented. Therefore, any location-selective delay activity cannot reflect a planned response. As in Experiment 3-1, evoked CTFs only transiently tracked stimulus location (Figure 3-7a), while the topographic distribution of total alpha power tracked the remembered location, allowing for the reconstruction of reliable CTFs throughout the delay period (Figure 3-7b). Thus, findings from Experiment 3-3 solidify our conclusion that alpha activity tracks the contents of spatial memory, rather than the trajectory of a planned response.

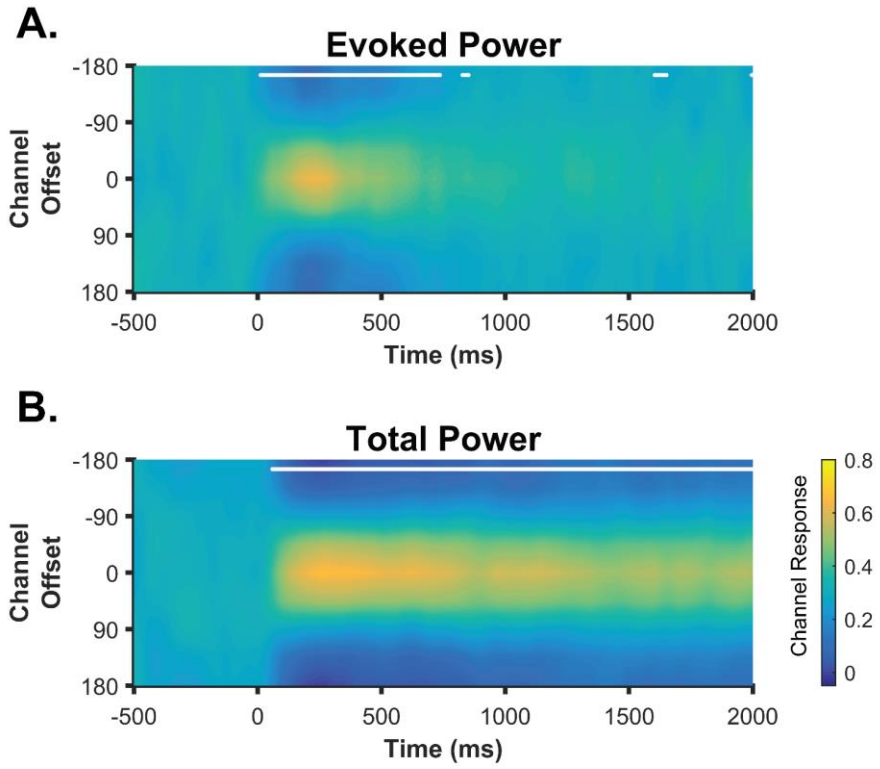

Figure 3-7. Ruling out response confounds.
Location CTFs reconstructed from evoked (a) and total (b) alpha (8–12 Hz) power during a spatial change detection task (Experiment 3-3). Because participants could not anticipate the response ("change" vs. "no change") during the delay period, this task eliminated any contribution of preparatory motor activity to the location CTFs. Nevertheless, total alpha power

Figure 3-7, continued..
tracked location held in spatial WM throughout the delay period, whereas evoked alpha power only transiently tracked stimulus position. White markers along the top of the panels indicate the points at which CTF slope was reliably above chance as determined by a permutation test.

Location-selective activity is specific to the alpha band: A cross-experiment analysis

In our analyses of experiments 3-2 and 3-3 so far, we focused exclusively on the alpha-band (8-12 Hz) because this frequency band was implicated in WM storage in Experiment 3-1. These analyses replicated the finding that the topographic distribution of alpha power tracks the content of spatial WM. However, because we focused on the alpha-band, these analyses did not replicate the finding that sustained, location-selective activity is *specific* to the alpha-band. Next, we sought to identify the frequency bands in which oscillatory power tracked stimulus location across all three experiments. To this end, we used the IEM to search time-frequency space for frequencies that carried location-specific information in Experiments 3-2 and 3-3 (as we previously reported for Experiment 3-1) to obtain maps of CTF slope across time and frequency (4-50 Hz, in increments of 1 Hz) for each experiment (Figure 3-8a). With these maps in hand, we then identified the points in time-frequency space that showed reliable CTF selectivity across all three experiments to create a cross-experiment map of location-selective oscillatory activity (Figure 3-8-b). This simple but conservative analysis revealed that evoked and total power across a range of frequency bands (~4-20 Hz) transiently tracked stimulus location, while only total alpha power tracked the remembered location throughout the delay period. The cross-experiment maps combined data from three independent experiments, with non-overlapping groups of

participants.[1] Therefore, this analysis provides clear evidence that location-selective oscillatory activity during WM maintenance, as reflected in the topographic distribution of oscillatory power, is specific to the alpha-band.

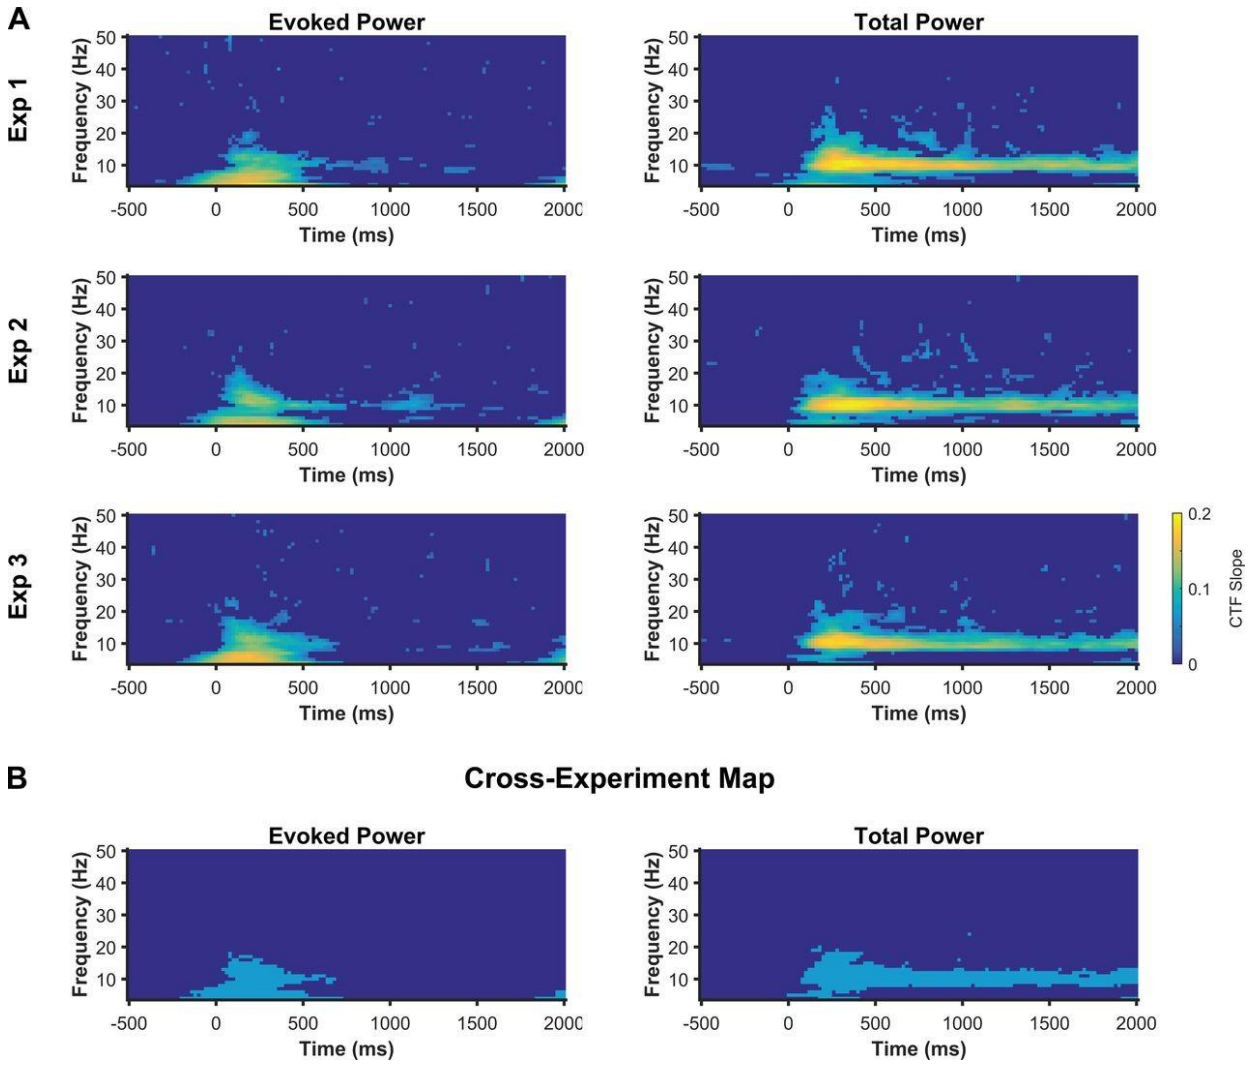

Figure 3-8. Location selectivity is specific to the alpha band.
(a) slope of CTFs reconstructed from the topographic distribution of evoked and total power across a broad range of frequencies (4–50 Hz, in increments of 1 Hz) and time for experiments

[1] One participant took part in both experiments 3-1 and 3-3. This subject was excluded from the Experiment 3-3 dataset for the purpose of the cross-experiment map to ensure that the data from each experiment were independent.

Figure 3-8, continued..
3-1, 3-2, and 3-3. Points at which CTF slope was not reliably above zero as determined by a permutation test are set to zero (dark blue). (b) cross-experiment map of the points for which CTF slope was reliably above chance for all three experiments. Reliable points are shown in light blue. Across three experiments, evoked and total power transiently tracked stimulus location after stimulus onset across a broad range of frequencies (4 to ~20 Hz) while only total alpha power tracked the content of spatial WM throughout the delay period.

Our results provide decisive evidence that the topography of alpha power tracks the location held in spatial WM, suggesting that alpha-band activity is related to the coding of spatial representations in WM. However, it is less clear whether other frequency bands play similar roles. Theta- and gamma-band activity (i.e., 4-7 Hz and ~30-100 Hz, respectively) are of particular interest because they have been proposed to play a central role in coordinating cellular assembles that code the content of WM (Lisman, 2010; Roux & Uhlhaas, 2013). Given that we observed no evidence of sustained location-selectivity in the topographic distribution of theta or gamma power, it is tempting to conclude that the frequency bands do not contribute to the coding of spatial representations in WM. However, it is possible that these frequency bands contribute to the coding of spatial WM representations in ways that do not result in location-specific patterns of EEG power across the scalp. For example, location-specific patterns of theta-band activity may exist in hippocampal local field potentials (e.g. Agarwal et al., 2014). However, such location-selectivity might not necessarily produce location-specific patterns of theta power on the scalp. Thus, our results do not rule out potential roles for oscillatory activity outside the alpha-band in coding the content of spatial WM.

Discussion

Previous work has demonstrated that alpha power is reduced contralateral to locations held in spatial WM (Medendorp et al., 2007; Van Der Werf et al., 2008; van Dijk et al., 2010).

However, those studies left it unanswered whether the topographic distribution of alpha activity tracked the precise location held in WM, or whether alpha activity reflected a lateralized memory process (e.g., a contralateral control signal) that is insensitive to the specific location that is stored. Here, we provide clear evidence that the topographic pattern of total alpha-band power precisely tracks the angular location maintained in WM. We used an inverted encoding model (IEM) to reconstruct CTFs from the pattern of EEG power across the scalp, which provided an assay of location-selective activity across the neuronal populations reflected in EEG activity. A clear pattern of results emerged over three independent experiments: while evoked and total power across a range of frequency bands transiently tracked the location of the sample stimulus, only total alpha power tracked the remembered location throughout the delay period. Furthermore, these location-specific patterns of alpha power showed the hallmark graded profile of a sensory population code. This result clearly demonstrates that alpha activity coordinates location-tuned neuronal populations rather than populations that code for more abstract variables, such as top-down control signals or categorical representations of location. Together, these results demonstrate that alpha activity is related to the sensory coding of spatial representations in WM.

Our findings build on a foundation of human neuroimaging work demonstrating that the feature content of visual WM can be recovered from voxel-wise patterns of activity (e.g. Christophel, Hebart, & Haynes, 2012; Emrich, Riggall, Larocque, & Postle, 2013; Ester, Anderson, Serences, & Awh, 2013; Ester et al., 2015; Harrison & Tong, 2009; Pratte & Tong, 2014; Riggall & Postle, 2012; Serences, Ester, Vogel, & Awh, 2009; Sprague et al., 2014). While stimulus-specific activity measured with functional magnetic resonance imaging (fMRI)

allows for reconstruction of feature-selective CTFs (e.g. Ester et al., 2013, 2015), the temporal precision of this approach is limited by the sluggish haemodynamic response. In a recent study, Garcia and colleagues reconstructed time-resolved CTFs that from patterns of power evoked by a flickering stimulus, highlighting the potential of combining IEMs with EEG recordings to obtained temporally-resolved stimulus-specific activity (Garcia et al., 2013). Here, we find that total alpha power tracks the content of spatial WM throughout a delay period, in the absence of rhythmic visual stimulation. Thus, the intrinsic role that alpha-band activity plays in spatial WM storage enables moment-by-moment tracking of location-specific activity, without the need for a flickering stimulus. Given that the topography of alpha power covaries with attended locations (Bahramisharif et al., 2011, 2010; Rihs et al., 2007; van Gerven & Jensen, 2009), we expect that this approach should also provide a powerful tool for tracking the dynamics of covert attention.

The IEM approach allowed us to reconstruct CTFs, providing an assay of location-specific activity across large populations of neurons (Serences & Saproo, 2012). While single neurons are the building block for sensory codes, it is the joint activity of a population of cells that guide behavior (Butts & Goldman, 2006; Jazayeri & Movshon, 2006). Thus, population-level encoding models provide a good approach for linking brain and behavior (Sprague et al., 2015). This approach has already enjoyed some success. For example, Sprague and colleagues (Sprague et al., 2014) used a spatial encoding model to obtain population-level reconstructions of stimulus location from patterns of activity measured with fMRI during a spatial WM task. Sprague and colleagues found that amplitude of spatial reconstructions tracked the decline in mnemonic precision that is seen with increasing set size in a WM tasks. The location-specific patterns of alpha power we report here provide a new window onto population-level coding.

Although we do not demonstrate it here, we expect that the topography of alpha activity should also track changes in the quality of population codes and predict behavior. Furthermore, because of the temporal resolution that EEG affords, this approach may be sensitive to rapid changes in the quality of population codes. Further work is necessary to test these predictions.

What are the cortical origins of the topographic pattern of alpha power that track the content of spatial WM? In studies of spatial attention, changes in topographic patterns of alpha power are thought to reflect synchronization of posterior visual areas tuned for unattended locations, reflecting suppression of processing in these regions (Kelly et al., 2006; Rihs et al., 2007; Thut et al., 2006). While it is tempting to conclude that the location-specific modulations of alpha activity reflect synchronization within visual cortex, it is difficult to infer the cortical source of oscillations based on EEG recordings alone. Neuroimaging studies have revealed location-specific activity in frontal, parietal, and occipital cortex during spatial attention and WM tasks (Silver & Kastner, 2009; Sprague et al., 2014; Sprague & Serences, 2013). Thus, it is possible that location-specific patterns of alpha power might reflect synchronization within or between any of the location-selective regions.

We found that alpha activity tracked stimulus location not only during the delay period, but also while the stimulus remained in view throughout an extended encoding period. Traditionally, visual WM has been characterized as a system for the maintenance of visual inputs that are no longer present. However, recent work has challenged this view, instead suggesting that WM also constrains representation of externally available stimuli (Chun, 2011; Tsubomi et al., 2013). For example, Tsubomi and colleagues demonstrated that memory-load dependent contralateral delay activity (CDA) – an electrophysiological marker of maintenance in visual

WM– showed a similar profile across both stimulus-absent and stimulus-present periods under the same task demands (Tsubomi et al., 2013). Here, we show that the alpha-band activity, like the CDA, plays a role in maintaining active spatial representations in both the presence and absence of visual input.

Our finding that the topography of alpha power tracks the content of spatial WM is consistent with the broad hypothesis that neuronal oscillations synchronize the cellular assemblies that code for mental representations (Fell & Axmacher, 2011; Hebb, 1949; Nicolelis, Fanselow, & Ghazanfar, 1997; Sejnowski & Paulsen, 2006; Singer, 1999; Singer & Gray, 1995; Watrous, Fell, Ekstrom, & Axmacher, 2015). Applied to WM, this view predicts that patterns of oscillatory activity should not only track spatial representations held in WM, but also non-spatial representations (e.g., color and orientation). Indeed, Salazar and colleagues found content-specific synchronization between frontal and parietal local field potentials in monkeys, peaking at 15 Hz. Fronto-parietal synchronization tracked both the location and identity (i.e., shape) of a remembered stimulus in a delayed-match to sample task, suggesting that both identity and location are encoded in patterns of synchronization (Salazar, Dotson, Bressler, & Gray, 2012). However, evidence for the role of oscillations in coordinating the code of non-spatial features is sparse, and further work is needed to examine whether neuronal synchrony plays a general role in coordinating feature-selective cellular assemblies that code the content of WM.

Conclusions

Using an inverted spatial encoding model, we demonstrated that the topographic distribution of alpha power tracks spatial representations held in WM. These findings show that

alpha-band activity plays a role in coding for spatial information held in WM, and this approach provides a time-resolved tool for tracking the content of spatial WM.

CHAPTER 4. ALPHA-BAND ACTIVITY REVEALS SPONTANEOUS REPRESENTATIONS OF SPATIAL POSITION IN VISUAL WORKING MEMORY

Summary

An emerging view suggests that spatial position is an integral component of working memory (WM), such that non-spatial features are bound to locations regardless of whether space is relevant (Rajsic & Wilson, 2014; Schneegans & Bays, 2017). For instance, past work has shown that stimulus position is spontaneously remembered when non-spatial features are stored. Item recognition is enhanced when memoranda appear at the same location where they were encoded (Dill & Fahle, 1998; D. H. Foster & Kahn, 1985; Zaksas, Bisley, & Pasternak, 2001), and accessing non-spatial information elicits shifts of spatial attention to the original position of the stimulus (Kuo et al., 2009; Theeuwes et al., 2011). However, these findings do not establish that a persistent, active representation of stimulus position is maintained in WM because similar effect have also been documented following storage in long-term memory (Martarelli & Mast, 2013; Nazir & O'Regan, 1990). Here, we show that the spatial position of the memorandum is actively coded by persistent neural activity during a non-spatial WM task. We used a spatial encoding model in conjunction with EEG measurements of oscillatory alpha-band (8–12 Hz) activity to track active representations of spatial position. The position of the stimulus varied trial-to-trial but was wholly irrelevant to the tasks. We nevertheless observed active neural representations of the original stimulus position that persisted throughout the retention interval. Further experiments established that these spatial representations are dependent on the volitional storage of non-spatial features rather than being a lingering effect of sensory energy or initial encoding demands. These findings provide strong evidence that online spatial representations are

spontaneously maintained in WM – regardless of task relevance – during the storage of non-spatial features.

Results

In Experiment 4-1, human observers performed a single-item color WM task (Figure 4-1a). Observers reported the color of a sample stimulus by clicking on a color wheel following a 1200-ms retention interval. Although the spatial position of the memorandum varied trial-to-trial, the position of the stimulus was wholly irrelevant to the task: observers never reported stimulus position, and the position of the test probe did not co-vary with the position of the memorandum. This design provided a strong test of whether observers spontaneously maintained an active representation of the irrelevant spatial position in WM. We assessed memory performance using the angular difference between the reported and studied color on the color wheel. We modeled these response errors as a mixture between a von Mises (circular normal) distribution and a uniform distribution (Zhang & Luck, 2008). The model fits established that observers rarely guessed, and precisely reported the color of the sample stimulus (Table 4-1).

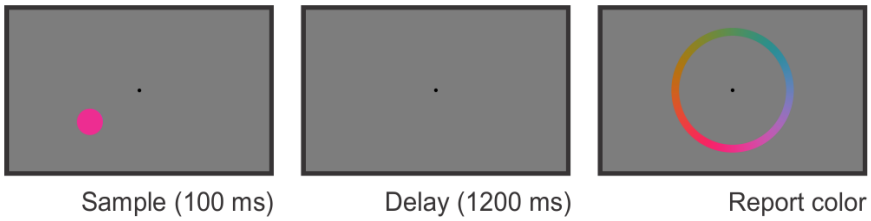

Figure 4-1. Color working memory task used in Experiment 4-1.
Observers in experiment 4-1 performed a color WM task. Observers saw a brief sample stimulus (100 ms). After a 1,200-ms retention interval, observers reported the color of the sample stimulus as precisely as possible by clicking on a color wheel. The position of the sample stimulus varied trial to trial but was wholly irrelevant to the task.

Table 4-1 Summary of parameter estimates (mean ± SEM) for mixture models fitted to response error distributions for each experiment.

	SD	P_g	P_s
Exp. 4-1 (360° color space)	13.5° ± 1.0	0.8% ± 0.3	N/A
Exp. 4-2a (360° color space)	16.5° ± 0.6	0.5% ± 0.2	0.3% ± 0.1
Exp. 4-2b (180° orientation space)	12.2° ± 0.6	2.6% ± 0.6	0.3% ± 0.1
Exp. 4-2c (180° orientation space)	10.1° ± 0.7	0.8% ± 0.3	N/A
Exp. 4-3 (360° color space)	15.0° ± 0.7	0.6% ± 0.1	1.2% ± 0.2

An emerging view suggests that spatial position is an integral component of WM representations (Rajsic & Wilson, 2014; Schneegans & Bays, 2017). Consistent with this view, behavioral studies have shown that observers spontaneously remember stimulus position when it is not task relevant (Dill & Fahle, 1998; D. H. Foster & Kahn, 1985; Kuo et al., 2009; Theeuwes et al., 2011; Zaksas et al., 2001). For example, item recognition is faster and more accurate when memoranda appear at the same location where they were encoded (Dill & Fahle, 1998; D. H. Foster & Kahn, 1985; Zaksas et al., 2001). However, behavioral signatures of spatial storage do not establish that a persistent, active representation of stimulus position is maintained in WM because they could instead reflect spatial representations that are stored in other memory systems (e.g., latent representations in episodic memory). Indeed, these behavioral signatures of storage of an irrelevant position are also seen following storage in long-term memory (Martarelli & Mast, 2013; Nazir & O'Regan, 1990). Therefore, it has remained unclear whether active representations of spatial position are spontaneously maintained in WM when position is irrelevant to the task. Here, we reasoned that persistent spatially selective neural activity tracking

the position of the memorandum could provide unambiguous evidence for online representations of space. Thus, we used ongoing oscillatory activity to test for active spatial representations during the delay period of non-spatial WM tasks.

We focused on oscillatory activity in the alpha-band (8–12 Hz) because past work has established that alpha-band activity encodes spatial positions that are actively maintained in WM (J. J. Foster et al., 2016). We used an inverted encoding model (IEM; J. J. Foster et al., 2016; Sprague et al., 2014; Sprague & Serences, 2013) to test for active alpha-band representations of the spatial position of the stimulus throughout the delay period (see Figure 4-2b-d and Materials and Methods). Our spatial encoding model assumed that the pattern of alpha-band power across the scalp reflects the activity of a number of spatially tuned channels (or neuronal populations; figures 4-2a and 4-2b). In a training phase (Figure 4-2c), we used a subset of the EEG data during the color WM task to estimate the relative contribution of these channels to each electrode on the scalp (the *channel weights*). Then, in a test phase (Figure 4-2d), using an independent subset of data, we *inverted* the model to estimate the response of the spatial channels from the pattern of alpha power across the scalp. This procedure produces a profile of responses across the spatial channels (termed channel-tuning functions or CTFs), which reflects the spatial selectivity of the large-scale neuronal populations that are measured by scalp EEG (J. J. Foster et al., 2016; Lopes da Silva, 2013; Nunez & Srinivasan, 2006). We performed this analysis at each time point throughout the trial, which allowed us to test whether active spatial representations were maintained throughout the retention interval.

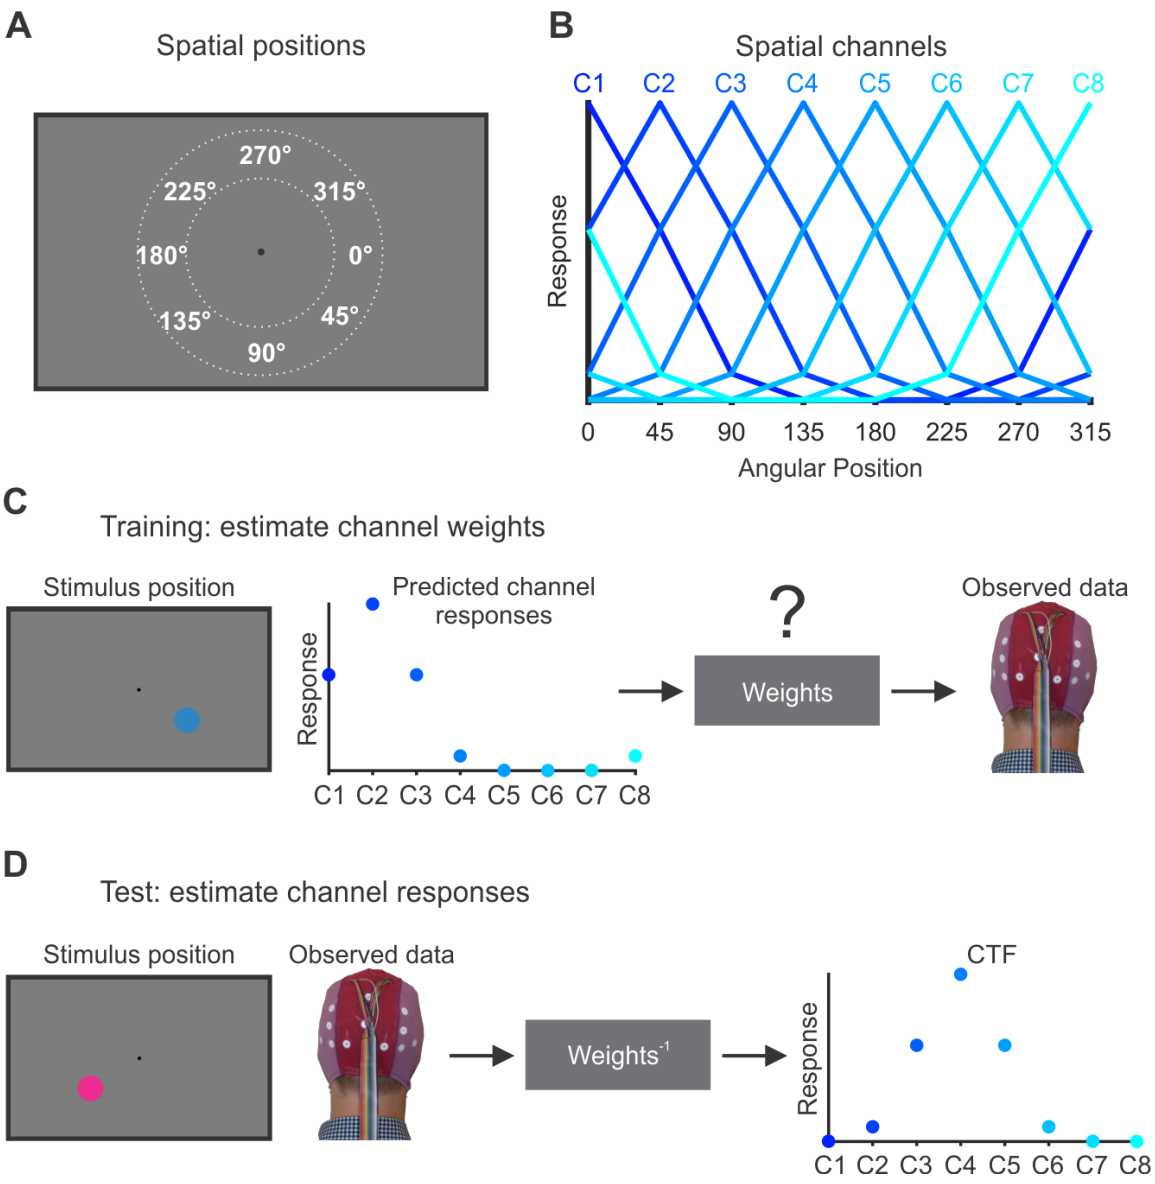

Figure 4-2. The inverted encoding model for reconstructing spatial channel-tuning functions.

(a) The sample stimulus could appear anywhere along an isoeccentric band around the fixation (dotted white lines). We categorized stimuli as belonging to one of eight positions bins centered at 0°, 45°, 90°, and so forth. Each position bin spanned a 45° wedge of positions (e.g., 22.5° to 67.5° for the bin centered at 45°). (b) We modeled oscillatory power at each electrode as the weighted sum of eight spatially selective channels (C1–C8), each tuned for the center of one of the eight position bins shown in (a). Each curve shows the predicted response of one of the channels across the eight positions bins (i.e., the "basis function"). (c) In the training phase, we used the predicted channel responses, determined by the basis functions shown in (b), to estimate a set of channel weights that specified the contribution of each spatial channel to the response measured at each electrode. The example shown here is for a stimulus presented at 45°. (d) In the

Figure 4-2, continued..
test phase, using an independent set of data, we used the channel weights obtained in the training phase to estimate the profile of channel responses given the observed pattern of activity across the scalp. The resulting CTF reflects the spatial selectivity of population-level oscillatory activity, as measured using EEG. The example shown here is for a stimulus presented at 135°. For more details, see Materials and Methods.

Our analysis revealed that precise representations of stimulus position were maintained in WM. Alpha-band CTFs revealed robust spatial selectivity, which persisted throughout the retention interval (Figure 4-3a). A permutation testing procedure (see Materials and Methods) confirmed that this spatial selectivity was reliably above chance throughout the retention interval (Figure 4-3b). In line with past work (J. J. Foster et al., 2016), we found that this spatially selective delay activity was restricted to the alpha-band (Figure 4-4). To examine the precision of the spatial representations encoded by alpha-band activity, we inspected the time-resolved channel-response profiles for each of the eight position bins separately (Figure 4-3c). For each position bin, the peak response was seen in the channel tuned for that position (i.e., a channel offset of 0°), which established that the recovered alpha-band CTFs tracked which of the eight positions that the stimulus had appeared in. Thus, stimulus position was precisely represented in alpha-band activity. In summary, Experiment 4-1 established that observers spontaneous maintained an active representation of stimulus position that persisted throughout the delay period of a color WM task, even though spatial position was completely irrelevant.

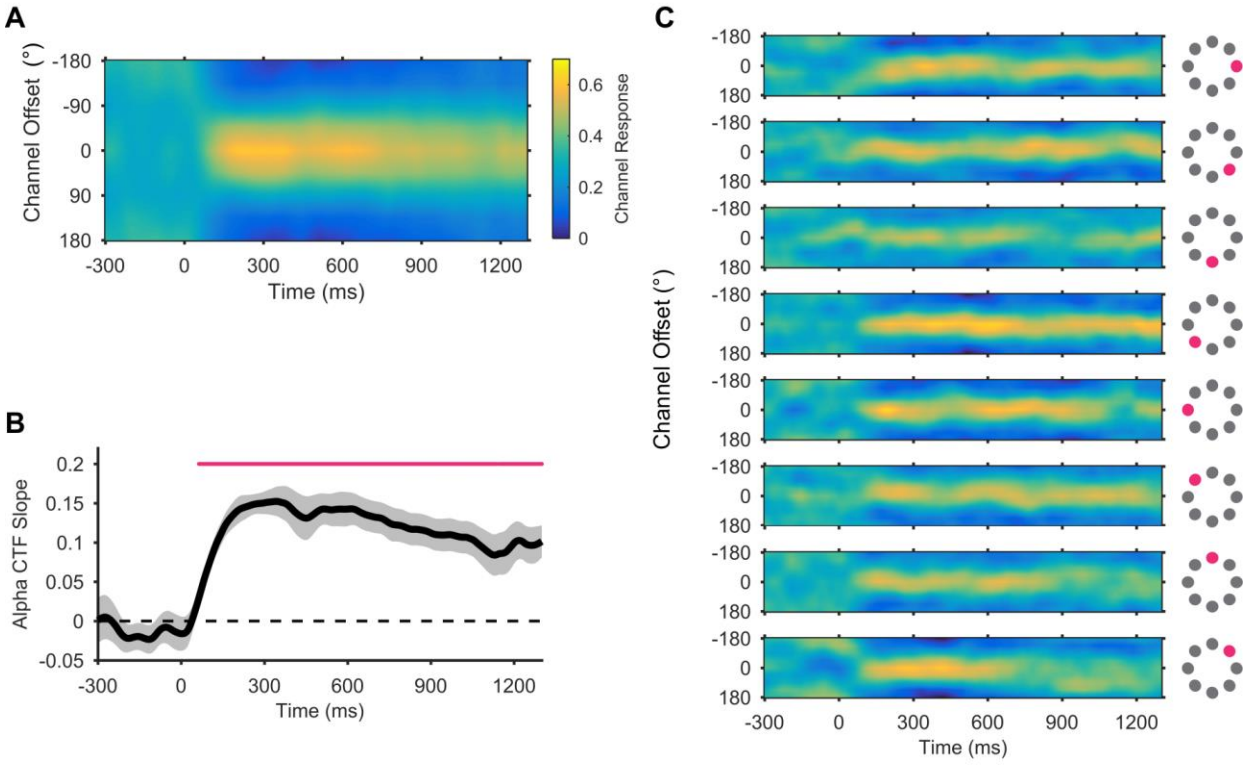

Figure 4-3. Spatial alpha-band CTFs during the storage of a color stimulus in working memory.
(a) Average alpha-band CTF in Experiment 4-1. Although the position of the stimulus was irrelevant to the task, we observed a robust, spatially selective alpha-band CTF throughout the delay period. (b) The selectivity of the alpha-band CTF across time (measured as CTF slope, see Materials and Methods). The magenta marker shows the period of reliable spatial selectivity. The shaded error bar reflects ±1 bootstrapped SEM across subjects. (c) Alpha-band CTFs for each position separately. The magenta markers show to which of the eight stimulus position bins each subplot corresponds.

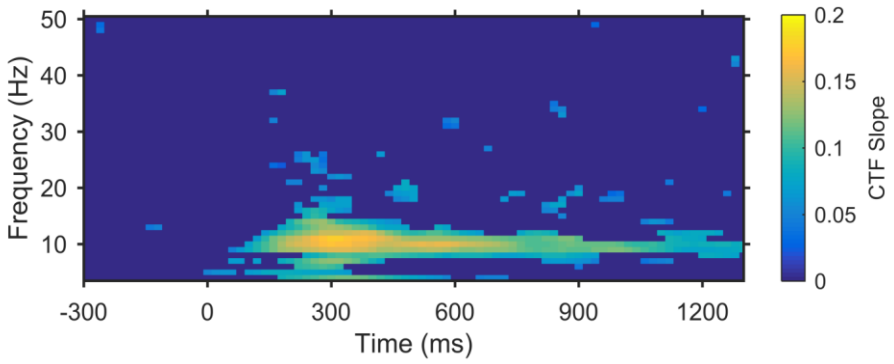

Figure 4-4. Spatially selective oscillatory activity is specific to the alpha band (8–12 Hz). CTF selectivity (measured as CTF slope) as a function of time and frequency (4–50 Hz, in 1-Hz bands) Points at which CTF selectivity was not reliable above chance, as determined by permutation testing, are set to zero (dark blue). To reduce computation time, we down-sampled power values from 250 Hz to 50 Hz (i.e., one sample every 20 ms). Note that we down-sampled power values after filtering and applying the Hilbert transform so that down-sampling did not affect how power values were obtained.

Another possibility is that the decoded spatial activity reflects lingering sensory activity that was evoked by the sample stimulus, rather than the online maintenance of the stimulus color in WM. In experiments 4-2a and 4-2b, we tested whether spatial alpha-band CTFs depend on the volitional storage goals of the observer. Observers performed a selective storage task in which they were instructed to store a non-spatial feature of a target that was presented alongside a distractor. In Experiment 4-2a, observers performed a color version of this task (Figure 4-5a): observers saw two shapes (a circle and a triangle) and were instructed to remember the color of the target shape, which was consistent throughout the experiment (see Materials and Methods). In Experiment 4-2b, observers performed an orientation version of the selective storage task (Figure 4-5b): observers saw two oriented gratings (one blue and one green) and were required to remember the orientation of the grating in the target color. Modeling of response errors revealed that observers rarely guessed or misreported the distractor item instead of the target (Table 4-1).

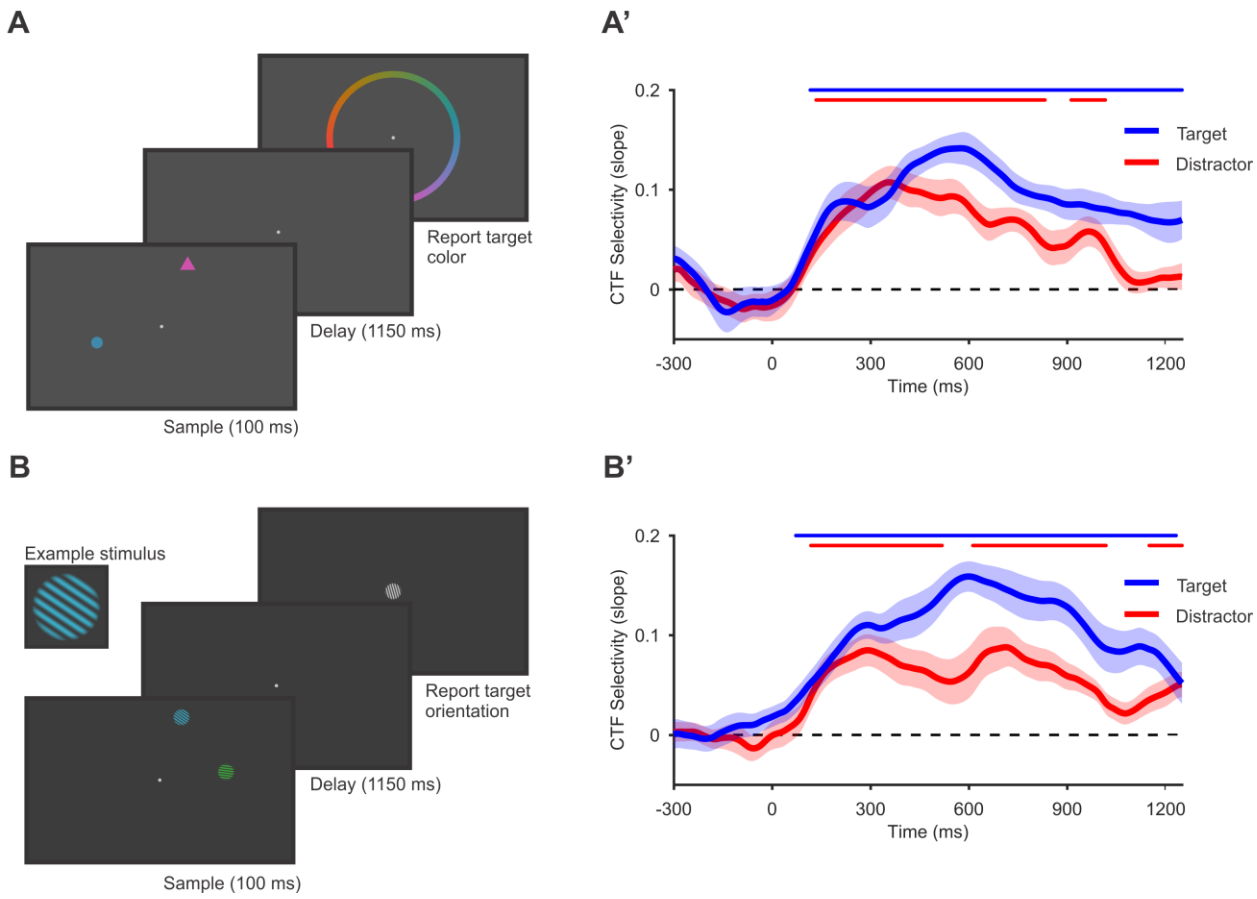

Figure 4-5. Tasks and results for experiments 4-2a and 4-2b.
(a) In experiment 4-2a, observers performed a selective storage task. The sample display contained two colored shapes (a circle and a triangle). One shape served as a target and the other as a distractor. We instructed observers to remember and report the color of the target and to ignore the distractor. The target shape did not change throughout the session and was counterbalanced across observers. (b) In experiment 4-2b, observers performed an orientation version of the selective storage task used in experiment 4-2a. In this experiment, observers saw two oriented gratings (one green and one blue) and were instructed to remember and report the orientation of the grating in the target color (counterbalanced across observers). (a'–b') Spatial selectivity of alpha-band CTFs in experiments 4-2a and 4-2b. In experiment 4-2a (a') and experiment 4-2b (b'), we examined the spatial selectivity (measured as CTF slope) for the target position (blue) and distractor position (red). In both experiments, spatial selectivity was higher for the target-related CTF than for the distractor-related CTF. This modulation of CTF selectivity by storage goals reveals a component of spatial selectivity that is related to the volitional storage of non-spatial features in WM. The markers at the top of each plot mark the periods of reliable spatial selectivity. All shaded error bars reflect ±1 SEM.

In both experiments, we varied the spatial positions of the target and distractor independently, which allowed us to compare target- and distractor-related alpha-band CTFs. By comparing target- and distractor-related CTFs, we were able to isolate spatially selective activity that must be related to the volitional storage of the target in WM. If sustained alpha-band CTFs are an automatic consequence of sensory activity evoked by the onset of a visual stimulus, then we should see identical alpha-band CTFs for the target and distractor items. Instead, we found that CTF selectivity was higher for the target location than for the distractor location (figures 4-5a' and 4-5b'). Bootstrap resampling tests confirmed that delay-period CTF selectivity (averaged from 100–1250 ms after stimulus onset) was reliably greater for the target location than for the distractor location in both experiments (Experiment 4-2a: $p < 0.001$; Experiment 4-2b: $p < 0.001$). While it is likely that there is a sensory contribution to the alpha-band CTFs, the amplified CTF selectivity observed for the target shows that these spatial representations cannot be wholly explained by stimulus-driven activity. Instead, they are strongly shaped by the observer's volitional storage goals.

Although we saw clear modulation of CTFs by the relevance of the stimulus, we also saw a spatial CTF for the distractor location that lasted for most of the retention interval (Figures 4-5a' and 4-5b'). In a further analysis, we found that the representation of the distractor position resembled that of the target position (Figure 4-6). We think it is unlikely that the sustained distractor-related CTF reflects lingering sensory activity alone. While observers have top-down control over the stimuli that are stored in WM, this control is imperfect, resulting in the unnecessary storage of irrelevant items (McNab & Klingberg, 2008; Vogel, McCollough, &

Machizawa, 2005). Therefore, the spatially specific distractor-related activity may in part reflect the unnecessary storage of the distractor in WM.

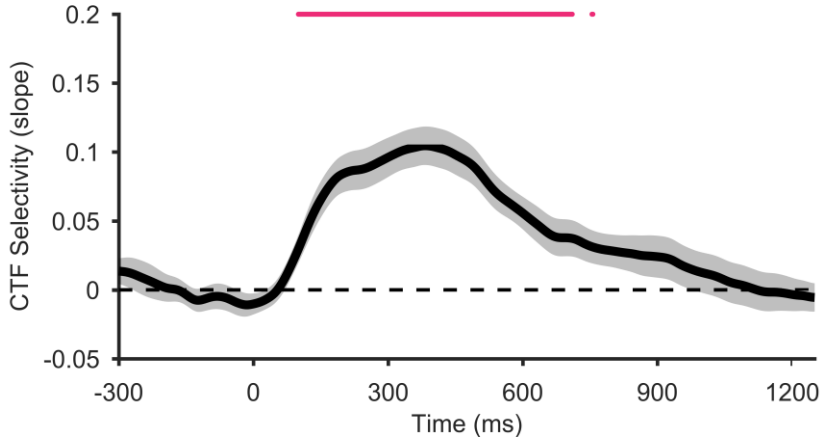

Figure 4-6. Spatial representations of the distractor position resemble the spatial representation of the target position.
In our standard IEM procedure, we obtained target-related CTFs by training and testing on the target position (see Materials and Methods). Similarly, we obtained distractor-related CTFs by training and testing on the distractor position. This analysis leaves open the possibility that the pattern of alpha-band activity that represents target and distractor positions is different. To test this possibility, we ran an analysis in which we trained the IEM on the target position and tested on the distractor position using the data from Experiment 4-2a (see Materials and Methods). The plot shows the selectivity of resulting distractor-related CTF across time (measured as CTF slope). We observed a reliable distractor-related CTF when training the IEM on the target position. This result demonstrates that the alpha-band representation of the distractor position resembles that of the target position. The magenta marker at the top of the plot shows the period of reliable spatial selectivity. The shaded error bar reflects ±1 bootstrapped SEM across subjects.

In Experiment 4-2c, we used a different approach to eliminate stimulus-driven activity as a source of spontaneous spatial representations during the delay period. We used a balanced visual display in which eight items were equally spaced around the fixation point (Figure 4-7a). We instructed observers to remember the orientation of the line in the target item (defined by color, with the relevant color varied across observers), and to reproduce the target orientation following the retention interval. Because the sample displays were visually balanced, stimulus-

driven activity was not spatially selective for the target position. Therefore, any spatially selective activity must be related to storage of the target item in WM. We observed a robust, target-related alpha-band CTF (Figure 4-7b), which emerged between 200 and 300 ms after onset of the stimulus array, and sustained throughout the retention interval. This observation provides a clean look at the spatial alpha-band representation when spatially selective stimulus-driven activity is eliminated. Taken together, experiments 4-2a-c provide clear evidence for a persistent active representation of the spatial position of memoranda stored in WM, even when stimulus-driven activity is controlled for.

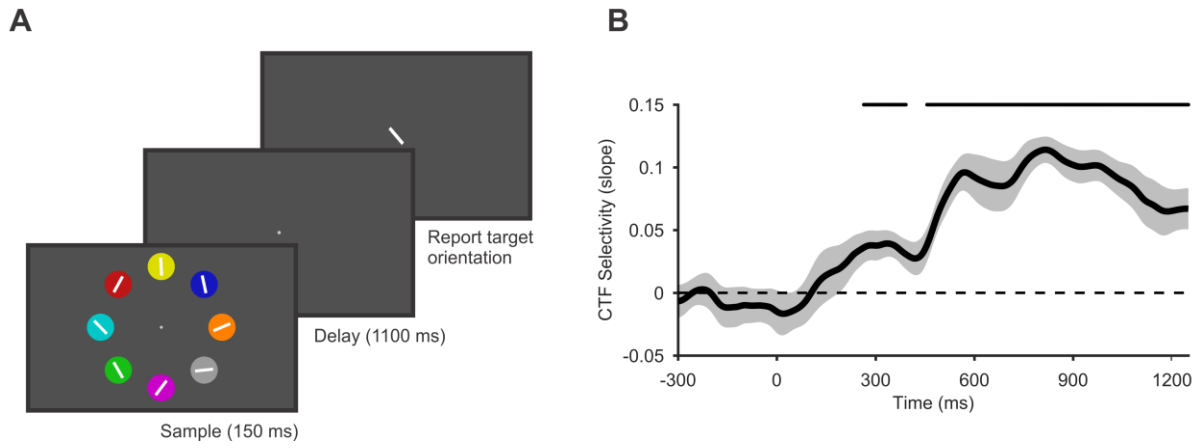

Figure 4-7. Task and results for Experiment 4-2c.
(a) Observers saw a balanced array of stimuli and were instructed to remember and report the orientation of the line in the target-colored circle. The target color did not change throughout the experiment and was varied across observers. The markers at the top of each plot mark the periods of reliable spatial selectivity. (b) We saw a robust alpha-band CTF that tracked the target position throughout the retention interval. In this experiment, we completely eliminated spatially specific stimulus-driven activity by presenting the target stimulus in a balanced visual display. Thus, the alpha-band CTF in this experiment must be related to the storage of the target orientation in WM. Shaded error bars reflect ±1 SEM.

Finally, we considered whether biased attention at the time of encoding could explain the stronger spatial selectivity observed for the target relative to the distractor items in experiments 2a-c. Given that the target-defining feature was constant during experiments 4-2a and 4-2b, it is possible that observers could have encoded the relevant feature of the target without fully directing attention towards the distractor; this may have yielded a lingering advantage for the target during the subsequent delay period. Experiment 4-3 tested this possibility with a task that required full attention towards both the target and distractor items. Observers saw two colored circles, each containing a digit, and the target was the circle that contained the larger digit (Figure 4-8a). Again, we instructed observers to remember the color of the target, and to disregard the distractor. Critically, this task forced subjects to attend both items to identify the target (Pashler & Badgio, 1985). Thus, any observed difference in alpha-band CTFs for the target and distractor cannot be attributed to biased attention toward the target during encoding. Indeed, we did not see any difference in CTF selectivity in the 500-ms period following onset of the sample display (0–500 ms, $p = 0.42$, bootstrap resampling test; Figure 4-8b), which far exceeds the time required to encode simple visual objects (Vogel, Woodman, & Luck, 2006). In contrast, when we examined the entire delay period (150–1500 ms), we found that alpha-band CTF selectivity was higher for the target than for the distractor throughout the delay period ($p < 0.01$, bootstrap resampling test). Together, our results provide clear evidence that alpha-band activity tracks active spatial representations that persist during the storage of non-spatial features in WM.

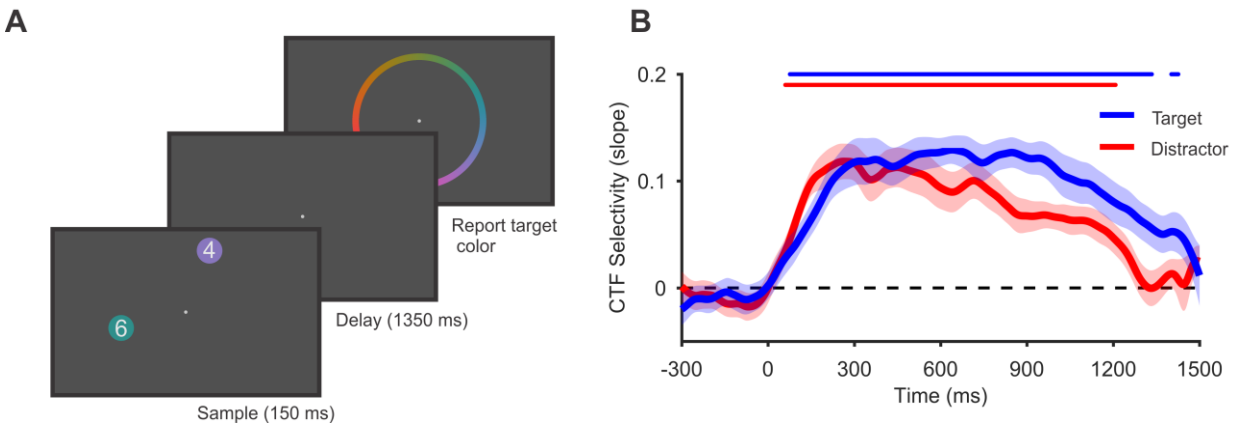

Figure 4-8. Task and results for Experiment 4-3.
(a) In Experiment 4-3, observers performed a variant on the selective storage task used in experiments 2a and 2b. Observers saw two colored circles, each containing a digit (between 0 and 7). The target was the circle that contained the larger digit. Observers were instructed to remember and report the color of the target and to disregard the distractor. Critically, this task forced observers to attend and encode both items in order to determine which one was the target. (b) Spatial selectivity (measured as CTF slope) for the target position (blue) and distractor position (red). Spatial selectivity was higher for the target-related CTF than for the distractor-related CTF. This modulation of CTF selectivity by storage goals cannot be explained by differential attention to the target during encoding, because our task forced subjects to attend and encode both items. Therefore, this modulation of CTF selectivity must reflect the continued maintenance of the target item in WM. The markers at the top of the plot mark the periods of reliable spatial selectivity. Shaded error bars reflect ±1 SEM.

Discussion

Past work has established that human observers can voluntarily control which visual features are stored in WM (Serences et al., 2009; Woodman & Vogel, 2008; Yu & Shim, 2017). For example, Serences and colleagues (2009) instructed observers to remember the color or orientation of a grating during a retention interval. Voxel-wise patterns of activity in visual cortex were measured with functional magnetic resonance imaging (fMRI) and revealed active coding of the relevant dimension but not the irrelevant dimension. Likewise, Woodman and Vogel (Woodman & Vogel, 2008) showed that contralateral delay activity, an electrophysiological marker of WM maintenance, was higher in amplitude for orientation stimuli

than for color stimuli. Critically, this "orientation bump" was only seen when observers voluntarily stored the orientation dimension of conjunction stimuli. These studies demonstrated that observers can control which aspects of the stimuli are held in visual WM, such that specific stimulus dimensions are excluded when they are behaviorally irrelevant. Here, we showed that stimulus position is a striking exception to this rule. We observed active representations of stimulus position that were spontaneously maintained throughout the delay period, even though location was never relevant to the task.

This finding lends support to models that posit a central role for space in the storage of non-spatial information in WM (Rajsic & Wilson, 2014; Schneegans & Bays, 2017). It has been a longstanding hypothesis that space is an integral component of an observer's representation of non-spatial information (Johnston & Pashler, 1990; Nissen, 1985; Tsal & Lavie, 1993). However, evidence for spontaneous representation of stimulus position in WM has been elusive. In non-human primates, there is evidence that spatial position is spontaneously represented when irrelevant (Salazar et al., 2012; Sereno & Amador, 2006). However, these studies found that non-spatial dimensions (e.g., shape) were also spontaneously represented when irrelevant (Sereno & Amador, 2006), raising the possibility that these findings might reflect a general failure of non-human primates to exclude any behaviorally irrelevant feature – spatial or non-spatial alike – from WM. In contrast, it has been well-established that human observers exclude irrelevant non-spatial features from storage in visual WM (Serences et al., 2009; Woodman & Vogel, 2008; Yu & Shim, 2017). Therefore, our finding that human observers spontaneously maintain spatial position in WM provides clear evidence that spatial position holds a special status in WM.

Our IEM approach also sheds light on the nature of spontaneous spatial representations in WM. By presenting stimuli at many spatial positions in conjunction with an IEM, we established that alpha-band activity precisely tracked the position of the stimulus (Figure 4-3c). This result rules out the possibility that spatially selective activity reflects an imprecise spatial signal that does not track the specific location of the stimulus (e.g., a signal that tracks hemifield). Furthermore, because scalp EEG activity reflects the synchronous activity of large neuronal populations (Lopes da Silva, 2013; Nunez & Srinivasan, 2006), our findings show that spatial position was robustly represented in a large-scale population code rather than an isolated group of spatially selective neuronal units.

Do these spontaneous spatial representations play a functional role in the online maintenance of non-spatial features? Past work has shown that observers recognize an item faster and more accurately when it appears at the location where it was encoded than when it appears elsewhere (Dill & Fahle, 1998; D. H. Foster & Kahn, 1985; Zaksas et al., 2001), thereby suggesting that access to a non-spatial feature is intertwined with spatial memory. Furthermore, Williams and colleagues (Williams et al., 2013) showed that color WM performance was impaired when observers were prevented from fixating or covertly attending the positions where the memoranda were presented. Thus, spatial attention toward the original locations of the items improved performance in a color WM task. That said, recognition memory in Williams and colleagues' study required knowledge of stimulus position. Thus, more work is needed to determine whether sustained spatial focus on an item's initial position will enhance the maintenance of non-spatial features alone.

A substantial body of work has linked spatially specific alpha-band activity with covert spatial attention (J. J. Foster, Sutterer, Serences, Vogel, & Awh, 2017; Foxe & Snyder, 2011; Jensen & Mazaheri, 2010; Thut et al., 2006), which raises the possibility that the spatial representations that we observed reflect sustained attention to the original location of the stimulus. If this activity does reflect covert spatial attention, does it qualify as a working memory for spatial position? Persistent stimulus-specific neural activity is an unambiguous signature of active maintenance in WM (Sreenivasan, Curtis, & D'Esposito, 2014), and the sustained spatially-selective activity that we observed satisfies this criterion. Furthermore, a broad array of evidence has shown that there is considerable overlap between attention and WM (Awh & Jonides, 2001; Awh et al., 2006; Gazzaley & Nobre, 2012), such that spatial attention supports maintenance of both locations and non-spatial features in WM (Awh, Anllo-Vento, & Hillyard, 2000; Awh, Jonides, & Reuter-Lorenz, 1998; Williams et al., 2013). Thus, while it is difficult to draw a precise line between spatial attention and spatial WM, our findings provide clear evidence that storing non-spatial features in WM elicits the spontaneous and sustained maintenance of an online spatial representation. Indeed, if the observed spatial representations are best interpreted as sustained covert attention, it is striking that this spatial focus was maintained throughout the delay period even though space was completely irrelevant to the task. This empirical pattern contrasts with findings from explicit studies of spatial attention, where behavioral relevance determines whether attention is sustained at a given location (Carrasco, 2011; Egeth & Yantis, 1997). Therefore, this perspective on our findings also highlights the special status of spatial position in visual WM.

In Summary, we found that human observers spontaneously maintained active neural representations of stimulus position that persisted during the storage of non-spatial features in WM. These spatial representations were precise and robustly coded by population-level alpha-band activity. Although human observers can exclude non-spatial features (e.g., orientation or color) from WM when they are irrelevant (Serences et al., 2009; Woodman & Vogel, 2008; Yu & Shim, 2017), our results show robust and sustained spatial representations despite hundreds of trials in which spatial position was behaviorally irrelevant.

Materials and Methods

Subjects

Subjects were human adults between 18 and 35 years old, who reported normal or corrected-to-normal vision. Twelve subjects (7 female, 5 male) participated in Experiment 4-1, 21 subjects (5 female, 16 male) participated in Experiment 4-2a, 20 subjects (14 female, 6 male) participated in Experiment 4-2b, 18 subjects (8 female, 10 male) participated in Experiment 4-2c, and 19 subjects (8 female, 11 male) participated in Experiment 4-3. Subjects in Experiment 4-1 provided informed consent according to procedures approved by the University of Oregon Review board. Subjects in all other experiments provided informed consent according to procedures approved by the University of Chicago Institutional Review Board.

Apparatus and Stimuli

We tested subjects in a dimly lit, electrically shielded chamber. Stimuli were generated using Matlab (The Mathworks, Natick, MA) and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). In Experiment 4-1, stimuli were presented on a 17-in CRT monitor (refresh rate = 60 Hz) at a viewing distance of 100 cm. In experiments 4-2a and 4-2b, stimuli were presented on

a 24-in LCD monitor (refresh rate = 120 Hz) at a viewing distance of 100 cm. In experiments 4-2c and 4-3, stimuli were presented on a 24-in LCD monitor (refresh rate = 120 Hz) at a viewing distance of 77cm.

Task Procedures

Experiment 4-1. Subjects performed a delayed color-estimation task (Figure 4-1). On each trial, subjects saw a sample stimulus and reported its color as precisely as possible following a retention interval. Subjects initiated each trial with a spacebar press. A fixation point (0.24° of visual angle) appeared for 800-1500 ms before a sample stimulus appeared for 100 ms. The sample stimulus was a circle (1.6° in diameter), centered 3.8° of visual angle from the fixation point. The angular position of the stimulus varied trial-to-trial, and was sampled from one of eight position bins around fixation (each bin spanned a 45° wedge of angular positions, bins were centered at 0°, 45°, 90°, and so forth). Critically, the position of the stimulus was irrelevant to the task, and subjects were told that this was the case at the start of the session. The color of the sample stimulus was chosen from a continuous color wheel that included 360 different colors. Following a 1200-ms retention interval during which only the fixation point remained on screen, a color wheel (8.0° in diameter, 0.5° thick) appeared around fixation, and subjects reported the color of the sample stimulus by clicking on the color wheel with a mouse. The orientation of the color wheel was randomized on each trial to ensure that subjects could not plan their response until the color wheel appeared. Before starting the task, subjects completed a brief set of practice trials to ensure that they understood the instructions. Subjects then completed 15 blocks of 64 trials (960 trials in total), or as many blocks as time-permitted (all subjects

completed at least 13 blocks; one subject completed 2 extra blocks, thus completed 17 blocks in total).

Experiments 4-2a and 4-2b. Subjects in Experiments 4-2a and 4-2b performed a selective storage version of the delayed-estimation task (figures 4-5a and 4-5b). In this procedure, two sample stimuli were presented on each trial (a *target* and a *distractor*). Subjects were instructed to remember the color (Experiment 4-2a) or orientation (Experiment 4-2b) of the target, and to ignore the distractor. As in Experiment 4-1, the angular positions of each stimulus around the fixation point were drawn from eight position bins, each spanning a 45° wedge of angular positions. The position bins that the target and distractor stimuli occupied were fully counterbalanced across trials for each subject. Thus, the position of one stimulus was random with respect to the other, allowing us to reconstruct spatial CTFs for the target and the distractor positions independently, disregarding the other. When the target and distractor stimuli occupied the same position bin, their exact position within the bin was constrained so that they did not overlap.

In Experiment 4-2a, subjects performed a color version of the selective storage task (Figure 4-5a). Subjects initiated each trial with a spacebar press. A fixation point (0.2°) appeared for 500-800 ms. Next, the sample display was presented for 100 ms. The sample display contained two stimuli: a circle (0.7° in diameter) and an equilateral triangle (length of edges were 0.94°, such that the circle and triangle were equated for area). The target shape was counterbalanced across subjects. Each shape was presented 4° of visual angle from the fixation point. Subjects were instructed to remember the color of the target shape (circle or triangle) as precisely as possible, and to ignore the distractor shape. Following an 1150-ms retention interval,

during which only the fixation point remained on the screen, a color wheel (8° in diameter, 0.4° thick) appeared centered around the fixation point, and subjects reported the color of the sample stimulus by clicking on the color wheel with a mouse. As in Experiment 4-1, the orientation of the color wheel was randomized on each trial to ensure that subjects could not plan their response until the color wheel appeared. Before starting the task, subjects completed a brief set of practice trials to ensure that they understood the instructions. Subjects then completed 15 blocks of 64 trials (960 trials in total), or as many blocks as time permitted (all subjects completed at least 11 blocks, one subject completed 2 extra blocks, thus completed 17 blocks in total).

In Experiment 4-2b, subjects performed an orientation version of the selective storage task (Figure 4-5b). The timing of the task was identical to Experiment 4-2a. Here, the sample display contained two oriented gratings (diameter, 1°; spatial frequency, 7 cycles/°), one green and one blue (equated for luminance), each presented 4° of visual angle from the fixation point. The spatial phase of each grating was randomized on each trial. Subjects were instructed to remember the orientation of the grating in the target color (blue or green) as precisely as possible, and to ignore the grating in the distractor color. The target color was counterbalanced across subjects. Following the retention interval, a white grating appeared at fixation, and subjects adjusted the orientation of this probe grating to match the remembered orientation of the target using a response dial (PowerMate USB Multimedia Controller, Griffin Technology, USA). The initial orientation of the probe grating was randomized on each trial to ensure that subjects could not plan their response until the probe appeared. Once subjects had finished adjusting the orientation of the probe grating, they registered their response by pressing the spacebar. Before

starting the task, subjects completed a brief set of practice trials to ensure that they understood the instructions. Subjects then completed 15 blocks of 64 trials (960 trials in total).

Experiment 4-2c. In Experiment 4-2c, a single target item was presented among seven distractors (Figure 4-7a). Subjects were instructed to remember the orientation of the line inside the circle in the target color (which was varied across subjects), and to ignore the distractor items. Subjects initiated each trial with a spacebar press. A fixation point (0.2°) appeared for 500-800 ms. Next, the sample display was presented for 150 ms. This display contained eight circles (1.8° in diameter), equally spaced around the fixation (each item was centered 4° of visual angle from the fixation point). In this experiment, the eight stimuli always occupied fixed locations at 0°, 45°, 90°, and so forth. Each circle contained a white line (1.2° long, 0.2° wide) presented at a randomly selected orientation. Each of the eight circles were presented in a different color (red, pink, yellow, green, light blue, dark blue, orange, and grey). These colors were chosen because they were easily discriminated. The target (i.e., the circle in the target color) appeared in each of the eight positions equally often, and the positions of the remaining seven colors were randomized. Following an 1100-ms retention interval, a white line appeared at fixation and subjects adjusted the orientation to match the remembered target orientation using a response dial. The initial orientation of the probe was randomized on each trial to ensure that subjects could not plan their response until the probe appeared. Once subjects had finished adjusting the orientation of the probe, they registered their response by pressing the spacebar. Immediately following their response, subjects were shown their response error (i.e., the angular difference between the target orientation and reported orientation) for 500 ms. Before starting the

task, subjects completed a practice block of 64 trials to ensure that they understood the instructions. Subjects then completed 16 blocks of 64 trials (1024 trials in total).

Experiment 4-3. Subjects in Experiment 4-3 performed a variant of the two-item selective storage task used in experiments 4-2a and 4-2b (Figure 4-8a). In this experiment, both stimuli (the target and distractor) were colored circles (1.8° in diameter) containing a digit (1.6° tall and 0.8° wide) between 0 and 7. Subjects were instructed to remember and report the color of circle that contained the larger digit (i.e., the target). On each trial, the target digit was randomly sampled from the digits 1-7, and the distractor was randomly sampled from the subset of digits that were less than that the target. Critically, this task ensured that subjects had to attend and encode both stimuli to the extent that they could identify the digit in each of the items, in order to determine which item was the target (Pashler & Badgio, 1985). The sample display was presented for 150 ms, and was followed by a 1350-ms retention interval. The trial timing was otherwise identical to that in experiments 2a and 2b. In this experiment, subjects were shown their response error for 500 ms immediately following their response. Before starting the task, subjects completed a practice block of 64 trials to ensure that they understood the instructions. Subjects then completed 16 blocks of 64 trials (1024 trials in total).

EEG Recording

Experiment 4-1. We recorded EEG using 20 tin electrodes mounted in an elastic cap (Electro-Cap International, Eaton, OH). We recorded from International 10/20 sites: F3, Fz, F4, T3, C3, Cz, C4, T4, P3, Pz, P4, T5, T6, O1, and O2, along with five nonstandard sites: OL midway between T5 and O1, OR midway between T6 and O2, PO3 midway between P3 and OL, PO4 midway between P4 and OR, and POz midway between PO3 and PO4. All sites were

recorded with a left-mastoid reference, and were re-referenced offline to the algebraic average of the left and right mastoids. Eye movements and blinks were monitored using electrooculogram (EOG). To detect horizontal eye movements, horizontal EOG was recorded from a bipolar pair of electrodes placed ~1 cm from the external canthus of each eye. To detect blinks and vertical eye movements, vertical EOG was recorded from an electrode placed below the right eye and referenced to the left mastoid. The EEG and EOG were amplified with an SA Instrumentation amplifier with a bandpass of 0.01 to 80 Hz and were digitized at 250 Hz using LabVIEW 6.1 running on a PC. Impedances were kept below 5 kΩ.

All other experiments. We recorded EEG activity using 30 active Ag/AgCl electrodes mounted in an elastic cap (Brain Products actiCHamp, Munich, Germany). We recorded from International 10-20 sites: FP1, FP2, F7, F3, Fz, F4, F8, FC5, FC1, FC2, FC6, C3, Cz, C4, CP5, CP1, CP2, CP6, P7, P3, Pz, P4, P8, PO7, PO3, PO4, PO8, O1, Oz, and O2. Two additional electrodes were placed on the left and right mastoids, and a ground electrode was placed at position FPz. All sites were recorded with a right-mastoid reference, and were re-referenced offline to the algebraic average of the left and right mastoids. Eye movements and blinks were monitored using EOG, recorded with passive electrodes. Horizontal EOG was recorded from a bipolar pair of electrodes placed ~1 cm from the external canthus of each eye. Vertical EOG was recorded from a bipolar pair of electrodes placed above and below the right eye. Data were filtered online (low cut-off = .01 Hz, high cut-off = 80 Hz, slope from low- to high-cutoff = 12 dB/octave), and were digitized at 500 Hz using BrainVision Recorder (Brain Products, Munich, Germany) running on a PC. For three subjects in Experiment 4-2b, data were digitized at 1000 Hz because of an experimenter error. The data for these subjects were down-sampled to 500 Hz

offline. Impedances were kept below 10 kΩ. For one subject in Experiment 3c, data from one electrode (F8) was discarded because of excessive noise.

Eye Tracking

In experiments 4-2a-c and 4-3, we monitored gaze position using a desk-mounted infrared eye tracking system (EyeLink 1000 Plus, SR Research, Ontario, Canada). According to the manufacturer, this system provides spatial resolution of 0.01° of visual angle, and average accuracy of 0.25-0.50º of visual angle. In Experiments 2a and 2b, gaze position was sampled at 500 Hz, and data were obtained in remote mode (without a chin rest). In Experiments 4-2c and 4-3, gaze position was sampled at 1000 Hz, and head position was stabilized with a chin rest. We obtained usable eye-tracking data for 11, 11, 13, and 15 subjects in experiments 4-2a, 4-2b, 4-2c, and 4-3, respectively.

Artifact Rejection

We visually inspected EEG for recording artifacts (amplifier saturation, excessive muscle noise, and skin potentials), and EOG for ocular artifacts (blinks and eye movements). For subjects with usable eye tracking data, we also inspected the gaze data for ocular artifacts. We discarded trials contaminated by artifacts. Subjects were excluded from the final samples if fewer than 600 trials remained after discarding trials contaminated by recording or ocular artifacts. For the analyses of gaze position, we further excluded any trials in which the eye tracker was unable to detect the pupil, operationalized as any trial in which gaze position was more than 15° of visual angle from the fixation point.

Subject Exclusions

Experiment 4-1. Two subjects were excluded because of excessive artifacts (see Artifact Rejection). The final sample included 10 subjects with an average of 785 (*SD* = 92) artifact-free trials.

Experiment 4-2a. Four subjects were excluded because of excessive artifacts, and one subject was excluded due to poor task performance (prevalence of guesses and swaps ~14%, estimated by fitting a mixture model to response errors). The final sample included 16 subjects with an average of 762 (*SD* = 99) artifact-free trials.

Experiment 4-2b. One subject was excluded because of excessive artifacts. Data collection was terminated early for two subjects because of excessive artifacts. Finally, one subject was excluded because of poor task performance (prevalence of guesses and swaps ~23%). The final sample included 16 subjects with an average of 775 (*SD* = 92) artifact-free trials.

Experiment 4-2c. One subject was excluded because of excessive artifacts. Data collection was terminated for two subjects because of excessive eye movements. Data collection was terminated for one subject because of an equipment failure. The final sample included 14 subjects with an average of 823 (*SD* = 92) artifact-free trials.

Experiment 4-3. Two subjects were excluded because of excessive artifacts. The final sample included 17 subjects with an average of 835 (*SD* = 103) artifact-free trials.

Modeling Response Error Distributions

In all experiments, response errors were calculated as the angular difference between the reported and presented color or orientation. Response errors could range between -180° and 180°

for the 360°-color space, and between -90° and 90° for the 180°-orientation space. To quantify performance in experiments 4-1 and 4-2c, we fitted a mixture model to the distribution of response errors for each subject using MemToolbox (Suchow et al., 2013). We modeled the distribution of response errors as the mixture of a von Mises distribution centered on the correct value (i.e., a response error of 0°), corresponding to trials in which sample color or orientation was remembered, and a uniform distribution, corresponding to guesses in which the reported color was random with respect to the sample stimulus (Zhang & Luck, 2008). We obtained maximum likelihood estimates for two parameters: (1) the dispersion of the von Mises distribution (*SD*), which reflects response precision; and (2) the height of the uniform distribution (P_g), which reflects the probability of guessing. For experiments 4-2a, 4-2b, and 4-3 (the experiments with one target and one distractor) we fitted a mixture model that also included an additional von Mises component centered on the color/orientation value of the distractor, corresponding to trials in which subjects mistakenly report the value of the distractor stimulus instead of the target stimulus (i.e., 'swaps', Bays, Catalao, & Husain, 2009). We obtained maximum likelihood estimates for the same parameters as in the previous model, with one additional parameter (P_s), which reflects the probability of swaps. The parameter estimates for each experiment are summarized in Table 4-1.

Time-Frequency Analysis

Time-frequency analyses were performed using the Signal Processing toolbox and EEGLAB toolbox (Delorme & Makeig, 2004) for MATLAB (The Mathworks, Natick, MA). To isolate frequency-specific activity, we band-pass-filtered the raw EEG signal using a two-way least-squares finite-impulse-response filter ("eegfilt.m" from EEGLAB Toolbox; Delorme &

Makeig, 2004). This filtering method used a zero-phase forward and reverse operation, which ensured that phase values were not distorted, as can occur with forward-only filtering methods. A Hilbert transform (MATLAB Signal Processing Toolbox) was applied to the band-pass-filtered data, producing the complex analytic signal, $z(t)$, of the filtered EEG, $f(t)$:

$$z(t) = f(t) + i\tilde{f}(t)$$

where $\tilde{f}(t)$ is the Hilbert transform of $f(t)$, and $i = \sqrt{-1}$. The complex analytic signal was extracted for each electrode using the following MATLAB syntax:

hilbert(eegfilt(data,F,f1,f2)')'

In this syntax, *data* is a 2-D matrix of raw EEG (number of trials × number of samples), *F* is the sampling frequency (250 Hz in Experiment 4-1, 500 Hz in all other experiments), *f*1 is the lower bound of the filtered frequency band, and *f*2 is the upper bound of the filtered frequency band. For alpha-band analyses, we used an 8- to 12-Hz band-pass filter; thus, *f*1 and *f*2 were 8 and 12, respectively. For the time-frequency analysis (Figure 4-4), we searched a broad range of frequencies (4–50 Hz, in increments of 1 Hz with a 1-Hz band pass). For these analyses, *f*1 and *f*2 were 4 and 5 to isolate 4- to 5-Hz activity, 5 and 6 to isolate 5- to 6-Hz activity, and so forth. Instantaneous power was computed by squaring the complex magnitude of the complex analytic signal.

Inverted Encoding Model

In keeping with our previous work on spatial working memory and spatial attention (J. J. Foster et al., 2016; J. J. Foster, Sutterer, et al., 2017), we used an IEM to reconstruct spatially selective CTFs from the topographic distribution of oscillatory power across electrodes (Figure 4-2). We assumed that power measured at each electrode reflected the weighted sum of eight

spatial channels (i.e., neuronal populations), each tuned for a different angular position (Figure 4-2b). We modeled the response profile of each spatial channel across angular positions as a half sinusoid raised to the seventh power:

$$R = sin(0.5\theta)^7$$

where θ is the angular position (ranging from 0° to 359°) and R is the response of the spatial channel in arbitrary units. This response profile was shifted circularly for each channel such that the peak response of each spatial channel was centered over one of the eight positions (corresponding to the centers of eight position bins: 0°, 45°, 90°, etc., see Figure 4-2b).

An IEM routine was applied to each time point in the alpha-band analyses and each time-frequency point in the time-frequency analysis. We partitioned our data into independent sets of training data and test data (for details, see the Training and Test Data section). The routine proceeded in two stages (training and test). In the training stage (Figure 4-2c), the training data (B_1) were used to estimate weights that approximated the relative contributions of the eight spatial channels to the observed response (i.e., oscillatory power) measured at each electrode. We define B_1 (m electrodes × n_1 measurements) as a matrix of the power at each electrode for each measurement in the training set, C_1 (k channels × n_1 measurements) as a matrix of the predicted response of each spatial channel (specified by the basis function for that channel; in experiments 4-2a, 4-2b, and 4-3, the predicted channel responses are determined by the target position bin when reconstructing target-related CTFs, and by the distractor position bin when reconstructing distractor-related CTFs) for each measurement, and W (m electrodes × k channels) as a weight matrix that characterizes a linear mapping from *channel space* to *electrode space*.

The relationships among B_1, C_1, and W can be described by a general linear model of the following form:

$$B_1 = WC_1$$

The weight matrix was obtained via least squares estimation as follows:

$$\widehat{W} = B_1 {C_1}^T \left(C_1 {C_1}^T\right)^{-1}$$

In the test stage (Figure 4-2d), we inverted the model to transform the test data, B_2 (m electrodes $\times$ n_2 measurements), into estimated channel responses, $\widehat{C_2}$ (k channels $\times$ n_2 measurements), using the estimated weight matrix, $\widehat{W}$, that we obtained in the training phase:

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

Each estimated channel-response function was circularly shifted to a common center, so that the center channel was the channel tuned for the position of the stimulus of interest (i.e., 0° on the channel offset axes of figures 4-3a and 4-3c). We then averaged these shifted channel-response functions to obtain the CTF averaged across the eight stimulus position bins. The IEM routine was performed separately for each time point.

Finally, because the exact contributions of the spatial channels to electrode responses (i.e., the channel weights, *W*) were expected to vary by subject, we applied the IEM routine to each subject separately. This approach allowed us to disregard differences in how spatially selective activity was mapped to scalp-distributed patterns of power across subjects and instead focus on the profile of activity in the common *stimulus*, or *information*, space (J. J. Foster et al., 2016; Sprague et al., 2015).

Training and Test Data

For the IEM procedure, we partitioned artifact-free trials for each subject into independent sets of training data and test data. Specifically, we divided the trials into three sets. For each of these sets, we averaged power across trials in which the relevant stimulus (see below) appeared in the same position bin. For example, trials in which the sample stimulus was presented at 32° and 60° were averaged because these positions both belonged the position bin centered at 45° (which spanned 22.5°–67.5°). For each set of trials, we obtained an *i* (electrodes) × 8 (position bins) matrix of power values, one for each set (note that the number of electrodes *i* was not the same in all experiments). We used a leave-one-out cross-validation routine such that two of these matrices served as the training data (B_1, *i* electrodes × 16 measurements), and the remaining matrix served as the test data (B_2, *i* electrodes × 8 measurements). Because no trial belonged to more than one of the three sets, the training and test data were always independent. We applied the IEM routine using each of the three matrices as the test data, and the remaining two matrices as the training data. The resulting CTFs were averaged across the three test sets.

When we partitioned the trials into three sets, we constrained the assignment of trials to the sets so that the number of trials was equal for all eight position bins within each set. To ensure that the number of trials per position would be equal within each of the three sets, we calculated the minimum number of trials per subject for a given position bin, *n*, and assigned *n*/3 trials for that position bin to each set. For example, if *n* was 100, we assigned 33 trials for each position bin to each set. Because of this constraint, some excess trials did not belong to any block.

We used an iterative approach to make use of all available trials. For each iteration, we randomly partitioned the trials into three sets (as just described) and performed the IEM routine on the resulting training and test data. We repeated this process of partitioning trials into sets 10 times for the alpha-band analyses, and 5 times for the full time-frequency analysis. For each iteration, the subset of trials that were assigned to blocks was randomly selected. Therefore, the trials that were not included in any block were different for each iteration. We averaged the resulting channel-response profiles across iterations. This iterative approach reduced noise in the resulting CTFs by minimizing the influence of idiosyncrasies that were specific to any given assignment of trials to blocks.

Note that when reconstructing target-related CTFs, we organized trials based on the position bin that the target stimulus appeared in, and when reconstructing distractor-related CTFs, we organized trials based on the position bin that the distractor stimulus appeared in. Thus, we trained and tested on the target's position bin to obtain target-related CTFs, and trained and tested on distractor's position bin to obtain distractor-related CTFs. In a supplemental analysis, we trained on the target position and tested on the distractor position (Figure 4-6). In this analysis, we assigned training blocks as in our standard analysis, and assigned all remaining trials to the test block.

CTF Selectivity

To quantify the spatial selectivity of alpha-band CTFs, we used linear regression to estimate CTF slope. Specifically, we calculated the slope of the channel responses as a function of spatial channels after collapsing across channels that were equidistant from the channel tuned for the position of the stimulus. Higher CTF slope indicates greater spatial selectivity.

Permutation Tests

To determine whether CTF selectivity was reliably above chance, we tested whether CTF slope was greater than zero using a one-sample *t* test. Because mean CTF slope may not be normally distributed under the null hypothesis, we employed a Monte Carlo randomization procedure to empirically approximate the null distribution of the *t* statistic. Specifically, we implemented the IEM as described above but randomized the position labels within each block so that the labels were random with respect to the observed responses in each electrode. This randomization procedure was repeated 1000 times to obtain a null distribution of *t* statistics. To test whether the observed CTF selectivity was reliably above chance, we calculated the probability of obtaining a *t* statistic from the surrogate null distribution greater than or equal to the observed *t* statistic (i.e., the probability of a Type 1 Error). Our permutation test was therefore a one-tailed test. CTF selectivity was deemed reliably above chance if the probability of a Type 1 Error was less than .01.

Bootstrap Resampling Tests

We used a subject-level bootstrap resampling procedure (Efron & Tibshirani, 1993) to test for differences in spatial selectivity (measured as CTF slope) of target- and distractor-related CTFs. We drew 10,000 bootstrap samples, each containing *N*-many subjects sampled with replacement, where *N* is the sample size. For each bootstrap sample, we calculated the mean difference in CTF slope (target – distractor), yielding a distribution of 10,000 mean difference values. We tested whether these difference distributions significantly differenced from zero in either direction, by calculating the proportion of values > or < 0. We doubled the smaller value to obtain a 2-sided *p* value.

Eye Movement Controls

To check that removal of ocular artifacts was effective, we examined baselined HEOG (baseline period: -300 to 0 ms, relative to onset of the sample display). Variation in the grand-averaged HEOG waveforms as a function of stimulus position was < 3 μV in all experiments. Given that eye movements of about 1° of visual angle produce a deflection in the HEOG of ~16 μV (Lins et al., 1993a), the residual variation in the average HEOG corresponds to variations in eye position of < 0.2° of visual angle.

In experiments 4-2a-c and 4-3, we also inspected gaze position (averaged from stimulus onset to the end of the delay period) as a function of the target's position bin. We drift-corrected gaze position data by subtracting the mean gaze position measured during a pre-stimulus window (-300 to -100 ms, relative to onset of the sample display) to achieve optimal sensitivity to changes in eye position relative to the pre-stimulus period (Cornelissen, Peters, & Palmer, 2002). This analysis revealed remarkable little variation in gaze position (< 0.05° of visual angle) as a function of target position in all four experiments with eye tracking data (Figure 4-9a-d), showing we achieved an extremely high standard of fixation compliance once trials with artifacts were discarded.

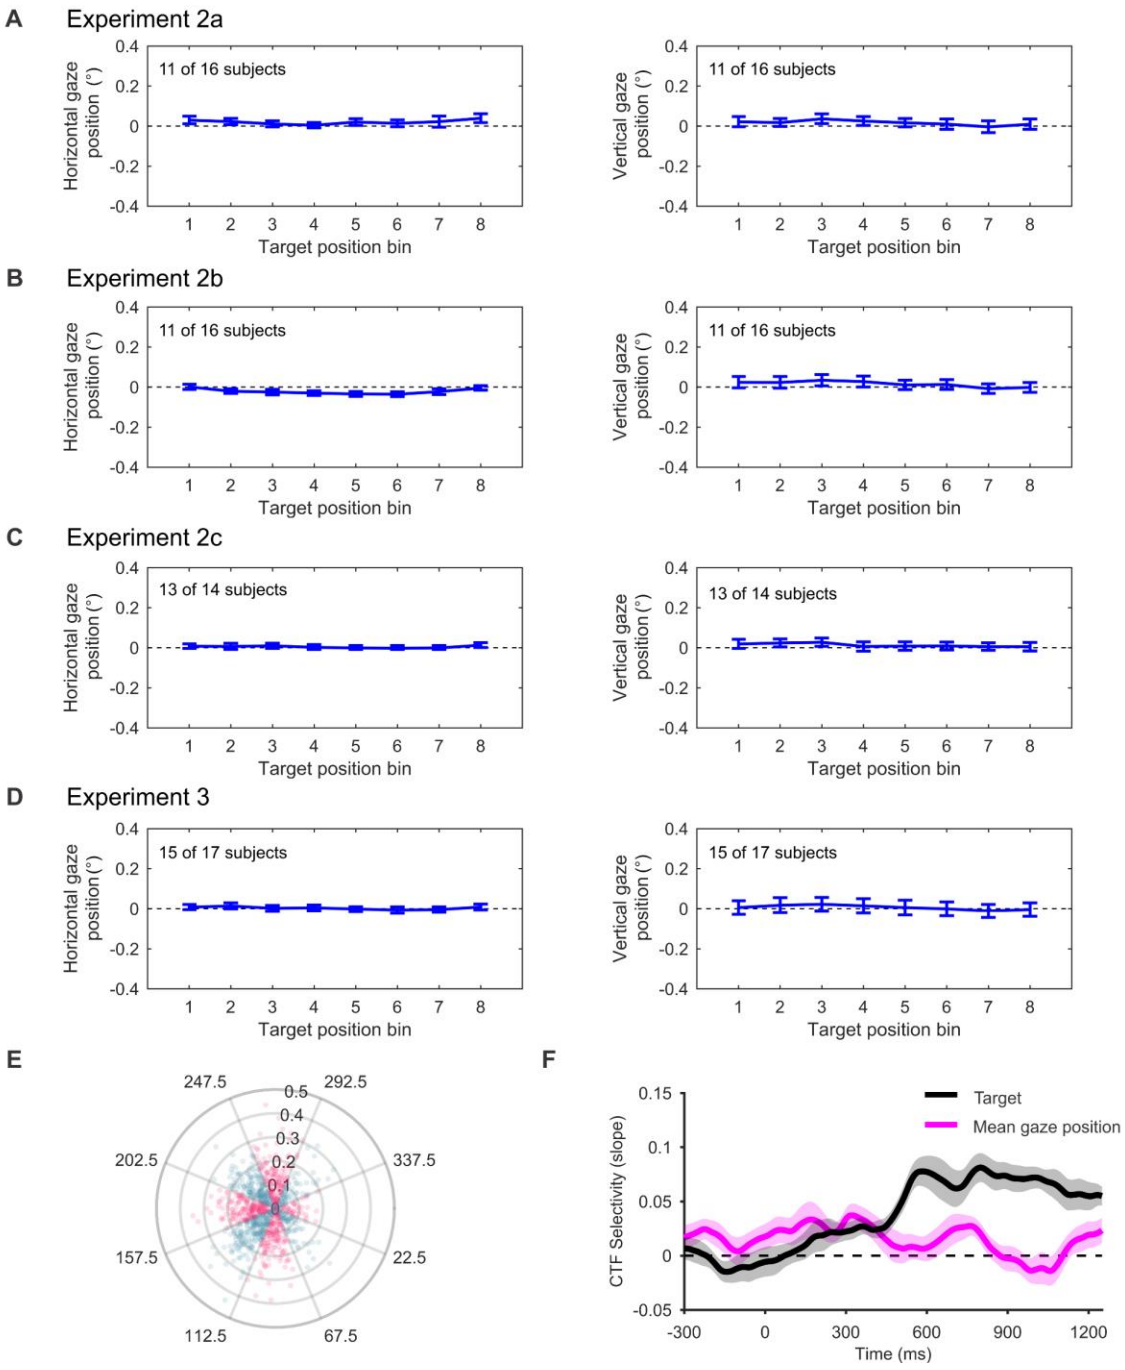

Figure 4-9. Eye movement controls.
(a-d) Mean horizontal (left) and vertical (right) gaze position (following the onset of the sample array) as a function of the position bin that the target appeared in. Gaze position was drift corrected to maximize sensitivity to any small changes in gaze position from the pre-stimulus period. We observed remarkable little variation in mean gaze position as a function of the position of the target position (less than 0.05° of visual angle), showing that we achieved an extremely high standard of fixation compliance. Each panel indicates the number of subjects that

Figure 4-9, continued..
had usable eye tracking data out of the total sample in each experiment. Error bars represent ±1 SEM across subjects. (e) Mean gaze position for each trial sorted in eight gaze position bins for a sample subject. The radial axis is in degrees of visual angle. (f) Alpha-band CTFs reconstructed on the basis of the target position bin (black line) and gaze position bin (pink line) in Experiment 4-2c. We observed a clear target-related CTF but did not observe clear evidence for a gaze-position-related CTF. This analysis shows that alpha-band activity tracks the location of the target position rather than small variations in gaze position. Shaded error bars represent ±1 bootstrapped SEM across subjects.

Although residual bias in gaze position is very small, it is possible that alpha-band activity might track small biases in gaze position rather than spatial representation in working memory. We ran a control analysis to test this possibility. We reasoned that if eye movements drive spatially specific alpha-band activity, then we should observe more robust alpha-band CTFs when we reconstruct CTFs on the basis of gaze position than on the basis of the stimulus position. We sorted trials for each subject in Experiment 4-2c into eight "gaze position bins" on the basis of the mean gaze position (drift corrected) during the post-stimulus period (0-1250 ms; see Figure 4-9e). We then used these gaze position bins to perform the CTF analysis. The number of trials available in each position bin limits the number of trials that can be assigned to the training/test sets. Two subjects were excluded from this analysis because they had fewer than 30 trials in a given gaze position bin. We also re-ran the target-related CTF analysis, this time equating the number of trials that were included in the training/test sets with the gaze position analysis. We did not see clear evidence for a gaze-position-related CTF (Figure 4-9f). However, we did observed a clear target-related CTF, confirming that we had enough trials to detect spatially selective alpha-band activity. This analysis shows that alpha-band activity tracks the location of the target position rather than small variations in gaze position.

CHAPTER 5. ATTENTION ENHANCES THE SPATIAL TUNING OF STIMULUS-EVOKED POPULATION RESPONSES

Introduction

Covert spatial attention improves visual processing at attended locations via enhanced neural representations of attended stimuli (Maunsell, 2015; Sprague et al., 2015). One important line of work has measured how attention impacts the responses of single cells, showing that the amplitude of neural responses to attended stimuli is increased (Luck, Chelazzi, Hillyard, & Desimone, 1997; McAdams & Maunsell, 1999; Treue & Martinez-Trujillo, 1999), and that spatial receptive fields are shifted towards attended stimuli (Anton-Erxleben & Carrasco, 2013; Anton-Erxleben, Stephan, & Treue, 2009; Connor, Preddie, Gallant, & Van Essen, 1997; Womelsdorf, Anton-Erxleben, Pieper, & Treue, 2006; Womelsdorf, Anton-Erxleben, & Treue, 2008). Ultimately, however, perception relies on information coded across large populations of cells (Pouget et al., 2000; Saxena & Cunningham, 2019), and the perceptual benefits of attention are thought to be mediated by changes in population codes (Cohen & Maunsell, 2009; Sprague et al., 2015). Thus, there is strong motivation to understand how attention modulates population-level neural responses (Sprague et al., 2015; Sprague & Serences, 2013).

In recent years, researchers have used functional MRI (fMRI) to examine how attention influences population-level representations of spatial position (Fischer & Whitney, 2009; Itthipuripat, Sprague, & Serences, 2018; Kay et al., 2015; Sprague, Itthipuripat, Vo, & Serences, 2018; Sprague & Serences, 2013; Vo, Sprague, & Serences, 2017). Sprague and Serences (2013), for example, used an inverted encoding model (IEM; Brouwer & Heeger, 2011, 2009) to reconstruct a population-level, spatial representation of a visual stimulus. This approach assumed

that the response of each voxel (in a visually responsive region) reflects the summed activity of a number of spatially tuned channels. In a training phase, the authors estimated the relative contributions of each of these channels (called "weights") to each voxel. Then, once the weights were estimated, they *inverted* the model to estimate the response of each of the spatially tuned channels given the pattern of activity across all voxels in each condition. Using this approach, Sprague and Serences (2013) found that attention to the stimulus increased the amplitude of spatial representations in visual areas without reliably changing their size, a finding that has been replicated by subsequent work (Itthipuripat et al., 2018; Vo et al., 2017; but see Fischer & Whitney, 2009).

At first glance, this finding seems to provide clear evidence that attention increases the gain of population-level representations. However, emerging evidence suggests that the effect of attention on the blood-oxygen-level-dependent (BOLD) signal that is measured with fMRI does not reflect an amplification of the stimulus-driven response, but instead reflects a shift in baseline activity in cortex tuned for the attended location that is *additive* with the stimulus-driven response (Buracas & Boynton, 2007; Murray, 2008). If the effect of attention on the BOLD response reflects a straightforward amplification of the stimulus-driven response, then the effect of attention should be greater for higher-contrast stimuli that drive a stronger neural response (Hillyard et al., 1998). Instead, attending a stimulus increases the BOLD response in visual cortex by the same amount, regardless of stimulus contrast (Buracas & Boynton, 2007; Itthipuripat et al., 2018; Murray, 2008; Pestilli, Carrasco, Heeger, & Gardner, 2011; Sprague, Itthipuripat, et al., 2018; but see Li, Lu, Tjan, Dosher, & Chu, 2008). This finding is consistent with earlier work that found that an increased BOLD response in regions of visual cortex tuned

for an attended location in balanced visual displays (Brefczynski & DeYoe, 1999; Kastner, Pinsk, De Weerd, Desimone, & Ungerleider, 1999; Tootell et al., 1998). Taken together, these studies suggest that the effect of attention on the BOLD signal reflects a stimulus-independent shift in baseline activity that does not reveal the impact of attention on the stimulus-evoked response.

Although fMRI may not be well suited to determine how attention modulates the stimulus-driven response, visually evoked potentials measured with EEG show more promise. The effect of attention on visually evoked potentials scales with stimulus contrast (Itthipuripat, Ester, Deering, & Serences, 2014; Itthipuripat et al., 2018), suggesting that they provide an opportunity to examine how attention influences stimulus-driven responses. Thus, the goal of the present work was to combine the EEG method with an IEM analytic approach that enabled us to characterize how attention modulates the spatial tuning of stimulus-evoked population activity. Furthermore, the temporally resolved nature of the EEG signal enabled a clear examination of how attention shapes the earliest stages of visual processing, occurring within the first few hundred milliseconds following stimulus onset.

We focused our analysis on *stimulus-evoked* EEG activity (i.e., activity phase-locked to stimulus onset), which provided a means of separating stimulus-driven activity from ongoing neural activity that is independent of the stimulus. In addition, we used the IEM approach to reconstruct spatially selective channel-tuning functions (CTFs) from stimulus-evoked EEG activity. In two experiments, we show that attention increases the amplitude of spatially tuned, stimulus-evoked responses within ~100 ms after stimulus onset, providing clear evidence that

attention enhances the selectivity of spatial population codes during the first wave of visual sensory activity.

Materials and Methods

Participants

Volunteers recruited from the University of Chicago and surrounding community participated for monetary compensation ($15/hr). Participants were between the ages of 18 and 35, reported normal or corrected-to-normal visual acuity, and provided informed consent according to procedures approved by the University of Chicago Institutional Review Board.

Experiment 5-1. Twenty-one volunteers participated in Experiment 5-1 (8 male, 13 female; mean age = 22.7 years, *SD* = 3.2). Four participants were excluded from the final sample for the following reasons: we were unable to prep the participant for EEG (n = 1); we were unable to obtain eye tracking data (n = 1); the participant did not complete enough blocks of the task (n = 1); and residual bias in eye position (see Artifact Rejection) was too large (n = 1). The final sample size was 17 (6 male, 11 female; mean age = 22.7 years, *SD* = 3.4).

Experiment 5-2. Twenty-four volunteers participated in Experiment 5-2 (6 male, 16 female; mean age = 24.0 years, *SD* = 3.0). For four participants, we terminated data collection and excluded the participant from the final sample for the following reasons: we were unable to obtain eye tracking data (n = 1); the participant had difficulty performing the task (n = 1); the participant made too many eye movements (n = 2). The final sample size was 20 (5 male, 15 female; mean age = 24.0 years, *SD* = 2.8).

Task and Stimuli

Participants were tested in a dimly lit, electrically shielded chamber. Stimuli were generated using Matlab (MathWorks, Natick, MA) and the Psychophysics toolbox (Brainard, 1997; Pelli, 1997). Participants viewed the stimuli on a contrast-linearized 24" LCD monitor (refresh rate: 120 Hz, resolution 1080 x 1920 pixels) with their chin rested on a padded chin rest (viewing distance: 76 cm in Experiment 5-1, 75 cm in Experiment 5-2). Stimuli were presented against a mid-gray background (~61 cd/m^2).

On each trial, observers viewed a sequence of four bullseye stimuli at a fixed location (Figure 5-1). Across blocks, we manipulated whether observers attended the bullseye stimuli (*attend-stimulus* condition) or attended the central fixation dot (*attend-fixation* condition). In the attend-stimulus condition, observers monitored the sequence of bullseyes for a bullseye that was lower contrast than the rest (while maintained fixation). In the attend-fixation condition, observers monitored the fixation dot for a brief reduction in contrast. In both conditions, changes in contrast occurred in both the attended and unattended stimulus, and we instructed participants to disregard changes in the unattended stimulus. Contrast changes occurred in each stimulus (bullseye or fixation) on 50% of trials. The trials that contained changes in the bullseye contrast and changes in the fixation contrast were determined independently.

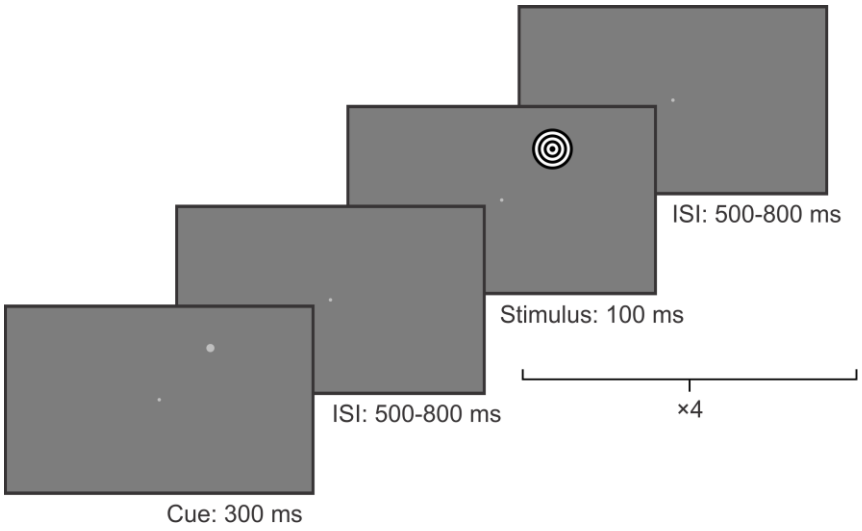

Figure 5-1. Illustration of task.
(a) On each trial, a series of bullseye stimuli appeared at a fixed location. The trial began with a cue that indicated where the bullseyes would appear. In the attend-stimulus condition, observers covertly monitored the bullseyes for one bullseye that was lower contrast than the rest. In the attend-fixation condition, observers monitored the fixation dot for a brief dimming.

Observers fixated the fixation dot (0.1° in diameter, 56.3% contrast) before pressing spacebar to initiate each trial. Each trial began with a 400 ms fixation display. A cue (0.25° in diameter, 32.8% contrast) presented for 300 ms, indicated where the bullseyes would appear. On each trial, the bullseye stimuli appeared at one of eight positions 4° from fixation (Figure 5-2a). Each bullseye (1.6° in diameter, 0.12 cycles/°) was presented for 100 ms, and was preceded by a jittered inter-stimulus interval (ISI) between 500 and 800 ms. The final bullseye was followed by a 500 ms ISI. The lower-contrast bullseye was never the first stimulus in the sequence, so that the first stimulus established the pedestal contrast for the trial. Decrements in the contrast of the fixation dot lasted 100 ms, and co-occurred with one of the bullseye stimuli (but never the first in the sequence). Each trial ended with a response screen that prompted observers to report whether the attended stimulus changed in contrast. Observers responded using the numberpad ("1" =

change, “2” = no change). The observer’s response appeared above the fixation dot, and the observer could correct their response if they pressed the wrong key. Finally, observers confirmed their response by pressing the spacebar. In Experiment 5-2, the experimenter manually provided feedback to the observer to indicate whether they had noticed blinks or eye movements during the trial by pressing a key outside the recording chamber. The feedback, “blink” or “eye movement” was presented in red for 500 ms after the observer had made their response.

In Experiment 5-1 the pedestal contrast of the bullseye was always 90%. Participants completed a 3.5-hour EEG session. The session began with a staircase procedure to adjust task difficulty (see Staircase Procedures). Participants then completed 12-20 blocks (40 trials each) during which we recorded EEG. Thus, participants completed between 480 and 800 trials (1920-3200 stimulus presentations). The blocks alternated between the attend-stimulus and attend-fixation conditions, and we counterbalanced task order across participants.

In Experiment 5-2, we parametrically manipulated pedestal contrast. We included 5 pedestal contrasts (6.25, 12.5, 25.0, 50.0, and 90.6%). Thus, there were 10 conditions in total (2 attention conditions × 5 pedestal contrasts). Participants completed three sessions: a 2.5-hour behavior session to adjust task difficulty in each condition (see Staircase Procedures), followed by two 3.5-hour EEG sessions. All sessions were completed within 10 days (M = 6 days, SD = 2). Across the EEG session, participants completed 20 blocks. Each block consisted of 104 trials: eight trials for each of the 10 conditions, and an additional 12 trials in each condition at the highest pedestal contrast (90.6% contrast) for the purpose of training the encoding model (see Training and Test Data). Each block included a break at the halfway point. As in Experiment 5-1, the blocks alternated between the attend-stimulus and attend-fixation conditions, and we

counterbalanced task order across participants. We aimed to run each participant through 20 blocks, to obtain 160 testing trials for each condition (640 stimulus presentations), and 480 training trials (1920 stimulus presentations). All participants completed 20 blocks with the following exceptions: three participants completed 18 blocks, and one participant completed 24 blocks.

Staircase Procedures

Experiment 5-1. We used a staircase procedure to match difficulty across conditions. We adjusted difficulty by adjusting the size of the contrast decrement for each condition (attend-stimulus and attend-fixation) independently. In Experiment 5-1, participants completed six staircase blocks of 40 trials (three blocks for each condition) before we started the EEG blocks of the task. Thus, participants completed 120 staircase trials for each condition. We used a 3-down-1-up procedure to adjust task difficulty: after three correct responses in a row, we reduced the size of the contrast decrement by 2%; after an incorrect response, we increased the size of the contrast decrement by 2%. This procedure was designed to hold accuracy at ~80% correct (García-Pérez, 1998). The final size of the contrast decrements in the staircase blocks were used for the EEG blocks of the task. Every four blocks (two blocks of each condition), we examined accuracy in each condition, and manually adjusted the size of the contrast decrements to hold accuracy as close to 80% as possible.

Experiment 5-2. In Experiment 5-2, participants completed a 2.5-hour staircase session prior to the EEG sessions. We adjusted difficulty for each of the 10 conditions independently (2 attention conditions × 5 pedestal contrast). Participants completed 16 blocks of 40 trials, alternating between the attend-fixation and attend-stimulus conditions. The five contrast levels

were randomized within each block. Thus, observers completed 64 staircase trials for each of the 10 conditions. In Experiment 5-2, we used a weighted up/down procedure to adjust task difficulty: after a correct response, we reduced the size of the contrast decrement by 5%; after an incorrect response, we increased the size of the contrast decrement by 17.6%. This procedure held accuracy fixed at ~76%. The staircase procedure continued to operate throughout the EEG sessions.

Eye Tracking

We monitored gaze position using a desk-mounted EyeLink 1000 Plus infrared eye-tracking camera (SR Research, Ontario, Canada). Gaze position was sampled at 1000 Hz. Head position was stabilized with a chin rest. According to the manufacturer, this system provides spatial resolution of 0.01° of visual angle, and average accuracy of 0.25-0.50° of visual angle. We calibrated the eye tracker every 1-2 blocks of the task. We re-calibrated they eye tracker during the blocks of the task (between trials) if necessary.

EEG Recording

We recorded EEG activity from 30 active Ag/AgCl electrodes mounted in an elastic cap (Brain Products actiCHamp, Munich, Germany). We recorded from International 10-20 sites: FP1, FP2, F7, F3, Fz, F4, F8, FT9, FC5, FC1, FC2, FC6, FT10, T7, C3, Cz, C4, T8, CP5, CP1, CP2, CP6, P7, P3, Pz, P4, P8, O1, Oz, O2. Two additional electrodes were affixed with stickers to the left and right mastoids, and a ground electrode was placed in the elastic cap at position FPz. Eye movements and blinks were also monitored using electrooculogram (EOG), recorded with passive Ag/AgCl electrodes. Horizontal EOG was recorded from a bipolar pair of electrodes placed ~1 cm from the external canthus of each eye. Vertical EOG was recorded from a bipolar

pair of electrodes placed above and below the right eye. Data were filtered online (low cut-off = .01 Hz, high cut-off = 80 Hz, slope from low- to high-cutoff = 12 dB/octave), and were digitized at 500 Hz using BrainVision Recorder (Brain Products, Munich, German) running on a PC. Impedance values were kept below 10 kΩ.

Preprocessing

EEG and eye tracking data were preprocessed using custom Matlab scripts. We drift-corrected the eye tracking data for each trial by subtracting the mean gaze position measured during a 200 ms window immediately before the onset of the cue. We re-referenced EEG data to the algebraic mean of the two mastoids. We excluded data from some electrodes for some participants because of low quality data (excessive high-frequency noise or sudden steps in voltage). We excluded no more than three electrodes for any given participant. In Experiment 5-2, we excluded electrodes FP1 and FP2 for all participants because we obtained poor-quality data (high-frequency noise and slow drifts) at these sites for most participants.

The data were segmented into epochs time-locked to the onset of each bullseye stimulus (starting 300-ms before stimulus onset, ending 500 ms after stimulus onset). We visually inspected the segmented EEG data for artifacts (amplifier saturation, excessive muscle noise, and skin potentials), and the eye tracking data for ocular artifacts (blinks, eye movements, and deviations in eye position from fixation), and discarded any epochs contaminated by artifacts. In Experiment 5-1, all subjects included in the final sample had at least 800 artifact-free epochs for each condition. In Experiment 5-2, all subjects included in the final sample had at least 450 artifact-epochs for testing the IEM in each condition, and at least 1500 artifact-free epochs for training the IEM (see Training and Test Data).

After artifact rejection, for each participant we inspected mean gaze position as a function of stimulus position for the attend-stimulus and attend-fixation conditions separately. For all participants in the final samples, mean gaze position varied by less than 0.2° of visual angle across stimulus positions. One subject in Experiment 5-1 was excluded from the final subject because they did not meet this criterion. To verify that removal of ocular artifacts was effective, we inspected mean gaze position (during the 100ms presentation of each stimulus) as a function of stimulus position for the attend-stimulus and attend-fixation conditions separately. In both experiments, we observed very little variation in mean gaze position (across subjects) as a function of stimulus position (< 0.05° of visual angle) in both the attend-stimulus and attend-fixation conditions. Thus, we achieved an extremely high standard of fixation compliance after epochs with artifacts were removed.

Evoked Power

A Hilbert Transform (Matlab Signal Processing Toolbox) was applied to the segmented EEG data to obtain the complex analytic signal, $z(t)$, of the EEG, $f(t)$:

$$z(t) = f(t) + i\tilde{f}(t)$$

where $\tilde{f}(t)$ is the Hilbert Transform of $f(t)$, and $i = \sqrt{-1}$. The complex analytic signal was extracted for each electrode using the following Matlab syntax:

hilbert(data')'

where data is a 2D matrix of segmented EEG (number of trials × number of samples). We calculated *evoked* power by first averaging the complex analytic signals across trials, and then squaring the complex magnitude of the averaged analytic signal. Evoked power isolates activity phase-locked to stimulus onset because only activity with consistent phase across trials remains

after averaging the complex analytic signal across trials. Trial averaging was performed for each stimulus position separately within each block of training or test data (see Training and Test Data).

For some analyses, we high-pass filtered the data with a low-cutoff of 4-Hz before calculating power to remove low-frequency activity. We used EEGLAB's eegfilt() function (Delorme & Makieg, 2004), which implements a two-way least-squares finite impulse response filter. This filtering method uses a zero-phase forward and reverse operation, which ensures that phase values are not distorted, as can occur with forward-only filtering methods.

Inverted Encoding Model

We used an inverted encoding model (Brouwer & Heeger, 2009, 2011) to reconstruct spatially selective channel-tuning functions (CTFs) from the distribution of power across electrodes (also see J. J. Foster, Bsales, Jaffe, & Awh, 2017; J. J. Foster et al., 2016). We assumed that the power at each electrode reflects the weighted sum of eight spatially selective channels (i.e., neuronal populations), each tuned for a different angular position (figures 5-2A and 5-2B). We modeled the response profile of each spatial channel across angular locations as a half sinusoid raised to the twenty-fifth power:

$$R = sin(0.5\theta)^{25}$$

where θ is angular location (0–359°), and R is the response of the spatial channel in arbitrary units. This response profile was circularly shifted for each channel such that the peak response of each spatial channel was centered over one of the eight positions corresponding to the eight positions bins (0°, 45°, 90°, etc.).

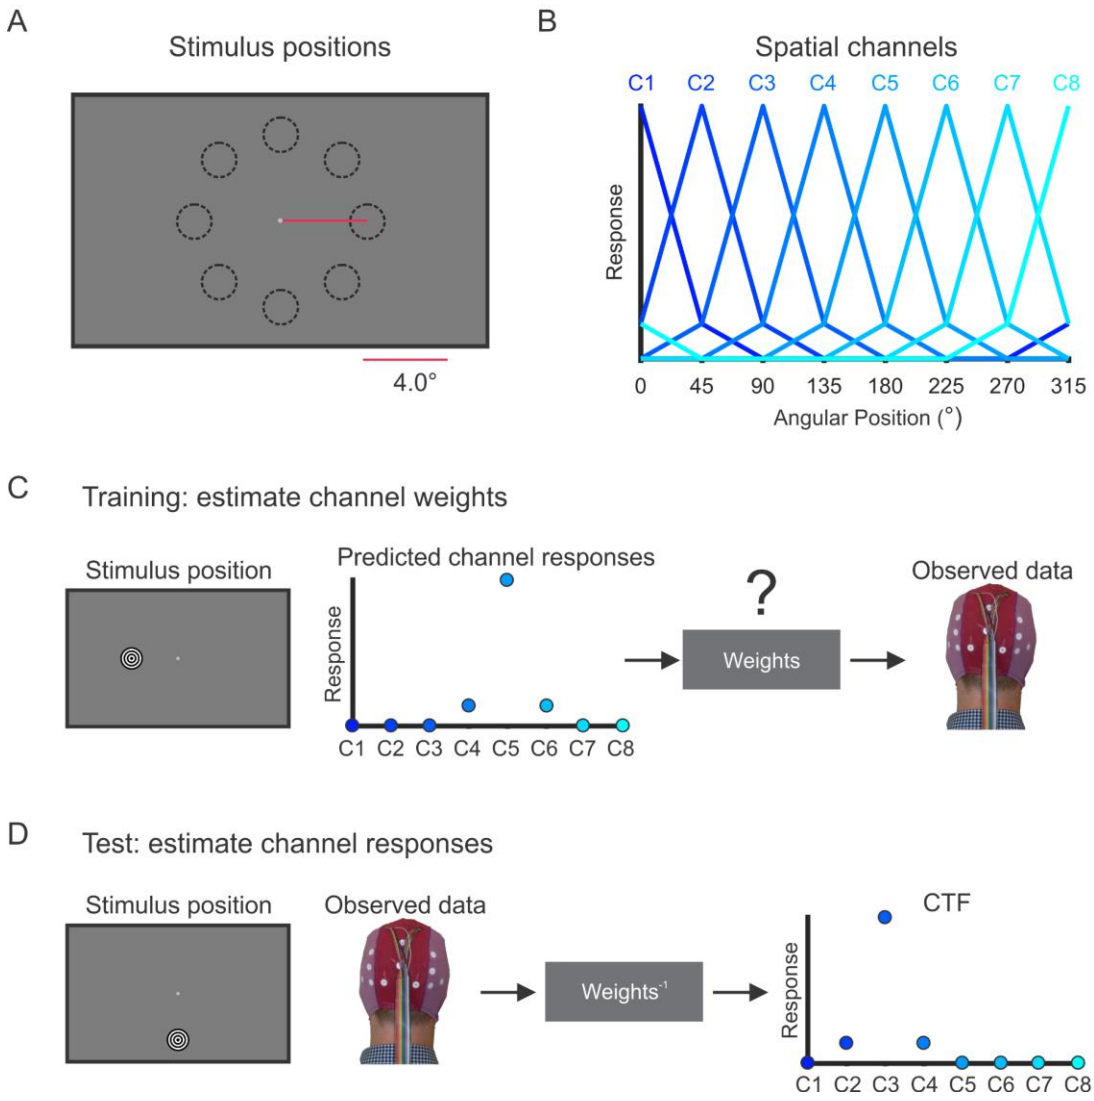

Figure 5-2. Inverted encoding model procedure.
(a) The eight spatial positions around the fixation dot where the bullseye stimuli could appear. (b) We modeled power at each electrode as the weighted sum of eight spatially selective channels (C1-C8), each tuned for the center of one of the eight stimulus positions. Each curve shows the predicted response of one of the channels across the eight stimulus positions (i.e., the "basis function"). (c) In the training phase, we used the predicted channel responses, determined by the basis functions shown in (b), to estimate a set of channel weights that specified the contribution of each spatial channel to the response measured at each electrode.(d) In the test phase, using an independent set of data, we used the channel weights obtained in the training phase to estimate the profile of channel responses given the pattern of activity across the scalp.

An IEM routine was applied to each time point. We partitioned our data into independent sets of training data and test data (see Training and Test Data). This routine proceeded in two stages (training and test). In the training stage (Figure 5-2C), training data (B_1) were used to estimate weights that approximate the relative contribution of the eight spatial channels to the observed response measured at each electrode. Let B_1 (m electrodes × n_1 measurements) be the power at each electrode for each measurement in the training set, C_1 (k channels × n_1 measurements) be the predicted response of each spatial channel (determined by the basis functions, see Figure 5-2B) for each measurement, and W (m electrodes × k channels) be a weight matrix that characterizes a linear mapping from "channel space" to "electrode space". The relationship between B_1, C_1, and W can be described by a general linear model of the form:

$$B_1 = WC_1$$

The weight matrix was obtained via least-squares estimation as follows:

$$\widehat{W} = B_1{C_1}^T\left(C_1{C_1}^T\right)^{-1}$$

In the test stage (Figure 5-2D), we inverted the model to transform the observed test data B_2 (m electrodes × n_2 measurements) into estimated channel responses, C_2 (k channels × n_2 measurements), using the estimated weight matrix, $\widehat{W}$, that we obtained in the training phase:

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

Each estimated channel response function was then circularly shifted to a common center, so the center channel was the channel tuned for the position of the probed stimulus (i.e., 0° on the "Channel Offset" axes), then averaged these shifted channel-response functions across the eight position bins to obtain a CTF. Finally, because the exact contributions of each spatial channel to each electrode (i.e., the channel weights, W) likely vary across participants, we applied the IEM

routine separately for each participant, and statistical analyses were performed on the reconstructed CTFs. This approach allowed us to disregard differences in the how location-selective activity is mapped to scalp-distributed patterns of power across participants, and instead focus on the profile of activity in the common stimulus or "information" space (Sprague et al., 2015).

Training and Test Data

For the IEM analysis, we partitioned artifact-free epochs into three independent sets: two training sets and one test set. Within each set, we calculated evoked power (see Evoked Power) across the epochs for each stimulus position to obtain a matrix of power values across all electrodes for each location bin (electrodes × stimulus positions, for each time point). We equated the number of epochs for each stimulus position in each set. Because of this constraint, some excess epochs were not assigned to any set. Thus, we used an interactive approach to make use of all available trials. For each of 500 iterations, we randomly partitioned the data into training and test data (see below for details of how data partitioned into training and test sets in each experiment). We averaged the resulting CTFs across iterations.

Experiment 5-1. When comparing CTF parameters across conditions, it is important to estimate a single encoding model that is then used to reconstruct CRFs for each condition separately (Sprague, Adam, et al., 2018). Thus, for Experiment 5-1, we estimated the encoding model using a training set consistent of equal numbers of trials from each condition. Specifically, we partitioned data for each condition (attend-stimulus and attend-fixation) into three sets (with the constraint that the number of trials per location in each set was also equated across conditions). We obtained condition-neutral training data by combining data across the two

conditions before calculating evoked power, resulting in two training sets that included equal numbers of trials from each condition. We then tested the model using the remaining set of data for each condition separately. Thus, we used the same training data to estimate a single encoding model, and varied only the test data that was used to reconstruct CTFs for each condition.

Experiment 5-2. In Experiment 5-2, we included additional trials in the highest contrast conditions (90.6% contrast) to train the encoding model (see Task and Stimuli). For each iteration of the analysis, we partitioned this data into two training sets, and generated a single test set for each of the 10 conditions separately. We equated the number of trials in included in each of the test sets. As described above, we averaged the resulting CTFs across 500 iterations of the analysis.

Statistical Analysis

Model fitting. To quantify how attention influenced stimulus-evoked CTFs, we fitted CTFS with an exponentiated cosine function (Figure 5-3) of the form:

$$f(x) = \beta + a\left(e^{k(\cos(\mu - x) - 1}\right)$$

where *x* is a vector of the angular offsets of each channel (-180°, -135°, -90° …, 135°). We fixed the μ parameter, which determines the center of the tuning function, at a channel offset of 0° (i.e., the peak of the function was fixed at the channel tuned for the stimulus position). The function had three free parameters: *baseline* (β) determines the vertical offset of the function from zero; *amplitude* (α) determines the height of the function above baseline function; *concentration* (*k*) determines the width of the function, with higher values corresponding to narrower tuning. We fitted the function with a general linear model combined with a grid search procedure (Ester et al., 2015). Rather than reporting *k* values (which are inversely related to CTF

width), we report the full-width half-maximum (fwhm) of the fitted function, measured as the width of the function in angular degrees halfway between baseline and the peak of the function.

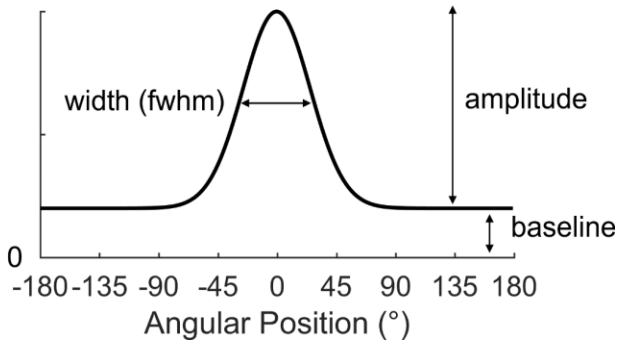

Figure 5-3. Parameters of the exponentiated cosine function.
The baseline parameter measured the overall offset of the function from zero The amplitude parameter measured the peak response of the CTF above baseline. The width parameter measured the width of the function halfway between baseline and the peak of the function (i.e. full-width half-maximum, or fwhm).

Resampling tests. We used a subject-level resampling procedure to test for differences in the parameters of the fitted function across conditions. We drew 10,000 bootstrap samples, each containing *N*-many subjects sampled with replacement, where *N* is the sample size. For each bootstrap sample, we fitted the exponentiated cosine function as described above to the mean CTF across subjects in the bootstrap sample. In Experiment 5-1, to test for differences between conditions for each parameter, we calculated the difference for the parameter between the attend-stimulus and attend-fixation conditions for each bootstrap sample, which yielded a distribution of 10,000 values. We tested whether these difference distributions significantly differed from zero in either direction, by calculating the proportion of values > or < 0. We doubled the smaller value to obtain a 2-sided *p* value.

In Experiment 5-2, we tested for main effects of attention and contrast, and for an attention × contrast interaction, for each parameter. To test for a main effect of attention, we

averaged parameter estimates across contrast levels for each bootstrap sample, and calculated the difference in each parameter estimate across attention conditions for each bootstrap sample. We tested whether these difference distributions significantly differed from zero in either direction, by calculating the proportion of values > or < 0. To test for a main effect of contrast, we averaged the parameter estimates across the attention conditions, and fitted a linear function to the parameter estimates as a function of contrast. For each bootstrap sample, we calculated the slope of the best-fit linear function. We tested whether the resulting distribution of slope values significantly differed from zero in either direction by calculating the proportion of values > or < 0. Finally, to test for an attention × contrast interaction, we fitted a linear function to the parameter estimates as a function of contrast for the attend-stimulus and attend-fixation conditions separately For each bootstrap sample, we calculated the difference in the slope of these functions between the attend-stimulus and attend-fixation conditions. We tested whether the resulting distribution of differences-in-slope values significantly differenced from zero differed from zero in either direction by calculating the proportion of values > or < 0. For both main effects and the interaction, we doubled the smaller *p* value to obtained a 2-sided *p* value.

Results

Experiment 5-1

In Experiment 5-1, we developed a paradigm to examine how spatial attention modulates the spatial tuning of the stimulus-evoked response. On each trial, observers saw a sequence of four bullseye stimuli were presented at one of eight positions around the fixation dot (Figure 5-1). Each trial began with a peripheral cue, which indicated where the bullseye stimuli would appear. The stimuli in the sequence were separated by a jittered (500-800-ms) inter-stimulus

interval so that we could isolate the visually evoked response to each stimulus. To examine how spatial attention modulates the visually evoked response, we manipulated whether observers attended the bullseye stimuli or the central fixation dot (c.f. Sprague & Serences, 2013). In the *attend-stimulus* condition, observers covertly monitored the sequence of bullseyes for one bullseye that was lower contrast than the rest. In the *attend-fixation* condition, observers covertly monitored the fixation dot for a brief decrement in contrast that co-occurred with one of the bullseye presentations. At the end of the trial, observers reported whether or not a contrast decrement occurred in the relevant stimulus (contrast changed occurred on 50% of trials). We closely matched difficulty in the attend-stimulus and attend-fixation conditions by adjusting the size of the contrast decrement in each task independently (see Materials and Methods). Accuracy was comparable in the attend-stimulus (M = 81.0%, SD = 3.7) and the attend-fixation (M = 80.0%, SD = 2.2) conditions.

We sought to characterize how attention modulates the stimulus-driven representations of the stimulus. To isolate stimulus-driven activity from activity independent of the stimulus, we measured the power of stimulus-evoked activity (i.e., activity phase-locked to stimulus onset; see Materials and Methods). We used an IEM (Brouwer & Heeger, 2009, 2011) to reconstruct spatially selective channel-tuning functions (CTFs) from the scalp distribution of stimulus-evoked power. This approach assumes that power measured at each electrode on the scalp reflects the underlying activity of a number of spatially tuned channels (or neuronal populations), each tuned for a different spatial position (Figure 5-2b). In a training phase, we estimated the relative contribution of each of these channels (called "channel weights") to the response measured at each electrode on the scalp (Figure 5-2c). Then, in a test phase, we used independent

set of data to estimate the response of the spatial channels given the pattern of activity across electrodes (Figure 5-2d). The resulting response channel-response profiles, called channel-tuning functions (CTFs) allowed us to measure the spatial selectivity of the population-level neural activity captured by EEG.

Human ERP studies have established that visually evoked responses are modulated by attention as early as 100 ms after stimulus onset (Hillyard & Anllo-Vento, 1998). Thus, we focused our analysis in an early window (80-130 ms after stimulus onset) to capture the early visually evoked response. We reconstructed spatial CTFs from stimulus-evoked activity for both the attend-stimulus and attend-fixation conditions, having estimated channel weights using a training set that included data from both the attend-stimulus and attend-fixation conditions. To quantify how attention modulates the CTFs, we fit CTFs with an exponentiated cosine function (see Materials and Methods), which enabled us to measure three parameters of the CTF: baseline, amplitude, and width (Figure 5-3). The *baseline* parameter measured the overall offset of the function from zero. An increase in baseline reflects an upward shift of entire CTF. Thus, changes in baseline do not reflect a change in the spatial selectivity of the CTF. The *amplitude* parameter measured the peak response of the CTF above baseline. An increase in amplitude reflects an increase in the spatial selectivity of the CTF. Finally, the *width* of the CTF (measured as full-width half-maximum, or fwhm) measured the width of the function halfway between baseline and the peak of the function, such that larger width values reflect broader spatial tuning. Figure 5-4a shows the reconstructed CTFs with the best fitting functions. We found that stimulus-evoked CTFs were higher amplitude ($p < .0001$) and more broadly tuned ($p < .0001$) in the attend-stimulus condition than in the attend fixation condition (Figure 5-4b). We observed no

difference in baseline between the conditions (p = .998). However, as we will see in the next section, the finding that CTFs were broader in the attend-stimulus condition than in the attend-stimulus appears to be an artifact of lingering activity from the preceding stimulus event. Furthermore, this effect did not replicate in Experiment 5-2. Thus, the primary effect of attention is to increase the spatial selectivity of population-level neural activity via an increase in the amplitude of spatially graded CTFs.

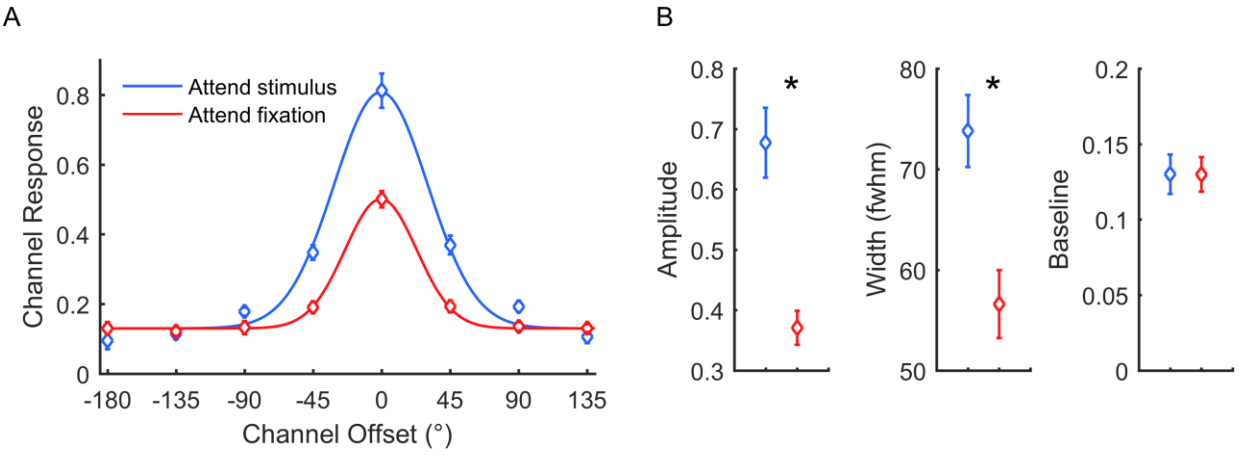

Figure 5-4. Stimulus-evoked CTFs and parameter estimates as a function of attention condition.
(a) Stimulus-evoked CTFs (measured 80-130 ms after stimulus onset) with best fitting functions. (b) Parameter estimates of the best fitting functions. Asterisks mark differences between the conditions at the .05 level. Error bars show ± 1 bootstrapped SEM.

We designed our task to measure activity evoked by each of the four stimuli presented within each trial. We jittered the inter-stimulus interval (ISI) between each stimulus so that activity evoked by each stimulus could be isolated from activity evoked by the stimulus before or after it in the sequence. A jittered ISI will ensure that activity evoked by the prior stimulus will not be phase-locked to the onset of the current stimulus. Therefore, by examining activity phase-locked to stimulus onset, we should eliminate activity evoked by the preceding stimulus in the

sequence. However, when we examined the amplitude of stimulus-evoked CTFs through time (Figure 5-5), we found pre-stimulus tuning (in the 300 ms preceding stimulus onset) that was higher amplitude in the attend-stimulus than attend-fixation condition ($p < .0001$). We hypothesized that this pre-stimulus activity reflects activity evoked by the preceding stimulus in the sequence that had a low enough frequency that it was not eliminated by the temporal jitter between stimulus events. Because this pre-stimulus activity was higher amplitude in the attend-stimulus condition than in the attend-fixation condition, it could have contaminated the apparent attentional modulations of stimulus-evoked activity that we observed 80-130 ms after stimulus onset. Thus, we assessed the effect of this lingering activity by examining CTFs as a function of position in the sequence of four stimuli within each trial. Within each trial, the second, third, and fourth stimuli were preceded by a high-contrast stimulus. In contrast, the first stimulus was preceded by the small, low-contrast cue (Figure 5-1) that would have had a much weaker effect on the stimulus-evoked signal to the first target in the sequence. Figure 5-6 shows the reconstructed CTFs from activity evoked by stimuli in each position on the sequence. For this analysis, we trained the IEM on all but the tested stimulus. For example, when testing on the first stimulus in the sequence, we training on stimuli in positions 2-4. We found a robust effect of attention on the amplitude of the stimulus-evoked CTFs across stimuli in all positions in the sequence (all ps $< .05$). In contrast, we found that the CTFs were broader in the attend-stimulus and attend-fixation conditions for the second, third, or fourth stimuli in the sequence (all ps $< .05$), but not for the first stimulus in the sequence ($p = .635$), when the influence of lingering stimulus-evoked activity should be greatly reduced. These results suggest that the increase in CTF width was driven by lingering activity evoked by the preceding stimulus in the sequence.

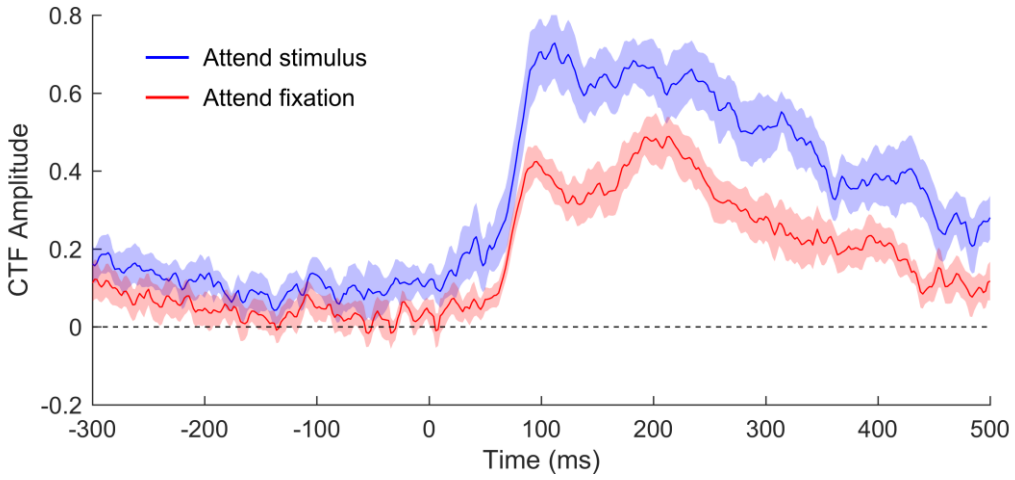

Figure 5-5. Amplitude of stimulus-evoked CTFs as a function of time in Experiment 5-1. Amplitude of stimulus-evoked CTFs as a function of time (0 ms indicates onset of the stimulus). Error bars reflect ±1 bootstrapped SEM.

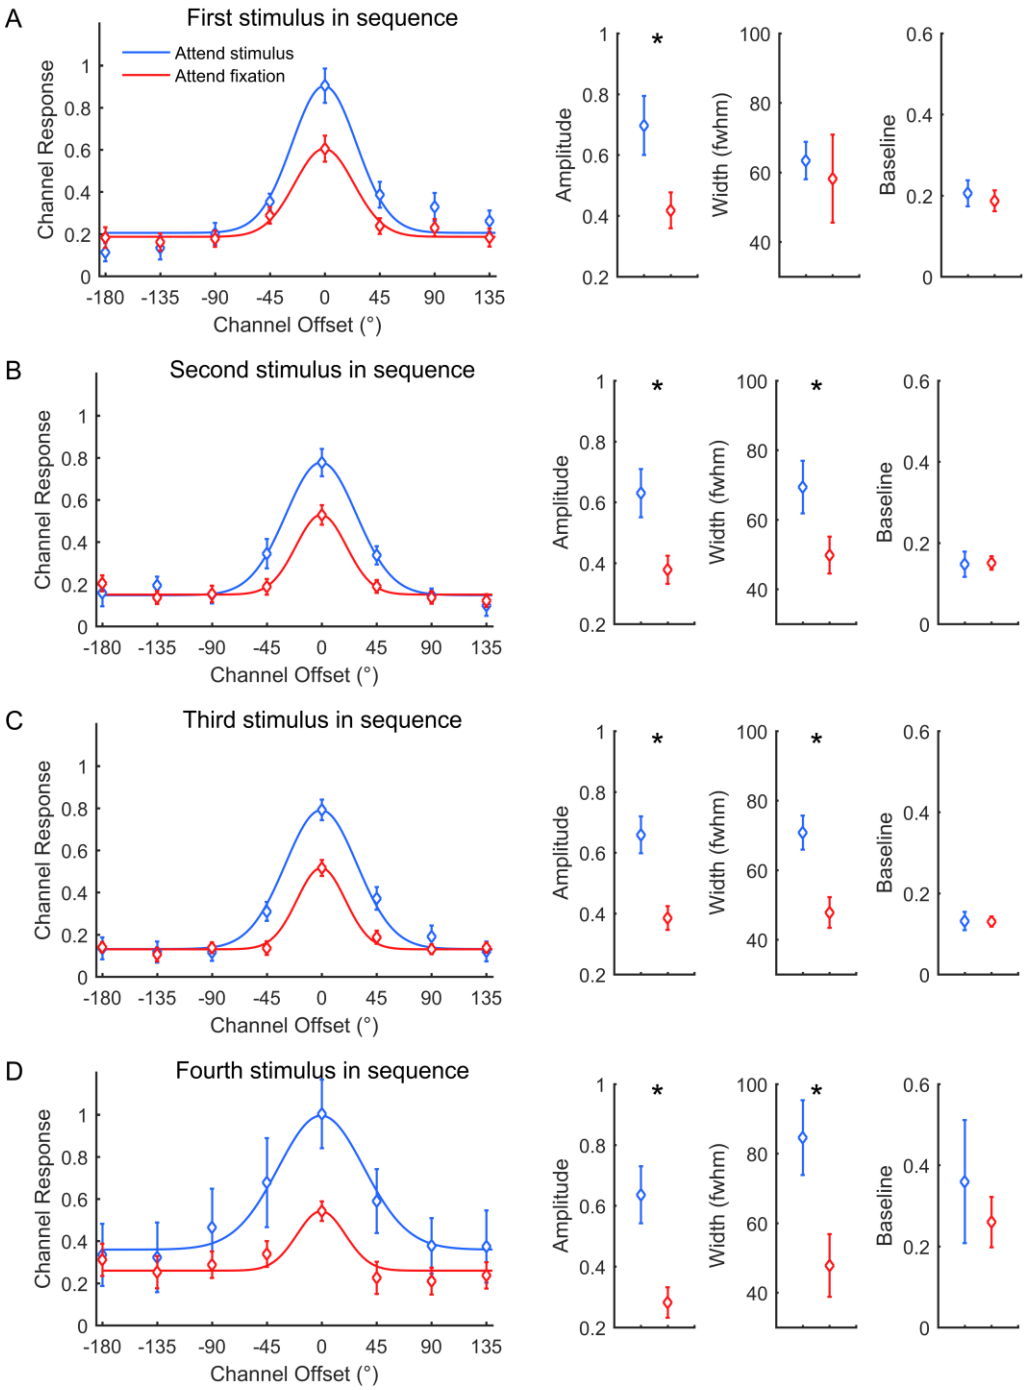

Figure 5-6. Stimulus-evoked CTFs as a function of position of the stimulus in the sequence in Experiment 5-1.
Stimulus-evoked CTFs (measured 80-130 ms after stimulus onset) with best fitting functions (left) and parameter estimates of the best fitting functions (right) for each stimulus in the sequence (a-d). Asterisks mark differences that are significant at the .05 level. Error bars show ± 1 bootstrapped SEM.

Next, to obtain converging evidence for this conclusion, we took a different approach to eliminate lingering activity evoked by the preceding stimulus while still collapsing across all stimulus positions in the sequence. It is primarily low-frequency components will survive temporal jitter. Thus, we reanalyzed the data, this time applying a 4-Hz high-pass filter to remove very low-frequency activity (Figure 5-7). We found that high-pass filtering the data eliminated the pre-stimulus difference in spatial selectivity between the attend-stimulus and attend-fixation conditions ($p = .286$), suggesting that the pre-stimulus activity was restricted to low frequencies. Having established that a high-pass filter eliminated pre-stimulus activity, we re-examined stimulus-evoked CTFs in our window of interest (80-130 ms) after high-pass filtering (Figure 5-8). Again, we found that the CTFs were higher amplitude when the stimulus was attended ($p < .0001$). We also found that CTFs were more broadly tuned when the stimulus was attended ($p < .01$). However, the width effect was considerably smaller after high-pass filtering the data than for the unfiltered data. As we will see in Experiment 5-2, this effect of attention on CTF width did not replicate in Experiment 5-2, suggesting that the primary effect of attention is to increase the amplitude of stimulus-evoked CTFs.

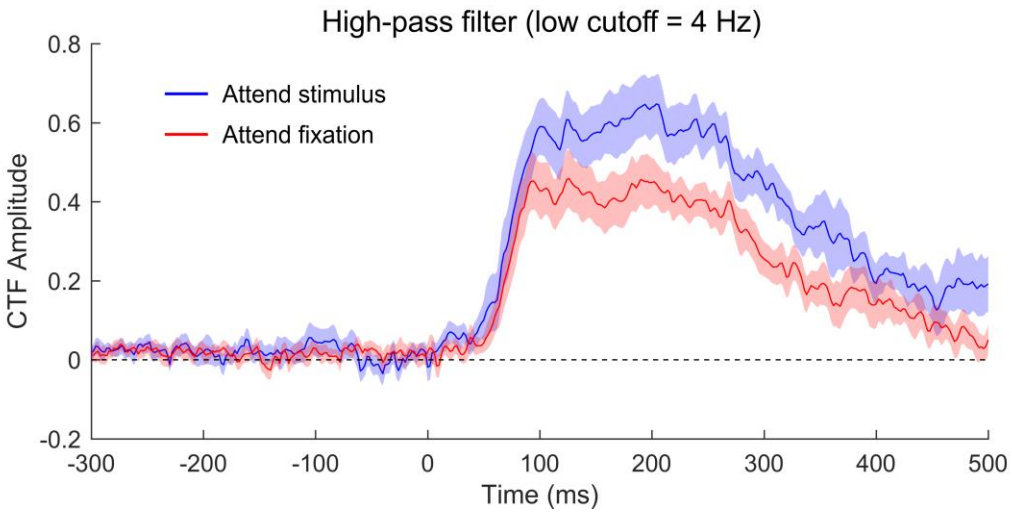

Figure 5-7. Amplitude of stimulus-evoked CTFs as a function of time after high-pass filtering (low cutoff = 4 Hz).
Amplitude of stimulus-evoked CTFs as a function of time (0 ms indicates stimulus onset) after high-pass filtering to remove lingering activity evoked by the preceding stimulus. Error bars reflect ±1 bootstrapped SEM.

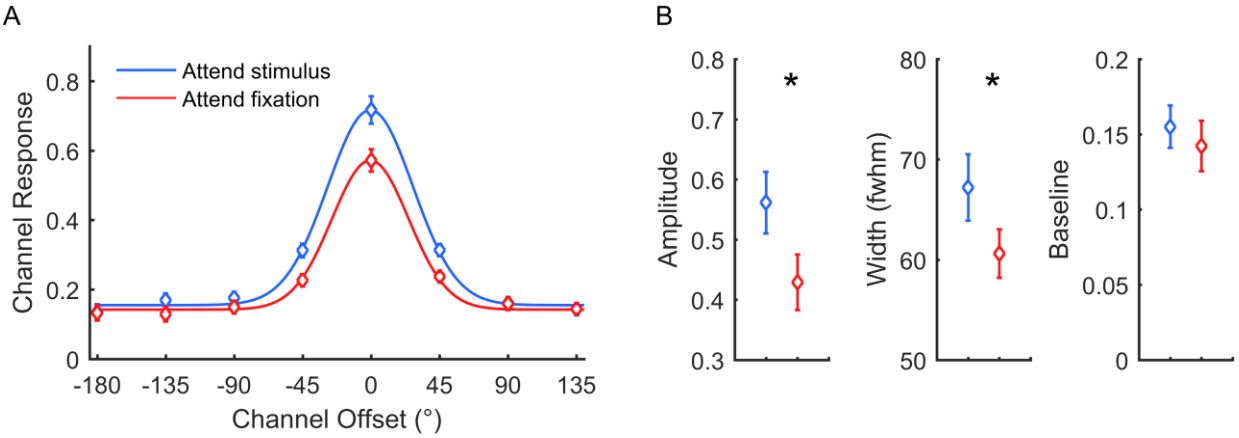

Figure 5-8. Stimulus-evoked CTFs and parameter estimates as a function of condition after high-pass filtering (low cutoff = 4 Hz)
(a) Stimulus-evoked CTFs (measured 80-130 ms after stimulus onset) with best fitting functions. (b) Parameter estimates of the best fitting functions. Error bars show ± 1 bootstrapped SEM.

Experiment 5-2

Past fMRI work has found that the effect of attention on the BOLD signal is stimulus-independent (Buracas & Boynton, 2007; Itthipuripat et al., 2018; Murray, 2008; Sprague, Itthipuripat, et al., 2018). This finding suggests that fMRI is not well suited to examine how attention modulates the stimulus-driven response. In contrast, the effect of attention on stimulus-evoked EEG signals scales with stimulus contrast (Itthipuripat et al., 2014, 2018), consistent with sensory gain models of attention, which hold that attention amplifies stimulus-driven responses (Hillyard et al., 1998). Thus, to replicate and further characterize the effect of attention on stimulus-evoked CTFs in Experiment 5-2, we manipulated stimulus contrast. Observers performed the same attend-stimulus and attend-fixation tasks as in Experiment 5-1, but we parametrically varied the pedestal contrast of the bullseye stimulus (five contrast levels ranging from 6.25 to 90.6%). We adjusted the size of the contrast decrement independently for each of the conditions using a staircase procedure designed to hold accuracy at ~76% correct (see Materials and Methods). Figure 5-9 shows accuracy as a function of pedestal contrast of the stimulus and attention condition. Accuracy was well matched across condition: mean accuracy across subjects did not deviate from 76% by more than 1% any condition.

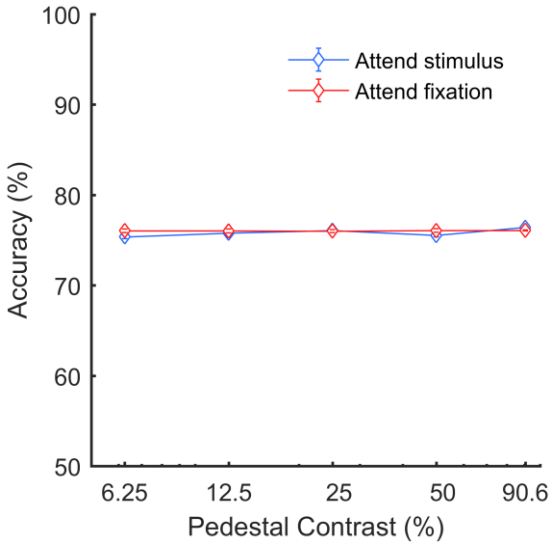

Figure 5-9. Task performance in Experiment 5-2.
Accuracy in Experiment 5-2 as a function of attention condition and pedestal contrast. Error bars reflects ±1 SEM across participants.

We reconstructed CTFs independently for each of the conditions, having estimated channel weights using additional trials (with a pedestal contrast of 90.6%) that were included in for this purpose (see Materials and Methods). In Experiment 5-2, we again used a high-pass filter to remove lingering activity evoked by the preceding stimulus in the sequence. Figures 5-10a and 5-10b show the reconstructed CTFs as a function of contrast for the attend-stimulus and attention-fixation conditions. We again fitted each CTF with an exponentiated cosine function to measure amplitude, width, and baseline of CTFs. For each of the three parameters, we performed a resampling test to test for a main effect of contrast, a main effect of attention, and an attention × contrast interaction. First, we examined CTF amplitude (Figure 5-10c). We found that CTF amplitude increased with stimulus contrast (main effect of contrast: $p < .0001$), and CTF amplitude was larger in the attend-stimulus condition than in the attend-fixation condition (main effect of attention: $p < .0001$). Critically, the effect of attention on CTF amplitude increased with

stimulus contrast (attention × contrast interaction, $p < .001$), as predicted by sensory gain accounts of attention (Hillyard et al., 1998). This finding provides clear evidence that the effect of attention on stimulus-evoked CTFs is not additive with stimulus contrast, as is the case with the BOLD signal (Buracas & Boynton, 2007; Itthipuripat et al., 2018; Murray, 2008; Sprague, Itthipuripat, et al., 2018). Next, we examined CTF width (Figure 5-10d). We found that estimates of CTF width were very noisy for the 6.25 and 12.5% contrast conditions because of the low amplitude of the CTFs in these conditions, precluding confidence in those estimates. Thus, we restricted our analysis to the higher contrast conditions (25.0, 50.0, and 90.6% contrast). We found no main effect of attention ($p = .824$), and no main effect of contrast ($p = .112$). However, we found a reliable attention × contrast interaction ($p = .031$), such that CTFs were narrower when the stimulus was attended for the 90.6% contrast condition and 50% contrast condition, and were broader for the 25% contrast condition, but none of these differences between the attend-stimulus and attend fixation conditions survived Bonferoni correction ($p = .031$, $p = .292$, and $p = .286$, respectively; $\alpha_{corrected} = .05/3 = .017$). Thus, we did not replicate the finding from Experiment 5-1 that stimulus-evoked CTFs were broader when the stimulus was attended. Finally, we examined CTF baseline (Figure 5-10e). We found that CTF baseline was higher in the attend-stimulus condition than in the attend-fixation condition (main effect of attention, $p < .01$), but found no main effect of contrast ($p = .678$) or attention × contrast interaction ($p = .132$)

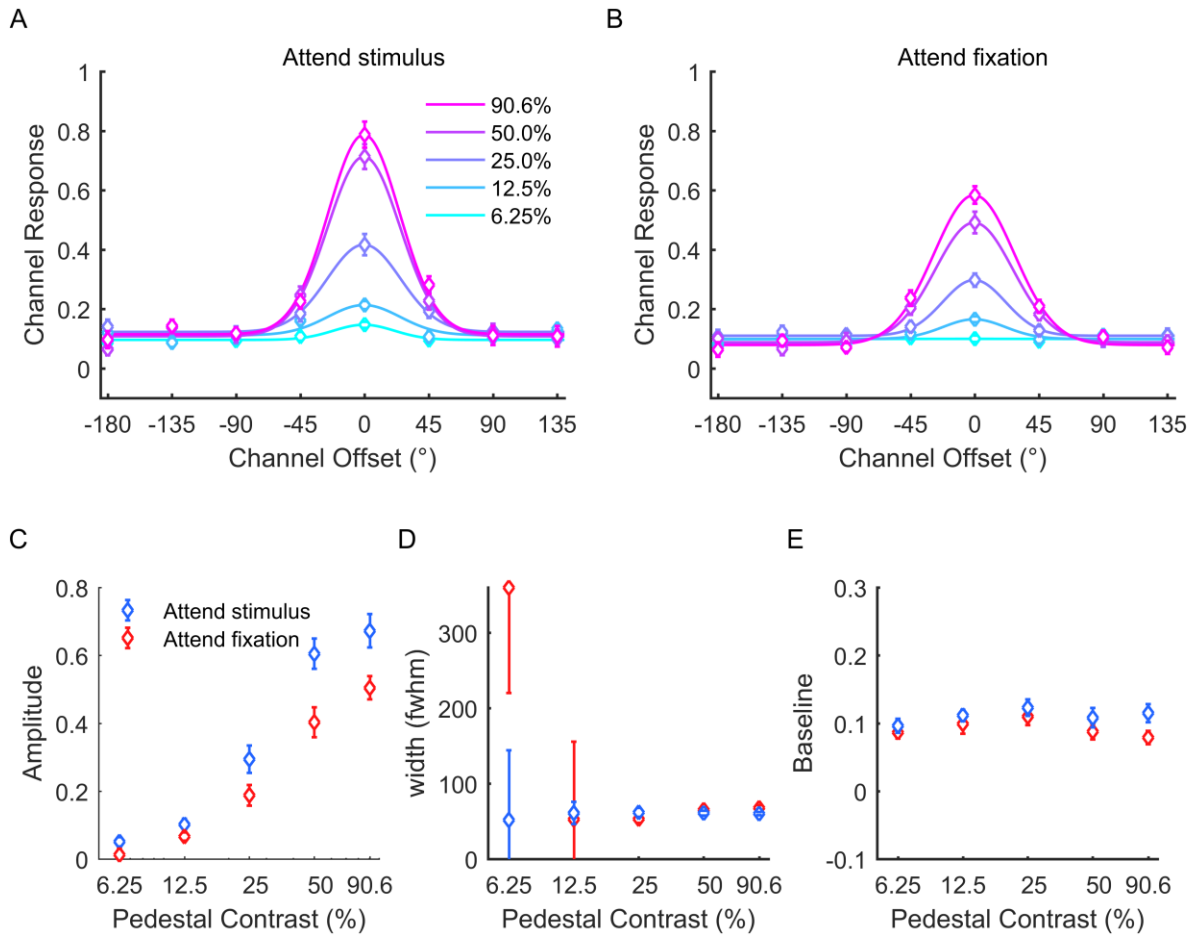

Figure 5-10. Stimulus-evoked CTFs and parameter estimates in Experiment 5-2. Stimulus-evoked CTFs as a function of stimulus contrast in the attend-stimulus (a) and attend-fixation (b) conditions. (c-e) parameter estimates as a function of task condition and stimulus contrast. Error bars reflect ±1 bootstrapped SEM across subjects.

Discussion

Past work that has examined the effect of attention on population-level representations of stimulus position has relied on BOLD activity measured with fMRI. These studies have found that attention primarily increases the amplitude of representations of stimulus position (Fischer & Whitney, 2009; Sprague, Itthipuripat, et al., 2018; Sprague & Serences, 2013; Vo et al., 2017). However, the effect of attention on the BOLD signal is best described as a stimulus-independent shift in baseline activity rather than reflecting a modulation of the stimulus-driven response

(Buracas & Boynton, 2007; Itthipuripat et al., 2018; Murray, 2008; Sprague, Itthipuripat, et al., 2018). Thus, it has remained unclear how attention modulates the spatial selectivity of stimulus-driven activity. Here, we provide compelling evidence regarding the effects of attention on stimulus-evoked activity. Across two experiments, we found clear evidence that attention increased the amplitude of stimulus-evoked CTFs. In Experiment 5-2, we parametrically varied stimulus contrast. Unlike with the BOLD signal, we found that the effect of attention on the amplitude of the stimulus-evoked CTF increased with stimulus contrast, confirming that this effect reflects an attentional modulation of the stimulus-driven signal rather than activity that is independent of the stimulus. Taken together, our results provide clear evidence that attention increases the amplitude of the spatially tuned stimulus-evoked representations.

The IEM approach also enabled us to examine how attention influenced width of stimulus-evoked CTFs. However, the evidence that attention changes the width of stimulus-evoked CTFs was equivocal. Stimulus-evoked CTFs were broader for attended stimuli than for unattended stimuli in Experiment 5-1, but this finding was largely eliminated when the influence of prior stimulus events was minimized, and the effect did not replicate in Experiment 5-2. It is tempting to conclude that attention does not change the width of stimulus-evoked CTFs. However, it is possible that the lack of a reliable effect of attention on the width of stimulus-evoked CTFs could reflect a lack of sensitivity of our approach to small changes in the width of spatially selective activity. Another possibility is that a different target discrimination task, one that requires precise encoding of the spatial position of the stimulus rather than of stimulus contrast (as in the current experiments), may be more conducive to a modulation of tuning width (e.g. Fischer & Whitney, 2009). Thus, more work is needed to assess the sensitivity of our

approach to changes in the width of tuning. One promising approach to address this issue is to parametrically increase the physical size of a stimulus, which should directly translate into broader CTFs. This is an important avenue for future work.

A core strength of our EEG-based approach is that it enabled temporally precise measurement of population-level CTFs. Past studies that have relied on the BOLD signal (e.g. Fischer & Whitney, 2009; Sprague & Serences, 2013; Vo et al., 2017) are limited by the sluggish hemodynamic response, which peaks several seconds after stimulus onset. The slow time course of the BOLD response makes it difficult to establish what stage of processing is influenced by attention. In contrast, the temporal resolution that EEG affords allowed us to isolate early visually evoked activity occurring ~100 ms after stimulus onset. Thus, our findings provide robust evidence that attention enhances the amplitude of population representations of stimulus position as early as early as 100 ms after stimulus onset.

CHAPTER 6. GENERAL DISCUSSION

Our capacity to process visual information is limited. Therefore, we must prioritize processing at relevant locations. In this dissertation, I have shown that alpha-band oscillations precisely track spatial priority in a range of contexts. To this end, I developed an encoding model approach to reconstruct spatially selective response profiles (called channel-tuning functions, or CTFs) from alpha-band power measured with EEG. These alpha-band CTFs reflect the spatial selectivity of the population-level activity that is measured with EEG. In Chapter 2, I showed that alpha oscillations precisely track where attention is deployed following central cues and during visual search. Furthermore, I showed that the time course of spatially selective alpha activity is sensitive to trial-by-trial variations in the latency of covert orienting. In Chapter 3, I showed that alpha activity precisely track locations held in working memory. In Chapter 4, I showed that alpha activity tracks the spatial position of memoranda held in working memory, even when spatial positon is completely irrelevant. Together, these studies show that-alpha band activity tracks a general selection mechanism that plays a general role in perception and in the maintenance of online representations in working memory. Finally, in Chapter 5, I leveraged the encoding model approach to examine the *consequences* of spatial attention. I showed that attention enhances population-level representations of stimulus position. In the following sections, I discuss open questions about the role of alpha oscillations in spatial selection.

Does alpha-band activity reflect signal enhancement or distractor suppression?

What aspect of visual selection does alpha activity reflect? The modal view is that alpha activity mediates the suppression or gating of irrelevant visual inputs (Clayton, Yeung, & Cohen Kadosh, 2015, 2017; Foxe & Snyder, 2011; Jensen & Mazaheri, 2010; Payne & Sekuler, 2014).

This view falls in line with the consensus that distractor exclusion is a critical component of visual attention (Desimone & Duncan, 1995). However, it is broadly acknowledged that improved perception at a relevant location can also occur via *signal enhancement*, which directly improves processing at attended locations (Carrasco, 2011; Hillyard et al., 1998). While many studies have shown that alpha activity tracks the attended positions in the visual field (Kelly et al., 2006; Thut et al., 2006; Worden et al., 2000), even in the absence of irrelevant distractors (Kelly, Foxe, Newman, & Edelman, 2010; Rihs et al., 2007; Thut et al., 2006), recent work has cast doubt on whether alpha activity tracks locations at which distractors are expected (Noonan et al., 2016). At first glance, this may seem to point towards a role in signal enhancement rather than distractor suppression. However, a neural signal that suppresses interference at all unattended locations may also enable precise tracking of target position. Thus, the fact that alpha activity precisely tracks the selected locations – but not the locations of distractors – still leaves open the question of *how* alpha activity supports selective attention.

In support of the distractor suppression account, past work has emphasized the finding in spatial-cueing tasks that alpha power *decreases* contralateral to the cued location and/or *increases* contralateral to the uncued location (Kelly et al., 2006; Rihs et al., 2007; Thut et al., 2006; Worden et al., 2000). This finding has been presented as suggesting that higher alpha power contralateral to the uncued location reflects increased suppression of irrelevant stimuli (Clayton et al., 2015, 2017; Foxe & Snyder, 2011; Jensen & Mazaheri, 2010; Payne & Sekuler, 2014). But this empirical pattern is equally compatible with the view that reduced alpha-band power reflects signal enhancement. Another finding that has motivated a suppression account of alpha activity is the inverse relationship between alpha power and other neural signals such as

spiking activity (Haegens, Nacher, Luna, Romo, & Jensen, 2011) and gamma-band oscillations (Spaak, Bonnefond, Maier, Leopold, & Jensen, 2012). However, these findings do not establish whether alpha activity influences the quality of *sensory representations* because the information content of neural activity can be disconnected from the overall amount of neural activity (e.g. Postle, 2015). For example, inhibition of specific neural units could improve the fidelity of a sensory representation. Thus, it is unclear whether these inverse relationships reflect an inverse relationship between alpha power and the strength of sensory representations. Even though there is some evidence for an inverse relationship between alpha power and the strength of sensory representations measured via BOLD responses in visual cortex (Zumer, Scheeringa, Schoffelen, Norris, & Jensen, 2014), these data are still compatible with the hypothesis that decreased alpha power reflects a relative increase in signal enhancement over the attended regions.

However, a more diagnostic approach for distinguishing between suppression and enhancement accounts is to use experimental designs that selectively manipulate the process of interest. If the degree of distractor suppression can be manipulated while signal enhancement is held constant, this provides an opportunity to link specific neural signals with suppression *per se*. For example, Serences and colleagues (Serences, Yantis, Culberson, & Awh, 2004) used this approach to test whether preparatory activity measured with fMRI reflected distractor exclusion. Spatial attention increases baseline activity measured with fMRI in visual cortex tuned for the attended location (Kastner et al., 1999; Tootell et al., 1998). To test whether this preparatory activity reflects distractor exclusion, Serences et al. varied the probability that distractors accompanied visual targets, a manipulation that has been shown to increase resistance to distractor interference without affecting performance with distractor-free displays (Awh,

Matsukura, & Serences, 2003). The selective effect of the probability manipulation suggests a specific effect on distractor exclusion, because changes in that process should not affect performance when there are no distractors to exclude. Critically, Serences et al. found that preparatory activity, measured via retinotopic changes in the amplitude of the BOLD signal, was greater when the probability of distractors was higher, suggesting that this preparatory activity plays a specific role in distractor exclusion. Therefore, this work provides a clear example of how preparatory activity can be unambiguously linked with distractor exclusion.

Of the substantial body of work that links alpha-band power with covert attention, only a few studies have attempted to selectively manipulate distractor suppression (Haegens, Luther, & Jensen, 2012; Janssens, De Loof, Boehler, Pourtois, & Verguts, 2018; Kelly et al., 2010; Noonan et al., 2016). In one study, Kelly and colleagues (Kelly et al., 2010) cued the location of an upcoming target. In some blocks, the target appeared alone, while in others a distractor appeared in the uncued hemifield. Kelly et al. reasoned that if spatially specific alpha-band power reflects distractor exclusion, then lateralized alpha-band should be stronger when observers expect a distractor than when no distractor is expected. Interestingly, lateralization of alpha-band power was *weaker* when distractors were expected (also see Haegens et al., 2012, which found no effect of the strength of distractors on the lateralization of alpha-band power in a somatosensory task). However, one caveat here is that there was no behavioral evidence that distractor exclusion was increased in blocks that contained distractors. Thus, the null effect on alpha laterality may not provide strong evidence against a distractor exclusion account.

Händel and colleagues (Händel, Haarmeier, & Jensen, 2011) took a different approach. Rather than manipulating the degree of distractor exclusion, they sought to test whether the

degree of lateralization of alpha-band activity predicted individual differences in distractor exclusion. The authors reasoned that if lateralized alpha-band activity predicted processing of the stimulus in the uncued hemifield (on invalidly cued trials) but not processing of the stimulus in the cued hemifield (on validly cued trials), then this would provide evidence in favor of the distractor exclusion account of alpha-band power. Their results followed this pattern, but there was considerably more variability in performance for the uncued stimulus than for the cued stimulus. Thus, the failure to detect a relationship between alpha-band lateralization and processing of the cued stimulus may reflect a restriction of range for performance at the cued position.

In another study, Noonan and colleagues (Noonan et al., 2016) successfully manipulated distractor exclusion. Observers responded to a target stimulus. Noonan et al. varied whether or not the target was accompanied by a distractor in different blocks of trials, and found that the presence of a distractor reliably slowed response times. In some blocks, Noonan et al. cued the location of the target or the distractor in advance. Unsurprisingly, a spatial cue indicating the target location speeded responses (compared to trials with no cue). Interestingly, a spatial cue indicating the *distractor* location also speeded responses (again, compared to trials with no cue). Critically, this distractor-cueing benefit was not seen when the distractor was absent, providing clear evidence that cues indicating the location of a distractor selectively enabled distractor exclusion. Strikingly, Noonan et al. found that alpha-band activity tracked the cued location when the target location was cued but not when the distractor location was cued, suggesting that alpha-band activity plays a role in signal enhancement but providing no evidence for a role in distractor exclusion.

To summarize, while it is clearly the modal view that alpha activity reflects the suppression of irrelevant visual information, the evidence is equivocal. Extant work linking alpha activity with spatial attention is compatible with an account in which alpha supports target selection via signal enhancement. Thus, there is strong motivation for further work in which alpha-band activity is assessed during selective manipulations of different aspects of spatial selection, preferably with analytic approaches that focus on spatially-selective alpha-band activity rather than overall power in that frequency band. For example, given prior work that has reported covariations of alpha power and BOLD activity (Hermes, Nguyen, & Winawer, 2017), it may be worthwhile to use an approach similar to that of Serences et al. (Serences et al., 2004) to examine how alpha activity is affected by selective changes in distractor exclusion. Until such work is conducted, the jury is still out regarding the computational role played by alpha activity during spatial selection.

Do alpha-band oscillations play a causal role in spatial selection?

The vast majority of evidence – including the work presented in this dissertation – is correlational. Although alpha activity precisely tracks selected locations, these findings do not provide evidence that alpha activity plays a causal role in spatial selection. To test for a causal relationship, one must manipulate alpha activity to and test whether this has consequences for spatial selection. Furthermore, the manipulation must be a *selective* manipulation of alpha activity: if the manipulation also influences other neural activity, then any effect of that manipulation on spatial selection could be mediated other neural consequences of the manipulation.

To date, just one study has sought to selectively manipulate alpha activity to test for a causal relationship between alpha activity and spatial selection. In this study, Romei and colleagues (Romei, Gross, & Thut, 2010) aimed to manipulate oscillatory activity in occipito-parietal cortex by stimulating cortex with rhythmic transcranial magnetic stimulus (TMS). Observers performed a visual detection task. On each trial, observers reported whether or not a near-threshold target appeared, which could appear in the left or right hemifield. Immediately before the target, Romei and colleagues applied a train of TMS pulses over occipito-parietal cortex, either contralateral or ipsilateral to the impending target. The authors applied 10-Hz stimulation to manipulate parieto-occipital alpha activity (Thut et al., 2011). Because alpha power is usually higher ipsilateral to an attended location, Romei and colleagues reasoned that stimulation (intended to *increase* alpha activity) should improve performance when it is applied *ipsilateral* to the target. Indeed, they found that target detection rate was improved following ipsilateral stimulation. Furthermore, this effect was not seen when stimulation was applied at 5-Hz or 20-Hz, showing that the effect was specific to 10-Hz stimulation. These findings led the authors to conclude that alpha oscillations play a causal role in modulating visual processing.

However, this finding must be interpreted with caution. The conclusion that alpha oscillations specifically play a causal role rest firmly on the assumption that the manipulation selectively influences alpha oscillations. While there is evidence that rhythmic TMS does drive neural oscillations (Thut et al., 2011), it remains unclear whether this kind of stimulation selectively influences alpha oscillations without changing other neuronal dynamics, leaving open the possibility that some other change in neural activity produced the change in behavior.

Therefore, whether alpha-band activity plays a causal role in spatial selection remains an important question for future work.

Conclusions

The work presented in this dissertation establishes that alpha-band oscillations are intimately linked with spatial priority in the human brain. I showed that alpha oscillations provide a spatially and temporally resolved index of spatial attention, establishing a powerful tool for precisely tracking spatial selection. Furthermore, I showed that this selection mechanism is recruited in a range of contexts – both during the selection of relevant locations during perception and during the maintenance of information in working memory. Although questions remain about the exact computational role that alpha-band activity plays in spatial selection, it is clear that alpha oscillations provide a powerful means of tracking spatial priority in the human brain.

REFERENCES

Agarwal, G., Stevenson, I. H., Berényi, A., Mizuseki, K., Buzsáki, G., & Sommer, F. T. (2014). Spatially distributed local fields in the hippocampus encode rat position. *Science*, *344*, 626–630.

Anton-Erxleben, K., & Carrasco, M. (2013). Attentional enhancement of spatial resolution: linking behavioural and neurophysiological evidence. *Nature Reviews. Neuroscience*, *14*(3), 188–200. https://doi.org/10.1038/nrn3443

Anton-Erxleben, K., Stephan, V. M., & Treue, S. (2009). Attention reshapes center-surround receptive field structure in macaque cortical area MT. *Cerebral Cortex*, *19*(10), 2466–2478. https://doi.org/10.1093/cercor/bhp002

Awh, E., Jonides, J., Smith, E. E., Buxton, R. B., Frank, L. R., Love, T., … Gmeindl, L. (1999). Rehearsal in spatial working memory: evidence from neuroimaging. *Psychological Science*, *10*(5), 433–437. https://doi.org/10.1111/1467-9280.00182

Awh, E, Anllo-Vento, L., & Hillyard, S. A. (2000). The role of spatial selective attention in working memory for locations: Evidence from event-related potentials. *Journal of Cognitive Neuroscience*, *12*(5), 840–847.

Awh, E, & Jonides, J. (2001). Overlapping mechanisms of attention and spatial working memory. *Trends in Cognitive Sciences*, *5*(3), 119–126.

Awh, E, Jonides, J., & Reuter-Lorenz, P. A. (1998). Rehearsal in spatial working memory. *Journal of Experimental Psychology. Human Perception and Performance*, *24*(3), 780–790.

Awh, E, & Pashler, H. (2000). Evidence for split attentional foci. *Journal of Experimental Psychology. Human Perception and Performance*, *26*(2), 834–846. https://doi.org/10.1037/0096-1523.26.2.834

Awh, E, Vogel, E. K., & Oh, S.-H. (2006). Interactions between attention and working memory. *Neuroscience*, *139*(1), 201–208.

Awh, Edward, Matsukura, M., & Serences, J. T. (2003). Top-down control over biased competition during covert spatial orienting. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(1), 52–63. https://doi.org/10.1167/2.7.15

Bahramisharif, A., Heskes, T., Jensen, O., & van Gerven, M. A. J. (2011). Lateralized responses during covert attention are modulated by target eccentricity. *Neuroscience Letters*, *491*(1), 35–39.

Bahramisharif, A., van Gerven, M. A. J., Heskes, T., & Jensen, O. (2010). Covert attention allows for continuous control of brain-computer interfaces. *European Journal of Neuroscience*, *31*(8), 1501–1508.

Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9*(10), 1–11.

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

Brefczynski, J. A., & DeYoe, E. A. (1999). A physiological correlate of the "spotlight" of visual attention. *Nature Neuroscience*, *2*(4), 370–374.

Brouwer, G. J., & Heeger, D. J. (2009). Decoding and reconstructing color from responses in human visual cortex. *Journal of Neuroscience*, *29*(44), 13992–14003.

Brouwer, G. J., & Heeger, D. J. (2011). Cross-orientation suppression in human visual cortex. *Journal of Neurophysiology*, *106*(5), 2108–2119. https://doi.org/10.1152/jn.00540.2011

Buracas, G. T., & Boynton, G. M. (2007). The effect of spatial attention on contrast response functions in human visual cortex. *Journal of Neuroscience*, *27*(1), 93–97. https://doi.org/10.1523/JNEUROSCI.3162-06.2007

Butts, D. A., & Goldman, M. S. (2006). Tuning curves, neuronal variability, and sensory coding. *PLOS Biology*, *4*(4), 639–646.

Canolty, R. T., & Knight, R. T. (2010). The functional role of cross-frequency coupling. *Trends in Cognitive Sciences*, *14*(11), 506–515.

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, *51*(13), 1484–1525.

Carrasco, M., & McElree, B. (2001). Covert attention accelerates the rate of visual information processing. *Proceedings of the National Academy of Sciences of the United States of America*, *98*(9), 5363–5367.

Cheal, M., & Lyon, D. R. (1991). Central and peripheral precuing of forced-choice discrimination. *Quarterly Journal of Experimental Psychology*, *43A*(4), 859–880.

Christophel, T. B., Hebart, M. N., & Haynes, J.-D. (2012). Decoding the contents of visual short-term memory from human visual and parietal cortex. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *32*(38), 12983–12989. https://doi.org/10.1523/JNEUROSCI.0184-12.2012

Chun, M. M. (2011). Visual working memory as visual attention sustained internally over time. *Neuropsychologia*, *49*(6), 1407–1409.

Clayton, M. S., Yeung, N., & Cohen Kadosh, R. (2015). The roles of cortical oscillations in sustained attention. *Trends in Cognitive Sciences*, *19*(4), 188–195.

Clayton, M. S., Yeung, N., & Cohen Kadosh, R. (2017). The many characters of visual alpha oscillations. *European Journal of Neuroscience*, *48*(7), 1–11.

Cohen, M. R., & Maunsell, J. H. R. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience*, *12*(12), 1594–1600. https://doi.org/10.1038/nn.2439

Connor, C. E., Preddie, D. C., Gallant, J. L., & Van Essen, D. C. (1997). Spatial Attention Effects in Macaque Area V4. *Journal of Neuroscience*, *17*(9), 3201–3214.

Cornelissen, F. W., Peters, E. M., & Palmer, J. (2002). The Eyelink Toolbox: eye tracking with MATLAB and the Psychophysics Toolbox. *Behavior Research Methods, Instruments, & Computers*, *34*(4), 613–617.

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.

Delorme, A., & Makieg, 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.

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

Dill, M., & Fahle, M. (1998). Limited translation invariance of human visual pattern recognition. *Perception & Psychophysics*, *60*(1), 65–81.

Downing, C. J. (1988). Expectancy and visual-spatial attention: effects on perceptual quality. *Journal of Experimental Psychology: Human Perception and Performance*, *14*(2), 188–202.

Drew, T., & Vogel, E. K. (2008). Neural measures of individual differences in selecting and tracking multiple moving objects. *Journal of Neuroscience*, *28*(16), 4183–4191.

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

Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap*. New York, NY: Chapman and Hall.

Egeth, H. E., & Yantis, S. (1997). Visual attention: Control, representation, and time course. *Annual Review of Psychology*, *48*, 269–297.

Emrich, S. M., Riggall, A. C., Larocque, J. J., & Postle, B. R. (2013). Distributed patterns of activity in sensory cortex reflect the precision of multiple items maintained in visual short-term memory. *Journal of Neuroscience*, *33*(15), 6516–6523.

Eriksen, C. W., & Collins, J. F. (1969). Temporal course of selective attention. *Journal of Experimental Psychology*, *80*(2), 254–261. https://doi.org/10.1037/h0027268

Eriksen, C. W., & Hoffman, J. E. (1974). Selective attention: Noise suppression or signal enhancement? *Bulletin of the Psychonomic Society*, *4*, 587–589.

Ester, E. F., Anderson, D. E., Serences, J. T., & Awh, E. (2013). A neural measure of precision in visual working memory. *Journal of Cognitive Neuroscience*, *25*(5), 754–761.

Ester, E. F., Sprague, T. C., & Serences, J. T. (2015). Parietal and frontal cortex encode stimulus-specific mnemonic representations during visual working memory. *Neuron*, *87*, 1–13. https://doi.org/10.1016/j.neuron.2015.07.013

Fell, J., & Axmacher, N. (2011). The role of phase synchronization in memory processes. *Nature Reviews Neuroscience*, *12*(2), 105–118.

Fischer, J., & Whitney, D. (2009). Attention narrows position tuning of population responses in V1. *Current Biology*, *19*(16), 1356–1361. https://doi.org/10.1016/j.cub.2009.06.059

Foster, D. H., & Kahn, J. I. (1985). Internal representations and operations in the visual comparison of transformed patterns: Effects of pattern point-inversion, positional symmetry, and separation. *Biological Cybernetics*, *51*(5), 305–312.

Foster, J. J., Bsales, E. M., Jaffe, R. J., & Awh, E. (2017). Alpha-band activity reveals spontaneous representations of spatial position in visual working memory. *Current Biology*, *27*(20), 3216–3223. https://doi.org/10.1016/j.cub.2017.09.031

Foster, J. J., Sutterer, D. W., Serences, J. T., Vogel, E. K., & Awh, E. (2016). The topography of alpha-band activity tracks the content of spatial working memory. *Journal of Neurophysiology*, *115*(1), 168–177. https://doi.org/10.1152/jn.00860.2015

Foster, J. J., Sutterer, D. W., Serences, J. T., Vogel, E. K., & Awh, E. (2017). Alpha-band oscillations enable spatially and temporally resolved tracking of covert spatial attention. *Psychological Science*, *28*(9), 929–941. https://doi.org/10.1177/0956797617699167

Foxe, J. J., & Snyder, A. C. (2011). The role of alpha-band brain oscillations as a sensory suppression mechanism during selective attention. *Frontiers in Psychology*, *2*, 154.

Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, *9*(10), 474–480.

García-Pérez, M. A. (1998). Forced-choice staircases with fixed step sizes: Asymptotic and small-sample properties. *Vision Research*, *38*(12), 1861–1881. https://doi.org/10.1016/S0042-6989(97)00340-4

Garcia, J. O., Srinivasan, R., & Serences, J. T. (2013). Near-real-time feature-selective modulations in human cortex. *Current Biology*, *23*(6), 515–522. https://doi.org/10.1016/j.cub.2013.02.013

Gazzaley, A., & Nobre, A. C. (2012). Top-down modulation: Bridging selective attention and working memory. *Trends in Cognitive Sciences*, *16*(2), 129–135.

Gould, I. C., Rushworth, M. F., & Nobre, A. C. (2011). Indexing the graded allocation of visuospatial attention using anticipatory alpha oscillations. *Journal of Neurophysiology*, *105*(3), 1318–1326. https://doi.org/10.1152/jn.00653.2010

Gruber, W. R., Klimesch, W., Sauseng, P., & Doppelmayr, M. (2005). Alpha phase synchronization predicts P1 end N1 latency and amplitude size. *Cerebral Cortex*, *15*(4), 371–377. https://doi.org/10.1093/cercor/bhh139

Haegens, S., Nacher, 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.

Haegens, Saskia, Luther, L., & Jensen, O. (2012). Somatosensory anticipatory alpha activity increases to suppress distracting input. *Journal of Cognitive Neuroscience*, *24*(3), 677–685.

Händel, B. F., Haarmeier, T., & Jensen, O. (2011). Alpha oscillations correlate with the successful inhibition of unattended stimuli. *Journal of Cognitive Neuroscience*, *23*(9), 2494–2502.

Hanslmayr, S., Klimesch, W., Sauseng, P., Gruber, W., Doppelmayr, M., Freunberger, R., … Birbaumer, N. (2007). Alpha phase reset contributes to the generation of ERPs. *Cerebral Cortex*, *17*(1), 1–8. https://doi.org/10.1093/cercor/bhj129

Harrison, S. A., & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, *458*(7238), 632–635.

Hebb, D. O. (1949). *The organization of behavior*. New York: Wiley.

Hermes, D., Nguyen, M., & Winawer, J. (2017). Neuronal synchrony and the relation between the BOLD response and the local field potential. *PLoS Biology*, *15*(7), e2001461.

Hillyard, S. A., & Anllo-Vento, L. (1998). Event-related brain potentials in the study of visual selective attention. *Proceedings of the National Academy of Sciences of the United States of America*, *95*(3), 781–787. https://doi.org/10.1073/pnas.95.3.781

Hillyard, S. A., Vogel, E. K., & Luck, S. J. (1998). Sensory gain control (amplification) as a mechanism of selective attention: electrophysiological and neuroimaging evidence. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *353*(1373), 1257–1270.

Hopfinger, J. B., & Mangun, G. R. (1998). Reflexive attention modulates processing of visual stimuli in human extrastriate cortex. *Psychological Science*, *9*(6), 441–447. https://doi.org/10.1111/1467-9280.00083

Itthipuripat, S., Ester, E. F., Deering, S., & Serences, J. T. (2014). Sensory gain outperforms efficient readout mechanisms in predicting attention-related improvements in behavior. *Journal of Neuroscience*, *34*(40), 13384–13398. https://doi.org/10.1523/JNEUROSCI.2277-14.2014

Itthipuripat, S., Sprague, T. C., & Serences, J. T. (2018). Reconciling fMRI and EEG indices of attentional modulations in human visual cortex. *BioRxiv*. Retrieved from https://www.biorxiv.org/content/early/2018/08/13/391193.full.pdf

Janssens, C., De Loof, E., Boehler, C. N., Pourtois, G., & Verguts, T. (2018). Occipital alpha power reveals fast attentional inhibition of incongruent distractors. *Psychophysiology*, *55*(3), 1–11.

Jazayeri, M., & Movshon, J. A. (2006). Optimal representation of sensory information by neural populations. *Nature Neuroscience*, *9*(5), 690–696. https://doi.org/10.1038/nn1691

Jensen, O., & Mazaheri, A. (2010). Shaping functional architecture by oscillatory alpha activity: Gating by inhibition. *Frontiers in Human Neuroscience*, *4*, 186.

Johnston, J. C., & Pashler, H. (1990). Close binding of identity and location in visual feature perception. *Journal of Experimental Psychology. Human Perception and Performance*, *16*(4), 843–856. https://doi.org/10.1037/0096-1523.16.4.843

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

Kay, K. N., Weiner, K. S., & Grill-Spector, K. (2015). Attention Reduces Spatial Uncertainty in Human Ventral Temporal Cortex. *Current Biology*, *25*(5), 595–600. https://doi.org/10.1016/j.cub.2014.12.050

Kelly, S. P., Foxe, J. J., Newman, G., & Edelman, J. A. (2010). Prepare for conflict: EEG correlates of the anticipation of target competition during overt and covert shifts of visual attention. *European Journal of Neuroscience*, *31*(9), 1690–1700.

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.

Kim, M., & Cave, K. R. (1995). Spatial attention in visual search for features and features conjunctions. *Psychological Science*, *6*(6), 376–380.

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

Kuo, B.-C. C., Rao, A., Lepsien, J., & Nobre, A. C. (2009). Searching for targets within the spatial layout of visual short-term memory. *Journal of Neuroscience*, *29*(25), 8032–8038. https://doi.org/10.1523/JNEUROSCI.0952-09.2009

Lavie, N. (2005). Distracted and confused?: Selective attention under load. *Trends in Cognitive Sciences*, *9*(2), 75–82. https://doi.org/10.1016/j.tics.2004.12.004

Li, X., Lu, Z.-L., Tjan, B. S., Dosher, B. A., & Chu, W. (2008). Blood oxygenation level-dependent contrast response functions identify mechanisms of covert attention in early visual areas. *Proceedings of the National Academy of Sciences*, *105*(16), 6202–6207. https://doi.org/10.1073/pnas.0801390105

Lins, O. G., Picton, T. W., Berg, P., & Scherg, M. (1993a). Ocular artifacts in EEG and event-related potentials. I: Scalp topography. *Brain Topography*, *6*(1), 51–63.

Lins, O. G., Picton, T. W., Berg, P., & Scherg, M. (1993b). Ocular artifacts in recording EEGs and event-related potentials. II: Source dipoles and source components. *Brain Topography*, *6*(1), 65–78. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/8260328

Lisman, J. (2010). Working memory: the importance of theta and gamma oscillations. *Current Biology*, *20*(11), R490-2. https://doi.org/10.1016/j.cub.2010.04.011

Liu, T., Stevens, S. T., & Carrasco, M. (2007). Comparing the time course and efficacy of spatial and feature-based attention. *Vision Research*, *47*(1), 108–113. https://doi.org/10.1016/j.visres.2006.09.017

Lopes da Silva, F. (2013). EEG and MEG: Relevance to neuroscience. *Neuron*, *80*(5), 1112–1128.

Luck, S J, Chelazzi, L., Hillyard, S. A., & Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology*, *77*(1), 24–42.

Luck, Steven J, Fan, S., & Hillyard, S. A. (1993). Attention-related modulation of sensory-evoked brain activity in a visual search task. *Journal of Cognitive Neuroscience*, *5*(2), 188–195. https://doi.org/10.1162/jocn.1993.5.2.188

Mack, A., & Rock, I. (1998). *Inattentional blindness*. Cambridge, MA: The MIT Press.

Martarelli, C. S., & Mast, F. W. (2013). Eye movements during long-term pictorial recall. *Psychological Research*, *77*(3), 303–309. https://doi.org/10.1007/s00426-012-0439-7

Maunsell, J. H. R. (2015). Neuronal Mechanisms of Visual Attention. *Annual Review of Vision Science*, *1*(1), 373–391. https://doi.org/10.1146/annurev-vision-082114-035431

McAdams, C. J., & Maunsell, J. H. (1999). Effects of attention on orientation-tuning functions of

single neurons in macaque cortical area V4. *Journal of Neuroscience*, *19*(1), 431–441.

McNab, F., & Klingberg, T. (2008). Prefrontal cortex and basal ganglia control access to working memory. *Nature Neuroscience*, *11*(1), 103–107. https://doi.org/10.1038/nn2024

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*, *17*(10), 2364–2374. https://doi.org/10.1093/cercor/bhl145

Miller, J., Patterson, T., & Ulrich, R. (1998). Jackknife-based method for measuring LRP onset latency differences. *Psychophysiology*, *35*(1), 99–115. https://doi.org/10.1111/1469-8986.3510099

Morgan, S. T., Hansen, J. C., & Hillyard, S. A. (1996). Selective attention to stimulus location modulates the steady-state visual evoked potential. *Proceedings of the National Academy of Sciences*, *93*(10), 4770–4774. https://doi.org/10.1073/pnas.93.10.4770

Müller, H. J., & Rabbitt, P. M. A. (1989). Reflexive and voluntary orienting of visual attention: time course of activation and resistance to interruption. *Journal of Experimental Psychology: Human Perception and Performance*, *15*(2), 315–330. https://doi.org/10.1037/0096-1523.15.2.315

Müller, M. M., Malinowski, P., Gruber, T., & Hillyard, S. A. (2003). Sustained division of the attentional spotlight. *Nature*, *242*, 309–312. https://doi.org/10.1038/nature01744.1.

Müller, M. M., Teder-Sälejärvi, W., & Hillyard, S. A. (1998). The time course of cortical facilitation during cued shifts of spatial attention. *Nature Neuroscience*, *1*(7), 631–634. https://doi.org/10.1038/2865

Murray, S. O. (2008). The effects of spatial attention in early human visual cortex are stimulus independent. *Journal of Vision*, *8*(10), 2.1-11. https://doi.org/10.1167/8.10.2.Introduction

Nakayama, K., & Mackeben, M. (1989). Sustained and transient components of focal visual attention. *Vision Research*, *29*(11), 1631–1647.

Nazir, T. A., & O'Regan, J. K. (1990). Some results on translation invariance in the human visual system. *Spatial Vision*, *5*(2), 81–100.

Nicolelis, M. A. L., Fanselow, E. E., & Ghazanfar, A. A. (1997). Hebb's dream: The resurgence of cell assemblies. *Neuron*, *19*(2), 219–221. https://doi.org/10.1016/S0896-6273(00)80932-0

Nissen, M. J. (1985). Accessing features and objects: Is location special? In M.I. Posner & O. S. M. Marin (Eds.), *Attention and Performance XI* (pp. 205–219). Hillsdale, NJ: Erlbaum.

Noonan, M. P., Adamian, N., Pike, X. A., Printzlau, F., Crittenden, B. M., & Stokes, M. G. (2016). Distinct mechanisms for distractor suppression and target facilitation. *Journal of Neuroscience*, *36*(6), 1797–1807.

Nunez, P. L., & Srinivasan, R. (2006). *Electric fields of the brain: The neurophysics of EEG*. New York, NY: Oxford University Press.

Pashler, H., & Badgio, P. C. (1985). Visual attention and stimulus identification. *Journal of Experimental Psychology: Human Perception & Performance*, *11*(2), 105–121.

Payne, L., & Sekuler, R. (2014). The importance of ignoring: alpha oscillations protect selectivity. *Current Directions in Psychological Science*, *23*(3), 171–177.

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

Pestilli, F., Carrasco, M., Heeger, D. J., & Gardner, J. L. (2011). Attentional enhancement via selection and pooling of early sensory responses in human visual cortex. *Neuron*, *72*(5), 832–846. https://doi.org/10.1016/j.neuron.2011.09.025

Posner, Michael I. (1980). Orienting of attention. *The Quarterly Journal of Experimental Psychology*, *32*(1), 3–25.

Posner, Michael I, Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General;Journal of Experimental Psychology: General*, *109*(2), 160–174. https://doi.org/10.1037/0096-3445.109.2.160

Postle, B. R. (2015). The cognitive neuroscience of visual short-term memory. *Current Opinion in Behavioral Sciences*, *1*, 40–46. https://doi.org/10.1016/j.cobeha.2014.08.004

Pouget, A., Pouget, A., Dayan, P., Dayan, P., Zemel, R., & Zemel, R. (2000). Information processing with population codes. *Nature Reviews Neuroscience*, *1*(2), 125–132. https://doi.org/10.1038/35039062

Pratte, M. S., & Tong, F. (2014). Spatial specificity of working memory representations in the early visual cortex. *Journal of Vision*, *14*(3), 1–12. https://doi.org/10.1167/14.3.22.doi

Rajsic, J., & Wilson, D. E. (2014). Asymmetrical access to color and location in visual working memory. *Attention, Perception, & Psychophysics*, *76*(7), 1902–1913. https://doi.org/10.3758/s13414-014-0723-2

Rensink, R. A., Regan, J. K. O., & Clark, J. J. (2003). To see or not to see: the need for attention to perceive changes in scenes. *Psychological Science*, *8*(5), 368–373.

Riggall, A. C., & Postle, B. R. (2012). The relationship between working memory storage and elevated activity as measured with functional magnetic resonance imaging. *Journal of*

Neuroscience, *32*(38), 12990–12998. https://doi.org/10.1523/JNEUROSCI.1892-12.2012

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.

Romei, V., Gross, J., & Thut, G. (2010). On the role of prestimulus alpha rhythms over occipito-parietal areas in visual input regulation: correlation or causation? *Journal of Neuroscience*, *30*(25), 8692–8697. https://doi.org/10.1523/JNEUROSCI.0160-10.2010

Roux, F., & Uhlhaas, P. J. (2013). Working memory and neural oscillations: alpha-gamma versus theta-gamma codes for distinct WM information? *Trends in Cognitive Sciences*, *18*(1). https://doi.org/10.1016/j.tics.2013.10.010

Salazar, R. F., Dotson, N. M., Bressler, S. L., & Gray, C. M. (2012). Content-specific fronto-parietal synchronization during visual working memory. *Science*, *338*, 1097–1100.

Samaha, J., Sprague, T. C., & Postle, B. R. (2016). Decoding and reconstructing the focus of spatial attention from the topography of alpha-band oscillations. *Journal of Cognitive Neuroscience*, *28*(8), 1090–1097.

Saproo, S., & Serences, J. T. (2014). Attention improves transfer of motion information between V1 and MT. *Journal of Neuroscience*, *34*(10), 3586–3596. https://doi.org/10.1523/JNEUROSCI.3484-13.2014

Sauseng, P., Klimesch, W., Stadler, W., Schabus, M., Doppelmayr, M., Hanslmayr, S., … Birbaumer, N. (2005). A shift of visual spatial attention is selectively associated with human EEG alpha activity. *The European Journal of Neuroscience*, *22*(11), 2917–2926. https://doi.org/10.1111/j.1460-9568.2005.04482.x

Saxena, S., & Cunningham, J. P. (2019). Towards the neural population doctrine. *Current Opinion in Neurobiology*, *55*, 103–111. https://doi.org/10.1016/j.conb.2019.02.002

Schneegans, S., & Bays, P. M. (2017). Neural architecture for binding in visual working memory. *Journal of Neuroscience*, *37*(14), 3913–3925. https://doi.org/10.1523/JNEUROSCI.3493-16.2017

Sejnowski, T. J., & Paulsen, O. (2006). Network oscillations: emerging computational principles. *Journal of Neuroscience*, *26*(6), 1673–1676. https://doi.org/10.1523/JNEUROSCI.3737-05d.2006

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

Serences, John T, Ester, E. F., Vogel, E. K., & Awh, E. (2009). Stimulus-specific delay activity

in human primary visual cortex. *Psychological Science*, *20*(2), 207–214. https://doi.org/10.1111/j.1467-9280.2009.02276.x

Serences, John T, & Saproo, S. (2012). Computational advances towards linking BOLD and behavior. *Neuropsychologia*, *50*(4), 435–446.

Sereno, A. B., & Amador, S. C. (2006). Attention and memory-related responses of neurons in the lateral intraparietal area during spatial and shape-delayed match-to-sample tasks. *Journal of Neurophysiology*, *95*, 1078–1098. https://doi.org/10.1152/jn.00431.2005

Silver, M. A., & Kastner, S. (2009). Topographic maps in human frontal and parietal cortex. *Trends in Cognitive Sciences*, *13*(11), 488–495. https://doi.org/10.1016/j.tics.2009.08.005

Singer, W. (1999). Neuronal synchrony: a versatile code for the definition of relations? *Neuron*, *24*, 49–65. https://doi.org/10.1016/B978-012370880-9.00287-5

Singer, W., & Gray, C. M. (1995). Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience*, *18*, 555–586.

Spaak, E., Bonnefond, M., Maier, A., Leopold, D. A., & Jensen, O. (2012). Layer-specific entrainment of gamma-band neural activity by the alpha rhythm in monkey visual cortex. *Current Biology*, *22*(24), 2313–2318. https://doi.org/10.1016/j.cub.2012.10.020

Sprague, T. C., Adam, K. C. S., Foster, J. J., Rahmati, M., Sutterer, D. W., & Vo, V. A. (2018). Inverted encoding models assay population-level stimulus representations, not single-unit neural tuning. *Eneuro*, *5*(June). https://doi.org/10.1523/ENEURO.0098-18.2018

Sprague, T. C., Ester, E. F., & Serences, J. T. (2014). Reconstructions of information in visual spatial working memory degrade with memory load. *Current Biology*, *24*(18), 1–7. https://doi.org/10.1016/j.cub.2014.07.066

Sprague, T. C., Itthipuripat, S., Vo, V., & Serences, J. T. (2018). Dissociable signatures of visual salience and behavioral relevance across attentional priority maps in human cortex. *Journal of Neurophysiology*, 1–26. https://doi.org/10.1152/jn.00059.2018

Sprague, T. C., Saproo, S., & Serences, J. T. (2015). Visual attention mitigates information loss in small- and large-scale neural codes. *Trends in Cognitive Sciences*, *19*(4), 215–226.

Sprague, T. C., & Serences, J. T. (2013). Attention modulates spatial priority maps in the human occipital, parietal and frontal cortices. *Nature Neuroscience*, *16*(12), 1879–1887. https://doi.org/10.1038/nn.3574

Sprague, T. C., & Serences, J. T. (2015). Using human neuroimaging to examine top-down modulation of visual perception. In B. U. Forstmann & E.-J. Wagenmakers (Eds.), *An Introduction to Model-Based Cognitive Neuroscience* (pp. 245–274). Springer. https://doi.org/10.1007/978-1-4939-2236-9_12

Sreenivasan, K. K., Curtis, C. E., & D'Esposito, M. (2014). Revisiting the role of persistent neural activity during working memory. *Trends in Cognitive Sciences*, *18*(2), 82–89. https://doi.org/10.1016/j.tics.2013.12.001

Suchow, J. W., Brady, T. F., Fougnie, D., & Alvarez, G. A. (2013). Modeling visual working memory with the MemToolbox. *Journal of Vision*, *13*, 1–8. https://doi.org/10.1167/13.10.9.doi

Theeuwes, J., Kramer, A. F., & Irwin, D. E. (2011). Attention on our mind: The role of spatial attention in visual working memory. *Acta Psychologica*, *137*(2), 248–251. https://doi.org/10.1016/j.actpsy.2010.06.011

Thut, G., Nietzel, A., Brandt, S. A., & Pascual-Leone, A. (2006). α-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *Journal of Neuroscience*, *26*(37), 9494–9502.

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

Tootell, R. B., Hadjikhani, N., Hall, E. K., Marrett, S., Vanduffel, W., Vaughan, J. T., & Dale, A. M. (1998). The retinotopy of visual spatial attention. *Neuron*, *21*(6), 1409–1422.

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

Tsal, Y., & Lavie, N. (1993). Location dominance in attending to color and shape. *Journal of Experimental Psychology: Human Perception and Performance*, *19*(1), 131–139. https://doi.org/10.1037/0096-1523.19.1.131

Tsubomi, H., Fukuda, K., Watanabe, K., & Vogel, E. K. (2013). Neural limits to representing objects still within view. *Journal of Neuroscience*, *33*(19), 8257–8263. https://doi.org/10.1523/JNEUROSCI.5348-12.2013

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. *Journal of Neuroscience*, *28*(34), 8397–8405. https://doi.org/10.1523/JNEUROSCI.0630-08.2008

van Dijk, H., Van Der Werf, J., Mazaheri, A., Medendorp, W. P., & Jensen, O. (2010). Modulations in oscillatory activity with amplitude asymmetry can produce cognitively relevant event-related responses. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(2), 900–905. https://doi.org/10.1073/pnas.0908821107

van Gerven, M. A. J., & Jensen, O. (2009). Attention modulations of posterior alpha as a control signal for two-dimensional brain-computer interfaces. *Journal of Neuroscience Methods*,

179(1), 78–84.

Vo, V., Sprague, T. C., & Serences, J. T. (2017). Spatial tuning shifts increase the discriminability and fidelity of population codes in visual cortex. *Journal of Neuroscience*, *37*(12), 3386–3401.

Vogel, E. K., McCollough, A. W., & Machizawa, M. G. (2005). Neural measures reveal individual differences in controlling access to working memory. *Nature*, *438*(7067), 500–503. https://doi.org/10.1038/nature04171

Vogel, E. K., Woodman, G. F., & Luck, S. J. (2006). The time course of consolidation in visual working memory. *Journal of Experimental Psychology. Human Perception and Performance*, *32*(6), 1436–1451. https://doi.org/10.1037/0096-1523.32.6.1436

Watrous, A. J., Fell, J., Ekstrom, A. D., & Axmacher, N. (2015). More than spikes: Common oscillatory mechanisms for content specific neural representations during perception and memory. *Current Opinion in Neurobiology*, *31*, 33–39. https://doi.org/10.1016/j.conb.2014.07.024

Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, *4*(12), 1120–1135. https://doi.org/10.1167/4.12.11

Williams, M., Pouget, P., Boucher, L., & Woodman, G. F. (2013). Visual–spatial attention aids the maintenance of object representations in visual working memory. *Memory & Cognition*, *41*(5), 698–715. https://doi.org/10.3758/s13421-013-0296-7

Womelsdorf, T., Anton-Erxleben, K., Pieper, F., & Treue, S. (2006). Dynamic shifts of visual receptive fields in cortical area MT by spatial attention. *Nature Neuroscience*, *9*(9), 1156–1160. https://doi.org/10.1038/nn1748

Womelsdorf, T., Anton-Erxleben, K., & Treue, S. (2008). Receptive field shift and shrinkage in macaque middle temporal area through attentional gain modulation. *Journal of Neuroscience*, *28*(36), 8934–8944. https://doi.org/10.1523/JNEUROSCI.4030-07.2008

Woodman, G. F., & Vogel, E. K. (2008). Selective storage and maintenance of an object's features in visual working memory. *Psychonomic Bulletin & Review*, *15*(1), 223–229. https://doi.org/10.3758/PBR.15.1.223

Worden, M. S., Foxe, J. J., Wang, N., & Simpson, G. V. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific α-band electroencephalography increases over occipital cortex. *Journal of Neuroscience*, *20*(RC63), 1–6.

Yu, Q., & Shim, W. M. (2017). Occipital, parietal, and frontal cortices selectively maintain task-relevant features of multi-feature objects in visual working memory. *NeuroImage*, *157*(May), 97–107. https://doi.org/10.1016/j.neuroimage.2017.05.055

Zaksas, D., Bisley, J. W., & Pasternak, T. (2001). Motion information is spatially localized in a visual working-memory task. *Journal of Neurophysiology*, *86*(2), 912–921.

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

Zumer, J. M., Scheeringa, R., Schoffelen, J. M., Norris, D. G., & Jensen, O. (2014). Occipital alpha activity during stimulus processing gates the information flow to object-selective cortex. *PLoS Biology*, *12*(10), e1001965.